

Optimal Taxation and Human Capital Policies over the Life Cycle

Stefanie Stantcheva

Harvard University and National Bureau of Economic Research

This paper derives optimal income tax and human capital policies in a life cycle model with risky human capital. The government faces asymmetric information regarding agents' ability, its evolution, and labor supply. When the wage elasticity with respect to ability is increasing in human capital, the optimal subsidy involves less than full deductibility of human capital expenses on the tax base and falls with age. Income-contingent loans or a deferred deductibility scheme can implement the optimum. Numerical results suggest that full deductibility of expenses is close to optimal and that simple linear age-dependent policies perform very well.

I. Introduction

Investments in human capital, in the form of both time and money, play a key role in most people's lives. Children and young adults acquire education, and human capital accumulation continues throughout life through job training. There is a two-way interaction between human capital and the tax system. On the one hand, investments in human capital are influenced by tax policy—a point recognized early on by Schultz

I want to thank James Poterba and Ivan Werning for invaluable advice, guidance, and encouragement throughout this project. Esther Duflo and Robert Townsend provided helpful and generous support. I benefited very much from Emmanuel Saez's insightful comments. I also thank seminar participants at Berkeley, Booth Chicago, Harvard, Michigan, Massachusetts Institute of Technology, Penn, Princeton, Santa Barbara, Stanford, Stanford Graduate School of Business, Wharton, and Yale for their useful feedback and comments. Data are provided as supplementary material online.

Electronically published October 31, 2017
[*Journal of Political Economy*, 2017, vol. 125, no. 6]
© 2017 by The University of Chicago. All rights reserved. 0022-3808/2017/12506-0005\$10.00

(1961).¹ Taxes on labor income discourage investment by capturing part of the return to human capital yet also help insure against earnings risk, thereby encouraging investment in risky human capital. Capital taxes affect the choice between physical and human capital. On the other hand, investments in human capital directly affect the available tax base and are a major determinant of the pretax income distribution. Policies to stimulate human capital acquisition, which vary greatly across countries, shape the skill distribution of workers—a crucial input into optimal income taxation models.

This two-way feedback calls for a joint analysis of optimal income taxes and optimal human capital policies over the life cycle, which is the goal of this paper. The vast majority of optimal tax research assumes that productivity is exogenously determined instead of being the product of investment decisions made throughout life. Therefore, this paper addresses the following questions. First, how, if at all, should the tax and social insurance system take into account human capital acquisition? Should human capital expenses be tax deductible? Second, what parameters are important for setting optimal human capital policies, such as subsidies, and how do optimal policies evolve over time? Finally, what combination of policy instruments implements the optimum? Can simple policies yield a level of welfare close to that achieved with complex systems?

Specifically, this paper jointly determines optimal tax and human capital policies over the life cycle and incorporates essential characteristics of the human capital acquisition process. First, human capital pays off over long periods of time and thus returns are inherently uncertain: skills can be rendered obsolete by unpredictable changes in technology, industry shocks, or macroeconomic contractions. Yet, private markets for insurance against personal productivity shocks are limited. Second, there are important and growing financial costs to human capital acquisition, which can be deterrents to an efficient investment in human capital. Finally, individuals have heterogeneous intrinsic abilities, which may differentially affect their returns to human capital investment.²

Accordingly, in the model, each individual's wage is a function of endogenous human capital and stochastic "ability." Ability, as in the standard Mirrlees (1971) income taxation model, is a comprehensive measure of the exogenous component of productivity. Agents have heterogeneous innate abilities, which are subject to persistent and privately uninsurable shocks. Throughout their lives, they can invest in human capital with risky returns by spending money. The government maximizes a standard so-

¹ "Our tax laws everywhere discriminate against human capital. Although the stock of such capital has become large and even though it is obvious that human capital, like other forms of reproducible capital, depreciates, becomes obsolete and entails maintenance, our tax laws are all but blind on these matters" (Schultz 1961, 17).

² The empirical evidence on this issue is reviewed in Sec. V.A.

cial welfare function under asymmetric information about any agent's ability—both its initial level and its evolution over life—as well as labor effort. This requires the imposition of incentive compatibility constraints in the dynamic mechanism designed by the government. To describe the distortions in the resulting constrained efficient allocations, the wedges, or implicit taxes and subsidies, are analyzed.

Despite the complexity of the model, a very simple relation between the optimal human capital and labor wedges is derived. The implicit subsidy for human capital expenses is determined by three goals. The first is to counterbalance the distortions to human capital indirectly stemming from the labor and savings distortions. When these distortions are perfectly counterbalanced, the tax system is neutral with respect to human capital investments. I introduce the notion of a full dynamic risk-adjusted deductibility that ensures this neutrality. The second goal is to stimulate labor supply by increasing the wage, that is, the returns to labor. The third is to redistribute and provide insurance, taking into account the differential effect of human capital on the pretax income of high- and low-ability people.

When the percentage (or proportional) change in the wage of high-ability agents from human capital is not larger than that of low-ability agents, human capital has a positive redistributive effect on after-tax income and a positive insurance value. It is then optimal to subsidize human capital expenses beyond simply ensuring a neutral tax system with respect to human capital expenses, that is, beyond making human capital expenses fully tax deductible in a dynamic, risk-adjusted fashion. In this case, the human capital wedge also drifts up with age. The persistence of ability shocks directly translates into a persistence of the optimal human capital wedge over time. While the sign of the human capital subsidy is exclusively determined by the complementarity between human capital and ability, its magnitude is modulated by the strength of the insurance and redistributive motives and the persistence of ability shocks.

This paper considers two ways of implementing the constrained efficient allocations: income-contingent loans (ICLs) and a “deferred deductibility” scheme. For ICLs, the loan repayment schedules are contingent on the past history of earnings and human capital investments. In the deferred deductibility scheme, only part of current investment in human capital can be deducted from current taxable income. The remainder is deducted from future taxable income, to account for the risk and the nonlinearity of the tax schedule.

I calibrate the model on the basis of US data to illustrate the optimal policies under different assumptions regarding the complementarity between human capital and ability. When human capital has a positive redistributive or insurance value, the net stimulus to human capital is small and positive and grows with age. It is not optimal to deviate much from a

neutral system with respect to human capital, a type of “production efficiency” result, and hence, full dynamic risk-adjusted deductibility is close to optimal. Simple linear age-dependent human capital subsidies, as well as income and savings taxes, achieve almost the entire welfare gain from the full second-best optimum for the calibrations studied.

Related literature.—The complex process of human capital acquisition has been studied in a long-standing literature, starting with Becker (1964), Ben-Porath (1967), and Heckman (1976). The model in this paper tries to adopt, in a stylized way, some of this literature’s main findings. The structural branch of the literature (Cunha et al. 2006; Cunha and Heckman 2007) emphasizes that human capital acquisition occurs throughout life, underscoring the need for a life cycle model. Both ex ante heterogeneity in the returns to human capital and uncertainty matter. A large body of empirical work documents the importance of human capital as a determinant of earnings (Card 1995; Goldin and Katz 2008) and the financial and other factors shaping individuals’ decisions to acquire human capital (Lochner and Monge-Naranjo 2011). The subset of this literature that studies the interaction between ability and schooling for earnings—a crucial consideration for optimal policies in this paper—is reviewed in detail in Section V.A.

On the other hand, the optimal taxation literature, dating back to Mirrlees (1971) and developed more recently by Saez (2001), Kocherlakota (2005), Albanesi and Sleet (2006), Golosov, Tsyvinski, and Werning (2006), Battaglini and Coate (2008), Farhi and Werning (2013), Golosov, Troshkin, and Tsyvinski (2013), and Scheuer (2014) typically assumes exogenous ability, thus abstracting from endogenous human capital investments. Therefore, this paper builds on the life cycle framework in Farhi and Werning (2013) and introduces endogenous stochastic productivity as the result of human capital acquisition by agents.

A series of papers, evolving from static to dynamic, have considered optimal taxation jointly with education policies. Bovenberg and Jacobs (2005), using a static taxation model, find that education subsidies and income taxes are “Siamese twins” and should always be set equal to each other, which is equivalent to making human capital expenses fully tax deductible. A few subsequent static papers emphasize the importance of the complementarity between intrinsic ability and human capital (Maldonado [2008] with two types and Bovenberg and Jacobs [2011] with a continuum of types) or between risk and human capital (Da Costa and Maestri 2007).

Several recent dynamic optimal tax papers examine the impact of taxation on human capital, with important differences from the current paper. Previous dynamic models allowed for heterogeneity across agents, but not uncertainty (Bohacek and Kapička 2008; Kapička 2013b), or uncertainty, but not heterogeneity (Anderberg 2009; Grochulski and

Piskorski 2010), which precludes a discussion of redistributive policies. Findeisen and Sachs (2012) include both heterogeneity and uncertainty but focus on a one-shot investment during “college,” before the work life of the agent starts, with a one-time realization of uncertainty. By contrast, this paper features life cycle investment in human capital and a progressive realization of uncertainty throughout life. Stantcheva (2015a) considers time investment in human capital. In a complementary analysis, Kapička and Neira (2014) posit a different human capital accumulation process with time investments and a fixed ability and consider the case in which effort spent to acquire human capital is unobservable. Also complementary is the work by Krueger and Ludwig (2013), who adopt a Ramsey approach by specifying ex ante the instruments available to the government, in contrast to the Mirrlees approach adopted here, which considers an unrestricted direct revelation mechanism. In their overlapping generations general equilibrium model, “education” is a binary decision that occurs exclusively before entry into the labor market. The life cycle analysis also addresses the issue of age-dependent taxation, as explored in Kremer (2002), Mirrlees et al. (2011), and Weinzierl (2011).³

The rest of the paper is organized as follows. Section II presents the dynamic life cycle model and the full information benchmark. Section III sets up a recursive mechanism design program using the first-order approach. Section IV solves for the optimal policies and interprets the results. Section V contains the numerical analysis. Section VI discusses the implementation of the optimal policies using ICLs and a deferred deductibility scheme. Section VII presents conclusions and discusses three alternative applications of the model: to intergenerational transfers and bequest taxation, to entrepreneurial taxation, and to health investments. All proofs are in Appendix B. The online appendix contains some of the lengthier proofs and additional figures from the simulations.

II. A Life Cycle Model of Human Capital Acquisition and Labor Supply

The economy consists of agents who live for T years, during which they work and acquire human capital. Agents who work $l_t \geq 0$ hours in period t at a wage rate w_t earn a gross income $y_t = w_t l_t$. Each period, agents can build their stock of human capital by spending money. A monetary investment of amount $M_t(e_t)$ generates an increase in human capital $e_t \geq 0$. The cost function satisfies $M'_t(e) > 0$ for all $e_t > 0$, $M'_t(0) = 0$, and $M''_t(e_t) \geq 0$ for all $e_t \geq 0$. These monetary investments add to a stock of hu-

³ This paper is more generally related to the dynamic mechanism design literature, as developed by, among many others, Fernandes and Phelan (2000), Doepke and Townsend (2006), and Pavan, Segal, and Toikka (2014).

man capital acquired by expenses (“expenses” for short), s_t , which evolves according to $s_t = s_{t-1} + e_t$.⁴ Expenses can be thought of as the necessary material inputs into the production of human capital, such as books, tuition fees, or living and board costs while at college, net of the cost of living elsewhere. The disutility cost to an agent of supplying labor effort l_t is $\phi_t(l_t)$; ϕ_t is strictly increasing and convex.

The wage rate w_t is determined by the stock of human capital built until time t and stochastic ability θ_t :

$$w_t = w_t(\theta_t, s_t),$$

where w_t is strictly increasing and concave in each of its arguments ($\partial w / \partial m > 0$, $\partial^2 w / \partial m^2 \leq 0$ for $m = \theta, s$). Importantly though, no restrictions are placed on the cross-partials. This formulation allows for human capital to affect the wage differently at different ages.⁵

Agents are born at time $t = 1$ with a heterogeneous earning ability θ_1 with distribution $f^1(\theta_1)$. Earning ability in each period is private information and evolves according to a Markov process with a time-varying transition function $f^t(\theta_t | \theta_{t-1})$, over a fixed support $\Theta \equiv [\underline{\theta}, \bar{\theta}]$. There are several possible interpretations for θ_t , such as stochastic productivity or stochastic returns to human capital. For example, with a separable wage form $w_t = \theta_t + h_t(s_t)$, for some increasing, concave function h_t , θ_t resembles a stochastic version of productivity from the static Mirrlees (1971) model. With a wage such as $w_t = \theta_t h_t(s_t)$, θ_t is perhaps more naturally interpreted as the stochastic return to human capital. To keep with the tradition in the literature, θ_t will be called “ability” throughout. Ability to earn income can be stochastic because of, among others, health shocks, individual labor market idiosyncrasies, or luck.

The agent’s per-period utility is separable in consumption and labor:

$$\tilde{u}_t(c_t, y_t, s_t; \theta_t) = u_t(c_t) - \phi_t\left(\frac{y_t}{w_t(\theta_t, s_t)}\right),$$

where u_t is increasing, twice continuously differentiable, and concave.

Denote by θ^t the history of ability shocks up to period t , by Θ^t the set of possible histories at t , and by $P(\theta^t)$ the probability of a history θ^t , $P(\theta^t) \equiv f^t(\theta_t | \theta_{t-1}) \cdots f^2(\theta_2 | \theta_1) f^1(\theta_1)$. An allocation $\{x_t\}_t$ specifies consumption, output, and expenses for each period t , conditional on the history θ^t , that is,

⁴ The agent cannot willfully destroy human capital, hence $e_t \geq 0$. The ability shock θ described right below can partially account for stochastic depreciation. Deterministic depreciation would enter as a scaling factor of the next period’s human capital stock (s_{t+1}) in all formulas.

⁵ Note that human capital yields an immediate benefit in the period in which it is acquired, as well as into the future. This reduces the uncertainty by one period and simplifies the optimal formulas below.

$x_t = \{x(\theta^t)\}_{\theta^t} = \{c(\theta^t), y(\theta^t), s(\theta^t)\}_{\theta^t}$. The expected lifetime utility from an allocation, discounted by a factor β , is given by

$$U(\{c(\theta^t), y(\theta^t), s(\theta^t)\}) = \sum_{t=1}^T \int \beta^{t-1} \left[u_t(c(\theta^t)) - \phi_t \left(\frac{y(\theta^t)}{w^t(\theta_t, s(\theta^t))} \right) \right] P(\theta^t) d\theta^t, \quad (1)$$

where, with some abuse of notation, $d\theta^t \equiv d\theta_t \cdots d\theta_1$.

Let $w_{m,t}$ denote the partial of the wage function with respect to argument m ($m \in \{\theta, s\}$) and $w_{mn,t}$ the second-order partial with respect to arguments $m, n \in \{\theta, s\} \times \{\theta, s\}$. One crucial parameter is the Hicksian coefficient of complementarity between ability and human capital in the wage function at time t (Hicks 1970; Samuelson 1974), denoted by $\rho_{\theta s,t}$:

$$\rho_{\theta s} \equiv \frac{w_{\theta s} w}{w_s w_\theta}. \quad (2)$$

A positive Hicksian complementarity between human capital s and ability θ means that higher-ability agents have a higher marginal benefit from human capital ($w_{\theta s} \geq 0$).⁶ Put differently, human capital compounds the exposure of the agent to stochastic ability and to risk. A Hicksian complementarity greater than one means that higher-ability agents have a higher proportional benefit from human capital; that is, the wage elasticity with respect to ability is increasing in human capital, that is,⁷

$$\frac{\partial}{\partial s} \left(\frac{\partial w}{\partial \theta} \frac{\theta}{w} \right) \geq 0.$$

A separable wage function of the form $w_t = \theta_t + h_t(s_t)$ for some function h_t implies that $\rho_{\theta s,t} = 0$. A multiplicative form $w_t = \theta_t h_t(s_t)$, the one typically used in the taxation literature, implies that $\rho_{\theta s,t} = 1$. Finally, with a constant elasticity of substitution (CES) wage function of the form

$$w_t = (\alpha_1 \theta^{1-\rho_t} + \alpha_2 s_t^{1-\rho_t})^{1/(1-\rho_t)}, \quad (3)$$

ability and human capital can be substituted one for the other at a fixed but potentially time-varying rate: $\rho_{\theta s,t} = \rho_t$.

III. The Planning Problem

Every period, the planner can observe an agent's choices of output y , consumption c , and human capital s . The informational problem is that

⁶ The term $\rho_{\theta s}$ is also the Hicksian complementarity coefficient between education and ability in earnings y .

⁷ Equivalently, the wage elasticity with respect to human capital is increasing in ability.

he cannot see ability θ_i in any period. This implies that, while the planner knows the wage function, he cannot know the wage realization $w_i(\theta_i, s_i)$ nor labor supply $l_i = y_i/w_i$ since those depend on the unobserved θ . Put differently, when seeing a low output produced by an agent, he cannot know whether it was due to the agent's low labor effort or to a bad ability (and, hence, wage) shock.

This technical section sets up the planning problem, starting from the sequential problem and defining incentive compatibility. It then goes through two steps to make this problem tractable, following the recent procedure proposed for dynamic Mirrlees models by Farhi and Werning (2013), augmented here with human capital. First, a relaxed problem based on the first-order approach is written out, which replaces the full set of incentive compatibility constraints by the agent's envelope condition. This relaxed program is then turned into a recursive dynamic programming problem through a suitable definition of state variables.

A. Incentive Compatibility

To solve for the constrained efficient allocations, suppose that the planner designs a direct revelation mechanism, in which, each period, agents have to report their ability θ_i . Denote a reporting strategy, specifying a reported type r_i after each history, by $r = \{r_i(\theta^t)\}_{i=1}^T$. Let \mathcal{R} be the set of all possible reporting strategies and $r^t = \{r_1(\theta^t), \dots, r_i(\theta^t)\}$ be the history of reports generated by reporting strategy r . Because output, savings, and human capital are observable, the planner can directly specify allocations as functions of the history of reports, according to some allocation rules $c(r^t)$, $y(r^t)$, and $s(r^t)$.⁸ Let the continuation value after history θ^t under a reporting strategy r , denoted by $\omega^r(\theta^t)$, be the solution to

$$\begin{aligned} \omega^r(\theta^t) = & u_i(c(r^t(\theta^t))) - \phi_i \left(\frac{y(r^t(\theta^t))}{w_i(\theta_i, s(r^t(\theta^t)))} \right) \\ & + \beta \int \omega^r(\theta^{t+1}) f^{t+1}(\theta_{t+1} | \theta_t) d\theta_{t+1}. \end{aligned}$$

The continuation value under truthful revelation, $\omega(\theta^t)$, is the unique solution to

$$\begin{aligned} \omega(\theta^t) = & u_i(c(\theta^t)) - \phi_i \left(\frac{y(\theta^t)}{w_i(\theta_i, s(\theta^t))} \right) \\ & + \beta \int \omega(\theta^{t+1}) f^{t+1}(\theta_{t+1} | \theta_t) d\theta_{t+1}. \end{aligned}$$

⁸ Hours of work are determined residually by $l(r^t) = y(r^t)/w(\theta_i, s(r^t))$.

Incentive compatibility requires that truth telling yields a weakly higher continuation utility than any reporting strategy r :

$$(IC): \quad \omega(\theta_1) \geq \omega^r(\theta_1) \quad \forall \theta_1, \forall r. \quad (4)$$

Denote by X^{IC} the set of allocations that satisfy incentive compatibility condition (4). To solve this dynamic problem, a version of the first-order approach is used, requiring the following assumptions.

ASSUMPTION 1. (i) $\tilde{u}_t(c, y, s; \theta)$ and

$$\frac{\partial \phi(l)}{\partial l} \frac{\partial w(\theta, s)}{\partial \theta} \frac{l}{w}$$

are bounded; (ii) $\partial f^t(\theta_t | \theta_{t-1}) / \partial \theta_{t-1}$ exists and is bounded;⁹ (iii) $f^t(\theta_t | \theta_{t-1})$ has full support on Θ .

Suppose that the agent has witnessed a history of shocks θ^t . Consider one particular deviation strategy \tilde{r}_t , under which he reports truthfully until period t ($\tilde{r}_s(\theta^s) = \theta_s$ for all $s \leq t-1$) and lies in period t by reporting $\tilde{r}_t(\theta^t) = \theta' \neq \theta_t$. The continuation utility under this strategy is the solution to

$$\begin{aligned} \omega^{\tilde{r}}(\theta^t) = & u_t(c(\theta^{t-1}, \theta^t)) - \phi_t \left(\frac{y(\theta^{t-1}, \theta^t)}{w_t(\theta_t, s(\theta^{t-1}, \theta^t))} \right) \\ & + \beta \int \omega^{\tilde{r}}(\theta^{t-1}, \theta', \theta_{t+1}) f^t(\theta_{t+1} | \theta_t) d\theta_{t+1}. \end{aligned}$$

Incentive compatibility in (4) implies that, after almost all θ^t , the temporal incentive constraint holds:

$$\omega(\theta^t) = \max_{\theta'} \omega^{\tilde{r}}(\theta^t). \quad (5)$$

Inversely, if (5) holds after all θ^{t-1} and for almost all θ_t , then (4) also holds (see Kapička 2013a, lemma 1). If we take the derivative of promised utility with respect to (true) ability, there are two *direct* effects, namely, on the wage (higher types have higher wages) and on the Markov transition $f^t(\theta_t | \theta_{t-1})$, and *indirect* effects on the allocation through the report. By the first-order condition of the agent, all indirect effects are jointly zero and only the two direct effects remain. This leads to the envelope condition of the agent:

⁹ For some distributions, this derivative is not bounded, and assumption 3 in Kapička (2013a) could be used instead; namely, that for $F^t(\theta_t | \theta_{t-1}) \equiv \int_{\underline{\theta}}^{\theta_t} f^t(\theta_s | \theta_{t-1}) d\theta_s$, we have $(\partial / \partial \theta_{t-1}) F^t(\theta_t | \theta_{t-1}) \leq 0$ and $F^t(\theta_t | \theta_{t-1})$ either concave or convex.

$$\begin{aligned} \dot{\omega}(\theta^t) &:= \frac{\partial \omega(\theta^t)}{\partial \theta_t} = \frac{w_{\theta,t}}{w_t} l(\theta^t) \phi_{l,t}(l(\theta^t)) \\ &+ \beta \int \omega(\theta^{t+1}) \frac{\partial f^{t+1}(\theta_{t+1} | \theta_t)}{\partial \theta_t} d\theta_{t+1}. \end{aligned} \quad (6)$$

The envelope condition tells us how promised utility changes with type at incentive compatible allocations, that is, how the informational rents evolve. The first term is the static rent, familiar from screening models (Mirrlees 1971), while the second is the dynamic rent that arises because the agent has some advance information about his type tomorrow (this second term disappears with independently and identically distributed [iid] shocks). Let X^{FOA} denote the set of allocations that satisfy the envelope condition (6). It can be shown that this is a necessary condition for incentive compatibility.¹⁰

The analysis is in partial equilibrium. There is a physical capital asset that yields a fixed gross interest rate R . Investments in this physical capital are called “savings.” The planner’s objective is to minimize the expected discounted cost of providing an allocation, subject to incentive compatibility as defined in (4) and to expected lifetime utility of each (initial) type θ being above a threshold $\underline{U}(\theta)$. Let $U(\{c, y, s\}; \theta)$ be lifetime utility as defined in (1) for agents with initial type θ . The relaxed planning problem, denoted by P^{FOA} , replaces the incentive constraint by the envelope condition and is given by

$$\begin{aligned} &\min_{\{c,y,s\}} \Pi(\{c, y, s\}; \underline{U}(\theta)_\theta) \\ &= \sum_{t=1}^T \left(\frac{1}{R}\right)^{t-1} \int_{\theta^t} [c(\theta^t) - y(\theta^t) + M_t(s(\theta^t) - s(\theta^{t-1}))] P(\theta^t) d\theta^t \\ &\quad \text{subject to } U(\{c, y, s\}; \theta) \geq \underline{U}(\theta), \\ &\quad y(\theta^t) \geq 0, s(\theta^t) \geq s(\theta^{t-1}), c(\theta^t) \geq 0, \\ &\quad \{c, y, s\} \in X^{\text{FOA}}. \end{aligned} \quad (7)$$

B. Recursive Formulation of the Relaxed Program

To write the problem recursively, let the future marginal rent (the second term in the envelope condition) be denoted by

$$\Delta(\theta^t) \equiv \int \omega(\theta^{t+1}) \frac{\partial f^{t+1}(\theta_{t+1} | \theta_t)}{\partial \theta_t} d\theta_{t+1}. \quad (8)$$

¹⁰ An application of theorem 2 in Milgrom and Segal (2002), under assumption 1.

The envelope condition can then be rewritten as

$$\dot{\omega}(\theta^t) = \frac{w_{b,t}}{w_t} l(\theta^t) \phi_{l,t}(l(\theta^t)) + \beta \Delta(\theta^t). \quad (9)$$

Let $v(\theta^t)$ be the expected future continuation utility:

$$v(\theta^t) \equiv \int \omega(\theta^{t+1}) f^{t+1}(\theta_{t+1} | \theta_t) d\theta_{t+1}. \quad (10)$$

Continuation utility $\omega(\theta^t)$ hence can be rewritten as

$$\omega(\theta^t) = u_t(c(\theta^t)) - \phi_t\left(\frac{y(\theta^t)}{w_t(\theta_t, s(\theta^t))}\right) + \beta v(\theta^t). \quad (11)$$

Define the expected continuation cost of the planner at time t , given v_{t-1} , Δ_{t-1} , θ_{t-1} , and s_{t-1} :

$$K(v_{t-1}, \Delta_{t-1}, \theta_{t-1}, s_{t-1}, t) = \min \left[\sum_{\tau=t}^T \left(\frac{1}{R}\right)^{\tau-t} \int [c_\tau(\theta^\tau) - y_\tau(\theta^\tau) + M_\tau(s_\tau(\theta^\tau) - s_\tau(\theta^{\tau-1}))] P(\theta^{\tau-t}) d\theta^{\tau-t} \right],$$

where, with some abuse of notation, $d\theta^{\tau-t} = d\theta_\tau d\theta_{\tau-1} \cdots d\theta_t$, and $P(\theta^{\tau-t}) = f^\tau(\theta_\tau | \theta_{\tau-1}) \cdots f^t(\theta_t | \theta_{t-1})$. A recursive formulation of the relaxed program is then for $t \geq 2$

$$K(v, \Delta, \theta_-, s_-, t) = \min \left[\int \left[c(\theta) + M_t(s(\theta) - s_-) - w_t(\theta, s(\theta)) l(\theta) + \frac{1}{R} K(v(\theta), \Delta(\theta), \theta, s(\theta), t+1) \right] f^t(\theta | \theta_-) d\theta \right] \quad (12)$$

subject to

$$\begin{aligned} \omega(\theta) &= u_t(c(\theta)) - \phi_t(l(\theta)) + \beta v(\theta), \\ \dot{\omega}(\theta) &= \frac{w_{b,t}}{w_t} l(\theta) \phi_{l,t}(l(\theta)) + \beta \Delta(\theta), \\ v &= \int \omega(\theta) f^t(\theta | \theta_-) d\theta, \\ \Delta &= \int \omega(\theta) \frac{\partial f^t(\theta | \theta_-)}{\partial \theta_-} d\theta, \end{aligned}$$

where the maximization is over the functions $(c(\theta), l(\theta), s(\theta), \omega(\theta), v(\theta), \Delta(\theta))$.

For period $t = 1$, the problem needs to be reformulated. Suppose that all agents have identical initial human capital levels s_0 . The problem for

$t = 1$ is then indexed by $(\underline{U}(\theta))_\theta$, the set of target lifetime utilities $\underline{U}(\theta)$ for each type θ :

$$K((\underline{U}(\theta))_\theta, 1) = \min \int \left[c(\theta) + M_1(s(\theta) - s_0) - w_1(\theta, s(\theta))l(\theta) + \frac{1}{R}K(v(\theta), \Delta(\theta), \theta, s(\theta), 2)) \right] f^1(\theta) d\theta$$

subject to

$$\omega(\theta) = u_1(c(\theta)) - \phi_1(l(\theta)) + \beta v(\theta),$$

$$\dot{\omega}(\theta) = \frac{w_{\theta,1}}{w_1} l(\theta) \phi_{l,1}(l(\theta)) + \beta \Delta(\theta),$$

$$\omega(\theta) \geq \underline{U}(\theta),$$

$$\Delta = \int \omega(\theta) \frac{\partial f^1(\theta|\theta_-)}{\partial \theta_-} d\theta,$$

where the maximization is now over $(c(\theta), l(\theta), \omega(\theta), s(\theta), v(\theta), \Delta(\theta), \Delta)$. Note that Δ is now a free variable that is chosen optimally. The set of constrained efficient allocations that solve the planning problem is indexed by the set of utilities $(\underline{U}(\theta))_\theta$ and denoted by $X^{*,\text{FOA}}((\underline{U}(\theta))_\theta)$.

The solution to the relaxed program might not be a solution to the full program, because the envelope condition is only a necessary condition. In the static taxation model (Mirrlees 1971), the validity of the first-order approach is guaranteed if the utility function satisfies the standard Spence-Mirrlees single-crossing property and a simple monotonicity condition on the allocation. However, in the dynamic case, the conditions imposed on the allocations are more involved (see Golosov et al. 2013; Pavan et al. 2014) and do not always provide much simplification. Hence, for the proposed calibrations in Section V, incentive compatibility of the candidate allocation, as well as any omitted nonnegativity constraints, are checked numerically, using a procedure in the spirit of Farhi and Werning (2013).

To reduce notational clutter throughout the paper, the dependence on the full history is often left implicit, for example, $c_t = c(\theta^t)$ and $\tau_{L,t} = \tau_{L,t}(\theta^t)$. Similarly, function arguments are sometimes left out, for example, $w_{s,t} = (\partial/\partial s)w_t(\theta_t, s(\theta^t))$ and $u(c_t) = u_t$. The term E_t denotes the expectation as of time t , conditional on θ_t .

IV. Optimal Human Capital Policies

This section characterizes the optimal allocations, obtained as solutions to the relaxed program P^{FOA} above, using wedges, or implicit taxes and subsidies.

A. *The Wedges or Implicit Taxes and Subsidies*

In the second best, marginal distortions in agents' choices can be described using "wedges." Since the agent has three possible choices (working, saving, or investing in human capital), there are three marginal distortions. For any allocation, define the intratemporal wedge on labor $\tau_L(\theta')$, the intertemporal wedge on savings (also called capital) $\tau_K(\theta')$, and the human capital wedge $\tau_S(\theta')$ as follows:

$$\tau_L(\theta') \equiv 1 - \frac{\phi_{l,t}(l_t)}{w_t u'_t(c_t)}, \quad (13)$$

$$\tau_K(\theta') \equiv 1 - \frac{1}{R\beta} \frac{u'_t(c_t)}{E_t(u'_t(c_{t+1}))}, \quad (14)$$

$$\begin{aligned} \tau_S(\theta') \equiv & -[1 - \tau_L(\theta')]w_{s,t}l_t + M'_t(e_t) \\ & - \beta E_t\left(\left[\frac{u'_{t+1}(c_{t+1})}{u'_t(c_t)}\right][M'_{t+1}(e_{t+1}) - \tau_{S,t+1}]\right). \end{aligned} \quad (15)$$

Wedges are akin to locally linear subsidies and taxes and would all be zero without government intervention. They define a measure of the amount and direction of distortion at an allocation relative to the laissez-faire allocation. The labor wedge, which is very standard in the dynamic taxation literature (Goloso et al. 2006), is defined as the gap between the marginal rate of substitution and the marginal rate of transformation between consumption and labor. In the laissez-faire, it would be zero since the agent would equate the marginal rates of substitution and transformation. On the other hand, imagine that the planner imposes a linear tax equal to $\tau_L(\theta')$ and lets the agent of type θ' choose his labor supply, conditional on human capital and savings in a neighborhood around $l(\theta')$. A necessary condition for the agent's labor supply choice would then be equation (13). Hence, a positive labor wedge means that labor is distorted downward. Similarly, the savings wedge τ_K is defined as the difference between the expected marginal rate of intertemporal substitution and the return on savings.

The implicit subsidy on human capital expenses τ_S is such that the agent's net marginal cost from investing in human capital is locally reduced to $M'_t(e_t) - \tau_{S,t}$. Like any subsidy, it is equal to the gap between marginal cost and marginal benefit. However, human capital yields benefits in all future periods, so the future discounted expected stream of mar-

ginal benefits is needed. I rewrite it here recursively, replacing the latter stream by the next period’s marginal cost. Solving this forward would yield the full future stream of marginal benefits.

The paper sometimes loosely refers to the wedges as taxes and subsidies and appeals to intuitions related to a standard tax system.¹¹ The relation between wedges and explicit taxes is studied in Section VI, which addresses the issue of implementation.

The following definitions will be needed for the formulas below. For any variable x , define the “insurance factor” of x , $\xi_{x,t+1}$:

$$\xi_{x,t+1} \equiv -\text{Cov}\left(\beta \frac{u'_{t+1}}{u'_t}, x_{t+1}\right) / \left[E_t\left(\beta \frac{u'_{t+1}}{u'_t}\right)E_t(x_{t+1})\right],$$

with $\xi_{x,t+1} \in [-1, 1]$. If x is a flow to the agent, it is a good hedge if $\xi < 0$ and a bad hedge otherwise. With some abuse of notation, define also

$$\xi'_{x,t+1} \equiv -\text{Cov}\left(\frac{\beta u'_{t+1}}{u'_t} - \frac{1}{R}, x_{t+1}\right) / \left[E_t\left(\frac{\beta u'_{t+1}}{u'_t} - \frac{1}{R}\right)E_t(x_{t+1})\right],$$

which, up to an additive constant, captures the same risk properties as $\xi_{x,t+1}$.

Denote by $\varepsilon_{xy,t}$ the elasticity of x_t to y_t , $\varepsilon_{xy,t} \equiv d\log(x_t)/d\log(y_t)$. Let ε_t^u and ε_t^c be the uncompensated and compensated labor supply elasticities to the net wage, holding savings fixed.¹²

B. Optimal Labor and Savings Wedges

Before characterizing the optimal human capital subsidy, it is worth mentioning how the two standard wedges—the labor and the savings wedges—are set in this model with human capital. Both results are derived in Appendix A.

The labor wedge (proposition 6 in App. A) looks as in the standard static and dynamic models without human capital cited in the introduc-

¹¹ The wedges can also have natural interpretations as marginal taxes or combinations of marginal taxes, for instance, in an implementation with a complex tax function that depends on the full history of savings, output, and human capital. The labor wedge would be the gradient of the tax function with respect to income. The map between marginal taxes and the human capital or savings wedges is more complex and is studied in detail in Sec. VI.C.

¹² That is, ε^c and ε^u are defined as in the static framework (Saez 2001) at constant savings:

$$\varepsilon^u = \frac{\phi_l(l)/l + \frac{\phi_l(l)^2}{u'(c)^2} u''(c)}{\phi_u(l) - \frac{\phi_l(l)^2}{u'(c)^2} u''(c)}, \quad \varepsilon^c = \frac{\phi_l(l)/l}{\phi_u(l) - \frac{\phi_l(l)^2}{u'(c)^2} u''(c)}.$$

With per-period utility separable in consumption and labor, $\varepsilon^c/(1 + \varepsilon_t^u - \varepsilon_t^c)$ is the Frisch elasticity of labor.

tion, with one exception: the wage elasticity with respect to ability $\varepsilon_{w\theta,t}$ is not constant at 1, but rather depends on human capital. For instance, with a CES wage as in (3), $\varepsilon_{w\theta} = [w_t(\theta_t, s_t)/\theta_t]^{p-1}$. A higher elasticity amplifies the labor wedge as it increases the value of insurance and redistribution. The standard zero distortion at the bottom and the top results from the static Mirrlees model continue to apply in the presence of human capital. The labor wedge at any age is inversely related to the elasticity of labor supply at that age.

The presence of observable human capital does not change a standard result in dynamic moral hazard models with observable savings and separable utility (Rogerson 1985). Appendix proposition 7 shows that, at the optimum, the inverse Euler equation holds, so that there is a positive savings wedge τ_k .

C. *The Net Human Capital Subsidy*

The previously defined wedges are in general sufficient to characterize an allocation. However, in the presence of human capital there are several simultaneous distortions: a zero human capital wedge, τ_s , does not imply that human capital is undistorted. For instance, if there is a positive labor wedge, human capital is indirectly distorted downward, since part of its return is taxed away. Similarly, if there is a positive savings wedge, human capital investments could be distorted upward since they allow the agent to transfer resources to the future without being subject to the savings tax.

Hence, part of the subsidy on human capital is simply undoing some of the effects of the labor and capital distortions on human capital. It would thus be useful to find a measure of the net distortion on human capital. A natural benchmark is the human capital subsidy that ensures that the tax system is neutral with respect to human capital. By neutrality, it is meant that conditional on the labor choice, the human capital decision is set according to the first-best rule (where consumption is also perfectly smooth). Formally, neutrality implies that

$$l_t w_{s_t}(s_t, \theta_t) = M'_t(s_t - s_{t-1}) - \frac{1}{R} E_t(M'_{t+1}(s_{t+1} - s_t)).$$

To build up the intuition, let us first think of a one-period version of the model, with $s = e$ and linear taxes and subsidies. An agent of type θ solves

$$\max_{s,l} u(w(s, \theta)l(1 - \tau_L) - M(s) + \tau_S s) - \phi(l). \quad (16)$$

The first-order condition with respect to human capital yields $w_s l(1 - \tau_L) - M'(s) + \tau_S = 0$. Imagine that we set the subsidy to be $\tau_S =$

$\tau_L M'(s)$. Then we obtain $(1 - \tau_L)[w, l - M'(s)] = 0$; that is, conditional on the labor choice, the tax system is neutral with respect to human capital or, put differently, guarantees a first-best efficient investment (conditional on the labor choice). Setting the subsidy equal to the marginal cost times the labor tax rate is the equivalent of making human capital expenses fully tax deductible; that is, taxable income is only $wl - M(s)$. The “net subsidy” on human capital, that is, the subsidy given in addition to full deductibility, is then (appropriately scaled)

$$t_{st} \equiv \frac{\tau_s - \tau_L M'(s)}{[M'(s) - \tau_s](1 - \tau_L)}.$$

As just shown, it is the subsidy that is zero when there is full deductibility and positive when human capital is encouraged more than at full deductibility.¹³

In this more complex multiperiod model, I introduce a similar concept of *full dynamic, risk-adjusted deductibility*. The standard contemporaneous full deductibility no longer guarantees neutrality of the tax system with respect to human capital as in the one-period model above. Instead, we need to account for four more elements.

First, marginal utility varies across states because of imperfect insurance and, thus, the agent does not value one dollar of transfer the same across states. Accordingly, any monetary amount has to be adjusted for risk using the insurance factors ξ . Second, the marginal benefits of human capital last into the future. An agent should be allowed to deduct from his taxable income today only his cost today minus the cost he is saving by not having to invest in the next period; that is, the deductible amount is

$$M_t(s_t - s_{t-1}) - \frac{1 - \xi_{M'}}{R(1 - \tau_K)} E_t(M_{t+1}(s_{t+1} - s_t))$$

for any choice of s_t rather than just $M_t(s_t - s_{t-1})$.

Third, there is a savings wedge that distorts intertemporal transfers. To avoid an arbitrage between transferring resources as savings versus as human capital, the agent pays a capital tax on the transfer of resources that would be needed for him to invest $s_{t+1} - s_t$ tomorrow, which is

¹³ The net subsidy is paid to the agent in addition to allowing him to fully deduct human capital expenses from taxable income. It is defined as a fraction of the net human capital benefit; i.e., for each unit of human capital investment, the agent receives (pretax) a net subsidy t_s as a fraction of $(1 - \tau_L)w(s, \theta)l$. The agent's consumption is then $c = (1 - \tau_L)[w(s, \theta)l - M(s) + t_s(1 - \tau_L)w(s, \theta)ls]$, which is exactly equivalent to consumption in (16) (expressed there without full deductibility and with the gross subsidy) using the definition of the net subsidy.

$$\frac{1}{R(1 - \tau_K)} (1 - \tau_{L_t})(1 - \xi'_{M'}) E_t(M_{t+1}(s_{t+1} - s_t)).$$

Finally, because there is also a human capital subsidy in the next period, $\tau_{s_{t+1}}$, the agent receives a compensation in this period for any investment s_t of the amount $(1 - \xi_{\tau_s})/R(1 - \tau_K)E_t(\tau_{s_{t+1}})(s_{t+1} - s_t)$. One can check that this ensures that, conditional on current labor supply l_t and future human capital s_{t+1} , the tax system is neutral with respect to human capital.¹⁴ Hence, I define the “net wedge” as the gross wedge from which we filter out all the parts just explained that only go toward compensating for the other distortions.

DEFINITION 1. Define the net wedge on human capital expenses, t_{st} , as

$$t_{st} \equiv \frac{\tau_{s_t}^d - \tau_{L_t} M_t'^d + P_t}{(M_t'^d - \tau_{s_t}^d)(1 - \tau_{L_t})}, \quad (17)$$

where

$$\tau_{s_t}^d \equiv \tau_{s_t} - \frac{1 - \xi_{\tau_s}}{R(1 - \tau_K)} E_t(\tau_{s_{t+1}})$$

is the dynamic risk-adjusted subsidy,

$$M_t'^d \equiv M_t' - \frac{1 - \xi_{M'}}{R(1 - \tau_K)} E_t(M_{t+1}')$$

denotes the dynamic, risk-adjusted cost, and

$$P_t \equiv \frac{\tau_K}{R(1 - \tau_K)} (1 - \tau_{L_t})(1 - \xi'_{M'}) E_t(M_{t+1}')$$

captures the risk-adjusted savings distortion.

A zero net wedge t_{st} again means that the tax system is neutral with respect to human capital. If

$$\tau_{s_t} = \tau_{L_t} M_t'^d - P_t + \frac{1 - \xi_{\tau_s}}{R(1 - \tau_K)} E_t(\tau_{s_{t+1}})$$

such that for every marginal investment e_t a locally linear subsidy $\tau_{s_t} e_t$ is received, there is full dynamic risk-adjusted deductibility.

D. The Optimal Human Capital Subsidy

Both the labor wedge and the net human capital wedge arise because of redistribution and insurance. At the optimum, they are constrained to

¹⁴ The insurance factors are evaluated at the efficient investment level and taken as given by the agent; this is the sense in which wedges are akin to only locally linear taxes and subsidies.

follow a specific relation, which is akin to a type of inverse elasticity rule. The next two propositions, which contain the two main theoretical results of the paper, describe the relation between these two wedges and determine the sign of the net human capital wedge.

PROPOSITION 1. At the optimum and at each history, the labor and human capital wedges need to satisfy the following relation:

$$t_{st}^* = \left(\frac{\tau_{Lt}^*}{1 - \tau_{Lt}^*} \right) \frac{\varepsilon_t^c}{1 + \varepsilon_t^u} (1 - \rho_{\theta_s,t}). \quad (18)$$

In addition, $t_{st}^*(\theta^{t-1}, \bar{\theta}) = t_{st}^*(\theta^{t-1}, \underline{\theta}) = 0$ for all t .

Despite the complexity of the model, formula (18) gives us a clear link between the labor wedge and the net human capital wedge. The two wedges need to comove if and only if $\rho_{\theta_s} < 1$. A particularly appealing relation arises if the wage is a CES function as in (3) with ρ_t constant and disutility is separable and isoelastic:

$$\phi(l) = \frac{1}{\gamma} l^\gamma \quad (\gamma > 1). \quad (19)$$

In this special case, the ratio of the net human capital wedge and labor wedge has to be constant over time and over agents:

$$t_{st}^* / \left(\frac{\tau_{Lt}^*}{1 - \tau_{Lt}^*} \right) = \frac{1 - \rho}{\gamma}.$$

This relation can be used to simply check for the optimality of a given existing tax and subsidy system.

The zero distortion at the bottom and top result, familiar for the labor wedge, holds here for the net human capital wedge. It does not hold for the gross wedge τ_{st} , underscoring again that the true incentive effects are captured by t_{st} , not τ_{st} .

If assumption 3 holds, which entails that the optimal labor wedge is nonnegative at all histories (see the proof in App. B), the sign of the net human capital wedge is determined by the Hicksian coefficient of complementarity, ρ_{θ_s} : the net human capital wedge is positive if and only if $\rho_{\theta_s} < 1$.

PROPOSITION 2. If assumption 3 holds, then

$$t_{st}^*(\theta^t) \geq 0 \Leftrightarrow \rho_{\theta_s,t} \leq 1.$$

An alternative, less stringent condition than (3) would be that $\tau_{Lt}^*(\theta^t) \geq 0$.

The Hicksian coefficient of complementarity ρ_{θ_s} , hence plays a key role in determining the sign of the net human capital wedge and its positive or negative relation to the labor wedge in (18). This key parameter ρ_{θ_s} and the critical threshold of one are discussed next.

E. The Redistributive and Insurance Values of Human Capital

The optimal net wedge results from the balance of two effects that human capital has on social welfare. First, because it increases returns to work, it encourages labor supply, the “labor supply effect.” This is beneficial in the presence of a distortion in the labor decision.¹⁵ At the same time, if $\rho_{\theta_s} > 0$, which is equivalent to ability being complementary to human capital in the wage ($\partial^2 w / \partial \theta \partial s > 0$), human capital mostly benefits already able agents and hence compounds existing inequality due to intrinsic differences in θ_r . The opposite occurs if $\rho_{\theta_s} < 0$, in which case human capital reduces inequality. This effect will be labeled the “inequality effect.” Because ability is stochastic, it is equivalent to say that if $\rho_{\theta_s} > 0$, human capital increases exposure to risk, because it mostly benefits agents when they receive high productivity shocks. Put differently, if $\rho_{\theta_s} > 0$, a human capital subsidy increases pretax income inequality.

This does not mean that human capital investments are undesirable: The question is what happens to posttax inequality. The answer is that, when $\rho_{\theta_s} < 1$, the positive labor supply effect dominates any potential inequality effect and human capital reduces posttax inequality. In this case, high-ability agents do not disproportionately benefit from human capital. Hence, it will be beneficial to encourage human capital on net. In this case, human capital is said to have a positive insurance effect, or a positive redistributive effect on after-tax income inequality.¹⁶

Technically, the inequality effect is the result of a rent transfer, which arises from the need to satisfy agents’ incentive compatibility constraints. If high-productivity agents benefit more from a marginal increase in human capital than lower-productivity agents ($\rho_{\theta_s} > 0$), an increase in their human capital tightens their incentive constraints. What matters for social welfare is whether, when encouraging human capital, the increase in resources from more labor is completely absorbed by the compensation forfeited to high-productivity agents to satisfy their incentive compatibility constraints, or whether there are resources left over. The latter case happens when $\rho_{\theta_s} < 1$, so that human capital investments generate net resources to be used for redistribution and insurance of all agents.

A special case is the multiplicatively separable wage $w = \theta h(s)$ for some function h of human capital. This wage function entails $\rho_{\theta_s} = 1$

¹⁵ As is clear from formula (18), this effect would disappear if there were no distortion on labor. This is different from the direct effect of human capital on earnings, through the increase in wage, which would exist even with no distortion on labor, or with fixed labor supply, and which is filtered out from the net subsidy.

¹⁶ Importantly, the social objective assigns nonnegative weights to all agents, and hence all consumption gains arising from higher resources are positively weighted, even if they increase inequality. Any Pareto improving rise in human capital would be stimulated. But the subsidy does not encourage rises in human capital, which benefit some agents at the expense of having to draw resources from other agents, with a resulting negative change in total social welfare.

and, hence, a null net wedge at the optimum. This is an application of the Atkinson and Stiglitz (1976) result on the nonoptimality of differential commodity taxation if preferences satisfy a form of separability between goods and labor. In a one-period model, it would be the “Siamese twins” result in Bovenberg and Jacobs (2005). It is interesting to note that it does not apply to the gross human capital wedge τ_s since the latter does not capture the full distortion on human capital. That a zero net subsidy is optimal when $\rho_{\theta_s} = 1$ does not depend on the optimality of the labor or intertemporal wedges.¹⁷ Indeed, when $\rho_{\theta_s} = 1$, the choice of education does not reveal any additional information on ability, as all types benefit equally from it in proportional terms.

Returning to the relation between the labor and net human capital wedges in (18), while the labor wedge typically has a positive redistributive or insurance value, the net human capital wedge has a positive redistributive value if and only if $1 - \rho_{\theta_s} > 0$.¹⁸ The optimal policies must be consistent with each other: if the labor wedge is higher so as to provide more insurance, the net human capital wedge must also be higher if and only if $1 - \rho_{\theta_s} > 0$.

Finally, if we extend the analysis to several types of human capital, s_1, \dots, s_j , with different Hicksian coefficients of complementarity, denoted by $\rho_{\theta_{s_j}}$, $j = 1, \dots, J$, formula (18) would apply for each type of human capital, so that at the optimum

$$\frac{t_{s,t}^*}{1 - \rho_{\theta_{s,t}}} = \frac{t_{s,t}^*}{1 - \rho_{\theta_{s,t}}} \quad \forall (i, j). \quad (20)$$

All else equal, human capital types that have the highest redistributive and insurance effects would be more subsidized on net.

F. A Recursive Representation of the Human Capital Wedge

While formula (18) links the labor and human capital wedges, it is also informative to express the optimal net wedge as a function of more primitive factors, as is done in the next proposition.

PROPOSITION 3.

- i. At the optimum, the net wedge can be written as

$$t_{st}^*(\theta^t) = \frac{\mu(\theta^t) u'_i(c(\theta^t)) \varepsilon_{\theta_{s,t}}}{f'(\theta_t | \theta_{t-1}) \theta_t} (1 - \rho_{\theta_{s,t}}), \quad (21)$$

where the multiplier $\mu(\theta_t)$ on the envelope condition can be written as

¹⁷ As in the direct proofs of the Atkinson-Stiglitz result by Laroque (2005).

¹⁸ As long as $\tau_{L_t} \geq 0$, i.e., assumption 3 holds.

$$\mu(\theta^t) = \kappa(\theta^t) + \eta(\theta^t), \quad (22)$$

$$\kappa(\theta^t) = \int_{\theta_s}^{\bar{\theta}} (1 - g_s) \frac{1}{u'_t(c(\theta^{t-1}, \theta_s))} f(\theta_s | \theta_{t-1}) d\theta_s, \quad (23)$$

with

$$g_s = u'_t(c(\theta^{t-1}, \theta_s)) \lambda_{t-1}$$

and

$$\lambda_{t-1} = \int_{\theta}^{\bar{\theta}} \frac{1}{u'_t(c(\theta^{t-1}, \theta_m))} f(\theta_m | \theta_{t-1}) d\theta_m,$$

$$\eta(\theta^t) = i_{st}^*(\theta^{t-1}) \left\{ \frac{R\beta}{u'_{t-1}(c(\theta^{t-1}))} \frac{1}{1 - \rho_{\theta_s, t-1} \varepsilon_{u\theta, t-1}} \frac{\theta_{t-1}}{\int_{\theta_s}^{\bar{\theta}} \frac{\partial f(\theta_s | \theta_{t-1})}{\partial \theta_{t-1}} d\theta_s} \right\}. \quad (24)$$

ii. $i_{st}^*(\theta^t) = 0$ if $u'_t(c_t) = 1$ and θ_t is iid.

This representation makes it clear that, while the coefficient of complementarity ρ_{θ_s} determines the sign of the net wedge, the two terms κ and η modulate its amplitude. The insurance motive is captured in $\kappa(\theta^t)$ familiar from the static taxation literature. It would be zero with linear utility. The term g_s is the marginal social welfare weight on an agent of type θ_s , measuring the social value of one more dollar transferred to that individual, and $1/\lambda_{t-1}$ is the social cost of public funds.

The term $\eta(\theta^t)$ captures the previous period's net wedge, hence indirectly the previous period's insurance motive, weighted by a measure of ability persistence. Recall that there can be a redistributive motive in the first period if there is initial heterogeneity.¹⁹ This motive remains effective through $\eta(\theta^t)$, the more so if types are more persistent, but vanishes as skills become less persistent. In the limit, if θ_t is iid, only the contemporaneous insurance motive $\kappa(\theta^t)$ matters. If, in addition to iid shocks, utility is linear in consumption, the optimal net subsidy is zero.²⁰

¹⁹ In the first period, heterogeneity in θ_t leads to

$$\mu(\theta_1) = \int_{\theta_s}^{\bar{\theta}} \frac{1}{u'_1(c_1(\theta_s))} [1 - \lambda_0(\theta_s) u'_1(c_1(\theta_s))] f(\theta_s),$$

where $\lambda_0(\theta_s)$ is the multiplier (scaled by $f(\theta_s)$) on type θ_s target utility. With linear utility, $1 = \int_{\theta}^{\bar{\theta}} \lambda_0(\theta_s) f(\theta_s)$.

²⁰ Except in the first period with a different utility threshold for different agents.

Da Costa and Maestri (2007), Maldonado (2008), Anderberg (2009), and Bovenberg and Jacobs (2011) also emphasize the role of the complementarity between human capital and ability (or risk). This dynamic model nests their findings. Subject to the suitable redefinition of the net wedge, it is striking that the sign of the net wedge is still determined as it would be in a completely static model, by the sign of $1 - \rho_{\theta_s}$, which can now vary over the life cycle. All dynamic considerations affect the net wedge's amplitude through the multiplicative term $\mu (= \kappa + \eta)$.

G. Age Dependency

In equation (21), the optimal wedge was expressed recursively and pointwise as a function of the previous period's wedges. One can also rewrite the formulas in terms of a weighted expectation across types at time t , using some weighting function $\pi(\theta)$.²¹ For the sake of the exposition, ability is assumed to follow a log autoregressive process:

$$\log(\theta_t) = p \log(\theta_{t-1}) + \psi_t, \quad (25)$$

where ψ_t has density $f^\psi(\psi|\theta_{t-1})$, with $E(\psi|\theta_{t-1}) = 0$. The general formulas for any stochastic process for the labor wedge and the net human capital wedges are in Appendix A.

COROLLARY 1. The optimal net subsidy evolves over time according to

$$\begin{aligned} E_{t-1} \left(t_{st} \frac{\varepsilon_{w\theta,t-1}}{\varepsilon_{w\theta,t}} \frac{1 - \rho_{\theta_s,t-1}}{1 - \rho_{\theta_s,t}} \left(\frac{1}{R\beta} \frac{u'_{t-1}}{u'_t} \right) \right) \\ = \varepsilon_{w\theta,t-1} (1 - \rho_{\theta_s,t-1}) \text{Cov} \left(\frac{1}{R\beta} \frac{u'_{t-1}}{u'_t}, \log(\theta_t) \right) + p t_{st-1}. \end{aligned} \quad (26)$$

Formula (26) exhibits a drift and an autoregressive term, like the labor wedge in Farhi and Werning (2013). Dynamic incentive compatibility constraints cause a positive covariance between consumption growth and productivity: by promising them higher consumption growth, the government induces higher-ability agents to truthfully reveal their types. This, however, makes insurance valuable and is captured by the drift term. The insurance motive is magnified by the sensitivity of the wage to ability $\varepsilon_{w\theta,t}$ and the redistributive or insurance factor of human capital $(1 - \rho_{\theta_s})$. The drift term inherits the sign of the latter; when $\rho_{\theta_s} \leq 1$, human capital has a positive insurance effect, which caters well to the rising need for insurance. The Hicksian complementarity ρ_{θ_s} might vary over

²¹ Different weighting functions $\pi(\theta)$ lead to different recursive relations, which must hold at the optimum, and some weighting functions draw out particularly enlightening effects. Such a reformulation for the optimal labor wedge formula was proposed by Farhi and Werning (2013) for a convenient weighting function $\pi(\theta) = 1$.

life. If it is decreasing faster, the net subsidy will rise faster or fall slower over the life cycle. The sensitivity of the wage amplifies the effect of the Hicksian complementarity coefficient in either direction. The persistence of the shock p translates into a persistence for the labor wedge. Over time, in the optimal mechanism, there will be “subsidy smoothing”: the net subsidy becomes more strongly correlated over time as age increases, because the variance of consumption growth falls to zero, which makes the drift term in the formula above vanish.

Finally, in the real world, education policies are often set independently from taxes. If we suppose that the tax system is fixed and only human capital policies can be optimized, the net subsidy will be set as before, but with additional terms that capture the indirect effect of human capital on labor supply (see formula [A8] in App. A).

V. Numerical Analysis

The numerical analysis has four goals: first, to highlight the quantitative importance of the Hicksian coefficient of complementarity ρ_{θ} ; second, to compare the labor and capital wedges to those arising in a standard model without human capital; third, to highlight the phenomenon of subsidy smoothing over the life cycle; and, fourth, to study the progressivity of the human capital subsidy and the labor wedge. Computational procedure, calibration details, and additional results for alternative calibrations are in the online computational appendix.

Before I present simulation results, the empirical evidence on the crucial parameter of the model, namely, the complementarity between human capital and ability in the wage, is discussed.

A. *Empirical Evidence on the Complementarity between Human Capital and Ability*

The formulas for the optimal net wedge and its evolution (see [17] and [26]) highlighted the importance of the complementarity between ability and human capital. Unfortunately, the evidence to date does not yield a conclusive answer regarding its magnitude, and further empirical work on this key parameter seems needed.²²

²² There is some evidence for early childhood investments. Ashenfelter and Rouse (1998) suggest that lower-ability children benefit more from schooling, a finding also in line with Cunha et al. (2006) for early childhood interventions. This is consistent with $\rho_{\theta} \leq 1$ for childhood investments. There is evidence that the Hicksian complementarity can change over life, as suggested by the structural literature on human capital formation (e.g., Cunha and Heckman 2007), so that estimates for primary and secondary schooling are of only limited use for the analysis of higher education or job training (which are the focus of this paper).

Several studies show that college education might mostly benefit already able students, implying that $\rho_{\theta_s} > 0$ and that $\rho_{\theta_s} > 1$ is possible. Cunha et al. (2006) estimate that the return to 1 year of college is around 16 percent at the 5th percentile of the math test scores distribution, as opposed to 26 percent at the 95th percentile. There is only scarce evidence on the complementarity between on-the-job training and ability, but the same authors show that on-the-job training is mostly taken up by those with higher Armed Forces Qualification Test scores, which might, all else equal, signal that they have a higher marginal return from it.²³ The OECD (2004) reports that training mostly benefits skilled workers in terms of higher wages but benefits low-skilled workers in terms of job security. Huggett, Ventura, and Yaron (2011) use a multiplicatively separable functional form for the wage in their structural model of time investments in human capital (implying $\rho_{\theta_s} = 1$), which generates a life cycle path of earnings that matches the data well.

B. Calibration

Calibration to US data.—To calibrate the model, I construct a “baseline economy,” which has the same primitives as in Section II but no social planner and, hence, no optimal tax system. Instead, the linear labor taxes, capital (here, savings) taxes, and human capital subsidies are set to their current averages in the United States. Public and private subsidies in the United States cover around 50 percent of total resource costs of formal higher education. However, in the model, human capital expenses are a comprehensive measure, including all types of formal and informal investments, and some mostly unsubsidized expenses (e.g., textbooks, computers). In addition, there are practically no subsidies beyond the initial 4 years of higher education. Hence, the linear subsidy in the baseline model, applicable to all expenses, is reduced to 35 percent for the first two periods and to 0 thereafter.²⁴ The linear labor tax rate is set to 13 percent and the savings tax to 25 percent.²⁵ Agents can borrow and save at a constant interest rate and start with zero asset holdings.²⁶ In this baseline economy, some parameters are set exogenously on the basis of the existing literature, while others are set to endogenously match two key moments from the data, namely, a wage premium and a ratio of

²³ It is also not evident that the test scores used as measures of “ability” are themselves exogenous, especially at later ages.

²⁴ See indicators B in OECD (2013).

²⁵ Only interest income and short-term capital gains are taxed at ordinary income rates. Taxes on a lot of dividends and long-term capital gains stop at 20 percent (plus possibly 2.9 percent for the new Medicare tax above \$250,000 of total income and various state taxes).

²⁶ Because of a natural debt limit consideration, agents do not choose to hold negative assets.

human capital expenses to lifetime income, as explained in more detail next.

Functional forms.—Agents live for $T = 30$ periods, each representing roughly 2 years of life. They work for 20 periods and spend 10 years in retirement. Preferences during working years are given by

$$\tilde{u}(c_t, y_t, s_t, \theta_t) = \log(c_t) - \frac{\kappa}{\gamma} \left[\frac{y_t}{w_t(\theta, s)} \right]^\gamma, \quad \kappa > 0, \gamma > 1.$$

During retirement, utility is simply $\tilde{u}^R(c_t) = \log(c_t)$. The aforementioned literature seemed to indicate that $\rho_{\theta_s} > 0$. The wage is assumed to be CES, which allows us to cleanly focus on the (constant) Hicksian coefficient of complementarity $\rho_{\theta_s} = \rho$:

$$w_t(\theta, s) = (\theta^{1-\rho} + b_s s^{1-\rho})^{1/(1-\rho)}, \quad (27)$$

where $b_s > 0$ is a scaling factor. Two values of ρ are studied in the main text, namely, $\rho = 0.2$ and $\rho = 1.2$, with additional values considered in the online computational appendix.

The cost function of human capital contains an adjustment term and takes the form

$$M_t(s_{t-1}, e_t) = b_l e_t + b_a \left(\frac{e_t}{s_{t-1}} \right)^2,$$

with b_l and b_a the linear and the adjustment components of cost. Consistent with the high persistence in earnings documented in Storesletten, Telmer, and Yaron (2004), ability is assumed to follow a geometric random walk:

$$\log \theta_t = \log \theta_{t-1} + \psi_t,$$

with $\psi_t \sim N(-1/2\sigma_\psi^2, \sigma_\psi^2)$.

Exogenously calibrated parameters.—The baseline model has $\gamma = 3$ and $\kappa = 1$, which implies a Frisch elasticity of 0.5 (Chetty 2012). The discount factor is set to $\beta = 0.95$ and the net interest rate to 5 percent. The adjustment cost is normalized to $b_a = 2$. There is large variation across empirical estimates of the variance of ability σ_ψ^2 . A medium-range value of 0.0095 (Heathcote, Storesletten, and Violante 2005) is adopted, with several alternative values for this parameter explored in the online computational appendix.

Endogenously matched parameters.—The scaling factor for human capital b_s and the linear cost parameter b_l are set to match two statistics in the data: a wage premium and a ratio of human capital expenses to income. One complication that arises is that the model does not a priori restrict investments to occur only during the traditional “college” years. Indeed,

one of the motivations for this study is the lifelong nature of human capital investments. However, most available estimates for wage premiums in the literature are for college education, with scarce evidence on job training. This difficulty can be overcome by redefining college appropriately for the model. Following Autor, Katz, and Krueger (1998), who find 42.7 percent full-time college equivalents, the top 42.7 percent in the population of the baseline economy, ranked by educational expenses, are assumed to represent the real-life college goers. Their average wage relative to the bottom 42.7 percent is set to match the wage premium for college estimated in the literature.²⁷ These estimates range from 1.58 in Murphy and Welch (1992), to between 1.66 and 1.73 for Autor et al. (1998), and above 1.80 in Heathcote, Perri, and Violante (2010). The calibration targets a midrange value of 1.7.

Turning to the second target, the net present value of higher education expenses over the net present value of lifetime income in the data is 13 percent.²⁸ However, since agents invest beyond traditional college years in the model, there needs to be an allowance for later-in-life investments. It is assumed that college costs represent two-thirds of all lifetime investments in human capital, so that the target ratio of the net present value of lifetime human capital expenses and the net present value of lifetime income is 19 percent. Table 1 summarizes the resulting parameters.

Using the computed policy functions, a Monte Carlo simulation with 100,000 draws is performed for each value of ρ_{θ} . The initial states are set to yield a zero present value resource cost for the allocation. This ensures comparability across simulations and gives a sense of outcomes achievable without outside government revenue.

To judge how well this calibration fits the data, online appendix B.1 contains graphs of the implied mean consumption, income, human capital expenses, and assets, as well as the variances of wages, income, and consumption over the life cycle, in the baseline economy. The processes for earnings and wages match the data well to a first order. In particular, earnings and the wage follow a near-random walk (as reported in Heathcote et al. [2005]) and the variances of log wages and earnings match those in Heathcote et al. (2010).

C. Results

The labor and savings wedges with and without human capital.—Figure 1 compares the labor and savings wedges that arise at the optimum for each

²⁷ The middle 14.6 percent are omitted to clearly delineate between college goers and others. Indeed, because of the continuous investments in the model, there is no sharp distinction between “college” and “no college” and, in particular, no notion of “a degree.”

²⁸ Computed using data on attendance at different types of colleges and costs (Wei 2010; OECD 2013). See the detailed calculations in the online computational appendix.

TABLE 1
CALIBRATION

Parameter	Definition	Simulation		Source/Target
		1	2	
Exogenously Calibrated or Normalized				
ρ	Hicksian complementarity	.2	1.2	Cunha et al. (2006)
κ	Disutility of work scale	1	1	
γ	Disutility elasticity	3	3	Chetty (2012)
σ_v^2	Variance of productivity	.0095	.0095	Heathcote (2005)
T	Working periods	20	20	
T_r	Retirement periods	10	10	
β	Discount factor	.95	.95	
R	Gross interest rate	1.053	1.053	
b_a	Scale of adjustment cost	2	2	
Endogenously Calibrated in Baseline Economy				
b_s	Scale of human capital in wage	.09	.1	Wage premium (Autor et al. 1998)
b_t	Linear cost	.5	.5	Wei (2010); OECD (2013)

NOTE.—Simulation 1 refers to the simulation for $\rho_{\theta_s} = 0.2$, while simulation 2 refers to the simulation for $\rho_{\theta_s} = 1.2$. The last column shows the source of the parameter for exogenously calibrated parameters and the target moment (and its source) for parameters endogenously calibrated in the baseline economy. If the source/target is blank, the parameter is normalized.

value of ρ_{θ_s} to those that arise in a model without human capital, where the wage is equal to exogenous ability, $w_t = \theta_t$. The process for θ is rescaled so that both models generate the same variance of wages over the life cycle (which, as noted, are consistent with the paths in Heathcote et al. [2010]). Hence, the exercise consists in holding observed wage volatility fixed and asking how optimal policies would change if that path of wages was generated by endogenous human capital or by a purely exogenous process.

The labor wedge is smaller and grows slower when $\rho_{\theta_s} < 1$. The reason is that human capital fulfills part of the redistribution and insurance and takes some of the burden away from the labor wedge. On the other hand, when $\rho_{\theta_s} > 1$, human capital does not fulfill any insurance or redistributive role and the full burden is again on the labor wedge, which is almost unaffected.²⁹

Panel B plots the savings wedge over time. For $\rho_{\theta_s} = 0.2$, it starts at 0.5 percent of the gross interest on savings, which corresponds to a

²⁹ It is clear from formula (A1) in App. A that the effect of human capital on the labor wedge is ambiguous.

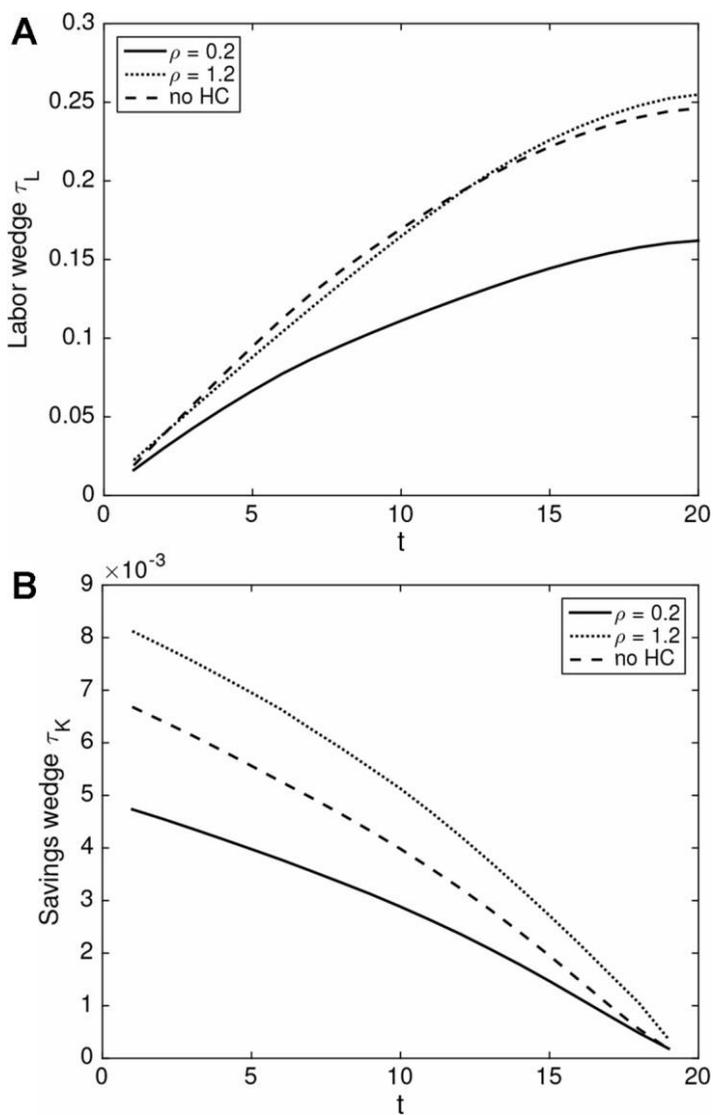


FIG. 1.—Average labor and savings wedges over time. “No HC” denotes the case without human capital. *A*, The labor wedge is lower and grows slower over time in the presence of human capital if human capital has positive redistributive and insurance effects ($\rho_{bs} < 1$). *B*, The savings wedge is also lower in the presence of human capital with a redistributive and insurance value ($\rho_{bs} < 1$), but larger otherwise.

10 percent tax on net interest and declines to zero.³⁰ It is higher when $\rho_{\theta_s} > 1$. Relative to a case with no human capital, savings are less distorted in the presence of human capital with a redistributive value ($\rho_{\theta_s} < 1$) since then, again, human capital helps the planner to screen agents. On the other hand, savings are more distorted when human capital increases the informational rents of high-ability agents ($\rho_{\theta_s} > 1$).

A brief note on the presentation of the wedges is required. When agents are at the corner solution of no investment, which occurs in the baseline calibration for most simulated paths approximately after period 12 or 13, the subsidy is indeterminate, as long as it remains below an upper bound that does not induce agents to invest. Hence, the policy function for the gross wedge is set to zero for agents after they stop investing.³¹ To make it most comparable to an explicit subsidy, it is presented as a fraction of the marginal cost, that is, $\tau_{SI}/M'_t(e_t)$.

The role of the Hicksian complementarity ρ_{θ_s} .—Figure 2 presents the human capital wedges. For this figure only, the focus should be on periods before the dotted vertical lines, as many agents no longer invest in human capital later in life, and wedges are normalized thereafter, as just explained. First, panel A shows that the optimal wedge on human capital is higher and grows faster when human capital has a positive insurance or redistributive effect ($\rho_{\theta_s} \leq 1$). When $\rho_{\theta_s} = 0.2$, the wedge starts from 1 percent and grows to 19 percent; for $\rho_{\theta_s} = 1.2$, it instead starts slightly negative and grows to only 11 percent. Panel B illustrates that with $\rho_{\theta_s} = 1.2$, the net subsidy is negative, so that human capital expenses are made less than fully deductible. Conversely, when $\rho_{\theta_s} = 0.2$, human capital expenses are subsidized on net beyond pure deductibility. The comparison between panels A and B once more highlights that the true incentive effect is different from the gross wedge when there are several wedges present. Finally, the net subsidy is growing when $\rho_{\theta_s} \leq 1$ and declining otherwise, as seen in the drift term of formula (26).

However, the net wedges are very small. Hence, the overall system remains very close to neutrality with respect to human capital expenses. Put differently, full dynamic risk-adjusted deductibility is very close to optimal.³² This is akin to a “production efficiency” result: human capital is an intertemporal decision, with persistent effects, and distorting it for redistributive or insurance reasons is relatively costly unless the redistributive or insurance effects are very strong.

³⁰ The equivalent tax on net interest, $\tilde{\tau}_{KI}$, solves $1 + (R - 1)(1 - \tilde{\tau}_{KI}) = R(1 - \tau_{KI})$.

³¹ The net wedge is then set to the right-hand side in (17), a level that will not artificially induce agents to invest.

³² The gross wedge is correspondingly smaller than a 50 percent real-world subsidy. Bear in mind, however, that the gross wedge here is more comprehensive, as it covers all human capital expenses, even those unsubsidized in the real world, and is available throughout life, not exclusively for college years.

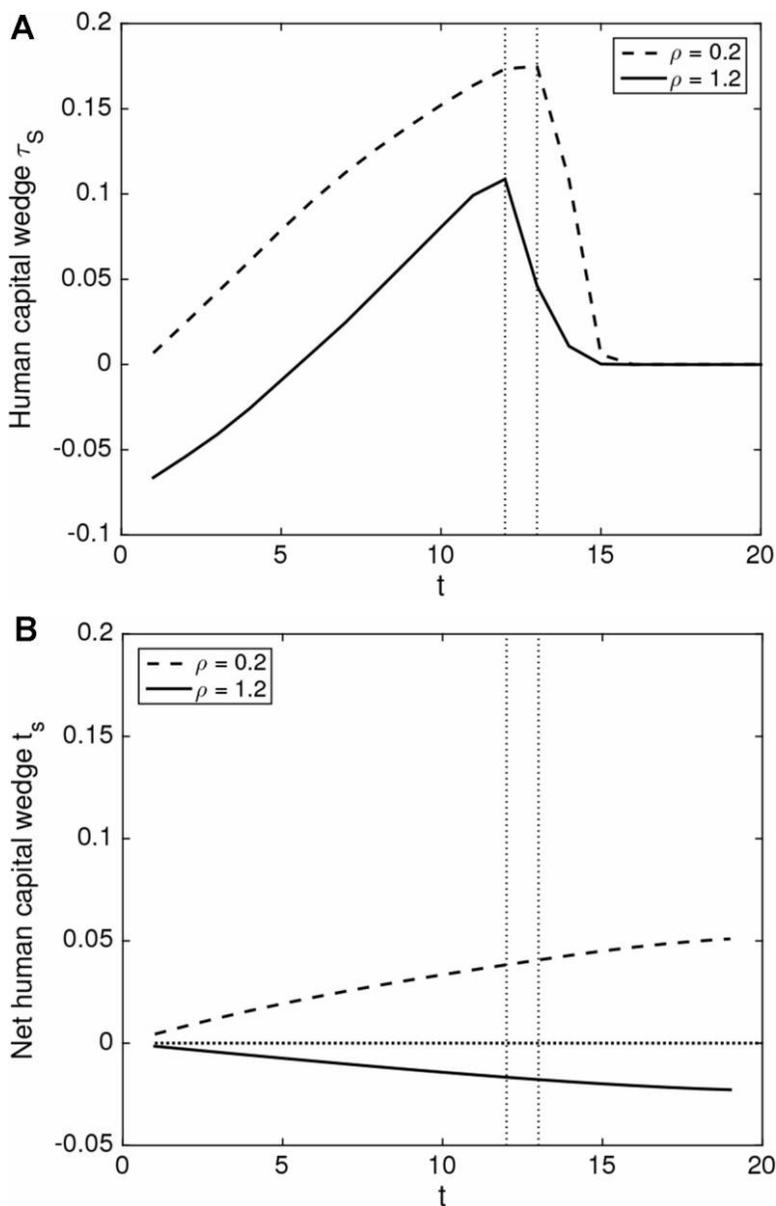


FIG. 2.—Average human capital wedges over time. Dotted vertical lines mark the time after which most agents no longer invest in human capital. *A*, The gross wedge is higher and grows faster when human capital has positive redistributive and insurance effects ($\rho_{\theta_s} < 1$). The wedge is normalized to zero for zero investment (a corner solution). *B*, If $\rho_{\theta_s} < 1$, human capital has positive redistributive and insurance values, and expenses are subsidized on net at a rising rate. Conversely, if $\rho_{\theta_s} > 1$, expenses are only partially deductible, and deductibility decreases over time.

The values of the complementarity between human capital and ability and the volatility of ability clearly matter for this result. Moving further away from a multiplicatively separable wage, that is, increasing $|\rho_{\theta_s} - 1|$, leads to larger net wedges in absolute value. Similarly, a higher volatility increases the value of insurance and yields a higher optimal net wedge if human capital has a positive insurance value ($\rho_{\theta_s} < 1$) and a lower optimal net wedge if not.

The production efficiency result hinges on the fact that the government jointly optimizes the full system, and insurance can also be achieved through the labor wedge. In particular, the planner is able through transfers to endogenously relieve “credit constraints,” which might be motivating high subsidies in the real world, without having to distort human capital acquisition at the margin.

Figure 1, panel A, shows that it is the optimal labor wedge that is large and rises over time to provide insurance against widening income dispersion. Hence, also, a large part of the growing human capital wedge merely goes toward compensating for this growing disincentive to accumulate human capital. But most of the redistribution and insurance go through the intratemporal labor wedge rather than the intertemporal net human capital wedge.³³

Subsidy smoothing over the life cycle.—Figure 3 plots the net subsidy in period t against the net subsidy at $t - 1$ for young adults ($t = 5$ in panel A) and for middle-aged workers ($t = 13$ in panel B). Earlier in life, the net wedge is more volatile from one period to the next but becomes more deterministic over time, leading to a “subsidy smoothing” result. The dynamic taxation literature has highlighted a similar “tax smoothing” result for the labor wedge (which continues to hold in the presence of human capital). The intuitions for these results are the same. A persistent productivity shock early in life has repercussions over many periods, leading to a larger present value change in the income flow than a later shock. Consumption in early years will react strongly to unexpected changes in ability, as the agent attempts to smooth out the shock. Accordingly, the variance of consumption growth is initially large but decreases to zero over time. The drift term in the net subsidy formula (26), which is proportional to the covariance between ability and consumption growth, tends to zero toward retirement. Then, only the autoregressive term of the random walk remains.

Progressivity or regressivity of the net human capital and labor wedges.—Figure 4 plots $\tau_{L,t}$ against the contemporaneous productivity shock, $\theta_{i,t}$, at $t = 19$ (arbitrarily chosen t), and figure 5 does the same for t_{st} . The net

³³ This model ignores the general equilibrium effects of human capital accumulation. Positive spillovers from education, through peer effects, innovation, or social and political channels, could lead to larger subsidies for human capital.

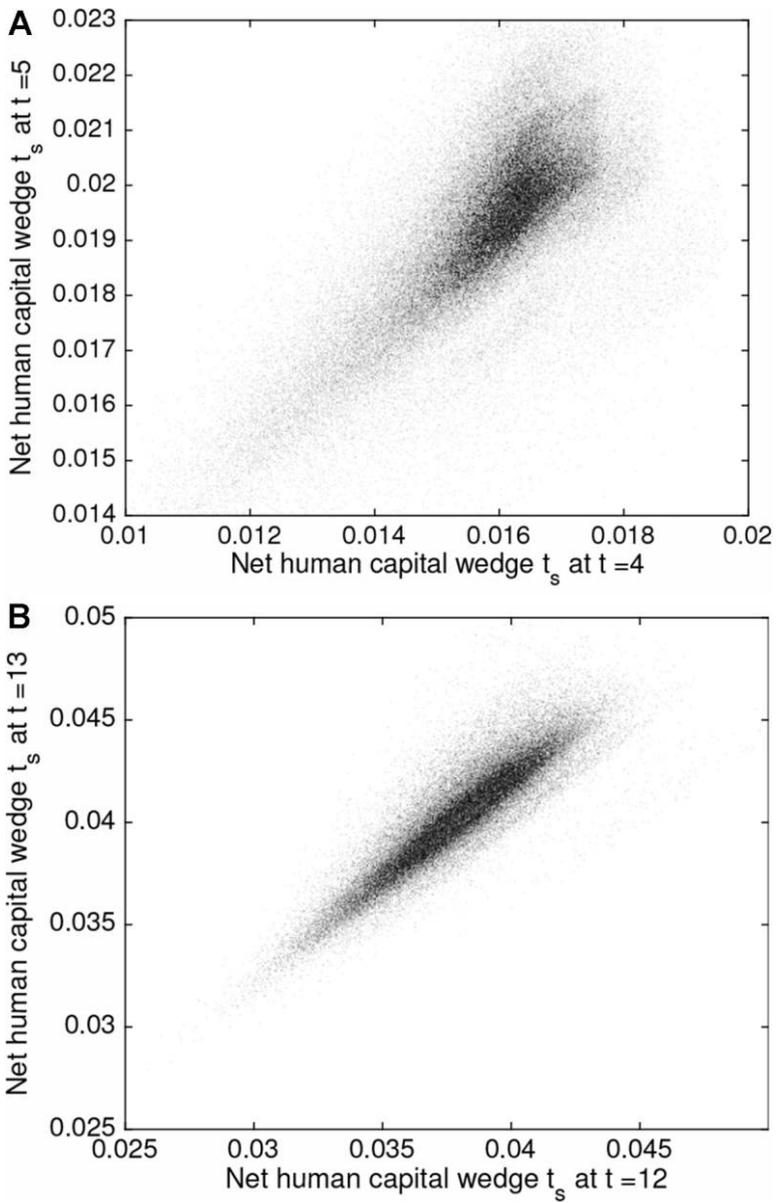


FIG. 3.—Subsidy smoothing over life. The net human capital wedge becomes more correlated from one period to the next as age increases, because the variance of consumption growth, which drives changes in the subsidy over time, vanishes. Figures are for $\rho_{\theta_s} = 0.2$.

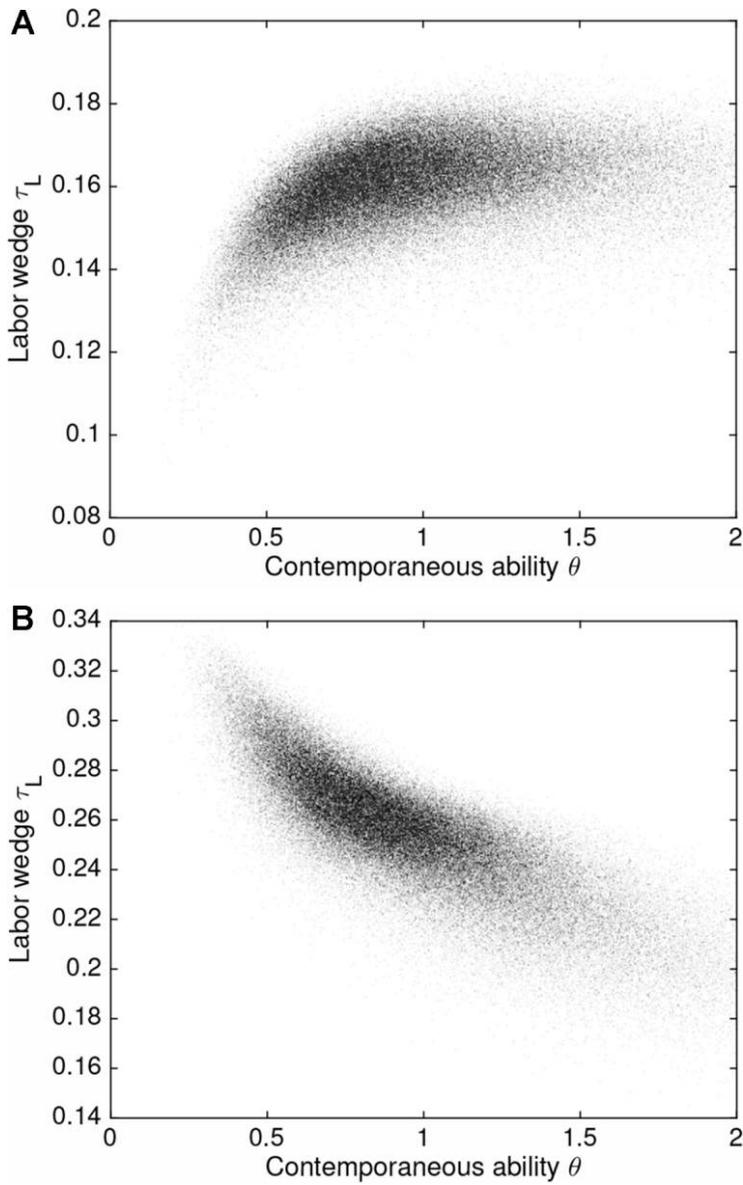


FIG. 4.—Progressivity and regressivity of the labor wedge. The labor wedge exhibits short-run progressivity when $\rho_{\theta_s} < 1$ but short-run regressivity when $\rho_{\theta_s} > 1$.

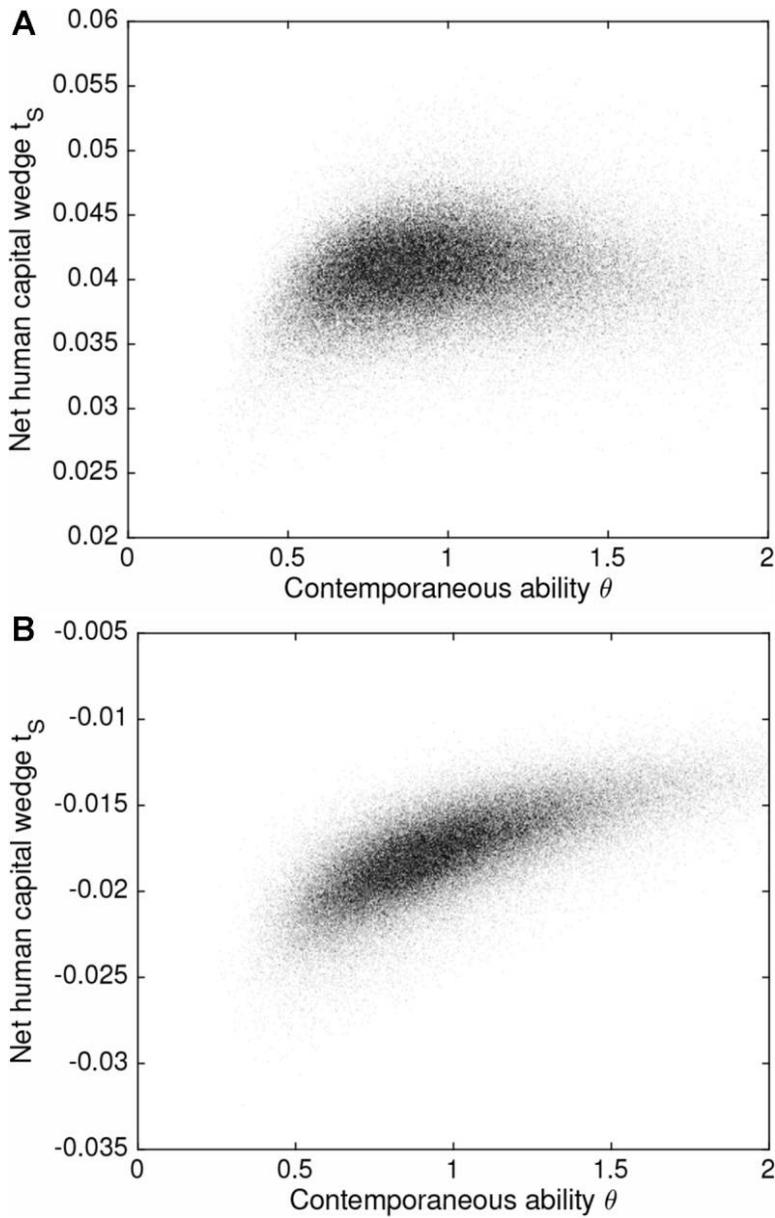


FIG. 5.—Regressivity of the net human capital wedge. The net wedge is always regressive in the short run, but more so when $\rho_{\theta_s} > 1$.

human capital subsidy is always regressive as long as $\rho_{\theta_s} > 0$ because higher-ability people benefit more from human capital. When $\rho_{\theta_s} > 1$, the labor wedge is regressive in the short run, which is true for a similar parameterization of the problem without human capital. On the other hand, when $\rho_{\theta_s} < 1$, the labor wedge exhibits a short-run progressivity. The reason for this reverse pattern is that both the labor wedge and the net subsidy are tools to insure against earnings risk. Along the optimal path, they need to evolve consistently, according to the key relation in (18). The labor wedge always has positive insurance and redistributive effects. The same is true for the net subsidy only if $\rho_{\theta_s} < 1$. Accordingly, the two instruments comove positively when $\rho_{\theta_s} < 1$ and negatively when $\rho_{\theta_s} > 1$.

VI. Implementation and Policy Implications

In this section, we step away from a direct revelation mechanism in which agents report their types and instead consider what policy tools can implement those same allocations. There are many possible implementations, and the theory does not provide guidance as to which one to choose, unless we include additional considerations such as political or administrative constraints. I present two implementations that are particularly appealing because they are variations on policies that already exist.

A. Income-Contingent Loans

Before presenting the ICLs, the decentralized economy is described and some notation introduced. In the decentralized economy, agents choose their human capital expenses e_t , income y_t , and savings b_t in a risk-free account at a gross rate R . Initial wealth is zero and initial human capital is s_0 .³⁴ The government can observe and keep record of the histories of consumption, output, human capital, and wealth.

Denote by $m_t^*(\theta^t)$ the optimal allocation of the social planner's problem after history θ^t for any choice variable $m \in \{c, y, b, e\}$.³⁵ For any history θ^t and subset of variables $\mathbf{m} \subset \{c, y, b, e\}$, let $Q_{\mathbf{m}}^t(\theta^{t-1})$ be the set of values for these variables at time t , which could arise in the planner's problem after history θ^{t-1} , that is, such that for some $\theta \in \Theta$, $\mathbf{m}_t = \mathbf{m}_t^*(\theta^{t-1}, \theta)$. For a history of observed choices \mathbf{m}^t , denote by $\Theta^t(\mathbf{m}^t)$ the set of all histories θ^t consistent with these choices, that is, all θ^t such that $\mathbf{m}_s = \mathbf{m}_s^*(\theta^s)$ for all $s \leq t$. Assumption 2 guarantees that in the planner's

³⁴ Initial wealth and human capital can be heterogeneous as long as they are observable and will enter the proposed repayment schedule as additional conditioning variables.

³⁵ Online app. A shows how to construct them from the recursive allocations derived above.

problem, the histories (y^t, e^t) can be uniquely inverted to identify the history of abilities, θ^t .

ASSUMPTION 2. $\Theta^t(y^t, e^t)$ is either the empty set or a singleton for all histories (y^t, e^t) .

In the proposed ICL scheme, loans are combined with a standard income tax based on contemporaneous income $\tilde{T}_Y(y_t)$ and a history-independent savings tax $\tilde{T}_K(b_t)$. In each period, the agent is offered a government loan $L_t(e_t)$ as a function of his human capital expenses and is required to make a history-contingent repayment $D_t(L^{t-1}, y^{t-1}, e_t, y_t)$, as a function of the full history of past loans and earnings, as well as current income and human capital expenses. Note that, without loss of generality, we could set $\tilde{T}_Y(y_t) = 0$ and adjust the repayment schedule D_t accordingly. The agent's problem is then to select the supremum over $\{c_t(\theta^t), y_t(\theta^t), b_t(\theta^t), e_t(\theta^t)\}_{\theta^t}$ in

$$V_1(b_0, \theta_0) = \sup \sum_{t=1}^T \int \left[u_t(c_t(\theta^t)) - \phi_t \left(\frac{y_t(\theta^t)}{w_t(\theta_t, s_{t-1}(\theta^{t-1}) + e_t(\theta^t))} \right) \right] P(\theta^t) d\theta^t \tag{28}$$

subject to

$$\begin{aligned} & c_t(\theta^t) + \frac{1}{R} b_t(\theta^t) + M_t(e_t(\theta^t)) - b_{t-1}(\theta^{t-1}) - L_t(e_t(\theta^t)) \\ & \leq y_t(\theta^t) - D_t(L^{t-1}(\theta^{t-1}), y^{t-1}(\theta^{t-1}), e_t(\theta^t), y_t(\theta^t)) - \tilde{T}_Y(y_t(\theta^t)) - \tilde{T}_K(b_t(\theta^t)), \\ & s_t(\theta^t) = s_{t-1}(\theta^{t-1}) + e_t(\theta^t), s_0 \text{ given}, e_t(\theta^t) \geq 0, b_0 = 0, b_T \geq 0. \end{aligned}$$

PROPOSITION 4. The optimum can be implemented through human capital loans $L_t(e_t)$, with repayments $D_t(L^{t-1}, y^{t-1}, e_t, y_t)$, contingent on the history of loans and earnings, current income, and human capital expenses, together with a history-independent savings tax $\tilde{T}_K(b_t)$, and an income tax on contemporaneous income $\tilde{T}_Y(y_t)$.

1. Features of the ICL System

Figure 6 illustrates the implementation through ICLs, by plotting the average loan received and average consolidated payment made as a fraction of contemporaneous income. The loan received naturally declines over life, as less human capital investments are needed. The magnitude of the loan depends on the extent of credit constraints in the private market: More stringent credit constraints would naturally lead to larger loans. The repayment is very mildly hump-shaped initially increasing with age, illustrating the insurance provided by the contingent repayment schedule and the increasing ability to pay over life.

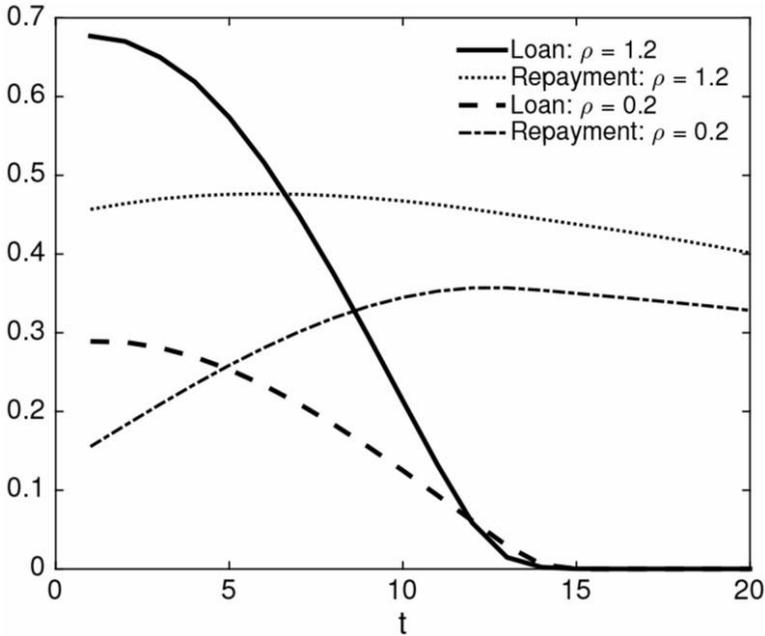


FIG. 6.—Income history–contingent loans: loans and repayments as a fraction of contemporaneous income. Loans are high early in life, while repayments increase to provide insurance.

Figure 7 highlights the insurance role of the repayment schedule. It shows that repayments, as a fraction of income, are increasing in the contemporaneous income realization. The income history–contingent nature of repayments is clearly seen in their large dispersion at a given contemporaneous income: repayments depend on the full past, not only on current income. Repayments are larger when $\rho_{\theta_s} > 1$, highlighting the need to provide more insurance when human capital does not have a redistributive and insurance value. Repayments are also higher and steeper in income when volatility is higher, as the need for insurance is then increased (not shown here, but clear in the formulas in the terms κ).

Figure 8 shows that the implied interest rate, defined as the ratio of the repayment and the outstanding loan balance, is also increasing in the contemporaneous income realization for both values of ρ_{θ_s} . Clearly, the interest rate is not fixed throughout life or across agents: it is history dependent. In the online appendix, figures 9 and 10 show, respectively, that the interest rate is declining in both the current loan size and the outstanding loan balance, so that there is a “quantity discount.” Similar to repayments, a larger variance of ability would make the interest rate steeper in income.

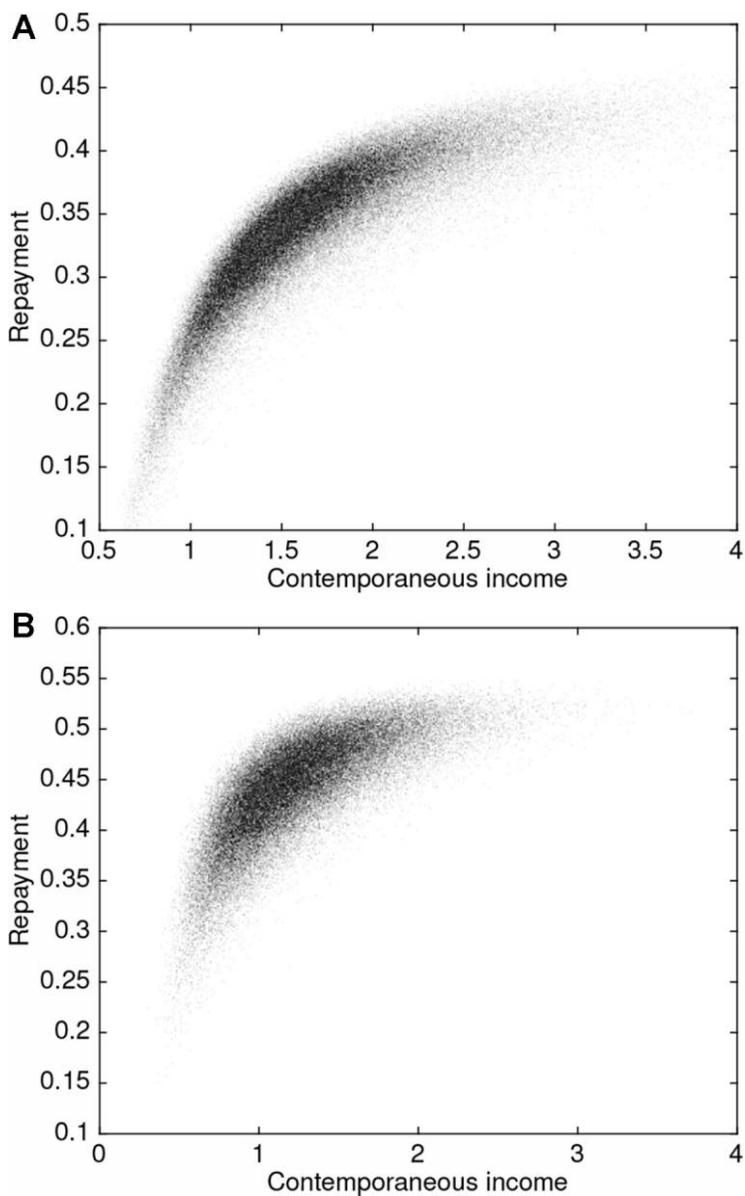


FIG. 7.—Insurance through contingent repayments. Repayment is plotted against contemporaneous income for $t = 15$. The repayment schedule provides insurance: repayments are higher when income is higher. The history-contingent nature is seen in the dispersion of the repayment at a given contemporaneous income.

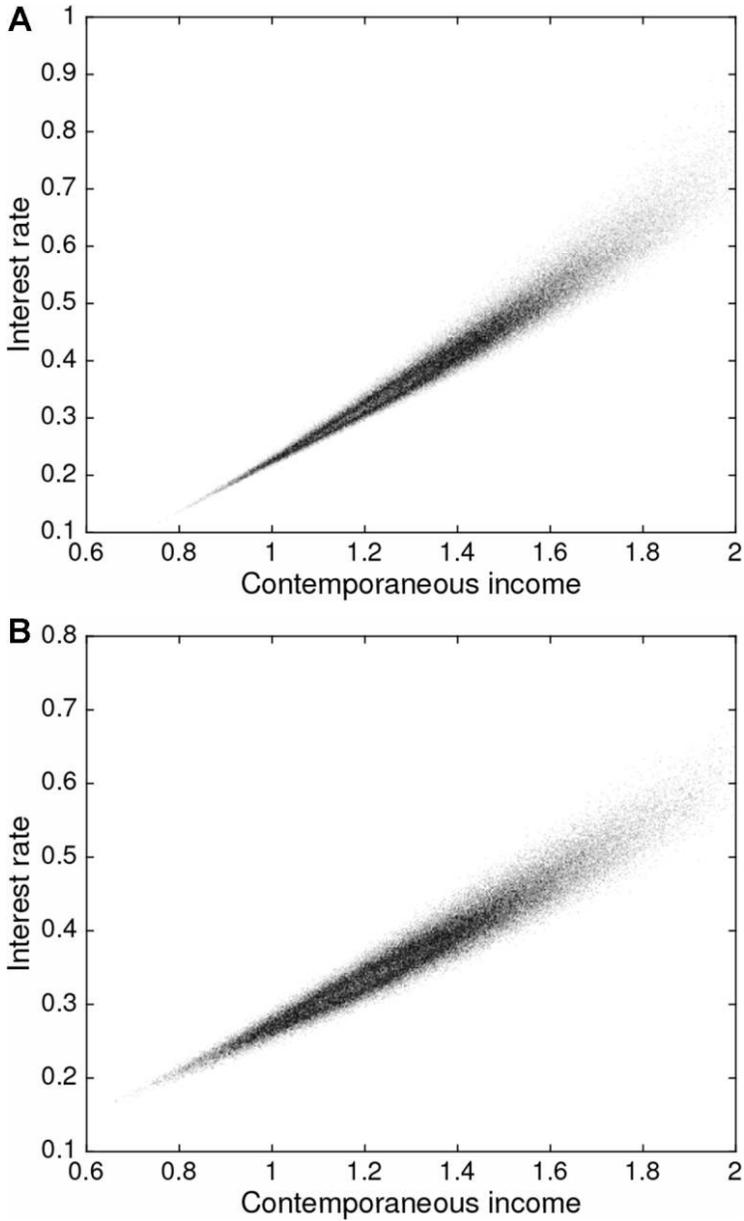


FIG. 8.—Interest rate on income-contingent loans. The implicit interest rate is defined as the ratio of the repayment and the outstanding loan balance. The interest rate provides insurance: interest paid is higher when income is higher.

2. Comparison of the Proposed Implementation to Existing ICLs

Certain types of ICLs for college education are used in several countries, including the United States, New Zealand, Australia, the United Kingdom, Chile, South Africa, Sweden, and Thailand, and have been growing in popularity as a tool to reduce public spending on education, while guaranteeing equality of access, and providing partial insurance in economic hardship (see Chapman [2006] for an overview).³⁶ The loans sometimes depend on the level of education acquired and the type of degree or field and are indexed to the costs of education, which mirrors the proposed loan above.

Several features of this optimal scheme have counterparts in existing policies. First, in some countries (e.g., Australia and New Zealand), repayments are directly collected through the tax authority and integrated with the tax system. The coercive tax power of the government is required to prevent agents from dropping out after the realization of their incomes, because more successful earners *ex post* cross-subsidize less successful ones. This is made clear by the prominent failure of the so-called “Yale plan,” an attempt at risk pooling within cohorts of students by Yale University in the 1970s. The plan suffered from a typical adverse selection problem: students with the best earnings prospects did not join or dropped out (Palacios 2004). Second, in the optimal scheme, repayments are consolidated repayments for all past loans and need not, in any way, be equal to the total loan amount for a given agent; there is an implicit subsidy or tax fully integrated into the system for insurance and redistribution purposes. Many existing systems also have a substantial insurance component (see, e.g., Dearden et al. 2008), although their repayments are typically more tightly linked to the actual amount borrowed and do not redistribute much. Third, until recently in ICL schemes observed in practice, the focus was almost exclusively on the downside, so that repayments could be deferred or forgiven in times of economic hardship, but repayments were not necessarily increasing after a good stream of earnings. In Australia, however, there is a clear income contingency, with different repayment schemes at different income threshold levels (Chapman 2006). The United Kingdom has been moving closer to the Australian scheme, and while the repayments were previously fixed for any in-

³⁶ In the United States, an important rationale seems to have been the fear that fixed repayment loans would discourage students from careers in the lower-paying public sector (Brody 1994). The first schemes introduced in 1994 were income-contingent repayments for public-sector jobs. In 1997, the College Cost Reduction and Access Act introduced income-based repayment beyond public-sector jobs. Australia is one of the success stories since 1989 with its nationwide scheme (Higher Education Contribution Scheme) that automatically enrolls students in an ICL, with repayments collected directly through the tax system.

come above a threshold, since 2012, the loans are repaid with an income-contingent interest rate.

There are two differences between the optimal scheme proposed and its real-world counterparts. First, loans are optimally made available throughout life—not only for young adults in university—for instance, for expenses related to job training or continuing education.³⁷ Fourth, the optimal repayment schedules depend not only on current income and the outstanding loan balance but rather on the full history of earnings and past loans.³⁸ There is very little history dependence in real-world ICLs, possibly with the exception of Sweden, which uses a 2-year averaging of earnings to determine repayments. Note that other policies exhibit exactly the kind of history dependence that is required here, for instance, social security and some types of tax-free savings accounts. However, the numerical analysis below shows that the gain from history-dependent policies, relative to simpler history-independent (but age-dependent) policies, is not very large for the calibration chosen, implying that history-independent ICLs might be close to optimal provided that age can be conditioned on.

B. Implementation with IID Shocks and Wealth Dependence

A natural question is when the history dependence of the optimal policies proposed above can be reduced. In the special case of iid shocks, wealth and the starting stock of human capital each period can serve as sufficient statistics for the full past. A similar implementation for physical capital, in the absence of human capital, is studied in Albanesi and Sleet (2006).

PROPOSITION 5. If θ is iid, the optimum can be implemented in an economy with borrowing constraints, an initial assignment of wealth, and an income tax schedule $\tilde{T}_t(b_{t-1}, s_{t-1}, y_t, e_t)$ that depends on the beginning-of-period wealth and human capital stocks, as well as on contemporaneous income and human capital investment.

The intuition for this result is that, conditional on human capital s_{t-1} , there is a direct mapping between the social planner's cost of providing the optimal allocation to an agent with promised utility v_{t-1} and the beginning-of-period wealth b_{t-1} of the agent in the decentralized equilibrium.

This recursive implementation allows a relatively simple map between the derivatives of the proposed tax function \tilde{T}_t and the optimal wedges.

³⁷ This discrepancy could, however, be bridged if the cost function M is such that human capital investments occur only early in life.

³⁸ In particular, the sum of past loans is not a sufficient statistic for the full sequence of loans L^{t-1} .

While the labor wedge is equal to the marginal income tax, the human capital wedge is not in general equal to the marginal subsidy, that is, the expected reduction in tax from an incremental investment in human capital. The relation between the human capital wedge and the tax system is

$$\begin{aligned} \tau_{St}^d = & -\left(\frac{\partial \tilde{T}_t}{\partial e_t}\right) + E_t\left(\beta \frac{u'_{t+1}}{u'_t}\right) E_t\left(\frac{\partial \tilde{T}_{t+1}}{\partial e_{t+1}} - \frac{\partial \tilde{T}_{t+1}}{\partial s_t}\right) \\ & + \text{Cov}\left(\beta \frac{u'_{t+1}}{u'_t}, \frac{\partial \tilde{T}_{t+1}}{\partial e_{t+1}} - \frac{\partial \tilde{T}_{t+1}}{\partial s_t}\right), \end{aligned}$$

where recall that

$$\tau_{St}^d \equiv \tau_{St} - \frac{1 - \xi_{\tau_s}}{R(1 - \tau_K)} E_t(\tau_{St+1}).$$

A positive wedge does not necessarily imply a positive marginal subsidy: it can be engineered either directly through expected positive marginal subsidies or, instead, more indirectly through the risk properties of the optimal tax schedule. If the marginal tax reduction from human capital

$$\frac{\partial \tilde{T}_{t+1}}{\partial e_{t+1}} - \frac{\partial \tilde{T}_{t+1}}{\partial s_t}$$

is high when marginal utility of consumption is high, human capital is a good hedge, and the covariance term is positive. It is then possible in theory that the overall wedge is positive even if expected marginal subsidies are zero.³⁹

Note that this tax system can instead be reformulated as a means-tested grant G , such that

$$G_t(y_t, e_t | b_{t-1}, s_{t-1}) = -\tilde{T}_t(b_{t-1}, s_{t-1}, y_t, e_t).$$

Means-tested grants for higher education are very common in many countries, while wealth-contingent income taxes are not. In the United States, for instance, Pell grants take assets as well as contemporaneous income into account.⁴⁰

³⁹ See Kocherlakota (2005) for the case of savings, where the income tax depends on the full history of incomes, and wealth carries no additional information about the past. However, in general, expected marginal subsidies are not zero. Human capital, like wealth in Albanesi and Sleet (2006), carries information value about the past, which restricts the marginal subsidies at the optimum.

⁴⁰ Beyond higher education only, grants for job training exist for the unemployed (hence, somewhat means tested), for youths at risk (YouthBuild in the United States), or for difficult to employ seniors (the Senior Aide Program), most often in the form of a direct provision of training. Some programs do provide funds for training based on need, such as the Adult and Dislocated Worker Program or the Trade Adjustment Assistance.

With iid shocks, the previous implementation through ICLs can also be modified to use wealth instead of the full history of earnings and loans to determine the repayments. The wealth- and income-contingent repayment schedules $D_t(b_{t-1}, s_{t-1}, y_t, e_t)$ are such that

$$\tilde{T}_y(y_t) - L_t(e_t) + D_t(b_{t-1}, s_{t-1}, y_t, e_t) = \tilde{T}_t(b_{t-1}, s_{t-1}, y_t, e_t).$$

C. A Deferred, Risk-Adjusted Human Capital Expense Deductibility Scheme

This implementation directly addresses the debate about whether education expenses should be tax deductible. Boskin (1977) has argued that a true economic depreciation of educational expenses, for which the net present value of the deduction is equal to the expense, would recover neutrality of the tax system with respect to human capital. Even starker is the argument by Bovenberg and Jacobs (2005) that a purely contemporaneous deduction of education expenses from taxable income is sufficient. In this section, I illustrate expense deductibility in the iid case.

Assume that $\rho_{\theta_s} = 1$ so that the optimal net wedge is zero, and the focus is on the human capital subsidy that aims to neutralize the distortionary effects of income and savings taxes. If the proposed tax system $\tilde{T}_t(b_{t-1}, s_{t-1}, y_t, e_t)$ is differentiable in all its arguments (at least over the range of equilibrium path values), then it must satisfy⁴¹

$$\begin{aligned} & -\frac{\partial \tilde{T}_t}{\partial e_t} + \beta E_t \left(\frac{u'_{t+1}}{u'_t} \left(\frac{\partial \tilde{T}_{t+1}}{\partial e_{t+1}} - \frac{\partial \tilde{T}_{t+1}}{\partial s_t} \right) \right) \\ & = \frac{\partial \tilde{T}_t}{\partial y_t} \left[M'_t - \frac{1}{R} E_t(M'_{t+1}) \right] - \frac{1}{R} (1 - \xi'_{M', t+1}) E_t \left(\beta R \frac{u'_{t+1}}{u'_t} \frac{\partial \tilde{T}_{t+1}}{\partial b_t} \right) E_t(M'_{t+1}). \end{aligned} \quad (29)$$

On the other hand, the pure contemporaneous deductibility scheme proposed by Bovenberg and Jacobs (2005) would imply that at all dates

$$\frac{\partial \tilde{T}_t}{\partial y_t} = M'_t(e_t) \frac{\partial \tilde{T}_t}{\partial e_t}$$

for all t .⁴² In this dynamic model, this is true only for the last period T in which agents face a static problem. In all earlier periods, recall the four

⁴¹ Obtained by applying formula (21) to this tax system.

⁴² The discrepancy in this recommendation with the one in Bovenberg and Jacobs (2005) does not come from the restrictive wage function assumed there, because the argument made in this subsection is for the case in which $\rho_{\theta_s} = 1$. Adding back $\rho_{\theta_s} \neq 1$ would push the optimum even further away from pure contemporaneous deductibility, as an additional net encouragement or discouragement of human capital would be desirable. It arises instead from the dynamics and risk.

additional effects discussed in Section IV.C that appeared in the definition of the net wedge t_t (the risk adjustment, the dynamic cost, the savings wedge, and the one-period-ahead subsidy). The policy counterpart to the concept of full risk-adjusted dynamic deductibility introduced in Section IV.C is a risk-adjusted deferred deductibility scheme.

For the sake of the exposition, assume that $\beta = 1/R$ and $M'_t(e_t) = 1$. Start from period T , in which a simple deductibility of expenses is sufficient with $-\partial\tilde{T}_T/\partial e_T = \partial\tilde{T}_T/\partial y_T$, and work backward in order to rewrite the total change in the tax burden from an incremental investment at t as

$$\begin{aligned}
 -\frac{\partial\tilde{T}_t}{\partial e_t} &= (1 - \beta) \sum_{j=1}^{T-t} \beta^{j-1} E_t \left(\frac{u'_{t+j-1}}{u'_t} \frac{\partial\tilde{T}_{t+j-1}}{\partial y_{t+j-1}} \right) \\
 &+ \beta^{T-t} E_t \left(\frac{u'_T}{u'_t} \left(\frac{\partial\tilde{T}_T}{\partial y_T} \right) \right) - \sum_{j=1}^{T-t} \beta^j E_t \left(\frac{u'_{t+j}}{u'_t} \left(\frac{\partial\tilde{T}_{t+j}}{\partial b_{t+j-1}} - \frac{\partial\tilde{T}_{t+j}}{\partial s_{t+j-1}} \right) \right).
 \end{aligned} \tag{30}$$

The optimal subsidy is hence equivalent to a deferred deductibility scheme, in which a fraction $(1 - \beta)$ of the human capital expenses of time t are deducted from taxable income in each subsequent period. Intuitively, with changing income tax rates $\partial\tilde{T}_t/\partial y_t$, a nondynamic deductibility scheme would mean that the expenses of time t would be deducted at time t 's marginal tax rate, but the returns to this investment would accrue in the future when the agent faces potentially different marginal tax rates. If income is growing over time and marginal tax rates are increasing, as in a progressive tax system, there would be insufficient incentives to invest in human capital. A poor student would see little benefit from deducting his tuition fees from his low income, only to pay high marginal tax rates in the future. In addition, there is a "no-arbitrage" term (the last term in [30]), which takes into account the relative shift in the future tax schedule from more savings versus more human capital stock. Since savings and human capital are two ways to transfer resources intertemporally, at the optimum there should be no incentive to substitute from one to the other because of a tax advantage. If full deductibility is not the target (i.e., $\rho_{\theta_s} \neq 1$), implementing the optimal allocation requires adding back the optimal net subsidy each period, on top of this scheme.⁴³

Tax incentives in the form of deduction schemes for higher-education expenses are common but are usually contemporaneous to the expense. In the United States, the American Opportunity Credit and the Lifetime Learning Credit allow families to claim a deduction up to a certain level

⁴³ The linear cost is for exposition only. See formula (A6) for the general case.

per student per year for college, as well as for books, supplies, and required equipment.

The deferred deductibility scheme sets the right incentives in utility terms. In monetary terms, figure 9 illustrates that the scheme is progressive and provides insurance. The figure plots the fraction of the net present value of human capital expenses that the agent cannot deduct against the net present value of lifetime income. Lower-income agents deduct more than they spent on human capital, while higher-income agents deduct less and hence implicitly cross-subsidize lower-income agents.

D. Welfare Gains and Simple Age-Dependent Policies

What are the welfare gains from the optimal mechanism, and how do they compare to the welfare gains from simpler, linear, but age-dependent policies? The first row of table 2 shows the welfare gains from the second-best relative to the *laissez-faire* economy, with no taxes or subsidies, in which agents are unconstrained to borrow and save at the gross interest rate R . Four cases are distinguished, according to the value of the complemen-

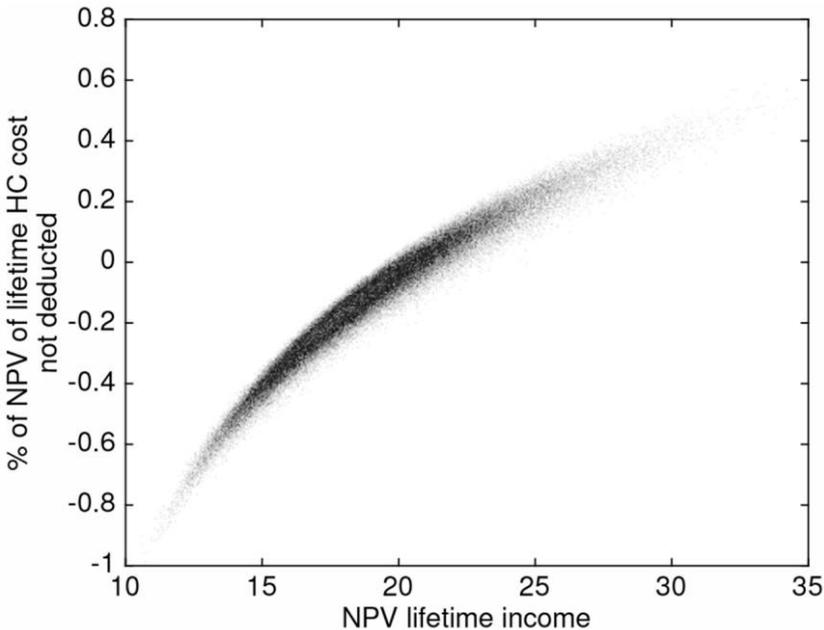


FIG. 9.—Progressivity of the deferred deductibility scheme. The figure shows the fraction of the net present value of human capital cost that agents cannot deduct as a function of the net present value of their lifetime income. Lower-lifetime-income agents deduct more than they spend on human capital, while higher-income agents deduct less.

TABLE 2
WELFARE GAINS FROM SIMPLER POLICIES

	$\rho_{\theta_s} = .2$		$\rho_{\theta_s} = 1.2$	
	Medium Volatility (%)	High Volatility (%)	Medium Volatility (%)	High Volatility (%)
Welfare gain from second best	.85	1.60	.98	1.76
Welfare gain from linear age-dependent policies	.79	1.53	.94	1.74
As % of second best	93	95.6	95.5	98.5

NOTE.—Medium volatility is 0.0095 and high volatility is 0.0161. Row 1 expresses the gain from the second best, relative to the laissez-faire economy, in terms of the equivalent increase in consumption after all histories. Welfare gains are higher when human capital has negative redistributive and insurance values ($\rho_{\theta} > 1$). Row 2 shows the gain from linear age-dependent policies relative to the laissez-faire, while row 3 expresses this gain as a fraction of the gain from the second best. Age-dependent linear policies achieve a very large fraction of the welfare gain from the second best.

tarity coefficient between human capital and ability ρ_{θ_s} , and the volatility of the productivity shock σ_{ψ}^2 .⁴⁴ Welfare gains are expressed as the percentage increase in consumption that, if received every year after all histories, would yield the same gain in lifetime utility.

Given the clear age trends in the above figures and in the optimal formulas, it is natural to compare the full optimum to simple age-dependent policies. The policy under consideration sets the linear human capital subsidy, the linear income tax rate, and the linear capital tax rate at each age equal to their cross-sectional averages at that age. It is numerically challenging to precisely optimize over age-dependent tax rates, given the number of periods and the presence of three instruments; hence, this procedure delivers a lower bound for the welfare gains. It turns out, however, that even this lower bound is very tight. Indeed, the third row in table 2 shows that welfare gains as a fraction of the second-best gains range from 93 percent for a low-volatility and low ρ_{θ_s} case to a surprising 98.5 percent for a high-volatility and high ρ_{θ_s} scenario. This suggests that—for these particular calibrations—the history-dependent policies can be informative about simpler, history-independent policies and that the bulk of the gain comes from the age trend of optimal policies.⁴⁵

These findings are reminiscent of Mirrlees's (1971) conclusion that static optimal income tax schedules appear close to linear. They also closely echo recent findings from the dynamic taxation literature (Farhi

⁴⁴ The high volatility (0.0161) is from Storesletten et al. (2004).

⁴⁵ A word of caution is needed. Given the order of magnitude of 10^{e-5} , it is actually very challenging to estimate these welfare comparisons between the second-best and age-dependent policies precisely, especially over longer horizons. The numbers should be taken as evidence only for small welfare gains, not precise welfare calculations.

and Werning 2013), suggesting that the addition of human capital per se does not change this result.

VII. Conclusion

This paper studies optimal dynamic taxation and human capital policies over the life cycle in a dynamic Mirrlees model with heterogeneous, stochastic, and persistent ability. The government aims to provide redistribution and insurance against adverse draws from the ability distribution. However, the government faces asymmetric information about agents' ability—both its initial level and its stochastic evolution over life—and about labor supply. The constrained efficient allocations were obtained using a dynamic first-order approach and are characterized by wedges or implicit taxes and subsidies. Formulas for the optimal labor and human capital wedges, as well as for their evolution over time, are derived.

A simple relation was derived between the optimal labor and human capital wedges. A crucial consideration for the design of optimal policies is whether human capital has overall positive redistributive and insurance effects. If human capital stimulates labor supply, and hence generates additional resources more than it amplifies existing pretax inequality, it reduces after-tax income inequality on balance. This occurs when the elasticity of the wage with respect to ability is decreasing in human capital or, put differently, when high-ability agents do not disproportionately benefit from human capital. In this case, the optimal net subsidy on human capital expenses is positive and increasing over time. The optimal allocations can be implemented with income-contingent loans, the repayment schedules of which depend on the full history of human capital investments and earnings. If shocks to ability are independently and identically distributed, a deferred deductibility scheme, in which part of current human capital expenses can be deducted from future years' incomes, can also implement the optimum.

The simulations reveal that the optimal net human capital wedges are small, which implies that neutrality of the tax system relative to human capital expenses is close to optimal for the proposed calibrations. In addition, simple age-dependent linear taxes and subsidies can achieve almost the entire welfare gain from the full second-best relative to the *laissez-faire* outcome. Further numerical work could shed light on whether this result remains true with different preferences, in particular with higher risk aversion.

There are three alternative questions for which this analysis can provide some answers. First, should the tax system preserve neutrality with respect to the choice between bequests and human capital spending, two important ways in which parents can transfer resources to their children? The life cycle can be reinterpreted as a dynastic household, in

which parents finance their children's human capital, with persistence in stochastic ability, and partial or full depreciation of human capital across generations, as in Stantcheva (2015b). Second, should productivity-enhancing investments by entrepreneurs or the self-employed be made tax deductible? Workers in the model can instead be viewed as entrepreneurs or self-employed, who can invest in their businesses' productivity through expenses for research, knowledge acquisition, or training of the workforce, generating risky and persistent profits.⁴⁶ If innovation and productivity expenses disproportionately increase risk, they should be made less than fully deductible.⁴⁷ Finally, the analysis could inform the study of optimal policies toward people's investments in health—another type of human capital—which also feature heterogeneity and uncertainty over life.

This theoretical research points to two important empirical explorations that could shed light on the mechanisms behind, and the magnitudes of, the optimal policies. First, how does the complementarity between ability and human capital change over life? In contrast to schooling or higher education, there is little evidence on this for human capital investments later in life, such as job training. Second, it has not been documented entirely yet how strongly people react to current and future expected taxes when making their human capital investment decisions. While challenging, estimating the long-term effects of taxation on human capital accumulation appears very important.

Appendix A

Additional Results

Technical Assumptions

The following assumptions are used as sufficient conditions only to determine the sign of the optimal wedges. All formulas derived still apply without them.⁴⁸

ASSUMPTION 3.

- i. $\int_{\theta}^{\theta'} f'(\theta|\theta_b)d\theta \leq \int_{\theta}^{\theta'} f'(\theta|\theta_s)d\theta$ for all $t, \theta',$ and $\theta_b > \theta_s$;
- ii.

$$\frac{\partial}{\partial \theta_t} \left[\frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial \theta_{t-1}} \frac{1}{f^t(\theta_t|\theta_{t-1})} \right] \geq 0 \quad \forall t, \forall \theta_{t-1};$$

- iii. $\partial v(\theta)/\partial \theta > 0$ for all θ ;
- iv. $(\partial/\partial v)K \geq 0$ and $(\partial^2/\partial v^2)K \geq 0$.

⁴⁶ For a study of optimal R&D policies for firms, see Akcigit, Hanley, and Stantcheva (2016).

⁴⁷ In the sense that the elasticity of business earnings to risk is increasing in innovative activity. In practice, many expenses for the self-employed can be deducted, but there is no special category for innovation or productivity-enhancing expenses.

⁴⁸ In addition, all theoretical results on the signs of the wedges are indeed satisfied in the simulations (Sec. V).

The first-order stochastic dominance assumption in assumption i ensures that, for any given future payoff function increasing in θ , higher types today have a higher expected payoff. Assumption ii introduces a form of monotone likelihood ratio property. Assumption iii guarantees that the expected continuation utility is increasing in the type. Assumption iv states that the resource cost is increasing and convex in promised utility.

PROPOSITION 6.

i. At the optimum, the labor wedge is equal to

$$\frac{\tau_{Lt}^*(\theta^t)}{1 - \tau_{Lt}^*(\theta^t)} = \frac{\mu(\theta^t) u'_t(c(\theta^t)) \varepsilon_{w\theta,t}}{f'(\theta_t|\theta_{t-1})} \frac{1 + \varepsilon_t^u}{\theta_t \varepsilon_t^c}, \tag{A1}$$

with $\mu(\theta^t) = \eta(\theta^t) + \kappa(\theta^t)$ as in (22), where $\eta(\theta^t)$ can be rewritten recursively as a function of the past labor wedge, $\tau_{L,t-1}$:

$$\eta(\theta^t) = \frac{\tau_{L,t-1}^*(\theta^{t-1})}{1 - \tau_{L,t-1}^*(\theta^{t-1})} \left[\frac{R\beta}{u'_{t-1}(c(\theta^{t-1}))} \frac{\varepsilon_{t-1}^c}{1 + \varepsilon_{t-1}^u} \frac{\theta_{t-1}}{\varepsilon_{w\theta,t-1}} \int_{\theta_t}^{\bar{\theta}} \frac{\partial f(\theta_s|\theta_{t-1})}{\partial \theta_{t-1}} d\theta_s \right];$$

ii. $\tau_{Lt}^*(\theta^{t-1}, \bar{\theta}) = \tau_{Lt}^*(\theta^{t-1}, \underline{\theta}) = 0$ for all t .

PROPOSITION 7. At the optimum, the inverse Euler equation holds:

$$\frac{R\beta}{u'_t(c(\theta^t))} = \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{u'_{t+1}(c(\theta^{t+1}))} f^{t+1}(\theta_{t+1}|\theta_t) d\theta_{t+1}. \tag{A2}$$

Note that both Anderberg (2009) and Da Costa and Maestri (2007) find versions of the zero distortion at the top and the inverse Euler equation.

COROLLARY 2. The labor wedge evolves over time according to

$$\begin{aligned} E_{t-1} \left(\frac{\tau_{Lt}}{1 - \tau_{Lt}} \frac{\varepsilon_{w\theta,t-1}}{\varepsilon_{w\theta,t}} \frac{\varepsilon_t^c}{1 + \varepsilon_t^u} \frac{1 + \varepsilon_{t-1}^u}{\varepsilon_{t-1}^c} \left(\frac{1}{R\beta} \frac{u'_{t-1}}{u'_t} \right) \right) \\ = \varepsilon_{w\theta,t-1} \frac{1 + \varepsilon_{t-1}^u}{\varepsilon_{t-1}^c} \text{Cov} \left(\frac{1}{R\beta} \frac{u'_{t-1}}{u'_t}, \log(\theta_t) \right) + p \frac{\tau_{L,t-1}}{1 - \tau_{L,t-1}}. \end{aligned} \tag{A3}$$

Construction of the ICL Schedule

The proof of proposition 4 is in online appendix A. First, the loan is set to exactly cover the cost of human capital:

$$L_t(e_t) = M_t(e_t) \quad \forall t, \forall e_t. \tag{A4}$$

The savings tax $\tilde{T}_k(b_t)$ is constructed to guarantee zero private wealth holdings.⁴⁹ The repayment schedule D and income tax \tilde{T}_y are such that, along the equilib-

⁴⁹ The construction builds on Werning (2011), which shows also that the savings tax can be redefined to implement nonzero savings, at the expense of modifying the repayment schedule. The repayment scheme could also allow for private savings and directly condition on their history (b^{t-1}). See the next implementation proposed, with nonzero private wealth holdings.

rium path, the optimal allocations from the social planner's problem are affordable for each agent after all histories, given zero asset holdings:

$$\begin{aligned} D_t(L^{t-1}, y^{t-1}, e_t^*(\theta^{t-1}, \theta), y_t^*(\theta^{t-1}, \theta)) + \tilde{T}_Y(y_t^*(\theta^{t-1}, \theta)) \\ = y_t^*(\theta^{t-1}, \theta) - c_t^*(\theta^{t-1}, \theta) \end{aligned}$$

for all (L^{t-1}, y^{t-1}) such that

$$\theta^{t-1} \in \Theta^{t-1}(\{M_1^{-1}(L_1), \dots, M_{t-1}^{-1}(L_{t-1})\}, y^{t-1}) \neq \emptyset,$$

and all $\theta \in \Theta$, where the history of education e^{t-1} is inverted from L^{t-1} using (A4). The repayment schedule on off-equilibrium allocations—those allocations that are not optimally assigned to any type in the social planner's program—is set to be sufficiently unattractive, to ensure that agents do not select them. Intuitively then, conditional on entering a period with no savings, and with a given history of loans and output, agents face only the choice of allocations available in the planner's problem after ability histories, which, up to this period, are consistent with the observed choices. By the temporal incentive compatibility of the constrained efficient allocation, they will choose the allocation designed for them.

This set of instruments defines a decentralized allocation rule, which, to an agent with past history (L^{t-1}, y^{t-1}) , assigns an allocation

$$\{\hat{c}_t(L^{t-1}, y^{t-1}, \theta_t), \hat{y}_t(L^{t-1}, y^{t-1}, \theta_t), \hat{b}_t(L^{t-1}, y^{t-1}, \theta_t), \hat{e}_t(L^{t-1}, y^{t-1}, \theta_t)\}_{t, \theta}.$$

The equilibrium allocation as a function of ability histories $\{\hat{c}(\theta^t), \hat{y}_t(\theta^t), \hat{b}(\theta^t), \hat{e}(\theta^t)\}$ can be deduced from the decentralization rule using the recursive relation

$$\hat{m}_t(\theta^t) = \hat{m}_t(L^{t-1}(\theta^{t-1}), y^{t-1}(\theta^{t-1}), \theta_t) \quad \text{for } m \in \{c, y, b, e\},$$

where

$$\theta^{t-1} \in \Theta^{t-1}(\{M_1^{-1}(L_1), \dots, M_{t-1}^{-1}(L_{t-1})\}, y^{t-1})$$

is unique by assumption 2. The decentralization rule is said to implement the optimum from the planner's problem for a given set of promised utilities $(\underline{U}(\theta))_\theta$ if, for all t and θ^t , the decentralized allocations under this rule coincide with the social planner's optimal allocations, that is, $\hat{m}_t(\theta^t) = m_t^*(\theta^t)$ for $m \in \{c, y, b, e\}$.⁵⁰

Implementation with IID Shocks

The recursive problem with iid shocks is nested in the formulation in Section II if the states θ_{t-1} and Δ_{t-1} , which account for persistence, are omitted and the distribution of shocks is $f(\theta)$ each period. Allocations can be expressed as functions of the reduced state space (v_{t-1}, s_{t-1}) for each θ_t , and the government's continuation cost is $K(v_{t-1}, s_{t-1}, t)$.

⁵⁰ In the planner's problem, savings are indeterminate when consumption is controlled, and without loss of generality, agents could be saving zero.

For this implementation, interpret the initial ability θ_1 as uncertainty, like all other shocks θ_t , rather than intrinsic heterogeneity. The government selects an initial promised utility U_1 . All agents again start with the same human capital s_0 and receive an initial wealth level b_0 assigned by the government. The government also sets borrowing limits \underline{b}_t , and wealth is constrained by $b_t \in B_t \equiv [\underline{b}_t, \infty)$. The proposed decentralization rule allocates $(\hat{c}_t, \hat{y}_t, \hat{b}_t, \hat{e}_t)$ to each agent type, following some mappings from observed initial wealth b_{t-1} and human capital s_{t-1} :

$$\hat{c}_t, \hat{y}_t, \hat{e}_t : B_{t-1} \times \mathbb{R}_+ \times \Theta \rightarrow \mathbb{R}_+ \quad \text{and} \quad \hat{b}_t : B_{t-1} \times \mathbb{R}_+ \times \Theta \rightarrow B_t.$$

Starting from an initial wealth level b_0 , the recursive decentralization rule can be mapped into a sequential allocation for all θ^t .

Let $V_t(b, s)$ denote the value of an agent with beginning-of-period wealth b and human capital s . A decentralization rule $(\hat{c}_t, \hat{y}_t, \hat{b}_t, \hat{e}_t)$, an initial assignment of wealth b_0 , a sequence of borrowing limits $\{\underline{b}_t\}_{t=1}^T$, and initial human capital level s_0 form a decentralized equilibrium if, in all periods, $(\hat{c}_t, \hat{y}_t, \hat{b}_t, \hat{e}_t)$ attains the supremum in the agent's problem in (A5):

$$V_t(b, s) = \sup_{c, y, b', e} \int \left[u_t(c(\theta)) - \phi_t \left(\frac{y(\theta)}{w_t(\theta, s + e(\theta))} \right) + \beta V_{t+1}(b'(\theta), s + e(\theta)) \right] f(\theta) d\theta \tag{A5}$$

subject to

$$c(\theta) + M_t(e(\theta)) + \frac{1}{R} b'(\theta) = y(\theta) - \tilde{T}_t(b, s, y(\theta), e(\theta)) + b \quad \forall \theta,$$

$$c, y, e : \Theta \rightarrow \mathbb{R}_+ \text{ and } b' : \Theta \rightarrow B_t = [\underline{b}_t, \infty) \text{ with } V_{T+1} \equiv 0, b_T \equiv 0.$$

A constrained efficient allocation from the planner's problem is implemented as a decentralized equilibrium if it arises as an equilibrium choice of agents in the above problem and delivers expected lifetime utility $V(b_0, s_0) = U_1$.

The link between the human capital wedge and the explicit tax system is

$$\tau_{S_t}(v_{t-1}, s_{t-1}, \theta_t) = \frac{-\partial \tilde{T}_t(K(v_{t-1}, s_{t-1}, t), s_{t-1}, y_t^*, e_t^*)}{\partial e_t} + \beta E_t \left(\frac{u_{t+1}(c_{t+1}^*)}{u_t'(c_t^*)} \left(\frac{\partial \tilde{T}_{t+1}}{\partial e_{t+1}} - \frac{\partial \tilde{T}_{t+1}}{\partial s_t} \right) \right),$$

where $c_t^* = c_t^*(v_{t-1}, s_{t-1}, \theta_t)$, $c_{t+1}^* = c_{t+1}^*(v_t^*, s_{t-1} + e_t^*, \theta_{t+1})$, and $\tilde{T}_{t+1} = \tilde{T}_{t+1}(b_t, s_t, y_{t+1}, e_{t+1})$ is evaluated at

$$v_t^* = v_t^*(v_{t-1}, s_{t-1}, \theta_t), \quad b_t = K(v_t^*, s_{t-1} + e_t^*, t + 1), \quad s_t = s_{t-1} + e_t^*, \\ y_{t+1} = y_{t+1}^*(v_t^*, s_{t-1} + e_t^*, \theta_{t+1}), \quad e_{t+1} = e_{t+1}^*(v_t^*, s_{t-1} + e_t^*, \theta_{t+1}).$$

General Deductibility Scheme

With nonlinear cost, the general expression for the deferred deductibility scheme is

$$\begin{aligned}
 -\frac{\partial \tilde{T}_t}{\partial e_t} &= \sum_{j=1}^{T-t} \beta^{j-1} E_t \left(\frac{u'_{t+j-1}}{u'_t} \frac{\partial \tilde{T}_{t+j-1}}{\partial y_{t+j-1}} \left(M'_{t+j-1} - \frac{1}{R} M'_{t+j} \right) \right) + \beta^{T-t} E_t \left(\frac{u'_T}{u'_t} \left(\frac{\partial \tilde{T}_T}{\partial y_T} \right) \right) \\
 &\quad - \sum_{j=1}^{T-t} \beta^j E_t \left(\frac{u'_{t+j}}{u'_t} \left[\left(1 - \xi'_{M',t+j} \right) E_{t+j-1} \left(M'_{t+j} \right) \frac{\partial \tilde{T}_{t+j}}{\partial b_{t+j-1}} - \frac{\partial \tilde{T}_{t+j}}{\partial s_{t+j-1}} \right] \right). \tag{A6}
 \end{aligned}$$

The first set of terms capture the deferred deductibility from the income tax base. Because the marginal cost is no longer constant, the deduction in period $t + j$ occurs at the dynamic marginal cost effective in that period $(M'_{t+j} - (1/R)M'_{t+j+1})$, not at the “historic” marginal cost faced by the agent at the time of the purchase M'_t ; that is, a purchase of Δe at time t is deducted as $[M'_{t+j} - (1/R)M'_{t+j+1}]\Delta e$ from y_{t+j} at $t + j$. Otherwise, there would be arbitrage possibilities.⁵¹ Similarly to the text, the “no-arbitrage” term takes into account the differential tax increases from physical capital versus human capital, except that now the nonlinear, risk-adjusted cost $(1 - \xi'_{M',t+j})E_{t+j-1}(M'_{t+j})$ enters the picture. The deduction is risk adjusted, as witnessed by the insurance factors, $\xi'_{M',t+1}$.

Optimal Human Capital Subsidy at a Given Labor Wedge

Suppose there is a linear tax on earnings τ_L . Define the consumption function $\tilde{c}(\theta) = c(\beta v(\theta) - \omega(\theta), \theta, l(\theta))$. The objective is now

$$\begin{aligned}
 K(v_-, \Delta_-, s_-, \theta_-, t) &= \min \left\{ \tilde{c}(\beta v(\theta) - \omega(\theta), \theta, l(\theta)) + w_l l(\theta) + M_t(s(\theta) - s_-) \right. \\
 &\quad \left. + \frac{1}{R} K(v(\theta), \Delta(\theta), s(\theta), \theta, t + 1) \right\}.
 \end{aligned}$$

The envelope condition is unchanged. Note that the first-order condition of the agent with respect to labor needs to hold as an additional constraint in the problem (since labor can now no longer be chosen directly):

$$w_t u'_l(\tilde{c}(\theta))(1 - \tau_L) = \phi_{l,t}(l(\theta)).$$

From this equation,

$$w_{s,t} u'_s(\tilde{c}(\theta))(1 - \tau_L) ds(\theta) + w_l u'_l(\tilde{c}(\theta))(1 - \tau_L) \frac{d\tilde{c}(\theta)}{dl(\theta)} dl(\theta) = \phi_{u,t}(l(\theta)) dl(\theta),$$

and hence

$$\frac{dl}{ds} = \frac{w_{s,t} u'_s(\tilde{c}(\theta))(1 - \tau_L)}{\phi_{u,t}(l(\theta)) - w_l u'_l(\tilde{c}(\theta))(1 - \tau_L) \frac{\phi_{l,t}(l(\theta))}{u'_l(\tilde{c}(\theta))}}. \tag{A7}$$

⁵¹ Note that with linear cost, as in the main text, this is just $1 - \beta$ with $\beta = 1/R$ for all $t < T$ and 1 for $t = T$.

The first-order condition with respect to education now yields

$$t_{st} = \frac{\mu_t}{f'(\theta_t|\theta_{t-1})} \frac{w_{\theta,t}}{w_t^2} \phi_{l,t}(l(\theta))(1 - \rho_{\theta,s,t}) - \left[\frac{1}{w_{s,t}l(\theta)} \tau_L w_t - \frac{\mu_t}{f'(\theta_t|\theta_{t-1})} \frac{w_{\theta,t}}{w_t^2} \phi_{l,t}(l(\theta)) \frac{w_t}{w_{s,t}l(\theta)} \left(\frac{1 + \varepsilon^u}{\varepsilon^c} \right) \right] \frac{dl}{ds}, \tag{A8}$$

with dl/ds as given by (A7). Of course, from the term multiplying dl/ds , we see that if the labor wedge is optimally set as in (A1), then the net human capital wedge is equal to the one in the text.

Appendix B

Derivations and Proofs

Additional proofs are in the online appendix.

Proof of Propositions 1, 3, and 6

The expenditure function $\tilde{c}(l, \omega - \beta v, \theta)$ defines consumption indirectly as a function of labor l , current-period utility ($\tilde{u} = \omega - \beta v$), and the current realization of the type (note that conditional on these variables, consumption does not depend on human capital s). Then $\omega(\theta) = u_t(c(\theta)) - \phi_t(l(\theta)) + \beta v(\theta)$ becomes redundant as a constraint, and the choice variables are $(l(\theta), s(\theta), \omega(\theta), v(\theta), \Delta(\theta))$. Let the multipliers in program (12) be (in the order of the constraints there) $\mu(\theta)$, λ_- , and γ_- . The problem is solved using the optimal control approach in which the “types” play the role of the running variable, $\omega(\theta)$ is the state (and $\dot{\omega}(\theta)$ its law of motion), and the controls are $l(\theta), v(\theta), s(\theta)$, and $\Delta(\theta)$. The Hamiltonian is

$$\begin{aligned} & (\tilde{c}(l(\theta), \omega(\theta) - \beta v(\theta), \theta) + M_t(s(\theta) - s_-) - w_t(\theta, s(\theta))l(\theta))f'(\theta|\theta_-) \\ & + \frac{1}{R} K(v(\theta), \Delta(\theta), \theta, s(\theta), t + 1)f'(\theta|\theta_-) + \lambda_- [v - \omega(\theta)f'(\theta|\theta_-)] \\ & + \gamma_- \left[\Delta - \omega(\theta) \frac{\partial f'(\theta|\theta_-)}{\partial \theta_-} \right] + \mu(\theta) \left[\frac{w_{\theta,t}}{w_t} l(\theta)\phi_{l,t}(l(\theta)) + \beta \Delta(\theta) \right] \end{aligned}$$

with boundary conditions

$$\lim_{\theta \rightarrow \underline{\theta}} \mu(\theta) = \lim_{\theta \rightarrow \underline{\theta}} \mu(\theta) = 0.$$

Taking the first-order conditions (FOCs) of the recursive planning problem yields (the variable with respect to which the FOC is taken appears in brackets)

$$[l(\theta)]: \frac{\tau_L^*(\theta)}{1 - \tau_L^*(\theta)} = \frac{\mu(\theta)}{f'(\theta|\theta_-)} \frac{w_{\theta,t}}{w_t} u'_t(c(\theta)) \left[1 + \frac{l(\theta)\phi_{ll,t}(l(\theta))}{\phi_{l,t}(l(\theta))} \right]$$

using the definitions of ε^c , ε^u , and ε in the text:

$$\frac{\tau_{L,t}^*(\theta)}{1 - \tau_{L,t}^*(\theta)} = \frac{\mu(\theta)}{f'(\theta|\theta_-)} \frac{\varepsilon_{w\theta}}{\theta} u'_t(c(\theta)) \frac{1 + \varepsilon^u}{\varepsilon^c},$$

$$[s(\theta)] : -M'_t(s(\theta) - s_-) + l(\theta)w_{s,t} - \frac{1}{R} \frac{\partial K(v(\theta), \Delta(\theta), \theta, s(\theta), t + 1)}{\partial s(\theta)}$$

$$= \frac{\mu(\theta)}{f'(\theta|\theta_-)} l(\theta)\phi_{l,t}(l(\theta)) \frac{1}{w_t^2} w_{\theta,t}w_{s,t}(\rho_{\theta s,t} - 1),$$

where (letting θ' and s' be the next period's type and human capital, respectively)

$$\frac{\partial K}{\partial s(\theta)} = - \int M'_{t+1}(s'(\theta') - s(\theta))f^{t+1}(\theta'|\theta)d\theta'$$

so that

$$-M'_t(s(\theta) - s_-) + l(\theta)w_{s,t} + \frac{1}{R} \int M'_{t+1}(s'(\theta') - s(\theta))f^{t+1}(\theta'|\theta)d\theta'$$

$$= \frac{\mu(\theta)}{f'(\theta|\theta_-)} l(\theta)\phi_{l,t}(l(\theta)) \frac{1}{w_t^2} w_{\theta,t}w_{s,t}(\rho_{\theta s,t} - 1).$$

Use the expression for $w_{s,t}l_t$ from the definition of the human capital wedge τ_{st} in (15) to write the FOC as a function of the modified wedge t_{st} :

$$t_{st}^*(\theta^t) = \frac{\mu(\theta^t)}{f'(\theta_t|\theta_{t-1})} u'(c_t(\theta^t)) \frac{\varepsilon_{w\theta,t}}{\theta_t} (1 - \rho_{\theta s,t}).$$

From this, we can immediately deduce the relation between the modified wedge and the tax rate in the text:

$$t_{st} = \frac{\tau_{L,t}}{1 - \tau_{L,t}} (1 - \rho_{\theta s,t}) \frac{\varepsilon_t^c}{1 + \varepsilon_t^u}.$$

The law of motion for the costate $\mu(\theta)$ comes from the FOC with respect to the state variable $\omega(\theta)$:

$$[\omega(\theta)] : \left[-\frac{1}{u'_t(c(\theta))} + (\lambda_-) + (\gamma_-) \frac{\partial f'(\theta|\theta_-)}{\partial \theta_-} \frac{1}{f'(\theta|\theta_-)} \right] f'(\theta|\theta_-) = \dot{\mu}(\theta). \quad (B1)$$

Integrating this and using the boundary condition $\mu(\bar{\theta}) = 0$ yields

$$\mu(\theta) = \int_{\theta}^{\bar{\theta}} \left[\frac{1}{u'_t(c(\theta))} - (\lambda_-) - (\gamma_-) \frac{\partial f'(\theta|\theta_-)}{\partial \theta_-} \frac{1}{f'(\theta|\theta_-)} \right] f'(\theta|\theta_-). \quad (B2)$$

Integrating and using both boundary conditions yields

$$\lambda_- = \int_{\theta}^{\bar{\theta}} \frac{1}{u'_t(c(\theta))} f'(\theta|\theta_-) d\theta. \quad (B3)$$

Using the envelope conditions

$$\frac{\partial K(v(\theta), \Delta(\theta), \theta, s(\theta), t + 1)}{\partial v(\theta)} = \lambda(\theta)$$

and

$$\frac{\partial K(v(\theta), \Delta(\theta), \theta, s(\theta), t + 1)}{\partial \Delta(\theta)} = \gamma(\theta),$$

the FOCs with respect to $v(\theta)$ and $\Delta(\theta)$, respectively, lead to

$$[v(\theta)] : \frac{1}{u'(c)} = \frac{\lambda(\theta)}{R\beta} \tag{B4}$$

and

$$[\Delta(\theta)] : -\frac{\gamma(\theta)}{R\beta} = \frac{\mu(\theta)}{f'(\theta|\theta_-)}. \tag{B5}$$

Using (B3) and (B5) in the expression for $\mu(\theta)$ from (B2) yields $\mu(\theta^t) = \kappa(\theta^t) + \eta(\theta^t)$, where

$$\begin{aligned} \kappa(\theta^t) &= \int_{\theta}^{\theta} \frac{1}{u'(c(\theta))} \left[1 - u'(c(\theta)) \int_{\theta}^{\theta} \frac{1}{u'(c(m))} f(m|\theta_-) dm \right] f'(\theta|\theta_-), \\ \eta(\theta^t) &= -(\gamma_-) \int_{\theta}^{\theta} \frac{\partial f'(\theta|\theta_-)}{\partial \theta_-} d\theta = \frac{\mu(\theta^{t-1})}{f(\theta_{t-1}|\theta_{t-2})} R\beta \int_{\theta}^{\theta} \frac{\partial f'(\theta|\theta_-)}{\partial \theta_-} d\theta, \end{aligned}$$

where the last equality uses the lag of (B5). The multiplier is replaced by the last period's t_{st-1}^* (respectively, τ_{Lt-1}^*) using the optimal formulas to obtain the expressions in proposition 3 (respectively, [6]).

The zero net wedge result at the top and bottom follows immediately from the boundary conditions $\mu(\underline{\theta}) = \mu(\bar{\theta}) = 0$.

Part ii of proposition 3: If θ is iid, $\gamma_- = 0$ and $\eta(\theta^t) = 0$ for all t . In addition, if $u'_i(c_i) = 1$ for all t , then $\kappa(\theta^t) = 0$ as well.

Proof of Proposition 2

From the expression of t , the proof is immediate by inspection as long as $\mu(\theta^t) \geq 0$ for all t , for all θ^t , which is now proved.

LEMMA 1. Under assumption 3, $\mu(\theta^t) \geq 0$ for all t , for all θ^t .

Proof. The proof is close to the one in Golosov et al. (2013), for a separable utility function and with human capital. From the envelope condition and the FOC for $v(\theta)$ in (B4),

$$\frac{\partial K}{\partial v} = \lambda(\theta) = \frac{R\beta}{u'_i(c(\theta))}.$$

Since by assumption $v(\theta)$ is increasing in θ and $K()$ is increasing and convex in v , it must be that $\partial K/\partial v$ is increasing in θ , so that $1/u'_i(c(\theta))$ as well is increasing in θ .

Start in period $t = 1$. In this case, since θ_0 has a degenerate distribution, $(\partial f^1/\partial \theta_0)(\theta_1|\theta_0) = 0$ and

$$\mu(\theta_1) = \int_{\theta_1}^{\theta} \left[\frac{1}{u'_i(c(\bar{\theta}_1))} - \lambda_- \right] f^1(\bar{\theta}_1) d\bar{\theta}_1.$$

Choose the θ' such that $1/u'(c(\theta')) = \lambda_-$. Since $1/u'(c(\theta))$ is increasing in θ , for $\theta \geq \theta'$, $\mu(\theta) \geq 0$ (integrating over nonnegative numbers only). Using the boundary condition $\mu(\underline{\theta}) = 0$, $\mu(\theta_1)$ can also be rewritten as

$$\mu(\theta_1) = \int_{\underline{\theta}}^{\theta_1} \left[-\frac{1}{u'_1(c(\tilde{\theta}_1))} + \lambda_- \right] f^1(\tilde{\theta}_1) d\tilde{\theta}_1.$$

Since for $\theta_1 \leq \theta'$, $1/u'(c(\theta_1)) \leq \lambda_-$, we again have $\mu(\theta_1) \geq 0$. Thus, for all θ_1 , $\mu(\theta_1) \geq 0$. By the FOC for Δ in (B5),

$$-\frac{\gamma(\theta_1)}{R\beta} = \frac{\mu(\theta_1)}{f(\theta_1)}$$

so that $\gamma(\theta_1) \leq 0$, for all θ_1 . Note that $\mu(\theta_2)$ is equal to

$$\mu(\theta_2) = \int_{\theta_2}^{\bar{\theta}} \left[\frac{1}{u'_2(c(\tilde{\theta}_2))} - (\lambda_-) - \frac{\partial f(\tilde{\theta}_2|\theta_1)}{\partial \theta_1} \frac{\gamma_-}{f(\tilde{\theta}_2|\theta_1)} \right] f(\tilde{\theta}_2|\theta_1).$$

Since by assumption 3(iii) $g(\tilde{\theta}_2|\theta_1)/f(\tilde{\theta}_2|\theta_1)$ is increasing in $\tilde{\theta}_2$ and we already showed that $1/u'(c(\tilde{\theta}_2))$ is increasing in θ_2 , there is a θ'_2 such that

$$\frac{1}{u'_2(c(\theta'_2))} - \frac{\partial f(\theta'_2|\theta_1)}{\partial \theta_1} \frac{\gamma_-}{f(\theta'_2|\theta_1)} = \lambda_-$$

and such that for $\theta_2 \geq \theta'_2$, $\mu(\theta^2) \geq 0$ (since integrating over nonnegative numbers only). Rewriting $\mu(\theta^2)$ as an integral from $\underline{\theta}$ to θ_2 and using the boundary condition $\underline{\theta} = 0$, we can again show that $\mu(\theta_2) \geq 0$ also for $\theta_2 \leq \theta'_2$. Proceeding in the same way for all periods up to T shows the result.

Proof That $t_{st} = 0$ When $\rho_{\theta_s} = 1$ Without Using the First-Order Approach

Consider a separable wage $w = \theta s$, a history θ^t , and a perturbation of the allocation for all $\tilde{\theta}^t$ in a neighborhood of θ^t , $|\theta^t - \tilde{\theta}^t| \leq \eta$ such that $s^{\delta}(\tilde{\theta}^t) = s(\tilde{\theta}^t) + \delta$ and $y^{\delta}(\tilde{\theta}^t) = y(\tilde{\theta}^t) + dy(\tilde{\theta}^t)$ such that $s^{\delta}(\tilde{\theta}^t)/y^{\delta}(\tilde{\theta}^t) = s(\tilde{\theta}^t)/y(\tilde{\theta}^t)$, that is, $dy(\tilde{\theta}^t) = \delta[y(\tilde{\theta}^t)/s(\tilde{\theta}^t)]$. This perturbation leaves utilities and incentive compatibility constraints unaffected. The change in the resource cost must hence be zero in this neighborhood, and letting $\eta \rightarrow 0$, we obtain

$$-\frac{y(\theta)}{s(\theta)} + \left[M'_t(e(\theta)) - \frac{1}{R} E_t(M'_{t+1}(e(\theta'))) \right] = 0,$$

which is equivalent to $t_{st} = 0$ for the multiplicative wage.

Proof of Proposition 7

Taking the integral of $\dot{\mu}(\theta)$ in equation (B1) between the two boundaries, $\bar{\theta}$ and $\underline{\theta}$, and using the boundary conditions $\mu(\bar{\theta}) = \mu(\underline{\theta}) = 0$, as well as the expression for λ_- from (B4), lagged by one period, yields the inverse Euler equation in (A2).

Proof of Corollary 2

The derivation of the time evolution of the labor wedge follows Farhi and Werning (2013). Take any weighting function $\pi(\theta) > 0$ and let $\Pi(\theta)$ denote a primitive of $\pi(\theta)/\theta$. Starting from the expression of the optimal labor wedge in (A1), multiply both sides of the expression by $\pi(\theta) > 0$. Integrating by parts yields

$$\begin{aligned} & \int \frac{\tau_{Lt}(\theta^t)}{1 - \tau_{Lt}(\theta^t)} f^t(\theta_t|\theta_{t-1}) \frac{\theta_t}{\varepsilon_{w\theta,t}} \frac{\varepsilon_t^c}{1 + \varepsilon_t^u} \frac{1}{u_t'(c(\theta^t))} \frac{\pi(\theta_t)}{\theta_t} d\theta_t = \int \mu(\theta^t) \frac{\pi(\theta_t)}{\theta_t} d\theta_t \\ & = - \int \dot{\mu}(\theta^t) \Pi(\theta_t) d\theta_t \\ & = \int \left[\frac{1}{u_t'(c(\theta^t))} - \frac{R\beta}{u_{t-1}'(c(\theta^{t-1}))} + \frac{R\beta\mu(\theta^{t-1})}{f(\theta_{t-1}|\theta_{t-2})} \frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial \theta_{t-1}} \frac{1}{f^t(\theta_t|\theta_{t-1})} \right] f^t(\theta_t|\theta_{t-1}) \Pi(\theta_t) d\theta_t \\ & = \int \left[\frac{1}{u_t'} - \frac{R\beta}{u_{t-1}'} + \frac{R\beta\tau_{Lt-1}}{1 - \tau_{Lt-1}} \frac{1}{u_{t-1}'\varepsilon_{w\theta,t-1}} \frac{\theta_{t-1}}{1 + \varepsilon_{t-1}^u} \frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial \theta_{t-1}} \frac{1}{f^t(\theta_t|\theta_{t-1})} \right] f^t(\theta_t|\theta_{t-1}) \Pi(\theta_t) d\theta_t d\theta_{t-1}, \end{aligned}$$

where the third line uses the expression for λ_- and γ_- from, respectively, (B4) and (B5) evaluated at $t - 1$. The fourth line uses the optimal wedge from (A1) at time $t - 1$ to substitute for the multiplier $\mu(\theta^{t-1})$. Using the inverse Euler equation yields a general formula for any stochastic process and weighting function:

$$\begin{aligned} & E_{t-1} \left(\frac{\tau_{Lt}(\theta^t)}{1 - \tau_{Lt}(\theta^t)} \frac{\theta_t}{\varepsilon_{w\theta,t}} \frac{\varepsilon_t^c}{1 + \varepsilon_t^u} \frac{u_{t-1}'(c(\theta^{t-1}))}{u_t'(c(\theta^t))} \frac{\pi(\theta_t)}{\theta_t} \right) \\ & = \text{Cov} \left(\frac{u_{t-1}'(c(\theta^{t-1}))}{u_t'(c(\theta^t))}, \Pi(\theta_t) \right) + \frac{R\beta\tau_{Lt-1}}{1 - \tau_{Lt-1}} \frac{\theta_{t-1}}{\varepsilon_{w\theta,t-1}} \frac{\varepsilon_{t-1}^c}{1 + \varepsilon_{t-1}^u} \int \frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial \theta_{t-1}} \Pi(\theta_t) d\theta_t. \end{aligned} \tag{B6}$$

For the particular weighting function $\pi(\theta_t) = 1$ (with $\Pi(\theta_t) = \log(\theta_t)$), and with the AR(1) process assumed for $\log(\theta_t)$, the formula becomes as in (A3).

Proof of Corollary 1

The net wedge on human capital can be rewritten similarly as in the proof of proposition 2, using the same weighting function $\pi(\theta) = 1$:

$$\begin{aligned} & \int \frac{t_{st}}{u_t'(c_t(\theta^t))\varepsilon_{w\theta,t}} \frac{1}{1 - \rho_{\theta_s,t}} f^t(\theta_t|\theta_{t-1}) \\ & = \int \left[\frac{1}{u_t'(c_t(\theta^t))} - \frac{R\beta}{u_{t-1}'(c(\theta^{t-1}))} \right] f^t(\theta_t|\theta_{t-1}) \log(\theta_t) d\theta_t \\ & \quad + R\beta \frac{\mu(\theta^{t-1})}{f(\theta_{t-1}|\theta_{t-2})} \int \frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial \theta_{t-1}} \log(\theta_t) d\theta_t, \end{aligned}$$

where the second line uses the expression for λ_- and γ_- from, respectively, (B4) and (B5) evaluated at $t - 1$. Using the optimal wedge from (21) at time $t - 1$ to

substitute for the multiplier $\mu(\theta^{-1})$, the boundary conditions for μ (and the resulting inverse Euler equation), and the log AR(1) process for θ_t , formula (26) in the text is obtained.

References

- Akcigit, Ufuk, Douglas Hanley, and Stefanie Stantcheva. 2016. "Optimal Taxation and R&D Policies." Working paper, Harvard Univ.
- Albanesi, Stefania, and Christopher Sleet. 2006. "Dynamic Optimal Taxation with Private Information." *Rev. Econ. Studies* 73 (1): 1–30.
- Anderberg, D. 2009. "Optimal Policy and the Risk Properties of Human Capital Reconsidered." *J. Public Econ.* 93 (9–10): 1017–26.
- Ashenfelter, Orley, and Cecilia Rouse. 1998. "Income, Schooling, and Ability: Evidence from a New Sample of Identical Twins." *Q.J.E.* 113 (1): 253–84.
- Atkinson, Anthony Barnes, and Joseph E. Stiglitz. 1976. "The Design of Tax Structure: Direct versus Indirect Taxation." *J. Public Econ.* 6 (1): 55–75.
- Autor, David, Lawrence Katz, and Alan Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *Q.J.E.* 113 (4): 1169–1213.
- Battaglini, Marco, and Stephen Coate. 2008. "Pareto Efficient Income Taxation with Stochastic Abilities." *J. Public Econ.* 92 (3): 844–68.
- Becker, Gary S. 1964. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. Chicago: Univ. Chicago Press.
- Ben-Porath, Yoram. 1967. "The Production of Human Capital and the Life Cycle of Earnings." *J.P.E.* 75 (4): 352–65.
- Bohacek, Radim, and Marek Kapička. 2008. "Optimal Human Capital Policies." *J. Monetary Econ.* 55 (1): 1–16.
- Boskin, Michael J. 1977. "Notes on the Tax Treatment of Human Capital." In *Proceedings of the Treasury Conference on Tax Policy*. Washington, DC: US Dept. Treasury.
- Bovenberg, Lans A., and Bas Jacobs. 2005. "Redistribution and Education Subsidies Are Siamese Twins." *J. Public Econ.* 89 (11–12): 2005–35.
- . 2011. "Optimal Taxation of Human Capital and the Earnings Function." *J. Public Econ. Theory* 13 (6): 957–71.
- Brody, Evelyn. 1994. "Paying Back Your Country through Income-Contingent Student Loans." *San Diego Law Rev.* 31:449–518.
- Card, David. 1995. "Earnings, Schooling, and Ability Revisited." *Res. Labor Econ.* 35:111–36.
- Chapman, Bruce. 2006. "Income Contingent Loans for Higher Education: International Reforms." In *Handbook of the Economics of Education*, vol. 2, edited by Eric A. Hanushek and Finis Welch, 1437–1503. Amsterdam: Elsevier.
- Chetty, Raj. 2012. "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply." *Econometrica* 80 (3): 969–1018.
- Cunha, Flavio, and James J. Heckman. 2007. "Identifying and Estimating the Distributions of Ex Post and Ex Ante Returns to Schooling: A Survey of Recent Developments." *Labour Econ.* 43 (4): 738–82.
- Cunha, Flavio, James J. Heckman, Lance Lochner, and Dimitriy V. Masterov. 2006. "Interpreting the Evidence on Life Cycle Skill Formation." In *Handbook of the Economics of Education*, vol. 1, edited by Eric A. Hanushek and Finis Welch, 697–812. Amsterdam: Elsevier.

- Da Costa, Carlos E., and Lucas J. Maestri. 2007. "The Risk Properties of Human Capital and the Design of Government Policies." *European Econ. Rev.* 51 (3): 695–713.
- Dearden, Lorraine, Alissa Goodman, Emla Fitzsimons, and Greg Kaplan. 2008. "Higher Education Funding Reforms in England: The Distributional Effects and the Shifting Balance of Costs." *Econ. J.* 118:100–125.
- Doepke, Matthias, and Robert M. Townsend. 2006. "Dynamic Mechanism Design with Hidden Income and Hidden Actions." *J. Econ. Theory* 126 (1): 235–85.
- Farhi, Emmanuel, and Iván Werning. 2013. "Insurance and Taxation over the Life Cycle." *Rev. Econ. Studies* 80 (2): 596–635.
- Fernandes, Ana, and Christopher Phelan. 2000. "A Recursive Formulation for Repeated Agency with History Dependence." *J. Econ. Theory* 19 (2): 223–47.
- Findeisen, Sebastian, and Dominik Sachs. 2012. "Education and Optimal Dynamic Taxation: The Role of Income-Contingent Student Loans." Technical report, Dept. Econ., Univ. Zurich.
- Goldin, Claudia Dale, and Lawrence F. Katz. 2008. *The Race between Education and Technology*. Cambridge, MA: Harvard Univ. Press.
- Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski. 2013. "Redistribution and Social Insurance." Working Paper no. 17642, NBER, Cambridge, MA.
- Golosov, Mikhail, Aleh Tsyvinski, and Iván Werning. 2006. "New Dynamic Public Finance: A User's Guide." In *NBER Macroeconomics Annual*, vol. 21, edited by Kenneth Rogoff, Daron Acemoglu, and Michael Woodford, 317–88. Cambridge, MA: MIT Press.
- Grochulski, Borys, and Tomasz Piskorski. 2010. "Risky Human Capital and Deferred Capital Income Taxation." *J. Econ. Theory* 145 (3): 908–43.
- Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006." *Rev. Econ. Dynamics* 13 (1): 15–51.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2005. "Two Views of Inequality over the Life Cycle." *J. European Econ. Assoc.* 3 (2–3): 765–75.
- Heckman, James J. 1976. "A Life-Cycle Model of Earnings, Learning, and Consumption." *J.P.E.* 84, no. 4, pt. 2 (August): S9–S44.
- Hicks, John. 1970. "Elasticity of Substitution Again: Substitutes and Complements." *Oxford Econ. Papers* 25:289–96.
- Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of Lifetime Inequality." *A.E.R.* 101 (7): 2923–54.
- Kapička, Marek. 2013a. "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach." *Rev. Econ. Studies* 80 (3): 1027–54.
- . 2013b. "Optimal Mirrleesian Taxation with Unobservable Human Capital Formation." Technical report, Univ. California Santa Barbara.
- Kapička, Marek, and Julian Neira. 2014. "Optimal Taxation in a Life-Cycle Economy with Endogenous Human Capital Formation." Technical report, Univ. California Santa Barbara.
- Kocherlakota, Narayana R. 2005. "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation." *Econometrica* 73 (5): 1587–1621.
- Kremer, Michael. 2002. "Should Taxes Be Independent of Age?" Manuscript, Harvard Univ.
- Krueger, Dirk, and Alexander Ludwig. 2013. "Optimal Progressive Labor Income Taxation and Education Subsidies When Education Decisions and In-

- tergenerational Transfers Are Endogenous." *A.E.R. Papers and Proc.* 103 (3): 496–501.
- Laroque, Guy R. 2005. "Indirect Taxation Is Superfluous under Separability and Taste Homogeneity: A Simple Proof." *Econ. Letters* 87 (1): 141–44.
- Lochner, Lance J., and Alexander Monge-Naranjo. 2011. "The Nature of Credit Constraints and Human Capital." *A.E.R.* 101 (6): 2487–2529.
- Maldonado, Darío. 2008. "Education Policies and Optimal Taxation." *Internat. Tax and Public Finance* 15 (2): 131–43.
- Milgrom, Paul, and Ilya Segal. 2002. "Envelope Theorems for Arbitrary Choice Sets." *Econometrica* 70 (2): 583–601.
- Mirrlees, James. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Rev. Econ. Studies* 38:175–208.
- Mirrlees, James, S. Adam, T. Besley, et al. 2011. *Tax by Design: The Mirrlees Review*. Oxford: Oxford Univ. Press.
- Murphy, Kevin M., and Finis Welch. 1992. "The Structure of Wages." *Q.J.E.* 107:285–326.
- OECD. 2004. "Improving Skills for More and Better Jobs: Does Training Make a Difference?" Technical report, OECD Employment Outlook, Paris.
- . 2013. "Education at a Glance." Technical report, OECD Employment Outlook, Paris.
- Palacios, M. 2004. *Investing in Human Capital: A Capital Markets Approach to Higher Education Funding*. Cambridge: Cambridge Univ. Press.
- Pavan, Alessandro, Ilya Segal, and Juuso Toikka. 2014. "Dynamic Mechanism Design: A Myersonian Approach." *Econometrica* 82 (2): 601–53.
- Rogerson, William P. 1985. "Repeated Moral Hazard." *Econometrica* 53:69–76.
- Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *Rev. Econ. Studies* 68 (1): 205–29.
- Samuelson, Paul A. 1974. "Complementarity: An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory." *J. Econ. Literature* 12 (4): 1255–89.
- Scheuer, Florian. 2014. "Entrepreneurial Taxation with Endogenous Entry." *American Econ. J.: Econ. Policy* 6 (2): 126–63.
- Schultz, Theodore W. 1961. "Investment in Human Capital." *A.E.R.* 51:1–17.
- Stantcheva, Stefanie. 2015a. "Learning and (or) Doing: Human Capital Investments and Optimal Taxation." Working Paper no. 21381, NBER, Cambridge, MA.
- . 2015b. "Optimal Income, Education, and Bequest Taxes in an Intergenerational Model." Working Paper no. 21177, NBER, Cambridge, MA.
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron. 2004. "Consumption and Risk Sharing over the Life Cycle." *J. Monetary Econ.* 51 (3): 609–33.
- Wei, C. Chang. 2010. "What Is the Price of College?" Technical Report no. NCES 2011-175, Nat. Center Educ. Statis., Dept. Educ., Washington, DC.
- Weinzierl, Matthew. 2011. "The Surprising Power of Age-Dependent Taxes." *Rev. Econ. Studies* 78 (4): 1490–1518.
- Werning, Iván. 2011. "Nonlinear Capital Taxation." Manuscript, Dept. Econ., Massachusetts Inst. Tech.