Optimal Taxation and Human Capital Policies over the Life Cycle*

Stefanie Stantcheva, MIT

June 19, 2014

Abstract

This paper derives optimal income tax and human capital policies in a dynamic life cycle model of labor supply and risky human capital formation. The wage is a function of both stochastic, persistent, and exogenous "ability" and endogenous human capital. Human capital is acquired throughout life through monetary expenses and training time. The government faces asymmetric information regarding the initial ability of agents and the lifetime evolution of ability, as well as the labor supply. The optimal subsidy on human capital expenses is determined by three considerations: counterbalancing distortions to human capital investment from the taxation of wage and capital income, encouraging labor supply, and providing insurance against adverse draws from the productivity distribution. 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. The optimal tax treatment of training time also depends on its interactions with contemporaneous and future labor supply. I consider two ways to implement the optimum: income contingent loans, and a tax scheme that allows for a deferred deductibility of human capital expenses. Numerical results are presented that suggest that full dynamic risk-adjusted deductibility of expenses might be close to optimal, and that simple linear age-dependent policies can achieve most of the welfare gain from the second best.

JEL classification: H21, H23, I21

Keywords: Income Taxation, Optimal Taxation, Lifecycle, Human Capital, Education, Subsidies

*Stefanie Stantcheva, MIT, stefanie@mit.edu; 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' insightful comments. I also thank seminar participants at Berkeley, Booth, Chicago, Harvard, Michigan, MIT, Penn, Princeton, Santa Barbara, Stanford, Stanford GSB, Wharton, and Yale for their useful feedback and comments.

1 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 (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 impact the available tax base and are a major determinant of the pre-tax 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.[2] 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, and how should the opportunity cost of time be treated? Second, how does the optimal tax system change when human capital acquisition is unobservable to the government? Third, 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.[3] Yet, private markets for insurance against personal productivity shocks are limited. Second, the investment has two types of costs, namely financial costs (e.g., tuition or books) and time costs (e.g., opportunity

[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, pg. 17).

[2] See Mirrlees (1971), Saez (2001), Golosov *et al.* (2006), Albanesi and Sleet (2006), Farhi and Werning (2013). Recent papers have relaxed the assumption of exogenous ability and are reviewed below.

[3] See Goldin and Katz (2008) and the literature reviewed later in the introduction.

cost of foregone wages and loss of leisure while studying or training),[4] the relative importance of which can change over the life cycle.[5] Finally, individuals have heterogeneous intrinsic abilities, which may differentially affect their returns to human capital investment.[6]

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, through either monetary expenses or time spent training at a disutility cost. The government maximizes a standard social 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.

The lifecycle model in the paper highlights the different forces at play. The implicit subsidy for human capital expenses is determined by three goals. The first is to counterbalance the distortions to human capital from implicit income and savings taxes. The second is to indirectly stimulate labor supply by increasing the wage, i.e., the returns to labor. The third is to redistribute, taking into account the differential effect of human capital on the pre-tax income of high and low ability people. When the wage elasticity with respect to ability is decreasing in human capital, 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 insuring a neutral tax system with respect to human capital expenses, i.e., beyond making human capital expenses fully tax deductible in a dynamic, risk-adjusted fashion. In this case, the human capital wedge drifts up with age, and the drift is amplified by the magnitude of the wage elasticity with respect to ability. The persistence of ability shocks directly translates into a persistence of the optimal human capital wedge over time.

When human capital can be acquired through training time at a disutility cost, as well as by spending money on tuition or other inputs, there is an additional, direct interaction of training with both contemporaneous and future labor supply, which affects optimal policies. Training may increase the marginal disutility cost from working, the case labeled "Learning-or-Doing," or decrease it, labeled

[4] OECD countries spend on average 6.3% of GDP on education; the US spends 7.3% (OECD, 2013). Heckman (1976a) indirectly infers that 30% of time in early working years is spent training. 40% of adults participate in formal and/or non-formal education in a given year in OECD countries, and receive on average 988 hours of instruction in non-formal education during life, out of which 715 are job-related (OECD, 2013).

[5] See Dynarski (2003) and Heckman et al. (2005) on, respectively, the effects of tuition and of disutility from education on school enrollment decisions.

[6] The empirical evidence on this issue is reviewed in section 5.1.

"Learning-And-Doing."

If, instead, human capital is unobservable to the government, the optimal labor and savings wedges need to indirectly provide the right incentives for human capital acquisition. The labor wedge may increase or decrease relative to the observable human capital case, depending, again, on whether human capital has a positive redistributive or insurance effect.

The 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 incomes, to account for the risk and the nonlinearity of the tax schedule.

I calibrate the model based on 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. The optimal labor wedge rises with age, but the age trend is less steep than in standard models with no human capital. Labor taxes are progressive in the short run when human capital has a positive insurance value, but regressive otherwise. 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.

1.1 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 (1976a,b). 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, 2008) 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. Disutility or psychic costs, above and beyond pure opportunity costs, are important in explaining human capital investment decisions (Heckman, Todd and Lochner, 2005). A large body of empirical work documents the importance of human capital as a determinant of earnings (Card, 1995a, Goldin and Katz, 2008, Chetty *et al.*, 2011, Huggett and Kaplan, 2011, Acemoglu and Autor, 2011), and the financial and other factors shaping individuals' decisions to acquire human capital (Lochner and Monge-Naranjo, 2011, Altonji, Blom and Meghir, 2012, Avery *et al.*, 2013). The subset of this literature which studies the interaction between ability and schooling for earnings – a crucial consideration for optimal policies in this paper – is reviewed in detail in section 5.1.

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 *et al.* (2006), Werning (2007a,b,c), Battaglini and Coate (2008), Scheuer (2013a), Golosov, Tsyvinski and Troshkin (2013), and Farhi and Werning (2013) typically assumes exogenous ability, thus abstracting from endogenous human capital investments. Therefore, this paper builds on the lifecycle 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.[7] A few subsequent static papers emphasize the importance of the complementarity between intrinsic ability and human capital (Maldonado, 2007, with two types, Bovenberg and Jacobs, 2011 with a continuum of types), or between risk and human capital (DaCosta and Maestri, 2007).

Several recent dynamic optimal tax papers examine the impact of taxation on human capital, with important differences to the current paper. Previous dynamic models allowed for heterogeneity across agents, but not uncertainty (Bohacek and Kapicka, 2008, Kapicka, 2013a), or uncertainty, but not heterogeneity (Anderberg, 2009, Grochulski and Piskorski, 2010), which precludes a discussion of redistributive policies. Findeisen and Sachs (2013) 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, through both expenses and time, and a progressive realization of uncertainty throughout life. A complementary analysis is Kapicka and Neira (2013), who 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

[7] Kaplow (1994) also considers the tax treatment of human capital.

binary decision that occurs exclusively before entry into the labor market.[8] The lifecycle analysis also addresses the issue of age-dependent taxation, as explored in Kremer (2002), Weinzierl (2011), and Mirrlees at al. (2011).[9]

The rest of the paper is organized as follows. Section 2 presents the dynamic lifecycle model and the full information benchmark. Section 3 sets up a recursive mechanism design program using the first-order approach. Section 4 solves for the optimal policies when human capital is observable and describes the forces shaping those policies, as well as their evolution over life. Section 5 contains the numerical analysis, which brings out the main mechanics behind the optimal policies, and illustrates their evolution over life. Section 6 discusses the implementation of the optimal policies using Income Contingent Loans (ICLs) and a Deferred Deductibility scheme. Section 7 considers the case with unobservable human capital. Section 8 concludes and discusses three alternative applications of the model: to intergenerational transfers and bequest taxation, to entrepreneurial taxation, and to health investments.

2 Lifecycle Model of Human Capital Acquisition and Labor Supply

This section starts with the setup of a lifecycle model with risky human capital.

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 in two ways. The first one is 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: $M_t'(e) > 0$, $\forall e_t > 0$; $M_t'(0) = 0$; $M_t''(e_t) \geq 0$, $\forall e_t \geq 0$. These monetary investments add to a stock of human capital acquired by expenses ("expenses" for short), s_t, which evolves according to $s_t = s_{t-1} + e_t$.[10] 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. Second, agents can spend time training. This may represent the time spent in formal classes or training courses, time spent reading, or on-the-job training. The disutility cost to an agent who provides l_t units of work and spends i_t units in training is $\phi_t(l_t, i_t)$, with

[8] Benabou (2002) jointly analyses taxes and education in a dynastic Ramsey model.

[9] This paper is more generally related to the dynamic mechanism design literature, as developed by, among many others, Baron and Besanko (1984), Spear and Srivastava (1987), Fernandes and Phelan (2000), and Doepke and Townsend (2006), which also underscores the challenges posed by general forms of persistence of shocks. The closest methodological link is to Farhi and Werning's (2013) first-order approach to such problems (see also Pavan, Segal and Toikka, 2013). The case with unobservable human capital builds on models with hidden savings (Cole and Kocherlakota, 2001, Werning, 2002, Kocherlakota, 2004, Abraham and Pavoni, 2008) or hidden trades (Golosov and Tsyvinski, 2007).

[10] The agent cannot wilfully destroy human capital, hence $e_t \geq 0$. There is no explicit depreciation of human capital, but the ability shock θ described right below can partially account for stochastic depreciation. Deterministic depreciation each period is easy to add and dropped for notational convenience.

$i_t \geq 0$. This second type of investment leads to a stock of training time z_t, which evolves according to $z_t = z_{t-1} + i_t$. ϕ_t is strictly increasing and convex in each of its arguments ($\frac{\partial \phi(l,i)}{\partial i}, \frac{\partial \phi(l,i)}{\partial l} > 0, \forall\, l, i > 0$; $\frac{\partial^2 \phi(l,i)}{\partial l^2}, \frac{\partial^2 \phi(l,i)}{\partial i^2} \geq 0, \forall l, i$); however, no assumption is yet made on the cross partial $\frac{\partial^2 \phi(l,i)}{\partial l \partial i}$.[11]

These two forms of human capital acquisition might vary in importance over the life cycle. For example, young people may face only low-wage opportunities, hence a low opportunity cost of time, but high tuition fees and financial costs. On the other hand, for working adults with college education and high opportunity cost of time, the costs of human capital investments in the form of on-the-job training might mostly be time and disutility costs.

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

$$w_t = w_t\left(\theta_t, s_t, z_t\right)$$

w_t is strictly increasing and concave in each of its arguments ($\frac{\partial w}{\partial m} > 0$, $\frac{\partial^2 w}{\partial m^2} \leq 0$ for $m = \theta, s, z$).[12] Importantly though, no restrictions are placed on the cross-partials. This formulation allows for different types of human capital to affect the wage differently over the lifecycle.

Agents are born at time $t = 1$ with a heterogeneous earning ability θ_1 with distribution $f^1\left(\theta_1\right)$. Earning ability in each period is private information, and evolves according to a Markov process with a time-varying transition function $f^t\left(\theta_t|\theta_{t-1}\right)$, over a fixed support $\Theta \equiv \left[\underline{\theta}, \bar{\theta}\right]$. 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 $w_t = \theta_t + h_t\left(s_t, z_t\right)$, for some function h_t, θ_t resembles a stochastic version of productivity from the static Mirrlees (1971) model. With a wage such as $w_t = \theta_t h_t\left(s_t, z_t\right)$, 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 among others because of health shocks, individual labor market idiosyncrasies or luck.

The agent's per period utility is separable in consumption and effort (both labor and training):

$$\tilde{u}_t\left(c_t, y_t, s_t, z_t; \theta_t, z_{t-1}\right) = u_t\left(c_t\right) - \phi_t\left(\frac{y_t}{w_t\left(\theta_t, s_t, z_t\right)}, z_t - z_{t-1}\right)$$

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 $P\left(\theta^t\right)$ the probability of a history θ^t, $P\left(\theta^t\right) \equiv f^t\left(\theta_t|\theta_{t-1}\right) ... f^2\left(\theta_2|\theta_1\right) f^1\left(\theta_1\right)$. An allocation $\{x_t\}_t$ specifies consumption, output, and expenses and training stocks for each period t, conditional

[11] Appendix A.1. links this model to the standard Ben-Porath model.

[12] Equivalently, we could define human capital as a composite function h_t of all expenses made and time spent training $h_t = h_t\left(s_t, z_t\right)$,with wage $\tilde{w}_t = w_t\left(\theta_t, h_t\right)$.

on the history θ^t, i.e., $x_t = \left\{x\left(\theta^t\right)\right\}_{\Theta^t} = \left\{c\left(\theta^t\right), y\left(\theta^t\right), s\left(\theta^t\right), z\left(\theta^t\right)\right\}_{\Theta^t}$. The expected lifetime utility from an allocation, discounted by a factor β, is given by:

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

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

Let $w_{m,t}$ denote the partial of the wage function with respect to argument m ($m \in \{\theta, s, z\}$), and $w_{mn,t}$ the second order partial with respect to arguments m,$n \in \{\theta, s, z\} \times \{\theta, s, z\}$. Two important parameters are the Hicksian coefficients of complementarity between ability and human capital (of type s or z) in the wage function at time t, denoted by $\rho_{\theta s,t}$ and $\rho_{\theta z,t}$:[13]

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

A positive Hicksian complementarity between human capital s and ability θ means that higher ability agents have a higher marginal benefit from human capital ($w_{\theta s} \geq 0$).[14] Put differently, human capital compounds the exposure of the agent to stochastic ability and to risk. A Hicksian complementarity greater than 1 means that higher ability agents have a higher proportional benefit from human capital, i.e., the wage elasticity with respect to ability is increasing in human capital, i.e., $\frac{\partial}{\partial s}\left(\frac{\partial w}{\partial \theta}\frac{\theta}{w}\right) \geq 0$.[15] The same applies to training time z.

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

$$w_t = \left[\alpha_{1t}\theta^{1-\rho_t} + \alpha_{2t}s_t^{1-\rho_t} + \alpha_{3t}z^{1-\rho_t}\right]^{\frac{1}{1-\rho_t}} \quad (3)$$

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

For training time, another relevant parameter is ρ_{lz}^{ϕ}, the Hicksian complementarity coefficient between labor and training in the disutility function $\phi_t\left(l_t, z_t - z_{t-1}\right)$. Let $\phi_{z,t} = \frac{\partial \phi_t}{\partial z_t}$, $\phi_{l,t} = \frac{\partial \phi_t}{\partial l_t}$, $\phi_{lz,t} = \frac{\partial^2 \phi}{\partial l \partial z}$. Then:

$$\rho_{lz}^{\phi} \equiv \frac{\phi_{lz}\phi}{\phi_l \phi_z} \quad (4)$$

[13] On studies of the Hicksian complementarity as see Hicks (1970), Samuelson (1973), Bovenberg and Jacobs (2011).

[14] $\rho_{\theta s}$ is also the Hicksian complementarity coefficient between education and ability in earnings y.

[15] Equivalently, the wage elasticity with respect to human capital is increasing in ability.

3 The Planning Problem

The informational problem that the planner faces is that he cannot see ability θ_t in any period. Hence, he also does not know an agent's wage $w_t(\theta_t, s_t, z_t)$, or labor supply $l_t = y_t/w_t$. Output y_t and consumption c_t on the other hand are observable. Human capital investments are also observable to the government but the case with unobservable human capital expenses is discussed in Section 7.

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 programing problem through a suitable definition of state variables. The results on optimal policies are in the subsequent section.

3.1 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_t. Denote a reporting strategy, specifying a reported type r_t after each history by $r = \left\{r_t\left(\theta^t\right)\right\}_{t=1}^{T}$. Let $\mathcal{R}$ be the set of all possible reporting strategies and $r^t = \left\{r_1(\theta_1), ..., r_t\left(\theta^t\right)\right\}$ 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 $c\left(r^t\right), y\left(r^t\right), s\left(r^t\right), z\left(r^t\right)$.[16] Let the continuation value after history θ^t under a reporting strategy r, denoted by $\omega^r\left(\theta^t\right)$, be the solution to:

$$\omega^r\left(\theta^t\right) = u_t\left(c\left(r^t\left(\theta^t\right)\right)\right) - \phi_t\left(\frac{y\left(r^t\left(\theta^t\right)\right)}{w_t\left(\theta_t, s\left(r^t\left(\theta^t\right)\right), z\left(r^t\left(\theta^t\right)\right)\right)}, z\left(r^t\left(\theta^t\right)\right) - z\left(r^{t-1}\left(\theta^{t-1}\right)\right)\right) + \beta\int \omega^r\left(\theta^{t+1}\right) f^{t+1}\left(\theta_{t+1}|\theta_t\right) d\theta_{t+1}$$

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

$$\omega\left(\theta^t\right) = u_t\left(c\left(\theta^t\right)\right) - \phi_t\left(\frac{y\left(\theta^t\right)}{w_t\left(\theta_t, s\left(\theta^t\right), z\left(\theta^t\right)\right)}, z\left(\theta^t\right) - z\left(\theta^{t-1}\right)\right) + \beta\int \omega\left(\theta^{t+1}\right) f^{t+1}\left(\theta_{t+1}|\theta_t\right) d\theta_{t+1}$$

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

$$(IC) : \omega(\theta_1) \geq \omega^r(\theta_1) \quad \forall\theta_1, \forall r \tag{5}$$

[16] Hours of work are determined residually by $l\left(r^t\right) = y\left(r^t\right)/w\left(\theta_t, s\left(r^t\right), z\left(r^t\right)\right)$ and hours of training by $i\left(r^t\right) = z\left(r^t\right) - z\left(r^{t-1}\right)$.

Denote by X^{IC} the set of allocations which satisfy incentive compatibility condition (5).

3.2 Planning Problem

The analysis is in partial equilibrium, or assuming a small open economy with fixed gross interest rate R. The planner's objective is to minimize the expected discounted cost of providing an allocation, subject to incentive compatibility as defined in (5), and to expected lifetime utility being above a threshold $\underline{U}$. The planning problem in sequential form is:

$$\min_{\{c,y,s,z\}} \Pi\left(\{c,y,s,z\}\right) = \left[\sum_{t=1}^{T}\left(\frac{1}{R}\right)^{t-1}\int_{\Theta^t}\left(c\left(\theta^t\right)-y\left(\theta^t\right)+M_t\left(s\left(\theta^t\right)-s\left(\theta^{t-1}\right)\right)\right)P\left(\theta^t\right)d\theta^t\right] \tag{6}$$

$$\begin{aligned}&\text{s.t.}: U\left(\{c,y,s,z\}\right)\geq \underline{U}\\ &y\left(\theta^t\right)\geq 0, s\left(\theta^t\right)\geq s\left(\theta^{t-1}\right), z\left(\theta^t\right)\geq z\left(\theta^{t-1}\right), c\left(\theta^t\right)\geq 0\\ &\{c,y,s,z\}\text{ is incentive compatible}\end{aligned}$$

First order approach: To solve this dynamic problem, a version of the first order approach is used. The following assumptions are needed and maintained throughout the paper.

Assumption 1 *i)* $\tilde{u}_t\left(c,y,s,z;\theta,z_-\right)$ *and* $\frac{\partial\phi(l,i)}{\partial l}\frac{\partial w(\theta,s,z)}{\partial\theta}\frac{l}{w}$ *are bounded. ii)* $\frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial\theta_{t-1}}$ *exists and is bounded.*[17] *iii)* $f^t\left(\theta_t|\theta_{t-1}\right)$ *has full support on* Θ*.*

Suppose 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 ($\tilde{r}_s\left(\theta^s\right)=\theta_s\ \forall s\leq t-1$), and lies in period t by reporting $\tilde{r}_t\left(\theta^t\right)=\theta'\neq\theta_t$. The continuation utility under this strategy is the solution to:

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

Incentive compatibility in (5) implies that, after almost all θ^t, the temporal incentive constraint holds:

$$\omega\left(\theta^t\right)=\max_{\theta'}\omega^{\tilde{r}}\left(\theta^t\right) \tag{7}$$

Inversely, if (7) holds after all θ^{t-1} and for almost al θ_t, then (5) also holds (see Kapicka (2013b), Lemma 1). The first order approach consists in replacing(7) by the envelope condition of the agent:

$$\dot{\omega}\left(\theta^t\right):=\frac{\partial\omega\left(\theta^t\right)}{\partial\theta_t}=\frac{w_{\theta,t}}{w_t}l\left(\theta^t\right)\phi_{l,t}\left(l\left(\theta^t\right),z_t\left(\theta^t\right)-z_{t-1}\left(\theta^{t-1}\right)\right)+\beta\int\omega\left(\theta^{t+1}\right)\frac{\partial f^{t+1}\left(\theta_{t+1}|\theta_t\right)}{\partial\theta_t}d\theta_{t+1} \tag{8}$$

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

or its integral version:

$$\omega\left(\theta^{t}\right)=\omega\left(\theta^{t-1},\underline{\theta}\right)+\int_{\underline{\theta}}^{\theta_{t}}(\frac{w_{\theta,t}}{w_{t}}l\left(\theta^{t-1},\theta_{s}\right)\phi_{l,t}\left(l\left(\theta^{t-1},\theta_{s}\right),z_{t}\left(\theta^{t-1},\theta_{s}\right)-z_{t-1}\left(\theta^{t-1}\right)\right)$$
$$+\beta\int\omega\left(\theta^{t-1},\theta_{s},\theta_{t+1}\right)\frac{\partial f^{t+1}\left(\theta_{t+1}|\theta_{s}\right)}{\partial\theta_{s}}d\theta_{t+1})d\theta_{s}\quad(9)$$

Let X^{FOA} denote the set of allocations which satisfy the envelope condition (9). It can be shown that this is a necessary condition for incentive compatibility, i.e., $X^{IC}\subseteq X^{FOA}$.[18] The relaxed planning problem, denoted by P^{FOA} is:

$$\min_{\{c,y,s,z\}}\Pi\left(\{c,y,s,z\}\right)$$

$$\text{s.t: } U\left(\{c,y,s,z\}\right)\geq\underline{U}$$
$$\{c,y,s,z\}\in X^{FOA}$$

3.3 Recursive Formulation of the Relaxed Program

To write the problem recursively, let the second term in the envelope condition in differential form – which governs the future incentives of the agent – be denoted by:

$$\Delta\left(\theta^{t}\right)\equiv\int\omega\left(\theta^{t+1}\right)\frac{\partial f^{t+1}\left(\theta_{t+1}|\theta_{t}\right)}{\partial\theta_{t}}d\theta_{t+1}\quad(10)$$

The envelope condition can then be rewritten as:

$$\dot{\omega}\left(\theta^{t}\right)=\frac{w_{\theta,t}}{w_{t}}l\left(\theta^{t}\right)\phi_{l,t}\left(l\left(\theta^{t}\right),z_{t}\left(\theta^{t}\right)-z_{t-1}\left(\theta^{t-1}\right)\right)+\beta\Delta\left(\theta^{t}\right)\quad(11)$$

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

$$v\left(\theta^{t}\right)\equiv\int\omega\left(\theta^{t+1}\right)f^{t+1}\left(\theta_{t+1}|\theta_{t}\right)d\theta_{t+1}\quad(12)$$

Continuation utility $\omega\left(\theta^{t}\right)$ can hence be rewritten as:

$$\omega\left(\theta^{t}\right)=u_{t}\left(c\left(\theta^{t}\right)\right)-\phi_{t}\left(\frac{y\left(\theta^{t}\right)}{w_{t}\left(\theta_{t},s\left(\theta^{t}\right),z\left(\theta^{t}\right)\right)},z\left(\theta^{t}\right)-z\left(\theta^{t-1}\right)\right)+\beta v\left(\theta^{t}\right)\quad(13)$$

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

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

[18] An application of Theorem 2 in Milgrom and Segal (2002), under assumption 1.

where, with some abuse of notation, $d\theta^{\tau-t} = d\theta_\tau d\theta_{\tau-1}...d\theta_t$, and $P\left(\theta^{\tau-t}\right) = f^\tau\left(\theta_\tau|\theta_{\tau-1}\right)...f^t\left(\theta_t|\theta_{t-1}\right)$. The minimization is over continuation allocations after period t, $\{c, y, s, z, \omega, v, \Delta\}_{\tau \geq t}$, and subject to the constraints (10), (11), (12), and (13), evaluated at $t-1$.

A recursive formulation of the relaxed program is then for $t \geq 2$:

$$K\left(v, \Delta, \theta_-, s_-, z_-, t\right) = \min \int \left(c\left(\theta\right) + M_t\left(s\left(\theta\right) - s_-\right) - w_t\left(\theta, s\left(\theta\right), z\left(\theta\right)\right) l\left(\theta\right)\right.$$
$$\left. + \frac{1}{R} K\left(v\left(\theta\right), \Delta\left(\theta\right), \theta, s\left(\theta\right), z\left(\theta\right), t+1\right)\right) f^t\left(\theta|\theta_-\right) d\theta \quad (14)$$

subject to:

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

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

For period $t = 1$, the problem needs to be reformulated only if the initial shock θ_1 is interpreted as heterogeneity, rather than uncertainty, and the objective is not utilitarian, so as to derive the full Pareto frontier. Suppose all agents have identical initial human capital levels s_0 and z_0. The problem for $t = 1$ is then indexed by $\left(\underline{U}\left(\theta\right)\right)_\Theta$, the set of target lifetime utilities $\underline{U}\left(\theta\right)$ for each type θ:

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

$$\begin{aligned}
\text{s.t.} \quad &: \quad \omega\left(\theta\right) = u_1\left(c\left(\theta\right)\right) - \phi_1\left(l\left(\theta\right), z\left(\theta\right) - z_0\right) + \beta v\left(\theta\right) \\
\dot{\omega}\left(\theta\right) &= \frac{w_{\theta,1}}{w_1} l\left(\theta\right) \phi_{l,1}\left(l\left(\theta\right), z\left(\theta\right)\right) + \beta \Delta\left(\theta\right) \\
\omega\left(\theta\right) &\geq \underline{U}\left(\theta\right)
\end{aligned}$$

where the maximization is again over $\left(c\left(\theta\right), l\left(\theta\right), \omega\left(\theta\right), s\left(\theta\right), z\left(\theta\right), v\left(\theta\right), \Delta\left(\theta\right)\right)$. The set of constrained efficient allocations that solve the Planning problem is indexed by the set of utilities $\left(\underline{U}\left(\theta\right)\right)_\Theta$ and denoted by $X^{*,FOA}\left(\left(\underline{U}\left(\theta\right)\right)_\Theta\right)$.

3.4 Validity of the First-order Approach and Assumptions

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 Battaglini and Lamba, 2013, Golosov, Troshkin and Tsyvinski, 2013 or Pavan, Segal and Toikka, 2013), and do not always provide much simplification. Hence, for the proposed calibrations in section 5, incentive compatibility of the candidate allocation, as well as any omitted non-negativity constraints, are checked numerically, using a procedure in the spirit of Werning (2007) and Farhi and Werning (2013).[19]

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.[20]

Assumption 2 *i)* $\frac{\partial v(\theta)}{\partial \theta} > 0$ *for all* θ*, ii)* $\int_{\underline{\theta}}^{\theta'} f^t(\theta|\theta_b)\, d\theta \leq \int_{\underline{\theta}}^{\theta'} f^t(\theta|\theta_s)\, d\theta$, $\forall t, \theta'$, *and* $\theta_b > \theta_s$*; iii)* $\frac{\partial}{\partial \theta_t}\left(\frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial \theta_{t-1}} \frac{1}{f^t(\theta_t|\theta_{t-1})}\right) \geq 0$, $\forall t$, $\forall \theta_{t-1}$*; iv)* $\frac{\partial}{\partial v} K \geq 0$ *and* $\frac{\partial^2}{\partial v^2} K \geq 0$.

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

To reduce notational clutter throughout the paper, the dependence on the full history is often left implicit, e.g.: $c_t = c\left(\theta^t\right)$ and $\tau_{Lt} = \tau_{Lt}\left(\theta^t\right)$. Similarly, function arguments are sometimes left out, e.g.: $w_{z,t} = \frac{\partial}{\partial z} w_t\left(\theta_t, s\left(\theta^t\right), z\left(\theta^t\right)\right)$. E_t denotes the expectation at time t, conditional on θ_t.

4 Optimal Human Capital Policies

This section characterizes the allocations, obtained as solutions to the relaxed program P^{FOA} above, using wedges, or implicit taxes and subsidies. Unless otherwise stated, "optimal" refers to the optimum of program P^{FOA}.

4.1 The Wedges or Implicit Taxes and Subsidies

In the second best, marginal distortions in agents' choices can be described using wedges which represent the implicit local marginal tax and subsidy rates. At a given allocation and history θ^t, define the

[19] An alternative could be the random contracts or "lotteries" approach (e.g.: Karaivanov and Townsend, 2013), which circumvents the sufficiency problem, and has explored increasingly sophisticated methods to increase computational efficiency, and counter the "curse of dimensionality," which arises from the exponential growth of incentive constraints when adding hidden states or actions. For optimal tax analysis, however, it is appealing to have analytical formulas, which the first order approach, when valid, can deliver, and which build the intuition for the solution.

[20] In addition, all theoretical results on the signs of the wedges are indeed satisfied in the simulations (section 5).

intratemporal wedge on labor $\tau_L(\theta^t)$, the intertemporal wedge on savings or capital $\tau_K(\theta^t)$, and the human capital wedges $\tau_S(\theta^t)$ and $\tau_Z(\theta^t)$ as follows:

$$\tau_L(\theta^t) \equiv 1 - \frac{\phi_{l,t}(l_t, z_t - z_{t-1})}{w_t(\theta_t, s_t, z_t)\, u'_t(c_t)} \tag{15}$$

$$\tau_K(\theta^t) \equiv 1 - \frac{1}{R\beta} \frac{u'_t(c_t)}{E_t(u'_t(c_{t+1}))} \tag{16}$$

$$\tau_S(\theta^t) \equiv -\left(1 - \tau_L(\theta^t)\right) w_{s,t} l_t + \left[M'_t(s_t - s_{t-1}) - \beta E_t\left(\frac{u'_{t+1}(c_{t+1})}{u'_t(c_t)} M'_{t+1}(s_{t+1} - s_t)\right)\right] \tag{17}$$

$$\tau_Z(\theta^t) \equiv -\left(1 - \tau_L(\theta^t)\right) w_{z,t} l_t + \left[\frac{\phi_{z,t}}{u'_t(c_t)} - \beta E_t\left(\frac{u'_{t+1}(c_{t+1})}{u'_t(c_t)} \frac{\phi_{z,t+1}}{u'_{t+1}(c_{t+1})}\right)\right] \tag{18}$$

The labor wedge τ_L is defined as the gap between the marginal rate of substitution and the marginal rate of transformation between consumption and labor. The savings or capital 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 can be thought of as the reduction in marginal cost for an additional dollar of spending on human capital: the agent's net marginal cost is $M'_t(e_t) - \tau_{St}$.[21] Similarly, the "bonus" for training τ_Z can be viewed as the incremental pay received by an agent for one more unit of training time.

Intuitively, conditional on a history, wedges are akin to locally linear subsidies and taxes, and would be zero absent government intervention. Indeed, the paper will sometimes loosely refer to the wedges as taxes and subsidies, and often appeal to intuitions related to a standard tax system. The relation between wedges and explicit taxes is studied in the section on implementation (section 6).

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

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

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:

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

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, $\varepsilon_{xyt} \equiv d\log(x_t)/d\log(y_t)$. Let ε^u_t and ε^c_t be the uncompensated and compensated labor supply elasticities to the net wage, holding both savings and training

[21] The wedges are here written with a recursive shorthand notation, replacing the future stream of marginal benefits by the future marginal cost, which holds at interior solutions.

fixed.[22]

4.2 Optimal Subsidy for Human Capital Expenses

4.2.1 The Net Subsidy

To achieve insurance and redistribution in the presence of incentive compatibility constraints, the optimal allocation features a labor wedge and an intertemporal wedge, both of which affect human capital investments. Intuitively, an increase in the implicit labor tax captures part of the return to the agent's investment, while the implicit savings tax encourages substitution from physical capital to human capital as an indirect way to transfer resources intertemporally. Hence, part of the implicit subsidy on human capital expenses merely acts to counterbalance the distortions in the human capital choice caused by other wedges. A useful benchmark is the subsidy which just ensures that the tax system is neutral with respect to human capital. Hence, I define a "net subsidy," as the gross subsidy from which we filter out all the parts that only go toward compensating for the other distortions.

Definition 1 *Define the net wedge on human capital expenses, t_{st}, as:*

$$t_{st} \equiv \frac{\tau_{St} - \tau_{Lt} M_t^{\prime d} + P_t}{\left(M_t^{\prime d} - \tau_{St}\right)(1 - \tau_{Lt})} \tag{19}$$

$M_t^{\prime d} \equiv M_t' - \frac{(1-\xi_{M'})}{R(1-\tau_K)} E_t\left(M_{t+1}'\right)$ *denotes the dynamic, risk-adjusted cost.*
$P_t \equiv \frac{\tau_K}{R(1-\tau_K)}(1-\tau_{Lt})\left(1-\xi_{M'}'\right) E_t\left(M_{t+1}'\right)$ *captures the risk-adjusted savings distortion.*

A zero net wedge t_{st} means that the tax system is neutral with respect to human capital. It is equivalent to full dynamic, risk-adjusted deductibility.[23] For instance, in the last period T, or in a static problem, $t_{st} = 0$ means that $\tau_{ST} = \tau_{LT} M_T'$, i.e., full local contemporaneous deductibility of expenses – akin to the type of deductions we are familiar with (see Bovenberg and Jacobs, 2005). In this dynamic setting, however, three additional effects need to be filtered out. First, the effective marginal cost of investing is dynamic ($M_t^{\prime d}$), since a higher investment at t means a lower cost of reaching any given human capital stock at $t+1$. Second, the intertemporal resource transfer from human capital investment is also taken into account in P_t, which would be zero with no savings distortion or with linear utility. Finally, uncertainty requires a risk-adjustment, embodied in the insurance factor $\xi_{M'}$,

[22] I.e., ε^c and ε^u are defined as in the static framework (e.g., Saez, 2001), at constant savings and training time:

$$\varepsilon^u = \frac{\phi_l(l,i)/l + \frac{\phi_l(l,i)^2}{u'(c)^2}u''(c)}{\phi_{ll}(l,i) - \frac{\phi_l(l)^2}{u'(c)^2}u''(c)} \qquad \varepsilon^c = \frac{\phi_l(l,i)/l}{\phi_{ll}(l,i) - \frac{\phi_l(l)^2}{u'(c)^2}u''(c)}$$

With per-period utility separable in consumption and labor, $\frac{\varepsilon_t^c}{1+\varepsilon_t^u-\varepsilon_t^c}$ is the Frisch elasticity of labor.

[23] It is not strictly speaking a standard tax deductibility, because wedges are not equivalent to explicit taxes, but this intuition is helpful to grasp the underlying logic, and will often be referred to.

since a given deduction is not worth the same in different states given that marginal utility is not constant.

The net wedge is positive if the government wants to encourage human capital acquisition on net, i.e., above and beyond simply ensuring a neutral tax system with respect to human capital. The next proposition shows how the net wedge is set in a recursive fashion at the optimum.

Proposition 1 *i) At the optimum, the net wedge is given by:*

$$t_{st}^{*}\left(\theta^{t}\right)=\frac{\mu\left(\theta^{t}\right)u_{t}'\left(c\left(\theta^{t}\right)\right)}{f^{t}\left(\theta_{t}|\theta_{t-1}\right)}\left(1-\rho_{\theta s,t}\right)\frac{\varepsilon_{w\theta,t}}{\theta_{t}} \tag{20}$$

where the multiplier $\mu\left(\theta_{t}\right)$ can be written recursively as:

$$\mu\left(\theta^{t}\right)=\kappa\left(\theta^{t}\right)+\eta\left(\theta^{t}\right) \tag{21}$$

$$\begin{aligned}
\kappa\left(\theta^{t}\right) &= \int_{\theta_{t}}^{\bar{\theta}}\left(1-g_{s}\right)\frac{1}{u_{t}'\left(c\left(\theta^{t-1},\theta_{s}\right)\right)}f\left(\theta_{s}|\theta_{t-1}\right)d\theta_{s} \qquad (22)\\
\textit{with } g_{s} &= u_{t}'\left(c\left(\theta^{t-1},\theta_{s}\right)\right)\lambda_{t-1} \textit{ and } \lambda_{t-1}=\int_{\underline{\theta}}^{\bar{\theta}}\frac{1}{u_{t}'\left(c\left(\theta^{t-1},\theta_{m}\right)\right)}f\left(\theta_{m}|\theta_{t-1}\right)d\theta_{m}\\
\eta\left(\theta^{t}\right) &= t_{st-1}^{*}\left(\theta^{t-1}\right)\left[\frac{R\beta}{u_{t-1}'\left(c\left(\theta^{t-1}\right)\right)}\frac{1}{\left(1-\rho_{\theta s,t-1}\right)}\frac{\theta_{t-1}}{\varepsilon_{w\theta,t-1}}\left(\int_{\theta_{t}}^{\bar{\theta}}\frac{\partial f\left(\theta_{s}|\theta_{t-1}\right)}{\partial\theta_{t-1}}d\theta_{s}\right)\right]
\end{aligned}$$

ii) If assumption (2) holds, then at each history θ^{t}:[24]

$$t_{st}^{*}\left(\theta^{t}\right)\geq0\Leftrightarrow\rho_{\theta s,t}\leq1$$

iii) $t_{st}^{*}\left(\theta^{t}\right)=0$ if $u_{t}'\left(c_{t}\right)=1$ and θ_{t} is iid.

iv) $t_{st}^{*}\left(\theta^{t-1},\bar{\theta}\right)=t_{st}^{*}\left(\theta^{t-1},\underline{\theta}\right)=0,\ \forall t.$

The multiplier $\mu\left(\theta^{t}\right)$ on the envelope condition is split into two parts. The insurance motive is captured in $\kappa\left(\theta^{t}\right)$, familiar from the static taxation literature. It would be zero with linear utility. 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 novel term $\eta\left(\theta^{t}\right)$ 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.[25] This motive remains effective

[24] Assumption 2 can be replaced by the condition $\tau_{Lt}\left(\theta^{t}\right)>0$. That $t_{st}=0$ when $\rho_{\theta s}=1$ and corollary 1 apply even when the first-order approach is not valid (see the Appendix).

[25] With a non-utilitarian objective, if θ_1 is interpreted as heterogeneity (see section 3.3):

$$\mu\left(\theta_{1}\right)=\int_{\theta_{1}}^{\bar{\theta}}\frac{1}{u_{1}'\left(c_{1}\left(\theta_{s}\right)\right)}\left(1-\lambda_{0}\left(\theta_{s}\right)u_{1}'\left(c_{1}\left(\theta_{s}\right)\right)\right)f\left(\theta_{s}\right)$$

where $\lambda_{0}\left(\theta_{s}\right)$ is the multiplier (scaled by $f\left(\theta_{s}\right)$) on type θ_s target utility. With linear utility: $1=\int_{\underline{\theta}}^{\bar{\theta}}\lambda_{0}\left(\theta_{s}\right)f\left(\theta_{s}\right)$.

through $\eta\left(\theta^t\right)$, the more so if types are more persistent, but vanishes as skills become less persistent. In the limit, if θ_t is identically and independently distributed (iid), only the contemporaneous insurance motive $\kappa\left(\theta_t\right)$ matters. If, in addition to iid shocks, utility is linear in consumption, the optimal net subsidy is 0.[26]

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

Importantly, the proposition underlines that the net wedge on human capital is positive if and only if $\rho_{\theta s,t} \leq 1$, a point discussed next.

4.2.2 The Redistributive and Insurance Values of Human Capital

Together with the aforementioned insurance and persistence channels, the Hicksian coefficient of complementarity $\rho_{\theta s}$ plays a crucial role. 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 makes leisure less attractive and encourages labor supply, the "Labor Supply Effect." This is a positive effect in the presence of a distortion in the labor decision.[27] At the same time, if $\rho_{\theta s} > 0$, which is equivalent to ability being complementary to human capital in the wage $\left(\frac{\partial^2 w}{\partial\theta\partial s} > 0\right)$, human capital mostly benefits already able agents, and hence compounds existing inequality due to intrinsic differences in θ_t. 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.

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 answer is that, when $\rho_{\theta s} < 1$, i.e., when high ability agents do not disproportionately benefit from human capital, the positive labor supply effect dominates any potential inequality effect, and generates net resources to be used for redistribution and insurance of all agents.

[26] Except at $t = 1$ with a non-utilitarian objective, see the previous footnote.

[27] This effect would disappear if there were no distortion on labor, ie., if $\mu\left(\theta^t\right) = 0$. 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.

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.[28] Put differently, if $\rho_{\theta s} > 0$, a human capital subsidy increases pre-tax income inequality, but as long as $\rho_{\theta s} \leq 1$, after-tax income inequality will be reduced thanks to the additional resources generated by human capital investment.

Although the human capital and labor literatures have long studied the interaction between human capital and ability (see the summary in section 5.1), the optimal taxation literature has mostly adopted the more restrictive wage form $w = \theta s$.[29] This special case entails $\rho_{\theta s} = 1$ and, hence, a null net wage at the optimum. This is reminiscent of the Atkinson and Stiglitz (1976) result on the non-optimality of differential commodity taxation if preferences satisfy a form of separability between goods and labor. Note that the zero net subsidy result for $\rho_{\theta s} = 1$ does not depend on the optimality of the labor or intertemporal wedges.[30] 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. This benchmark case is discussed next.

4.2.3 Full Risk-adjusted Dynamic Deductibility

With any multiplicatively separable wage of the form $w(\theta, s, z) = w^1(\theta)\, w^2(s, z)$, where w^1 and w^2 are increasing functions, the gross wedge τ_{St} only compensates for the distortion caused by income and savings taxes. The following corollary shows the implicit dynamic deductibility relation that the gross wedge, the labor tax, and the capital tax need to satisfy at the optimum.[31]

Corollary 1 *If $\rho_{\theta s} = 1$, the constrained efficient human capital wedge τ^*_{St}, labor wedge τ^*_{Lt}, and capital wedge τ^*_{Kt} must satisfy:*

$$\tau^*_{St} = \left(M'_t - \frac{1}{R} E_t\left(M'_{t+1}\right)\right)\tau^*_{Lt} - \frac{\tau^*_{Kt}}{R\left(1-\tau^*_{Kt}\right)}\left(1 - \xi'_{M',t+1}\right) E_t\left(M'_{t+1}\right) \tag{23}$$

Applying this corollary to a static model or to time $t = T$ shows why Bovenberg and Jacobs (2005) find that education subsidies and income taxes are "Siamese Twins," and that the optimal (linear) subsidy on education is equal to the optimal (linear) tax rate. In a lifecycle setting with uncertainty,

[28] Importantly, the social objective assigns non-negative 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.

[29] With the exception of the references in the introduction (Bovenberg and Jacobs, 2011, Maldonado, 2007, DaCosta and Maestri, 2007). This generalized dynamic model nests their findings. In particular, it emphasizes the role of the distribution of productivity and its persistence over time, which can modulate, amplify, or fully dampen the effects of the complementarity between human capital and ability. In addition, I show below how the *evolution* of the net wedge is affected by the coefficient of complementarity, something which could not be seen in a static model.

[30] As in the direct proofs of the Atkinson-Stiglitz result by Kaplow (2006) and Laroque (2005).

[31] See the general formula for $\rho_{\theta s} \neq 1$, as well as an alternative formulation in the Appendix.

additional considerations, namely the dynamic nature of the investment cost, the intertemporal wedge, and the insurance factor, complicate this simple contemporaneous deductibility. Section 6 maps this result into a deferred deductibility scheme with explicit tax instruments.

4.2.4 The Capital Wedge with Observable Human Capital

The presence of observable human capital does not change a standard result in dynamic moral hazard models with observable savings and separable utility, namely the Inverse Euler Equation.

Proposition 2 *At the optimum, the inverse Euler Equation holds*:

$$\frac{R\beta}{u_t'\left(c\left(\theta^t\right)\right)}=\int_{\underline{\theta}}^{\bar{\theta}}\frac{1}{u_{t+1}'\left(c\left(\theta^{t+1}\right)\right)}f^{t+1}\left(\theta_{t+1}|\theta_t\right)d\theta_{t+1} \tag{24}$$

Equation (24), combined with Jensen's inequality, implies that the savings wedge in (16) is positive, i.e., savings are discouraged. Indeed, saving and shirking in the future are complements. Human capital is an alternative way to transfer resources to the future, and a substitute to savings through physical capital (Heckman, 1976). Hence, in the absence of additional redistributive or insurance effects from human capital ($\rho_{\theta s}=1$), formula (23) shows that the human capital subsidy and the savings wedge co-move inversely.

4.3 The Optimal Labor Wedge with Observable Human Capital

The optimal labor wedge is very similar to the one in dynamic taxation models without human capital, except for the novel wage function.

Proposition 3 *i) At the optimum, the labor wedge is equal to:*

$$\frac{\tau_{Lt}^*\left(\theta^t\right)}{1-\tau_L^*\left(\theta^t\right)}=\frac{\mu\left(\theta^t\right)u_t'\left(c\left(\theta^t\right)\right)}{f^t\left(\theta_t|\theta_{t-1}\right)}\frac{\varepsilon_{w\theta,t}}{\theta_t}\frac{1+\varepsilon_t^u}{\varepsilon_t^c} \tag{25}$$

with $\mu\left(\theta^t\right)=\eta\left(\theta^t\right)+\kappa\left(\theta^t\right)$ *as in* (21)*, where* $\eta\left(\theta^t\right)$ *can be rewritten recursively as a function of the past labor wedge,* $\tau_{L,t-1}$*:*

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

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

Unlike in the standard taxation model, the wage elasticity with respect to ability $\varepsilon_{w\theta,t}$ is not constant at 1; a higher elasticity amplifies the labor wedge as it increases the value of insurance and redistribution.[32] However, the standard zero distortion at the bottom and the top results from the

[32] E.g.: with a CES wage as in (3), $\varepsilon_{w\theta}=\left(\frac{w_t(\theta_t,s_t,z_t)}{\theta_t}\right)^{\rho-1}$; human capital indirectly enters the optimal tax formula.

static Mirrlees model continue to apply in the presence of observable human capital. The labor wedge at any age is inversely related to the elasticity of labor supply that prevails at that time.

In this setting, both the labor wedge and the net human capital wedge are tools for redistribution and insurance. The following relation shows that they are set according to a type of inverse elasticity rule, i.e., inversely proportional to their efficiency costs:

Corollary 2 *At the optimum, the labor wedge and human capital wedge need to satisfy the following relation after each history:*

$$t_{st}^{*} = \left(\frac{\tau_{Lt}^{*}}{1-\tau_{Lt}^{*}}\right)\frac{\varepsilon_t^c}{1+\varepsilon_t^u}\left(1-\rho_{\theta s,t}\right) \tag{26}$$

While the labor wedge always has a positive redistributive or insurance value,[33] the net human capital wedge only has a positive redistributive value if and only if $(1-\rho_{\theta s})>0$. 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 $(1-\rho_{\theta s})>0$.

In the special case of a CES wage function as in (3) and a separable isoelastic disutility:

$$\phi\left(l,i\right)=\frac{1}{\gamma}l^{\gamma}+\frac{1}{\eta}i^{\eta}\quad\left(\gamma>1,\eta>1\right) \tag{27}$$

the ratio of the net human capital wedge and labor wedge is constant over time and across all agents, and equal to:

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

4.4 Age-dependency

This section focuses on the lifetime evolution of the optimal wedges. In formulas (20) and (25), the optimal wedges after each history were expressed recursively as a function of the previous period's wedges. Instead of these point-wise expressions, one can also rewrite the formulas in terms of a weighted expectation across types at time t, using some weighting function $\pi\left(\theta\right)$. Different weighting functions $\pi\left(\theta\right)$ lead to different recursive relations, which must hold at the optimum, and some weighting functions draw out particularly enlightening effects.[34]

For the exposition only, ability is assumed to follow a log autoregressive process:

$$\log\left(\theta_t\right)=p\log\left(\theta_{t-1}\right)+\psi_t \tag{28}$$

where ψ_t has density $f^{\psi}\left(\psi|\theta_{t-1}\right)$, with $E\left(\psi|\theta_{t-1}\right)=0$. The general formula without this assumption is in the Appendix (formula (56)):

[33] As long as $\tau_{Lt}\geq 0$, i.e., $\mu\left(\theta^t\right)\geq 0$, which is true under assumption (2).

[34] Such a reformulation for the optimal labor wedge formula was proposed by Farhi and Werning (2013) for a convenient weighting function $\pi\left(\theta\right)=1$.

Corollary 3 *The labor wedge evolves over time according to:*

$$E_{t-1}\left(\frac{\tau_{Lt}}{(1-\tau_{Lt})}\frac{\varepsilon_{w\theta,t-1}}{\varepsilon_{w\theta,t}}\frac{\varepsilon_t^c}{1+\varepsilon_t^u}\frac{1+\varepsilon_{t-1}^u}{\varepsilon_{t-1}^c}\left(\frac{1}{R\beta}\frac{u'_{t-1}}{u'_t}\right)\right)$$
$$=\varepsilon_{w\theta,t-1}\frac{1+\varepsilon_{t-1}^u}{\varepsilon_{t-1}^c}Cov\left(\frac{1}{R\beta}\frac{u'_{t-1}}{u'_t},\log\left(\theta_t\right)\right)+p\frac{\tau_{Lt-1}}{(1-\tau_{Lt-1})} \quad (29)$$

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 scaled down by the efficiency cost of the labor wedge, $\varepsilon_t^c/\left(1+\varepsilon_t^u\right)$, which makes provision of insurance more expensive, and magnified by the sensitivity of the wage to stochastic ability $\varepsilon_{w\theta,t}$, which increases the value of insurance. The persistence of the shock p translates into a persistence for the labor wedge. But the relation is not one-for-one: the autocorrelation of the labor wedge is modulated by the growth in efficiency costs and changes in the sensitivity of the wage to ability over time. Age-varying labor supply elasticities are a commonly stated argument in favor of age-dependent taxation. A time-varying elasticity of the wage to ability can provide a further rationale for age-contingent taxes.

A similar recursive formulation can be derived for the lifetime evolution of the net human capital wedge. Because of its tight link with the labor wedge (as explained in formula (26)), many of the same forces driving the evolution of the labor wedge are also important for the net subsidy.

Corollary 4 *The optimal net subsidy evolves over time according to:*

$$E_{t-1}\left(t_{st}\frac{\varepsilon_{w\theta,t-1}}{\varepsilon_{w\theta,t}}\frac{\left(1-\rho_{\theta s,t-1}\right)}{\left(1-\rho_{\theta s,t}\right)}\left(\frac{1}{R\beta}\frac{u'_{t-1}}{u'_t}\right)\right)$$
$$=\varepsilon_{w\theta,t-1}\left(1-\rho_{\theta s,t-1}\right)Cov\left(\frac{1}{R\beta}\frac{u'_{t-1}}{u'_t},\log\left(\theta_t\right)\right)+pt_{st-1} \quad (30)$$

Formula (30) again exhibits a drift term capturing insurance concerns, 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 simulations in section 5 show that, indeed, the net subsidy rises with age when human capital has a positive insurance value ($\rho_{\theta s}\leq 1$), and falls when it has a negative insurance value ($\rho_{\theta s}\geq 1$). They also highlight a "subsidy smoothing" result: 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.

The Hicksian complementarity $\rho_{\theta s}$ might vary over life, as suggested by the empirical literature in section 5.1. If it is decreasing faster, the net subsidy will rise faster or fall slower over the lifecycle.

The sensitivity of the wage amplifies the effect of the Hicksian complementarity coefficient in either direction.

Finally, note that all results from Propositions 1, 2, and 3, and Corollaries 1, 2, 3, and 4 still apply with a moving support $\left[\underline{\theta}_t\left(\theta_{t-1}\right), \bar{\theta}_t\left(\theta_{t-1}\right)\right] \subseteq \Theta$, except the zero distortions at the top and bottom results for τ_L and t_s.[35]

4.5 Optimal Bonus for Training Time

The foregoing results described the optimal subsidy for human capital expenses. Should training time be treated differently than expenses? This section analyzes the optimal bonus for training time, and focuses on the peculiarities of training relative to expenses, namely its direct interaction with labor supply.

As in subsection 4.2.1, a net wedge, the "net bonus," is defined. It measures the deviation away from a tax system that is neutral with respect to training time. However, the cost being deducted is agents' disutility cost, converted into a monetary cost ($\phi_{z,t}/u'$), and again accounting for the dynamic nature of investment.

Definition 2 *The net wedge or net bonus on training time t_{zt} is defined as:*

$$t_{zt} \equiv \frac{\tau_{Zt} - \tau_{Lt}\left(\frac{\phi_{z,t}}{u_t'(c_t)}\right)^d + P_{Zt}}{\left(\left(\frac{\phi_{z,t}}{u_t'(c_t)}\right)^d - \tau_{Zt}\right)\left(1-\tau_{Lt}\right)} \tag{31}$$

$\left(\frac{\phi_{z,t}}{u_t'(c_t)}\right)^d \equiv \frac{\phi_{z,t}}{u_t'(c_t)} - \frac{1}{R(1-\tau_K)}\left(1-\xi_{\phi'/u',t+1}\right)E_t\left(\frac{\phi_{z,t+1}}{u_{t+1}'(c_{t+1})}\right)$ *is the dynamic risk-adjusted disutility cost, converted into monetary units.*

$P_{Zt} \equiv \frac{\tau_K}{R(1-\tau_K)}\left(1-\xi'_{\phi'/u',t+1}\right)\left(1-\tau_{Lt}\right)E_t\left(\frac{\phi'_{z,t+1}}{u_{t+1}'(c_{t+1})}\right)$ *is the risk-adjusted savings distortion.*

The following proposition characterizes the net bonus at the optimum.

Proposition 4 *At the optimum, the net bonus is given by:*

$$t_{zt}^*\left(\theta^t\right) = \frac{\tau_{Lt}^*\left(\theta^t\right)}{1-\tau_{Lt}^*\left(\theta^t\right)}\frac{\varepsilon_t^c}{1+\varepsilon_t^u}\left(1-\rho_{\theta z,t}-\frac{\varepsilon_{\phi z,t}}{\varepsilon_{wz,t}}\rho_{lz,t}^{\phi}\right)+\frac{1}{R}E_t\left(\frac{\tau_{Lt+1}^*\left(\theta^{t+1}\right)}{1-\tau_{Lt}^*\left(\theta^t\right)}\frac{\varepsilon_{t+1}^c}{1+\varepsilon_{t+1}^u}\frac{w_{z,t+1}}{w_{z,t}}\frac{l_{t+1}}{l_t}\frac{\varepsilon_{\phi z,t+1}}{\varepsilon_{wz,t+1}}\rho_{lz,t+1}^{\phi}\right) \tag{32}$$

with $\varepsilon_{\phi z,t} = d\log\left(\phi_t\right)/d\log\left(z_t\right)$ the elasticity of disutility with respect to training time. $\rho_{lz,t}^{\phi}$ is the Hicksian coefficient of complementarity between l and z in the disutility ϕ_t. $\tau_{Lt}^*\left(\theta^t\right)$ and $\tau_{Lt+1}^*\left(\theta^{t+1}\right)$ are at their optimal levels from (25).

[35] See Farhi and Werning (2013) for the analysis of the optimal tax with moving support, which can easily be extended here to the model with human capital.

First, the Labor Supply Effect and the Inequality Effect are again present, with exactly the same interpretation as in subsection 4.2.1. But the bonus on training time has an additional direct interaction effect with current and future labor supply through the disutility function.

A first natural conjecture is that contemporaneous training and labor supply are substitutes: time spent acquiring human capital cannot be spent working. In addition, "fatigue" effects can set in. In this case, the Hicksian complementarity between training and labor in the disutility function is positive ($\rho^{\phi}_{lz,t} > 0$), a case labeled "Learning-or-Doing." However, because $\phi(l, i)$ is a disutility cost – not purely an opportunity cost of time – it might also be that labor supply and training are complements at least over some range ($\rho^{\phi}_{lz,t} < 0$). For example, the guarantee of regular training, which gives workers a prospective for progress, might boost motivation and make labor effort seem less painful, a situation called "Learning-and-Doing."

In the formula, $\left(1 - \frac{\varepsilon_{\phi z}}{\varepsilon_{wz}}\rho^{\phi}_{lz,t}\right)$ is the total effect of the bonus on labor, i.e., the sum of the Labor Supply Effect, which simulates labor because of the higher wage, and either one of the Learning-and-Doing or Learning-or-Doing effects, which can increase or decrease work. The total effect of a training subsidy on contemporaneous labor is positive if and only if

$$\varepsilon_{wz} > \varepsilon_{\phi_l z} \quad \text{with } \varepsilon_{\phi_l z} \equiv \partial \log(\phi_l) / \partial \log(z)$$

i.e., if and only if the wage is more sensitive to training than the marginal disutility of work is.[36]

Even if training diverts time away from contemporaneous labor supply, the effects on future labor supply can motivate a positive net subsidy. Whatever the pattern of complementarity or substitutability between contemporaneous labor supply and training, the relation between current training and future labor supply is its mirror image. If the former are complements, the latter are substitutes, and vice versa, because investing in human capital today means having to invest less tomorrow to reach any given level.[37]

If there is Learning-or-Doing, the net bonus co-moves positively with the future income tax rate τ_{Lt+1}, but negatively with the current tax rate τ_{Lt}. Intuitively, if there is a higher contemporaneous wedge on labor, there already is an indirect stimulus to training because training is a substitute for labor; hence the need for an additional stimulus through a net bonus is reduced. However, a higher future labor wedge is a stimulus for training in the future, which is a substitute for training today. Hence, to stimulate training today, there is a need for a higher net bonus.

Corollary 5 illustrates two cases, among other possible ones.

[36] Note that: $\left(\varepsilon_{\phi z,t}/\varepsilon_{wz,t}\right)\rho^{\phi}_{lz,t} = \varepsilon_{\phi_l z,t}/\varepsilon_{wz,t}$.

[37] This is because only the flow $i_t = z_t - z_{t-1}$, and not the stock z_{t-1} enters the disutility. This reasoning is also only valid at interior solutions for training.

Corollary 5 *Under assumption 2:*

i) The net bonus is positive if $\left(1 - \frac{\varepsilon_{\phi z}}{\varepsilon_{wz}} \rho_{lz,t}^{\phi}\right) \geq \rho_{\theta z,t}$ *and* $\rho_{lz,t}^{\phi} \geq 0$.

ii) The net bonus is negative if $\left(1 - \frac{\varepsilon_{\phi z}}{\varepsilon_{wz}} \rho_{lz,t}^{\phi}\right) \leq \rho_{\theta z,t}$ *and* $\rho_{lz,t}^{\phi} \leq 0$.

The sufficient conditions in i) guarantee that, although training diverts time away from contemporaneous labor effort, it stimulates labor supply in the future and has a sufficiently large positive redistributive or insurance effect $(1 - \rho_{\theta z})$. On the other hand, the sufficient conditions in ii) imply that despite stimulating contemporaneous labor supply through Learning-and-Doing, training disproportionately benefits high productivity workers, and has a negative redistributive or insurance effect.

What this analysis also highlights is that time and money costs are not equivalent in the presence of asymmetric information and incentive problems on unobservable labor. Unlike in the full information case, disutility costs cannot simply be converted into monetary units and treated as equivalent to expenses. An exception case is when the disutility function is separable in labor and human capital $\left(\frac{\partial^2 \phi}{\partial z \partial l} = 0\right)$. Then, as corollary (6) shows, the forces governing the net bonus on training are the same as those driving the net subsidy on human capital expenses.

Corollary 6 *i) If* $\frac{\partial^2 \phi_t}{\partial z \partial l} = 0$ *and assumption 2 holds:*

$$t_{zt}^*\left(\theta^t\right) \geq 0 \Leftrightarrow \rho_{\theta z,t} \leq 1$$

ii) The following relations between the wedges hold at the optimum at every history θ^t*:*

$$t_{zt}^* = \frac{\tau_{Lt}^*}{1 - \tau_{Lt}^*} \frac{\varepsilon_t^c}{1 + \varepsilon_t^u} \left(1 - \rho_{\theta z,t}\right) \tag{33}$$

$$\frac{t_{zt}^*}{t_{st}^*} = \frac{\left(1 - \rho_{\theta z,t}\right)}{\left(1 - \rho_{\theta s,t}\right)} \tag{34}$$

ii) If in addition the wage function is CES with parameter ρ *as in (3):*

$$t_{st}^* = t_{zt}^* \tag{35}$$

First, when there no longer is a direct interaction with labor supply, the net stimulus to training, i.e., the net bonus, is positive if and only if training has a positive insurance effect $(\rho_{\theta z} \leq 1)$. Formula (33) shows that in this case, the training wedge and the labor tax evolve simply according to a type of inverse elasticity rule. Formula (34) highlights that both types of human capital need to be encouraged proportionally to their redistributive and insurance effects $(1 - \rho_{\theta z})$ and $(1 - \rho_{\theta s})$. With a CES wage function, these redistributive effects are the same: there is then no reason to stimulate one type of human capital more than the other. At the optimum, the net wedges in (35) fully equate the incentives to invest in both types of human capital.

5 Numerical Analysis

In the numerical analysis in this Section, I take a middle stand between a simple illustration of the qualitative features of the optimal mechanism and a more careful calibration with quantitative implications for the optimal wedges and their lifecycle patterns. The focus is on human capital expenses. Additional details of the computational procedure and the calibration, as well as results for alternative calibrations, are in the online Computational Appendix. Before presenting 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.

5.1 Empirical Evidence on the Complementarity between Human Capital and Ability

The formulas for the optimal net wedge and its evolution (see (19) and (30)) highlighted the importance of the complementarity between ability and human capital. The benchmark Mincer model of returns to schooling yields a log-linear relation between the wage and schooling, with homogeneous returns across ability levels. The Becker (1964) model, reformulated and developed by Card (1995a), allows for returns to education to vary with ability. When put to work empirically though, the goal is frequently to circumvent the problem of unobserved ability. Instrumental variables related to institutional changes for schooling (Angrist and Krueger, 1991, Card, 1995b), or matching between siblings or twins (see Ashenfelter and Krueger, 1994, Ashenfelter and Rouse, 1998) have been exploited. 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 $\rho_{\theta s} \leq 1$ for childhood investments.[38]

Importantly, however, the Hicksian complementarity can change over life, as suggested by the structural literature on human capital formation (see among others Cunha *et al.*, 2005, Cunha and Heckman, 2007, 2008), so that estimates for primary and secondary schooling could be of limited use for the analysis of higher education or job training. 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) (hereafter, CHLM, 2006) estimate that the return to one year of college is around 16% at the 5^{th} percentile of the math test scores distribution, as opposed to 26% at the 95^{th} 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 AFQT scores, which

[38] However, since schooling – as well as college and job training – builds the two stocks of human capital z and s, this finding could be explained by several configurations of the parameters $\rho_{\theta s}$ and $\rho_{\theta z}$.

might, all else equal, signal that they have a higher marginal return from it.[39] 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 *et al.* (2011) use a multiplicatively separable functional form for the wage in their structural model of time investments in human capital (implying $\rho_{\theta z} = 1$), which generates a lifecycle path of earnings that matches the data well.

5.2 Calibration

Calibration to US data: To calibrate the model, I construct a "baseline economy." The baseline economy has the same primitives as in Section 2, but no social planner and, hence, no optimal tax system. Instead, the linear labor taxes, capital taxes, and human capital subsidies are set to their current averages in the US. Public and private subsidies in the US cover around 50% 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% for the first 2 periods and to 0 thereafter.[40] The linear labor tax rate is set to 13%, and the savings tax to 25%.[41] Agents can borrow and save at a constant interest rate, subject to a debt limit. In this baseline economy, some parameters are set exogenously based on the existing literature, while others are set to endogenously match two key moments from the data, namely, a wage premium and a ratio of 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 life are given by:

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

During retirement, utility is simply $\tilde{u}^R\left(c_t\right) = \log\left(c_t\right)$. The aforementioned literature seemed to indicate that $\rho_{\theta s} > 0$. Given the uncertainty about the magnitude of this parameter though, one of the goals in the simulations will be to explore different values, and, in particular, how the results depend on whether it is above or below its critical threshold of 1. To this end, the wage is assumed to be CES

$$w_t\left(\theta, s\right) = \left(\theta^{1-\rho} + c_s s^{1-\rho}\right)^{\frac{1}{1-\rho}} \tag{36}$$

[39] It is also not evident that the test scores used as measures of "ability" are themselves exogenous, especially at later ages.

[40] See indicators B, in OECD "Education At A Glance, 2013."

[41] 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% (plus possibly 2.9% for the new Medicare tax above $250K of total income, and various state taxes).

with a constant parameter $\rho_{\theta s} = \rho$, where c_s 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 Computational Appendix.[42]

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

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

with c_l and c_a the linear and the adjustment components of cost. Consistent with the high persistence in earnings documented in Storesletten *et al.* (2004), ability is assumed to follow a geometric random walk:

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

with $\psi_t \sim N\left(-\frac{1}{2}\sigma_\psi^2, \sigma_\psi^2\right)$.

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%. The adjustment cost is normalized to $c_a = 2$. There is large variation across empirical estimates of the variance of productivity σ_ψ^2. A medium-range value of 0.0095 (Heathcote *et al.*, 2005, hereafter HSV 2005) is adopted. Although the qualitative features of the solution are unaffected for a wide range of parameters, the quantitative results are. Therefore, several alternative calibrations are explored in the online Computational Appendix.

Endogenously matched parameters: The scaling factor for human capital c_s and the linear cost parameter c_l are set to match two statistics in the data: a wage premium and a ratio of human capital expenses to income. One complication which 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 *et al.* (hereafter AKK, 1998) who find 42.7% full-time college equivalents, the top 42.7% 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% is set to match the wage premium for college estimated in the literature.[43] These estimates range from 1.58 in Murphy and Welch (1992), to between 1.66 and 1.73 for AKK (1998), and above 1.80 in Heathcote *et al.* (2010). The calibration targets a mid-range value of 1.7.

[42] Parameter ρ affects the scale of the wage. To make the two values of ρ most comparable, for $\rho > 1$, the CES form used is instead: $w_t(\theta, s) = \left(\theta^{\rho-1} + c_s s^{\rho-1}\right)^{\frac{1}{\rho-1}}$ with $\rho = 1.45$, which implies $\rho_{\theta s} = \left(\frac{2-\rho}{\rho-1}\right) = 1.2$.

[43] The middle 14.6% 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."

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%.[44] 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 2/3 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%. Table 1 summarize the resulting parameters.

Using the computed policy functions, a Monte Carlo simulation with 100,000 draws is performed for each value of $\rho_{\theta s}$. 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.

	Definition	**Sim 1**	**Sim 2**	**Source/Target**
Exogenously calibrated or normalized				
ρ	Hicksian complementarity	0.2	1.2	CHLM (2006)
κ	Disutility of work scale	1	1	
γ	Disutility elasticity	3	3	Chetty (2012)
σ_ψ^2	Variance of productivity	0.0095	0.0095	HSV (2005)
T	Working periods	20	20	
T_r	Retirement periods	10	10	
β	Discount factor	0.95	0.95	
R	Gross interest rate	1.053	1.053	
Endogenously calibrated in baseline economy				
c_s	Scale of HC in wage	0.09	0.1	Wage premium (AKK, 1998)
c_l	Linear cost	0.5	0.5	(OECD, 2013, US Dept. Educ, 2010)

TABLE 1

5.3 Results

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 13 (after 26 years), 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.[45] To make it most comparable to an explicit subsidy, it is presented as a fraction of the marginal cost, i.e., $\tau_{St}/C_t'(e_t)$.

The optimal gross and net wedges as functions of $\rho_{\theta s}$: Figure 1 presents the human capital

[44] Computed using data on attendance at different types of colleges and costs (Chang Wei, 2010 US Department of Education, OECD, 2013). See the detailed calculations in the online Computational Appendix. For an earlier period, Caucutt and Kumar (2003) find a slightly lower number of 12%.

[45] The net wedge cannot be innocuously normalized to zero; hence, when agents no longer invest, it is set to the right hand side in (19), a level that will not artificially induce agents to invest.

wedges. For this figure only, the focus should be on periods $t \leq 13$, as many agents no longer invest in human capital later in life, and wedges are normalized thereafter, as just explained. The theoretical results from Section 4 are illustrated here: 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% and grows to 19%; for $\rho_{\theta s} = 1.2$, it instead starts at -2% and grows to only 13%. 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 (30).

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.[46] 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 (see also Diamond and Mirrlees, 1971, Chamley, 1986, and Judd, 1985) unless the redistributive or insurance effects are very strong.

Because of this, the values chosen for the complementarity between human capital and ability and for the volatility of ability clearly matter. Moving further away from a multiplicatively separable wage, i.e., increasing the complementarity coefficient further away from 1 in either direction, 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. Unfortunately, the estimates of these parameters from the literature are surrounded by large uncertainty. Therefore, the online Computational Appendix provides alternative calibrations.

The production efficiency result also stems from 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*.[47]

In addition, absent from this model are positive spillovers from human capital on growth, political

[46] The gross wedge is correspondingly smaller than 50% 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.

[47] A related reason is that, although it is optimal to invest more in human capital early in life in the model, there is no sharp exogenous constraint which forces investments to occur exclusively during college years. In practice, it seems that the signaling value of a college degree makes it critical to acquire human capital in a very concentrated manner at the beginning of working life.

stability, or social cohesion, often studied in the literature, and pointed to by policy makers as reasons to subsidize education. Finally, the belief about $\rho_{\theta s}$ held by policy makers, and ultimately, society, could be very different from the current estimates of $\rho_{\theta s}$ in the literature. In particular, one way to rationalize the higher subsidies in the real world is that human capital is believed to have very strong redistributive and insurance values, and strongly benefit lower ability people. If the model is taken literally, a very negative $\rho_{\theta s}$ would be needed. Related to this, education is often perceived as a basic human right.

Figure 1: Average human capital wedges over time

(a) Gross Wedge τ_{St} (b) Net Wedge t_{st}

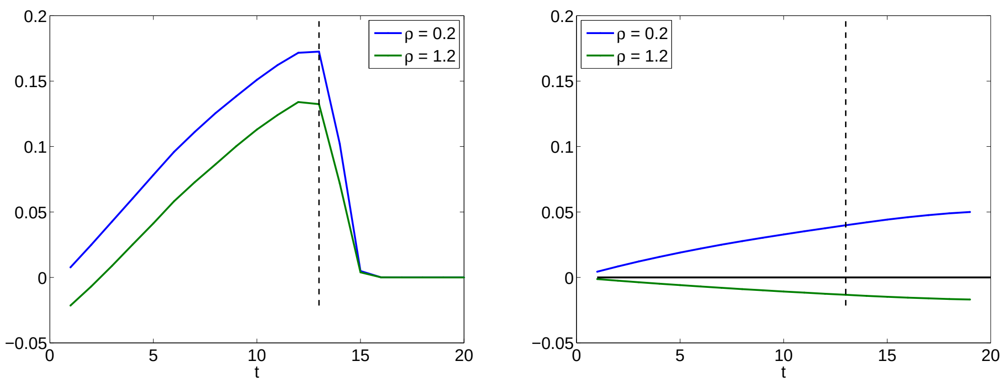

Dashed lines mark the time after which most agents no longer invest in human capital. (a) The gross wedge is higher and grows faster when human capital has positive redistributive and insurance effects ($\rho_{\theta s} < 1$). The wedge is normalized to zero for zero investment (a corner solution). (b) The provision of dynamic incentives also creates a value for insurance. 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.

Optimal labor wedges: Figure 2 panel (a) explains why the gross and net human capital wedges differ so starkly: the optimal labor wedge rises over time to provide insurance against widening income dispersion, and a large part of the growing human capital wedge merely goes towards compensating for this growing disincentive to accumulate human capital. The labor wedge is compared to the one in a standard dynamic taxation model without human capital, in which the wage is equal to exogenous ability, $w_t = \theta_t$. As is intuitive, the wedge grows slower in the presence of human capital, particularly when $\rho_{\theta s} \leq 1$. The labor wedge has a disincentive effect on human capital acquisition, which is undesirable, and the more so when human capital itself has positive insurance and redistributive effects ($\rho_{\theta s} \leq 1$).

Optimal capital wedges: Figure 2 panel (b) plots the capital wedge over time. It starts at 0.5% of the gross interest on savings, which corresponds to a 10% tax on net interest, and declines to zero.[48]

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

Figure 2: Average labor and capital wedges over time

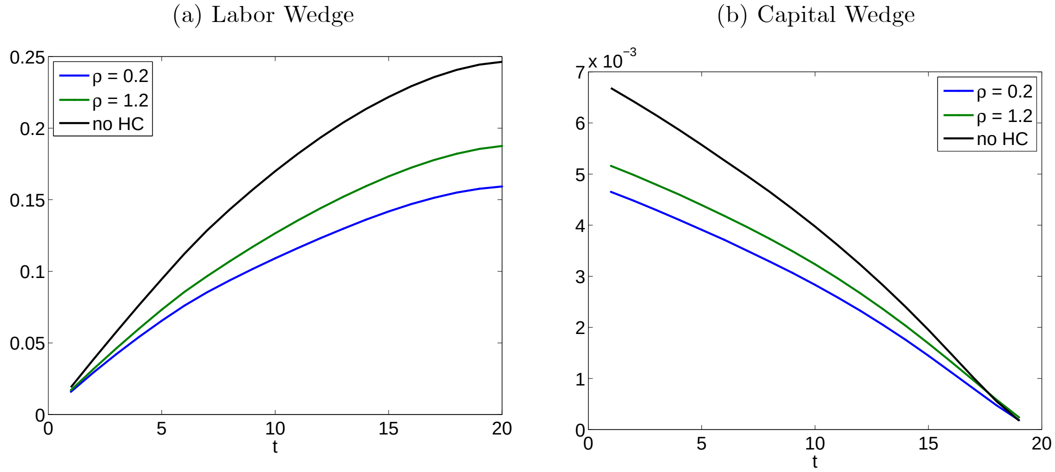

"No HC" denotes the case without human capital. (a) The labor wedge – which represents a disincentive for human capital investments but also insures agents against risky human capital returns – grows slower over time in the presence of human capital, the more so if human capital has positive redistributive and insurance effects ($\rho_{\theta s} < 1$). (b) The capital wedge is still positive, but lower in the presence of human capital.

The capital wedge arises at the optimum from the inverse Euler equation in (24), and is a standard feature in dynamic moral hazard models, in which savings are complementary to future shirking. However, savings are less distorted in the presence of human capital. One possible intuition for this is that human capital investments improve productivity in the future, which to some extent counters incentives to shirk. The capital wedge without human capital is 0.7% of gross interest, equivalent to 14% of net interest.

Subsidy Smoothing: 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 also applies in the presence of human capital (see the online Computational Appendix). 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 (see (30)), which is proportional to the covariance between ability and consumption growth, tends to zero towards retirement. Then, only the autoregressive term of the random walk remains.

Figure 3: Subsidy smoothing over life

(a) t_{st} against t_{st-1} at $t = 5$ (b) t_{st} against t_{st-1} at $t = 13$

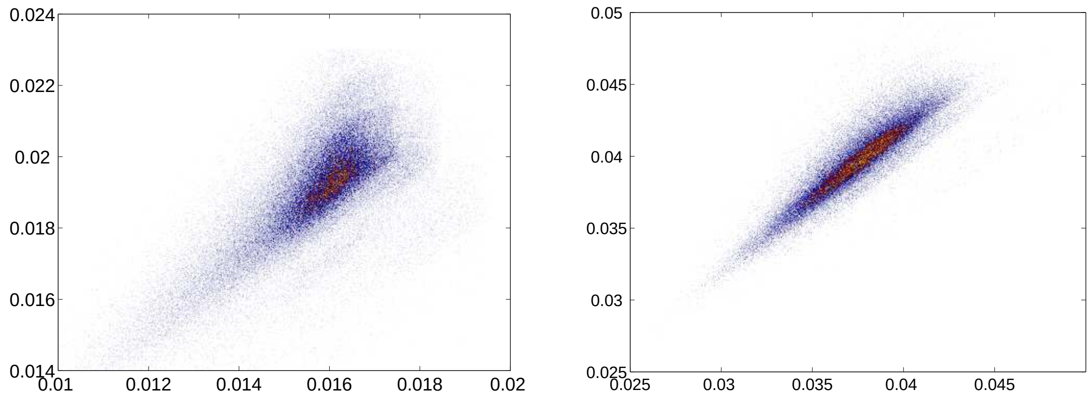

The net human capital wedge becomes more correlated from one period to the next as age increases, because the variance of consumption growth, which drives changes in the subsidy over time, vanishes. Figures are for $\rho_{\theta s} = 0.2$.

Allocations and Insurance: Figure 6 plots the average allocations over time. Average human capital investments are almost flat and highest early in life, before declining with age. Mean consumption is constant, a result due to the Inverse Euler equation in (24), and log utility with $\beta = \frac{1}{R}$, which imply that consumption is a martingale: $E_{t-1}(c_t) = c_{t-1}$. Mean output is increasing, despite the rising labor wedge, because of the growing productivity of agents driven by their endogenous human capital investments.

Figure 4 panel (a) shows that lifetime spending on human capital is more tightly linked to lifetime income when human capital disproportionately benefits high ability agents. The causality goes both ways: higher ability agents both acquire more human capital and earn more. In turn, human capital increases earnings even further. When $\rho_{\theta s} > 1$, both effects are amplified. However, the fact that higher ability people also acquire more human capital does not mean that there is no insurance. Panel (b) highlights the extent of lifecycle insurance at the optimum by plotting the net present value of lifetime spending (consumption plus human capital expenses) against the net present value of lifetime income. Clockwise pivots of the line represent higher insurance.

Finally, figure 5 describes the cross-sectional variances of output, human capital, consumption, and ability over time. The variance of output is driven not only by stochastic ability, but also by differential investments in human capital at different ability levels. Output is much more volatile than consumption. Hence, pre-tax income inequality grows at an increasing rate, but the provision of insurance prevents this from translating fully into consumption inequality. In addition, while consumption

Figure 4: Human Capital and insurance over the life cycle

(a) Human capital expenses and lifetime income (b) Consumption and education expenses

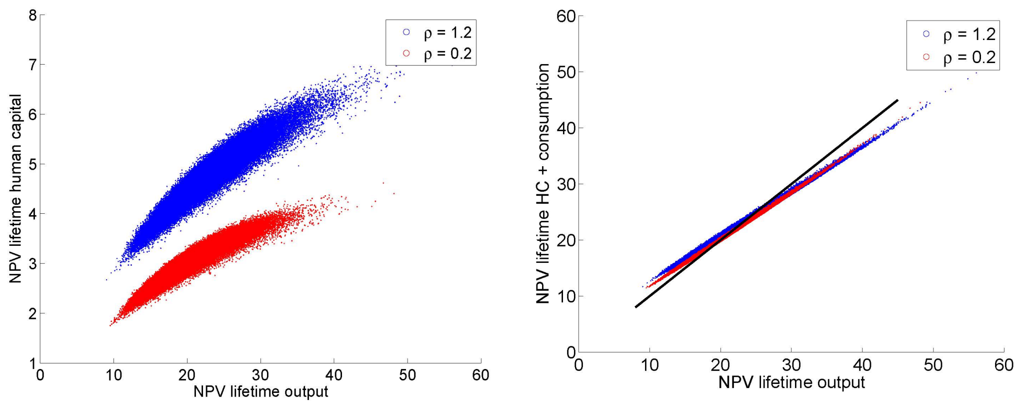

(a) Lifetime income is positively correlated with lifetime human capital expenses, the more so when $\rho_{\theta s} > 1$. There is a two-way causality: Higher ability people both acquire more human capital and have higher earnings potential. At the same time, human capital increases earnings.

(b) The figure shows the present value of consumption and human capital expenses against the present value of lifetime income. The laissez-faire outcome is represented by the 45 degree line. Clockwise pivots of the line represent more insurance.

Figure 5: Variance and Risk

(a) $\rho_{\theta s} = 0.2$ (b) $\rho_{\theta s} = 1.2$

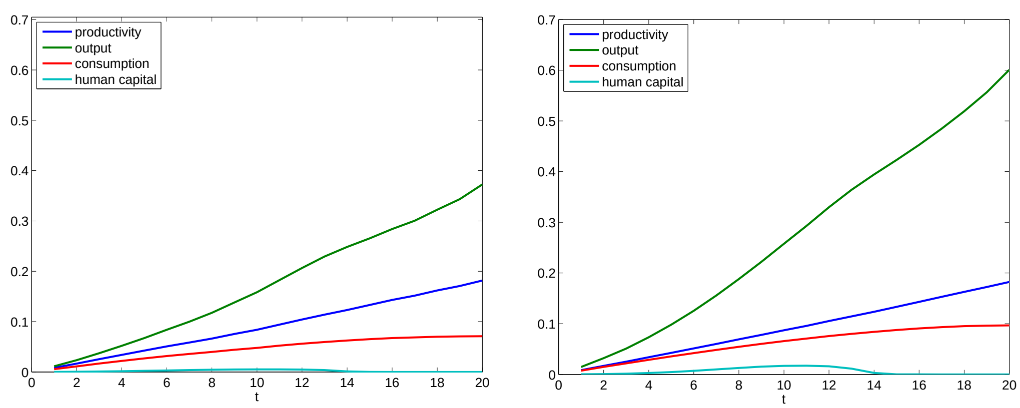

The figure shows cross-sectional variances over time. Output is more volatile than consumption. Its volatility grows at an increasing rate, driven both by ability shocks and differential investments in human capital. But pre-tax income inequality does not fully translate into consumption inequality. All outcomes are more volatile when human capital has a negative insurance value ($\rho_{\theta s} > 1$).

Figure 6: Allocations

(a) $\rho_{\theta s} = 0.2$ (b) $\rho_{\theta s} = 1.2$

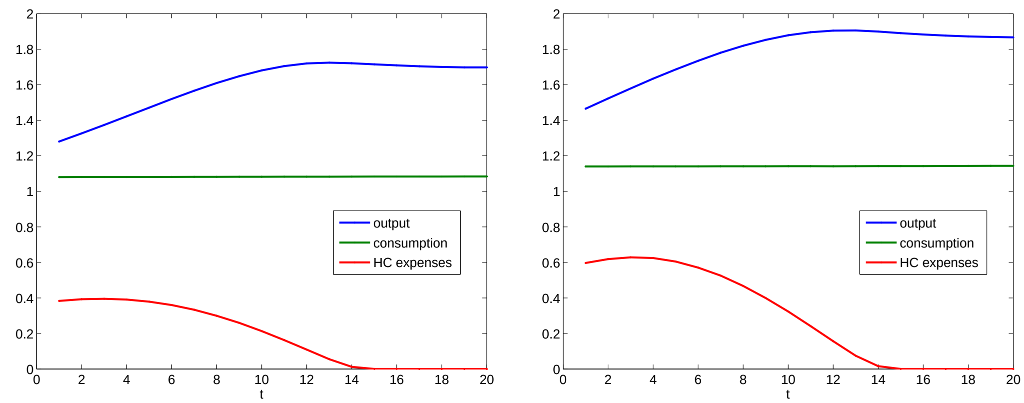

Agents optimally invest in human capital early in life. Consumption is a martingale; hence, average consumption is perfectly flat. Output and consumption are higher when human capital disproportionately benefits higher ability agents.

variance grows, it does so at a decreasing rate, echoing the tax and subsidy smoothing results described above.

Progressivity: Figure 7 show the implicit progressivity of the labor wedge, by plotting τ_{Lt} against the contemporaneous productivity shock, θ_t, at $t = 19$. 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 "modified inverse elasticity rule" in (26). 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 co-move positively when $\rho_{\theta s} < 1$ and negatively when $\rho_{\theta s} \geq 1$. The net subsidy is always regressive when $\rho_{\theta s} > 0$ because higher ability people benefit more from human capital (see figure 8 for $t = 13$). The labor wedge will hence exhibit the inverse relation when $\rho_{\theta s} \leq 1$ (which is equivalent to a progressive labor wedge). Despite any short-run regressivity, the system is progressive overall and does provide insurance, as was shown in figure 4.

6 Implementation

The wedges, as functions of histories of types, do not immediately translate into explicit taxes and subsidies on income, human capital or savings. This section considers the implementation of the optimal allocations using decentralized instruments. For clarity, the focus is on educational expenses only,

Figure 7: Progressivity and regressivity of the labor wedge

(a) τ_{Lt} against θ_t for $\rho_{\theta s} = 0.2$ (b) τ_{Lt} against θ_t for $\rho_{\theta s} = 1.2$

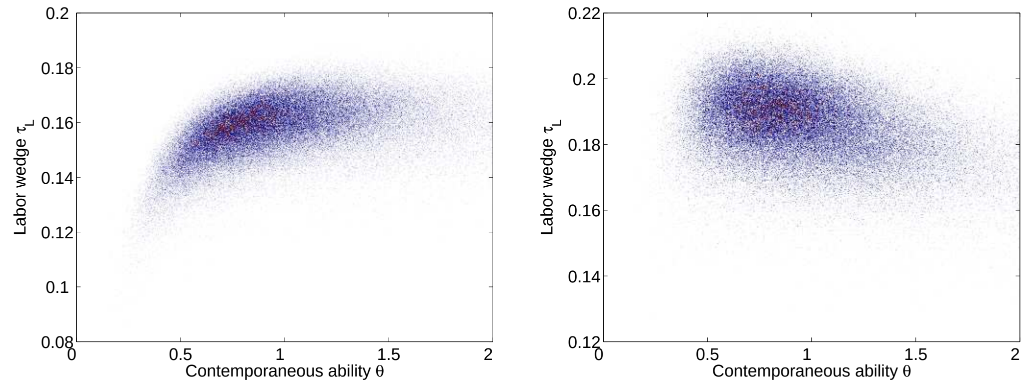

The labor wedge exhibits short-run progressivity when $\rho_{\theta s} < 1$, but short-run regressivity when $\rho_{\theta s} > 1$.

Figure 8: Regressivity of the net human capital wedge

(a) t_{st} against θ_t for $\rho_{\theta s} = 0.2$ (b) t_{st} against θ_t for $\rho_{\theta s} = 1.2$

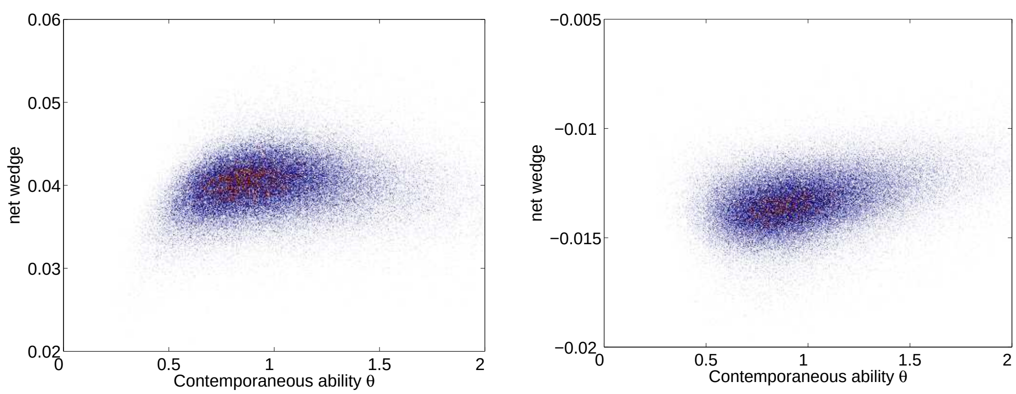

The net wedge is always regressive in the short run, but more so when $\rho_{\theta s} > 1$.

but observable training time can be added as an additional conditioning variable. The first implementation is through income contingent loans (hereafter, ICLs), which are compared and contrasted to certain types of ICLs currently used in several countries. The second is through a Deferred Deductibility scheme, which can be applied in the special case when shocks are independently and identically distributed.

6.1 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.[49] The government can observe and keep record of the histories of consumption, output, human capital, and wealth.

Denote by $m_t^*\left(\theta^t\right)$ the optimal allocation of the social planner's problem after history θ^t for any choice variable $m \in \{c, y, b, e\}$ (the Appendix shows how to construct them from the recursive allocations derived above). 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}, i.e., such that for some $\theta \in \Theta$, $\mathbf{m}_t = \mathbf{m}_t^*\left(\theta^{t-1}, \theta\right)$. For a history of observed choices $\mathbf{m}^t$, denote by $\Theta^t\left(\mathbf{m}^t\right)$ the set of all histories θ^t consistent with these choices, i.e., all θ^t such that $\mathbf{m}_s = \mathbf{m}_s^*\left(\theta^s\right)$ for all $s \leq t$. Assumption 3 guarantees that in the planner's problem, the histories $\left(y^t, e^t\right)$ can be uniquely inverted to identify the history of abilities, θ^t.

Assumption 3 *$\Theta^t\left(y^t, e^t\right)$ is either the empty set or a singleton for all histories $\left(y^t, e^t\right)$.*

In the proposed ICL scheme, loans are combined with a standard income tax based on contemporaneous income $T_Y\left(y_t\right)$, and a history-independent savings tax $T_K\left(b_t\right)$. In each period, the agent is offered a government loan $L_t\left(e_t\right)$ as a function of his human capital expenses, and is required to make a history contingent repayment $D_t\left(L^{t-1}, y^{t-1}, e_t, y_t\right)$, as a function of the full history of past loans and earnings, as well as current income, and human capital expenses. The agent's problem is then to select the supremum over $\left\{c_t\left(\theta^t\right), y_t\left(\theta^t\right), b_t\left(\theta^t\right), e_t\left(\theta^t\right)\right\}_{\theta^t}$ in:

[49] 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.

$$
\begin{aligned}
V_1\left(b_0,\theta_0\right) &= \sup\sum_{t=1}^{T}\int\left[u_t\left(c_t\left(\theta^t\right)\right)-\phi_t\left(\frac{y_t\left(\theta^t\right)}{w_t\left(\theta_t,s_{t-1}\left(\theta^{t-1}\right)+e_t\left(\theta^t\right)\right)}\right)\right]P\left(\theta^t\right)d\theta^t \qquad (37)\\
\text{s.t:} \quad & c_t\left(\theta^t\right)+\frac{1}{R}b_t\left(\theta^t\right)+M_t\left(e_t\left(\theta^t\right)\right)-b_{t-1}\left(\theta^{t-1}\right)-L_t\left(e_t\left(\theta^t\right)\right)\\
&\leq y_t\left(\theta^t\right)-D_t\left(L^{t-1}\left(\theta^{t-1}\right),y^{t-1}\left(\theta^{t-1}\right),e_t\left(\theta^t\right),y_t\left(\theta^t\right)\right)-T_Y\left(y_t\left(\theta^t\right)\right)-T_K\left(b_t\left(\theta^t\right)\right)\\
& s_t\left(\theta^t\right)=s_{t-1}\left(\theta^{t-1}\right)+e_t\left(\theta^t\right),\ \ s_0\text{ given, }e_t\left(\theta^t\right)\geq 0, b_0=0,\ b_T\geq 0
\end{aligned}
$$

The construction of the ICL schedule, explained formally in the Appendix, is intuitively as follows. First, the loan is set to exactly cover the cost of human capital:

$$
L_t\left(e_t\right)=M_t\left(e_t\right)\quad \forall t,\,\forall e_t \qquad (38)
$$

The savings tax $T_K\left(b_t\right)$ is constructed to guarantee zero private wealth holdings.[50] The repayment schedule D and income tax T_Y are such that, along the equilibrium path, the optimal allocations from the social planner's problem are affordable for each agent after all histories, given zero asset holdings:

$$
D_t\left(L^{t-1},y^{t-1},e_t^*\left(\theta^{t-1},\theta\right),y_t^*\left(\theta^{t-1},\theta\right)\right)+T_Y\left(y_t^*\left(\theta^{t-1},\theta\right)\right)=y_t^*\left(\theta^{t-1},\theta\right)-c_t^*\left(\theta^{t-1},\theta\right)
$$

for all $\left(L^{t-1},y^{t-1}\right)$ such that $\theta^{t-1}\in\Theta^{t-1}\left(\left\{M_1^{-1}\left(L_1\right),...,M_{t-1}^{-1}\left(L_{t-1}\right)\right\},y^{t-1}\right)\neq\emptyset$, and all $\theta\in\Theta$, where the history of education e^{t-1} is inverted from L^{t-1} using (38). The repayment schedule on off-equilibrium allocations – those allocations which 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 only face 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 $\left(L^{t-1},y^{t-1}\right)$, assigns an allocation:

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

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

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

[50] The construction in the Appendix builds on Werning (2007), who shows also that the savings tax can be redefined to implement non zero 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 non-zero private wealth holdings.

where $\theta^{t-1} \in \Theta^{t-1}\left(\left\{M_1^{-1}\left(L_1\right), ..., M_{t-1}^{-1}\left(L_{t-1}\right)\right\}, y^{t-1}\right)$ is unique by assumption 3. The decentralization rule is said to implement the optimum from the planner's problem for a given set of promised utilities $\left(\underline{U}\left(\theta\right)\right)_{\Theta}$ if, for all t and θ^t, the decentralized allocations under this rule coincide with the social planner's optimal allocations, i.e., $\hat{m}_t\left(\theta^t\right) = m_t^*\left(\theta^t\right)$ for $m \in \{c, y, b, e\}$.[51]

Proposition 5 *The optimum can be implemented through human capital loans $L_t\left(e_t\right)$, with repayments $D_t\left(L^{t-1}, y^{t-1}, e_t, y_t\right)$, contingent on the history of loans and earnings, current income, and human capital expenses, together with a history-independent savings tax $T_K\left(b_t\right)$, and an income tax on contemporaneous income $T_Y\left(y_t\right)$.*

Income Contingent Loans: Figure 9 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 repayment rises as a fraction of income until late in life, illustrating the insurance provided by the contingent repayment schedule and the increasing ability to pay over life, driven by human capital investments.

Figure 10 further highlights the insurance role of the repayment schedule, by showing a snapshot of repayment, as a fraction of income, against the contemporaneous realization of the productivity shock, at an arbitrarily chosen time ($t = 15$). Repayments increase on average in higher productivity states and decrease in lower productivity ones. The income-history contingent nature of repayments is clearly seen in their large dispersion at a given θ_t: repayments depend on the full past, not only on current productivity. Indeed, a positive productivity shock is correlated with higher repayments over several future periods.

6.1.1 Comparison of the Proposed Implementation to Existing ICLs

Certain types of ICLs for college education are used in several countries, including the US, New Zealand, Australia, the UK, 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, 2005, for an empirical overview).[52]

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

[52] In the US, 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 (ICRs) for public sector jobs. In 1997, the College Cost Reducation and Access Act (CCRAA), introduced Income Based Repayment (IBR) 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.

Figure 9: Income history contingent loans

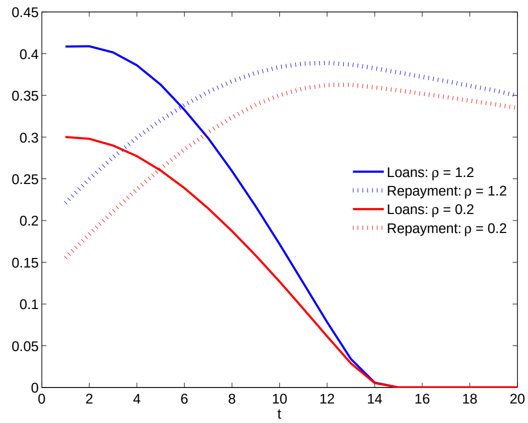

Loans and Repayments as a fraction of contemporaneous income.
Loans are high early in life, while repayments increase to provide insurance.

Figure 10: Insurance through contingent repayments

(a) repayment against θ_t for $\rho_{\theta s} = 0.2$

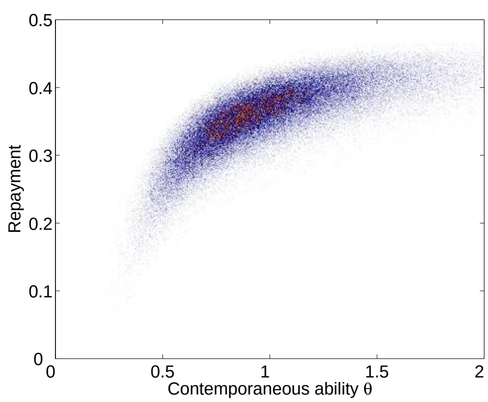

(b) repayment against θ_t for $\rho_{\theta s} = 1.2$

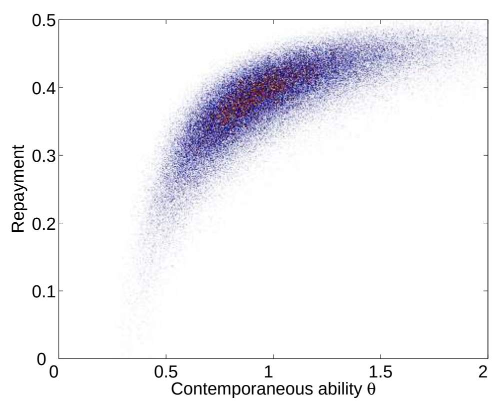

The repayment schedule provides insurance: repayments are higher when productivity realizations are higher. The history-contingent nature is seen in the dispersion of the repayment at a given contemporaneous ability realization.

The loans sometimes depend on the level of education acquired, the type of degree or field, and are indexed to the costs of education, which mirrors the proposed loan above. In several countries, such as Australia or New Zealand, repayments are directly collected through the tax system, as in the optimal integrated system suggested. The coercive tax power of the government is required, together with full-commitment to announced policies, to prevent agents from dropping out ex post after the realization of their incomes. 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. Because more successful earners would ex post face higher repayments, and hence cross-subsidize less successful ones, the plan suffered from a typical adverse selection problem: students with the best earnings prospects did not join or dropped out (Palacios, 2004).

There are four main differences between the ICL proposed here and existing ones. First, in the real world, the total present value of repayments is closely linked to the amount borrowed, except for subsidized interest rates or exceptional loan forgiveness. In the model, 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 optimal system. Second, in ICL schemes observed in practice, the focus is almost exclusively on the downside, so that repayments can be deferred or forgiven in times of economic hardship. The optimal scheme is also focused on the upside, with repayments potentially increasing after a good history of earnings. Third, 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. 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.[53] There is very little history-dependence in real-world ICLs, possibly with the exception of Sweden, which uses a two-year averaging of earnings to determine repayments. Interestingly, there are other real-world policies which exhibit exactly the kind of history dependence which is required, 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.

6.2 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 independently and identically distributed (iid) shocks, wealth and the

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

starting stock of human capital each period can serve as sufficient statistics for the full past. This yields history-independent policies, explored in this subsection. A similar implementation for physical capital, in the absence of human capital, is studied in Albanesi and Sleet (2006).

The recursive problem with iid shocks is nested in the formulation in section 2, 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 $K(v_{t-1}, s_{t-1}, t)$.

The decentralization rule with iid shocks and taxes: For this implementation, interpret the initial ability θ_1 as uncertainty, like all other shocks θ_t, rather than intrinsic heterogeneity.[54] 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.[55] 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 $\left(\hat{c}_t, \hat{y}_t, \hat{b}_t, \hat{e}_t\right)$ to each agent type, following some mappings from observed initial wealth b_{t-1} and human capital s_{t-1}:

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

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 $\left(\hat{c}_t, \hat{y}_t, \hat{b}_t, \hat{e}_t\right)$, 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, $\left(\hat{c}_t, \hat{y}_t, \hat{b}_t, \hat{e}_t\right)$ attains the supremum in the agent's problem in (39) :

$$\begin{aligned} V_t(b, s) \quad &= \quad \sup_{c, y, b', e} \int \left(u_t(c(\theta)) - \phi_t\left(\frac{y(\theta)}{w(\theta, s + e(\theta))}\right) + \beta V_{t+1}(b'(\theta), s + e(\theta)) \right) f(\theta)\, d\theta \qquad (39) \\ \text{s.t:} \quad & \quad c(\theta) + M_t(e(\theta)) + \frac{1}{R} b'(\theta) = y(\theta) - T_t(b, s, y(\theta), e(\theta)) + b \quad \forall \theta \\ c, y, e \quad &: \quad \Theta \to \mathbb{R}_+ \text{ and } b' : \Theta \to B_t = [\underline{b}_t, \infty) \text{ with } V_{T+1} \equiv 0,\ b_T \equiv 0. \end{aligned}$$

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 $V(b_0, s_0) = U_1$.

[54] This incidentally makes the model directly comparable to Albanesi and Sleet (2006). Unlike Atkeson and Lucas (1995) and Albanesi and Sleet (2006), the analysis is in partial equilibrium analysis, with agents, as well as the government facing a constant gross rate R.

[55] This distribution of initial wealth need not be degenerate if there is a non-degenerate distribution of initial utility promises, from which we abstract here.

Proposition 6 *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 $T_t(b_{t-1}, s_{t-1}, y_t, e_t)$ 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. That is, for all (v_{t-1}, s_{t-1}), there is an associated wealth level $b_{t-1} = K(v_{t-1}, s_{t-1}, t)$. Taxes are designed such that the problem of an agent starting with wealth level $b_{t-1} = K(v_{t-1}, s_{t-1}, t)$ and human capital s_{t-1} is the dual of the planner's problem with promised utility v_{t-1}, facing an agent with human capital s_{t-1}.

This recursive implementation allows a relatively simple map between the derivatives of the proposed tax function T_t and the optimal wedges. Let $\tau_{Lt}(v_{t-1}, s_{t-1}, \theta_t)$ and $\tau_{St}(v_{t-1}, s_{t-1}, \theta_t)$ stand for the labor wedge and the gross human capital wedge (as defined in (15) and (17), respectively) evaluated at the optimal allocation for $(v_{t-1}, s_{t-1}, \theta_t)$. The labor wedge is immediately linked to the marginal income tax through:

$$\tau_{Lt}(v_{t-1}, s_{t-1}, \theta_t) = \frac{\partial T_t(K(v_{t-1}, s_{t-1}, t), s_{t-1}, y_t^*, e_t^*)}{\partial y_t}$$

with $y_t^* = y_t^*(v_{t-1}, s_{t-1}, \theta_t)$, and $e_t^* = e_t^*(v_{t-1}, s_{t-1}, \theta_t)$ the optimal allocations from the planner's problem. The human capital wedge is linked to the marginal subsidy more indirectly:

$$\tau_{St}(v_{t-1}, s_{t-1}, \theta_t) = \frac{-\partial T_t(K(v_{t-1}, s_{t-1}, t), s_{t-1}, y_t^*, e_t^*)}{\partial e_t} + \beta E_t\left(\frac{u'_{t+1}(c_{t+1}^*)}{u'_t(c_t^*)}\left(\frac{\partial T_{t+1}}{\partial e_{t+1}} - \frac{\partial T_{t+1}}{\partial s_t}\right)\right)$$

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 $T_{t+1} = T_{t+1}(b_t, s_t, y_{t+1}, e_{t+1})$ is evaluated at:

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

The wedge is not in general equal to the marginal subsidy, that is, the expected reduction in tax from an incremental investment in human capital. A positive wedge does also not necessarily imply a positive marginal subsidy. This can be made clear by rewriting the wedge as:

$$\tau_{St} = -\left(\frac{\partial T_t}{\partial e_t}\right) + E_t\left(\beta\frac{u'_{t+1}}{u'_t}\right) E_t\left(\frac{\partial T_{t+1}}{\partial e_{t+1}} - \frac{\partial T_{t+1}}{\partial s_t}\right) + Cov\left(\beta\frac{u'_{t+1}}{u'_t}, \frac{\partial T_{t+1}}{\partial e_{t+1}} - \frac{\partial T_{t+1}}{\partial s_t}\right)$$

A positive wedge on human capital 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 $\left(\frac{\partial T_{t+1}}{\partial e_{t+1}} - \frac{\partial T_{t+1}}{\partial s_t}\right)$ 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.[56]

Means-tested Human Capital Grants with iid shocks: As an immediate corollary of Proposition 6, the tax system can instead be reformulated as a means-tested grant. In period t, the agent's assets and stock of human capital are verified, and he receives a grant such that:

$$G_t\left(y_t, e_t | b_{t-1}, s_{t-1}\right) = -T_t\left(b_{t-1}, s_{t-1}, y_t, e_t\right)$$

Means-testing based only on contemporaneous assets and income is hence optimal if shocks are iid. It is interesting that, although formally just a reformulation of the tax from Proposition 6, means-tested grants for higher education are very common in many countries, while wealth contingent income taxes are not. In the US for instance, Pell grants take assets as well as contemporaneous income into account.[57]

ICLs with iid shocks: The implementation through ICLs from subsection 6.1 can be modified to use wealth instead of the full history of earnings and loans to determine the repayments (hence abandoning the savings tax proposed above). In particular, the loans are again equal to the cost of human capital acquisition in each period, $L_t\left(e_t\right) = M_t\left(e_t\right)$, and the wealth and income contingent repayment schedules $D_t\left(b_{t-1}, s_{t-1}, y_t, e_t\right)$, are such that:

$$T_Y\left(y_t\right) - L_t\left(e_t\right) + D_t\left(b_{t-1}, s_{t-1}, y_t, e_t\right) = T_t\left(b_{t-1}, s_{t-1}, y_t, e_t\right)$$

This implementation is closer in spirit to the means-tested college loans, where the amount of funds provided is conditional on the contemporaneous resources of the child or her parents, and potentially on the highest education already attained, rather than on the history of earnings.

This subsection highlights that, while the optimal allocation derived in the direct revelation mechanism in section 4 is generally unique, there are many possible combinations of subsidies, loans, grants, and income taxes that can implement it.[58] If grants, repayments and subsidies were set at the appropriate levels, the difference between a grant-intensive and heavily subsidized education system such as

[56] See Kocherlakota (2005) for the case of savings. In Kocherlakota (2005), 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.

[57] Again, these grants are typically limited to tertiary education only. Grants for job training exist for the unemployed (hence, somewhat means-tested), for youth at risk ("YouthBuild" in the US), or for difficult to employ seniors (the "Senior Aide Program"). They are 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."

[58] On the tax incentives for higher education in the US, see for instance Hoxby (1998).

in Continental Europe, and a loan-intensive, high-tuition system as in the US might be more apparent than real.[59] The real focus should be on the parameters entering the optimal formulas, such as the Hicksian complementarity $\rho_{\theta s}$, which determines to what extent human capital should be subsidized on net. The exact mix of instruments used is more a matter of administrative capabilities and infrastructure, which might very well be country-specific.

6.3 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, Blomquist, 1982, Bovenberg and Jacobs, 2005). It has been 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.

Assume first 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 wedges, which is the case most analogous to full deductibility. In this case, equation (23) must hold at the optimum. Accordingly, if the proposed tax system $T_t(b_{t-1}, s_{t-1}, y_t, e_t)$ is differentiable in all its arguments (at least over the range of equilibrium path values), then it must satisfy:[60]

$$-\frac{\partial T_t}{\partial e_t} + \beta E_t\left(\frac{u'_{t+1}}{u'_t}\left(\frac{\partial T_{t+1}}{\partial e_{t+1}} - \frac{\partial T_{t+1}}{\partial s_t}\right)\right) = \frac{\partial T_t}{\partial y_t}\left(M'_t - \frac{1}{R}E_t\left(M'_{t+1}\right)\right) - \frac{1}{R}\left(1-\xi'_{M',t+1}\right)E_t\left(\beta R\frac{u'_{t+1}}{u'_t}\frac{\partial T_{t+1}}{\partial b_t}\right)E_t\left(M'_{t+1}\right) \quad (40)$$

On the other hand, the pure contemporaneous deductibility scheme proposed by Bovenberg and Jacobs (2005) would imply that at all dates $\frac{\partial T_t}{\partial y_t} = M'_t(e_t)\frac{\partial T_t}{\partial e_t}$, $\forall t$.[61] In this dynamic model, this is only true for the last period of investment, T, in which agents face a static problem. In all earlier periods, first, the true cost is not simply the static cost M'_t, but instead the dynamic expected cost $M'_t - \frac{1}{R}E_t\left(M'_{t+1}\right)$. Second, there is uncertainty and risk aversion; as already mentioned above, a positive wedge can arise if the tax burden on human capital is lower in states with high marginal utility of consumption. Third, there is also a (stochastic) distortion to physical capital accumulation,

[59] Given the discrepancies in the net cost of education in Continental Europe and the US, it is clear that the financing systems are far from equivalent at their current levels. The argument proposed here is that they *could* be made equivalent while still preserving their general structure (loans-based vs. subsidy and grants-based).

[60] Obtained by applying formula (23) to this tax system. See Appendix formula (58) for a rewriting.

[61] The discrepancy in this recommendation to 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 $\rho_{\theta s} = 1$. It arises instead from the dynamics and risk.

which needs to be taken into account. The (risk-adjusted) human capital subsidy and the (risk-adjusted) intertemporal wedge need to co-move inversely. This is because, if the tax on physical capital increases, there is a substitution towards human capital, and a lower subsidy is needed to stimulate it. Finally, 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.

The right way to implement full risk-adjusted dynamic deductibility, which is the optimal policy when $\rho_{\theta s} = 1$, is a risk-adjusted deferred deductibility scheme. For the sake of the exposition, assume that $\beta = \frac{1}{R}$ and $M_t'(e_t) = 1$. Start from period T, in which a simple deductibility of expenses is sufficient with $-\frac{\partial T_T}{\partial e_T} = \frac{\partial T_T}{\partial y_T}$, and work backwards in order to rewrite the total change in the tax burden from an incremental investment at t as:

$$-\frac{\partial T_t}{\partial e_t} = (1-\beta)\sum_{j=1}^{T-t}\beta^{j-1}E_t\left(\frac{u'_{t+j-1}}{u'_t}\frac{\partial T_{t+j-1}}{\partial y_{t+j-1}}\right) + \beta^{T-t}E_t\left(\frac{u'_T}{u'_t}\left(\frac{\partial T_T}{\partial y_T}\right)\right) - \sum_{j=1}^{T-t}\beta^{j}E_t\left(\frac{u'_{t+j}}{u'_t}\left(\frac{\partial T_{t+j}}{\partial b_{t+j-1}} - \frac{\partial T_{t+j}}{\partial s_{t+j-1}}\right)\right) \quad (41)$$

The optimal subsidy is hence equivalent to a deferred deductibility scheme, in which a fraction $(1-\beta)$ of the human capital expense of time t are deducted from taxable income in each subsequent period. Intuitively, with changing income tax rates $\frac{\partial T_t}{\partial y_t}$, a non-dynamic deductibility scheme would mean that the expense 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 (41)), which takes into account the relative shift in the future tax schedule from more physical capital stock versus more human capital stock. Since physical and human capital are two ways to transfer resources intertemporally, there should at the optimum 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.[62]

Tax incentives in the form of deduction schemes for higher education expenses are common, but are usually contemporaneous to the expense. In the US, the American Opportunity Credit and the

[62] The linear cost is for exposition only, since we want interior solutions in general. See Appendix formula (57) for the general case.

Figure 11: Progressivity of the deferred deductibility scheme

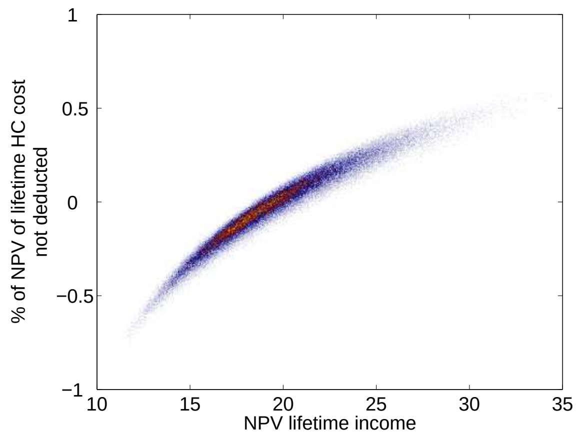

Lower income people can deduct more than they spend on human capital, while higher income people deduct less.

Lifetime Learning Credit allow families to claim a deduction up to a certain level per student per year for college, as well as for books, supplies, and required equipment.

The deferred deductibility scheme sets the right incentives in expected, discounted utility terms. Figure 11 illustrates that the scheme is naturally progressive and provides insurance, by plotting 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 hence end up deducting more than they actually spent on human capital, while higher income agents end up deducting less and hence implicitly cross-subsidizing lower income agents.

6.4 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 line 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 complementarity coefficient between human capital and ability $\rho_{\theta s}$ and the volatility of the productivity shock σ^2_ψ.[63] Welfare gains are expressed as the percentage increase in consumption which, if received every year after all histories, would yield the same gain in lifetime utility.

Although comparable, welfare gains are higher when human capital has negative insurance and redistributive effects ($\rho_{\theta s} > 1$). In this case, the laissez-faire economy features both more cross-sectional consumption inequality and higher consumption volatility over time, which makes the insurance from the constrained efficient mechanism more valuable. Naturally, welfare gains are higher when produc-

[63] The high volatility (0.0161) is from Storesletten *et al.* (2004).

tivity is more volatile.

Table 2: Welfare Gains

	$\rho_{\theta s} = 0.2$		$\rho_{\theta s} = 1.2$	
Volatility	Medium	High	Medium	High
Welfare gain from second best	0.85%	1.60%	0.98%	1.76%
Welfare gain from linear age-dependent policies	0.79%	1.53%	0.94%	1.74%
as % of second best	93%	95.6%	95.5%	98.5%

Medium volatility is 0.0095, high volatility is 0.0161. Line 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 ($\rho_{\theta s} > 1$). Line 2 shows the gain from linear age-dependent policies relative to the laissez-faire, while line 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.

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 line in table 2 shows that welfare gains as a fraction of the second best gains range from 93% for a low-volatility and low $\rho_{\theta s}$ case to a surprising 98.5% for a high volatility and high $\rho_{\theta s}$ 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.[64]

These findings are reminiscent of Mirrlees' (1971) conclusion that static optimal income tax schedules appear close to linear. They also echo closely recent findings from the dynamic taxation literature (Farhi and Werning, 2013), suggesting that the addition of human capital per se does not change this result.

7 Extension: Unobservable Human Capital

Until now, human capital was observable to the government, which might be a strong assumption under some circumstances. This section considers unobservable human capital investments, and focuses on how optimal policies are adjusted relative to the observable case. It starts with a simplified version of

[64] A word of caution is needed. Given the order of magnitude of 10e-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 only be taken as evidence for small welfare gains, not precise welfare calculations. The author is currently studying which other factors are most important for the welfare gains. For instance, stepping away from log utility and from a log-normal process for θ can increase welfare gains.

the model without training, but with unobservable human capital expenses. An augmented program is set up, in which the agent's human capital choice needs to be incentive compatible, in addition to his labor effort and type revelation. This combines dynamic moral hazard models with hidden savings and models with hidden persistent types,[65] both of which unobservably modify the agent's future response to incentives. Then, the complete model is analyzed, assuming that training time is not observable, but human capital expenses are. In this case, consumption is again controlled by the planner, and hidden training is akin to an unobservable effort with persistent effects over time.

7.1 Unobservable Monetary Investments in Human Capital

7.1.1 Planning Problem

When monetary investments in human capital are unobservable, the planner can no longer directly control consumption c_t. However, savings or physical capital investments remain observable, and hence the planner controls total financial resources transferred to an agent each period. The agent can both misreport his type and spend a different amount on human capital than the planner would choose. As in subsection 3.1, a reporting strategy $r = \left\{r_t\left(\theta^t\right)\right\}$ specifies a report after each history. In addition, a human capital strategy $\sigma = \left\{s_t\left(\theta^t\right)\right\}$ specifies human capital choices. The planner allocates output $y\left(r^t\right)$ and resources, denoted by $c^a\left(r^t\right)$, as functions of the history of reports by the agent. The agent then unobservably chooses his human capital investment $s\left(\theta^t\right)$. Hours of work are determined residually by $l\left(r^t\right) = y\left(r^t\right)/w\left(\theta_t, s\left(\theta^t\right)\right)$. The continuation value after history θ^t from reporting and human capital strategies r and σ can be written as:

$$\omega^{r,\sigma}\left(\theta^t\right) = u_t\left(c\left(r^t\left(\theta^t\right), s_t\left(\theta^t\right), s_{t-1}\left(\theta^{t-1}\right)\right)\right) - \phi_t\left(\frac{y\left(r^t\left(\theta^t\right)\right)}{w_t\left(\theta_t, s_t\left(\theta^t\right)\right)}\right) + \beta\int \omega^{r,\sigma}\left(\theta^{t+1}\right) f^{t+1}\left(\theta_{t+1}|\theta_t\right) d\theta_{t+1}$$

where actual consumption is equal to

$$c\left(r^t\left(\theta^t\right), s_t\left(\theta^t\right), s_{t-1}\left(\theta^{t-1}\right)\right) = c^a\left(r^t\left(\theta^t\right)\right) - M_t\left(s_t\left(\theta^t\right) - s_{t-1}\left(\theta^{t-1}\right)\right)$$

Similarly, let $\omega\left(\theta^t\right)$ be the continuation value after history θ^t from reporting truthfully and following the planner's recommended human capital strategy. Incentive compatibility requires that, after any history θ^t, for all reporting and human capital strategies r and σ :

$$(IC) : \omega\left(\theta^t\right) \geq \omega^{r,\sigma}\left(\theta^t\right) \quad \forall r, \sigma, \theta^t \tag{42}$$

[65] On hidden savings, see for instance Cole and Kocherlakota (2001), Werning (2002), Kocherlakota (2004), Abraham and Pavoni (2008). On hidden persistent types, see Farhi and Werning (2013), Kapicka (2013) among others.

The first-order approach replaces incentive constraint (42) by two local necessary conditions along the equilibrium path. Again, the agent's envelope condition must hold:

$$\frac{\partial \omega\left(\theta^{t}\right)}{\partial \theta_{t}}=\frac{w_{\theta,t}\left(\theta_{t},s_{t}\left(\theta^{t}\right)\right)}{w_{t}\left(\theta_{t},s_{t}\left(\theta^{t}\right)\right)^{2}}y\left(\theta^{t}\right)\phi_{t}^{\prime}\left(\frac{y\left(\theta^{t}\right)}{w_{t}\left(\theta_{t},s_{t}\left(\theta^{t}\right)\right)}\right)+\beta\int\omega\left(\theta^{t+1}\right)\frac{\partial f^{t+1}\left(\theta_{t+1}|\theta_{t}\right)}{\partial \theta_{t}}d\theta_{t+1} \tag{43}$$

(since there is no training in this subsection, $\phi_t^{\prime}\left(l_t\right)$ denotes the derivative with respect to labor). Second, since human capital is now unobservable, the agent's first order condition with respect to s_t, i.e., the "Euler Equation for human capital," must also hold, which is equivalent to a zero gross wedge τ_{St} in every period:

$$\begin{aligned}M_{t}^{\prime}\left(s_{t}\left(\theta^{t}\right)-s_{t-1}\left(\theta^{t-1}\right)\right)u_{t}^{\prime}\left(c\left(\theta^{t}\right)\right) &= \frac{w_{s,t}\left(\theta_{t},s_{t}\left(\theta^{t}\right)\right)}{w_{t}\left(\theta_{t},s_{t}\left(\theta^{t}\right)\right)^{2}}y\left(\theta^{t}\right)\phi_{t}^{\prime}\left(\frac{y\left(\theta^{t}\right)}{w_{t}\left(\theta_{t},s_{t}\left(\theta^{t}\right)\right)}\right) \\ &\quad +\beta\int M_{t+1}^{\prime}\left(s_{t+1}\left(\theta^{t+1}\right)-s_{t}\left(\theta^{t}\right)\right)u_{t+1}^{\prime}\left(c\left(\theta^{t+1}\right)\right)f^{t+1}\left(\theta_{t+1}|\theta_{t}\right)\end{aligned} \tag{44}$$

The focus is again on interior solutions (see the discussion on this in subsection 3.2).

Unobservable human capital raises the challenge of guaranteeing that the first-order approach is valid. In standard hidden savings problems, a sufficiently strong complementarity between shirking and savings can invalidate it (Kocherlakota, 2004). But there are still many cases in which it remains valid (see Abraham and Pavoni, 2008), and the gain in tractability from it is even larger relative to direct approaches, as the number of potential deviations grows. The goal of this section is to highlight analytically how optimal policies need to be adjusted relative to the observable human capital case, in circumstances under which the first-order approach can be applied. Studying general conditions for the validity of this approach with hidden human capital is left for future work.

The notation is as in subsection 3.3. In addition, a new endogenous state variable is introduced. Let $\Delta^{s}\left(\theta^{t}\right)$ be the negative of the expected marginal cost of education in utility units:

$$\Delta^{s}\left(\theta^{t}\right)=-\int M_{t+1}^{\prime}\left(s_{t+1}\left(\theta^{t+1}\right)-s_{t}\left(\theta^{t}\right)\right)u_{t+1}^{\prime}\left(c\left(\theta^{t+1}\right)\right)f^{t+1}\left(\theta_{t+1}|\theta_{t}\right) \tag{45}$$

$\Delta^{s}\left(\theta^{t}\right)$ captures the expected future cost reduction from an investment at time t, or, equivalently, at interior solutions, the future stream of benefits from investment. The Euler Equation from (44) can be rewritten as:

$$M_{t}^{\prime}\left(s_{t}\left(\theta^{t}\right)-s_{t-1}\left(\theta^{t-1}\right)\right)u_{t}^{\prime}\left(c\left(\theta^{t}\right)\right)=\frac{w_{s,t}}{w_{t}}l\left(\theta^{t}\right)\phi_{t}^{\prime}\left(l\left(\theta^{t}\right)\right)-\beta\Delta^{s}\left(\theta^{t}\right) \tag{46}$$

An additional "promise-keeping" constraint with respect to the future benefit of human capital needs to be added to the program

$$\Delta^{s}\left(\theta^{t-1}\right)=-\int M_{t}^{\prime}\left(s_{t}\left(\theta^{t}\right)-s_{t-1}\left(\theta^{t-1}\right)\right)u_{t}^{\prime}\left(c\left(\theta^{t}\right)\right)f^{t}\left(\theta_{t}|\theta_{t-1}\right)d\theta_{t}$$

The full recursive formulation of the program is in the Appendix.

7.1.2 The Optimal Labor Wedge with Unobservable Human Capital

With unobservable human capital, both the income tax and the savings tax will be adjusted to indirectly provide incentives for HC accumulation. A lower labor tax and a higher capital tax can stimulate human capital investments and mimic a positive net subsidy t_s, which is the optimal policy when $\rho_{\theta s} < 1$.

The optimal formula for the labor wedge is very complicated, because distorting labor downwards has several indirect feedback effects on unobservable human capital. First, with lower hours of work, there is less incentive for the agent to acquire (hidden) human capital. The wage is lower, which tends to reduce labor further. While this effect is also present with observable human capital, the planner can no longer directly counter it by controlling human capital incentives: the realized net incentive t_{st} will be different than the optimal t_{st}^* with observable human capital in (20). The effective labor distortion is hence larger than τ_{Lt} as defined in (15) if $t_{st} < 0$. Second, the income effect from a change in the labor distortion modifies the agent's incentive to invest through his Euler equation for human capital (in (46)). Finally, changing labor in one period has repercussions on future human capital choices, through the change in the contemporaneous human capital stock, because of the nonlinear cost of human capital M_t. All these additional feedback effects of the labor distortion are summed up for notational purposes in an adjusted labor wedge, $\tilde{\tau}_{Lt}$:[66]

Definition 3 *Define the adjusted labor wedge $\tilde{\tau}_{Lt}$:*

$$\frac{\tilde{\tau}_{Lt}}{1-\tilde{\tau}_{Lt}} \equiv \frac{\tau_{Lt}}{1-\tau_{Lt}} - \frac{1}{\left(1-\rho_{ss,t}\right)} \frac{1+\varepsilon_t^u}{\varepsilon_t^c} t_{st} + \left(\tilde{\gamma}_t^E \left(\theta^t\right) M_t' u_t'' - \frac{\left(\frac{1}{R} E_t \left(\tilde{\gamma}_{t+1}^E M_{t+1}'' u_{t+1}'\right) - \tilde{\gamma}_t^E M_t'' u_t'\right)}{w_{s,t} l_t \left(\theta^t\right) \left(1-\tau_{Lt}\right) \left(1-\rho_{ss,t}\right)} \frac{1+\varepsilon_t^u}{\varepsilon_t^c} \right) \quad (47)$$

where t_{st} is as defined in (19) evaluated at $\tau_{St} \equiv 0$; $\rho_{ss,t} \equiv \left(\frac{\partial^2 w_t}{\partial s_t} w_t\right) / \left(\frac{\partial w_t}{\partial s_t}\right)^2 < 0$; $\tilde{\gamma}_t^E \left(\theta^t\right) \equiv \tilde{\gamma}_t^E \left(\theta^t\right) / f^t \left(\theta_t | \theta_{t-1}\right)$ is the normalized multiplier on the Euler for human capital (46).

The following proposition shows how the adjusted labor wedge is set at the optimum.

Proposition 7 *i) The optimal adjusted labor wedge $\tilde{\tau}_{Lt}$ satisfies:*

$$\frac{\tilde{\tau}_{Lt}\left(\theta^t\right)}{1-\tilde{\tau}_{Lt}\left(\theta^t\right)} = \frac{\mu\left(\theta^t\right)}{f^t\left(\theta_t | \theta_{t-1}\right)} \frac{\varepsilon_{w\theta}}{\theta_t} u_t'\left(c\left(\theta^t\right)\right) \frac{1+\varepsilon_t^u}{\varepsilon_t^c} \left(1 - \frac{\left(1-\rho_{\theta s,t}\right)}{\left(1-\rho_{ss,t}\right)}\right) \quad (48)$$

where the multiplier $\mu\left(\theta_t\right)$ can be written recursively as:

$$\mu\left(\theta^t\right) = \kappa^u\left(\theta^t\right) + \eta^u\left(\theta^t\right)$$

[66] The term in brackets in the definition in (47) measures how the cost side of the Euler for human capital is relaxed or tightened. If the planner wants to stimulate human capital beyond what agents would themselves chose, then it can be shown that $\tilde{\gamma}_t^E(\theta) < 0$. The adjusted labor wedge is just a re-definition of a distortion that conveniently captures several terms. But the formulas can of course be re-stated in terms of the τ_{Lt} wedge, with the additional terms from $\tilde{\tau}_{Lt}$ appearing explicitly on the right hand side.

$$\begin{aligned}
\kappa^u\left(\theta^t\right) &= \int_{\theta_t}^{\bar{\theta}} (1-g_s) \frac{\left(1-\tilde{\gamma}^E\left(\theta_s\right) M_t'\left(e\left(\theta^{t-1},\theta_s\right)\right) u_t''\left(c\left(\theta^{t-1},\theta_s\right)\right)\right)}{u_t'\left(c\left(\theta^{t-1},\theta_s\right)\right)} f\left(\theta_s|\theta_{t-1}\right) d\theta \\
\textit{with } g_s &= \frac{u_t'\left(c\left(\theta^{t-1},\theta_s\right)\right)}{\left(1-\tilde{\gamma}^E\left(\theta_s\right) M_t'\left(e\left(\theta^{t-1},\theta_s\right)\right) u_t''\left(c\left(\theta^{t-1},\theta_s\right)\right)\right)} \lambda_- \ , \\
\lambda_- &= \int_{\underline{\theta}}^{\bar{\theta}} \frac{\left(1-\tilde{\gamma}^E\left(\theta_m\right) M_t'\left(e\left(\theta^{t-1},\theta_s\right)\right) u_t''\left(c\left(\theta^{t-1},\theta_m\right)\right)\right)}{u_t'\left(c\left(\theta^{t-1},\theta_m\right)\right)} f\left(\theta_m|\theta_{t-1}\right) d\theta_m \\
\eta^u\left(\theta^t\right) &= \frac{\mu\left(\theta^{t-1}\right)}{f\left(\theta_{t-1}|\theta_{t-2}\right)} \left[R\beta \left(\int_{\theta_t}^{\bar{\theta}} \frac{\partial f\left(\theta_s|\theta_{t-1}\right)}{\partial \theta_{t-1}} d\theta_s \right) \right]
\end{aligned}$$

ii) Distortions at the bottom and the top are not zero: $\tau_{Lt}\left(\theta^{t-1},\bar{\theta}\right) \neq 0$, $\tau_{Lt}\left(\theta^{t-1},\underline{\theta}\right) \neq 0$.

Formula (48) highlights the two roles that the income tax needs to fulfill with unobservable human capital, namely to insure and redistribute income through the standard channel, but also to partially substitute for the missing human capital subsidy. Because of the latter role, the factors which were previously entering the optimal net human capital wedge now affect the labor wedge formula. Among others, a new scaling factor appears, $\left(1-\frac{(1-\rho_{\theta s,t})}{(1-\rho_{ss,t})}\right)$, which can be greater than 1 when human capital has negative redistributive and insurance effects ($\rho_{\theta s} \geq 1$).[67] The scaling factor is dampened by the labor elasticity, since by manipulating the labor wedge to indirectly provide human capital incentives, the labor choice is also distorted, which is costlier if labor is more elastic. It is not always the case that the optimal labor wedge should be lower when human capital is unobservable. Intuitively, this is more likely to occur when $\rho_{\theta s}$ is small, in which case the planner would favor a larger net subsidy on human capital if it were observable, but agents do not factor in the full effect of their investment on social welfare.

The (unadjusted) labor wedge τ_{Lt} will no longer be zero at the bottom and at the top with unobservable human capital. Instead, it is the *adjusted* wedge $\tilde{\tau}_{Lt}$ which is zero.

The adjusted wedge can be rewritten recursively to study its life cycle evolution, again under the assumption that $\log(\theta_t)$ follows the auto-regressive process given by (28), as in subsection 4.4.[68]

$$\begin{aligned}
&E_{t-1}\left[\frac{\tilde{\tau}_L}{(1-\tilde{\tau}_{Lt})}\frac{1}{R\beta}\frac{\varepsilon_{w\theta,t-1}}{\varepsilon_{w\theta,t}}\left(\frac{\varepsilon_t^c}{1+\varepsilon_t^u}\right)\left(\frac{1+\varepsilon_{t-1}^u}{\varepsilon_{t-1}^c}\right)\frac{\left(\rho_{\theta s,t-1}-\rho_{ss,t-1}\right)}{\left(\rho_{\theta s,t}-\rho_{ss,t}\right)}\frac{\left(1-\rho_{ss,t}\right)}{\left(1-\rho_{ss,t-1}\right)}\frac{u_{t-1}'\left(c_{t-1}\right)}{u_t'\left(c_t\right)}\right] \\
&\quad = \frac{1}{R\beta}\varepsilon_{w\theta,t-1}\frac{1+\varepsilon_{t-1}^u}{\varepsilon_{t-1}^c}\frac{\left(\rho_{\theta s,t-1}-\rho_{ss,t-1}\right)}{\left(1-\rho_{ss,t-1}\right)}Cov\left(\frac{u_{t-1}'\left(c_{t-1}\right)}{u_t'\left(c_t\right)}\left(1-\tilde{\gamma}_t^E M_t' u_t''\right),\log\left(\theta_t\right)\right)+p\left(\frac{\tilde{\tau}_{Lt-1}}{1-\tilde{\tau}_{Lt-1}}\right)
\end{aligned}$$

In addition to the familiar factors, the labor wedge evolution is now also loaded with the factors which would determine the evolution of the net human capital wedge t_{st}, were human capital observable

[67] Because all terms in this formula are endogenous, it is not evident whether, after any given history, τ_{Lt} is higher or lower than in an equivalent world with observable human capital (i.e., a world with the same primitives, but not the same allocations).

[68] The proof is the same as for Corollary 3, except that τ_{Lt} is replaced by $\tilde{\tau}_{Lt}$, and the point-wise formula is (48).

(formula (30)). Notably, if the insurance and redistributive effects of hidden human capital grow over time ($\rho_{\theta s,t}$ decreases), the income tax will rise more slowly (or fall more quickly) to provide the needed stimulus for human capital. Unobservable human capital, with potentially varying redistributive implications throughout life, is another element to take into account when thinking about age-dependent taxation.

7.1.3 The Optimal Capital Wedge with Unobservable Human Capital Expenses

Unobservable human capital expenses are a hidden channel of intertemporal resource transfer, invalidating the standard Inverse Euler equation (in (24)). This is consistent with previous work with hidden consumption (e.g., Townsend, 1982, Cole and Kocherlakota, 2001), or hidden preferences (Atkeson and Lucas, 1992). The setup here is a hybrid of the observable and unobservable consumption cases, because the planner does control total financial resources allocated per period, but does not control how these resources get split by the agent between consumption and human capital expenses. Therefore, the agent's standard Euler for savings or physical capital does not need to hold, but the agent's Euler equation for human capital (in (44)) does. This imposes a restriction on the marginal utilities across periods. Attempting to manipulate assigned consumption in different periods carries an additional cost for satisfying the agent's Euler equation for human capital, because of an income effect. In this case, a modified Inverse Euler equation applies at the optimum:[69]

$$\frac{\beta R(1-\tilde{\gamma}_t^E\left(\theta^t\right)M_t'\left(e_t\right)u_t''\left(c_t\right))}{u_t'\left(c_t\right)}=\int_{\underline{\theta}}^{\bar{\theta}}\frac{\left(1-\tilde{\gamma}_{t+1}^E\left(\theta^{t+1}\right)M_{t+1}'\left(e_{t+1}\right)u_{t+1}''\left(c_{t+1}\right)\right)}{u_{t+1}'\left(c_{t+1}\right)}f^{t+1}\left(\theta_{t+1}|\theta_t\right) \tag{49}$$

7.2 Unobservable Training Time

This subsection reintroduces training time. Suppose that training time is unobservable to the government, but human capital expenses are observable. This is most consistent with the standard taxation framework, as training is another type of effort, such as labor, potentially difficult to keep track of. Expenses on tuition fees or formal education costs on the other hand are tangible, recordable transactions – similar to income. Under this scenario, the net subsidy on human capital expenses is another tool available to the government, in addition to the labor wedge and the savings wedge, to indirectly encourage or discourage training. The following assumption simplifies the analysis in this subsection.[70]

Assumption 4 $\phi_{lz,t}=0\ \forall t$.

[69] The capital wedge must satisfy a "no-arbitrage" condition:

$$\left(1-\tau_{Kt}\right)\beta RE_t\left(u_{t+1}'\right)=\frac{w_{s,t}}{w_t}l_t\frac{\phi'\left(l_t\right)}{M_t'}+\beta E_t\left(u_{t+1}'\frac{M_{t+1}'}{M_t'}\right)$$

[70] The setup of the program and the general derivations are in the Appendix.

Despite this assumption, training still affects labor supply through the wage. Recall that τ^*_{Lt} and t^*_{st} are the optimal labor and human capital wedges when training is observable, as given in (25) and (20) respectively.[71] The relation between the optimal labor tax and the net wedge on observable expenses is given by the following proposition.

Proposition 8 *At the optimum, the deviations of the labor wedge and net human capital wedge from their optimal levels when training is observable must satisfy the following relation:*

$$\left(\frac{\tau_{Lt}}{1-\tau_{Lt}} - \frac{\tau^*_{Lt}}{1-\tau^*_{Lt}}\right) = \frac{1}{\left(1-\rho_{zs,t}\right)} \frac{1+\varepsilon^u_t}{\varepsilon^c_t} \left(t_{st} - t^*_{st}\right) \tag{50}$$

where $\rho_{zs} \equiv \frac{w_{sz}w}{w_s w_z}$ *is the Hicksian coefficient of complementarity between human capital and training in the wage function.*

Hence, the labor wedge and the net subsidy will both be adjusted away from their optimal levels which would apply if training were observable. They will be modified in proportion to their effectiveness in affecting training time, that is, respectively, $1/\frac{\varepsilon^c_t}{1+\varepsilon^u_t}$ and $\left(1-\rho_{zs,t}\right)$. This leads to a variation on the inverse elasticity rule, in which the sharpest instrument needs to be used to indirectly provide incentives for an unobservable choice.

Should human capital expenses be distorted additionally to take into account hidden training? According to the above relation, in general, they should be, unless the elasticity of the wage to training does not depend on expenses $\left(\rho_{zs,t} = 1\right)$. Note once more that this is different from training and expenses being separable in the wage ($\rho_{zs,t} = 0$). The labor and human capital wedges will co-move if $\rho_{zs,t} \leq 1$, and move inversely otherwise. With a CES wage as in (3) with $\rho_{zs} = \rho_{\theta s} = \rho$ and isoelastic disutility as in (27):

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

With unobservable training, the standard Inverse Euler equation from (24) holds again, since consumption is controlled by the planner.

8 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. Agents invest in human capital throughout life, by either spending money or time. The government aims to provide redistribution and insurance against adverse draws from the ability distribution. However, the government faces

[71] These are endogenous functions of the allocations, hence their *values* here might not coincide with the optimal ones from the previous section.

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 crucial consideration for the design of optimal policies is whether human capital has overall positive redistributive and insurance effects. If human capital subsidies stimulate labor supply, and hence generate additional resources more than they amplify existing pre-tax inequality, they reduce after-tax income inequality on balance. This occurs when the elasticity of the wage with respect to ability is decreasing in human capital. In this case, the optimal net subsidy on human capital expenses is positive and increasing over time. When considering the optimal subsidies on training time, the additional interactions of training with both contemporaneous and future labor supply need to be taken into account. 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. A reinterpretation of the optimal formulas derived here shows that the optimal subsidy on parents' investments in children's human capital increases with their children's expected labor taxes, and decreases with the bequest tax. If the elasticity of the children's wage with respect to ability is decreasing in human capital, the positive redistributive and insurance effects of human capital transfers on the next generation push the optimal subsidy higher.[72] Second, should productivity-enhancing

[72] Higher than the subsidy needed to simply guarantee tax neutrality with respect to the choice between bequests and human capital transfers.

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.[73] If innovation and productivity expenses disproportionately increase risk, they should be made less than fully deductible.[74] Finally, the analysis could inform the study of optimal policies towards people's investments in health – another type of human capital – which also involves both time and monetary costs, as well as 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. This creates a fruitful link between optimal taxation and the long-standing empirical labor literature on this issue. Secondly, 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.

[73] These productivity investments, embodied in people, are distinct from investment in physical machines or financial assets.

[74] 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.

9 References

Abraham, A. and N. Pavoni. 2008. “Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending,” *Review of Economic Dynamics* 11, 781-803.

Acemoglu, D. and D. Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in O. Ashenfelter and D. Card (eds.) *Handbook of Labor Economics,* Vol. 4.

Albanesi, S. and C. Sleet. 2006. “Dynamic Optimal Taxation with Private Information,” *Review of Economic Studies* 73, 1-30.

Altonji, J., E. Blom and C. Meghir. 2012. “Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers,” *Annual Review of Economics*, Annual Reviews, vol. 4(1), 185-223.

Anderberg, D. 2009. “Optimal Policy and the Risk Properties of Human Capital Reconsidered,” *Journal of Public Economics* 93(9-10), 1017-1026.

Angrist, J., and A. Krueger 1991. “Does compulsory school attendance affect schooling and earnings,” *Quarterly Journal of Economics* 106, 979-1041.

Ashenfelter, O. and A. Krueger 1994. "Estimates of the economic return to schooling for a new sample of twins", *American Economic Review* 84:1157-1173.

Ashenfelter, O. and C. Rouse 1998. “Income, schooling, and ability: Evidence from a new sample of identical twins,” *Quarterly Journal of Economics* 113, 253-284.

Atkeson, A. and R. Lucas 1995. “Efficiency and Equality in a Simple Model of Efficient Unemployment Insurance,” *Journal of Economic Theory* 66, 64-88.

Atkinson, A. and J. Stiglitz 1976. “The Design of Tax Structure: Direct versus Indirect Taxation,” *Journal of Public Economics* 6, 55-75.

Autor, D.H., L. Katz and A. Krueger. 1998. “Computing inequality: have computers changed the labor market?,” *Quarterly Journal of Economics* 113, 1169-1213.

Avery, C., M. Glickman, C. Hoxby, and A. Metrick. 2013. “A Revealed Preference Ranking of American Colleges and Universities,” *Quarterly Journal of Economics* 128(1).

Baron, D. and D. Besanko. 1984. “Regulation and Information in a Continuing Relationship,” *Information Economics and Policy,* 1(3), 267-302.

Battaglini, M. and S. Coate. 2008. “Pareto Efficient Income Taxation with Stochastic Abilities,” *Journal of Public Economics* 92(3-4), 844-868.

Battaglini, M. and R. Lamba. 2013. “Optimal Dynamic Contracting: the First Order Approach and Beyond,” *Working Paper, Princeton University.*

Becker, G. 1964. *A Theoretical and Empirical Analysis, with Special Reference to Education*, The University of Chicago Press.

Benabou, R. 2002. "Tax and Education Policy in a Heterogeneous Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70(2), 481-517.

Ben-Porath, Y. 1967. “The Production of Human Capital and the Life Cycle of Earnings”, *Journal of Political Economy*, Vol. 75, No. 4, pp. 352-365.

Blomquist, Soren. 1982. “Should educational expenses be deductible?,” *Journal of Public Economics* 18, 65-83.

Bohacek, R. and M. Kapicka. 2008. “Optimal Human Capital Policies,” *Journal of Monetary Economics* 55(1), 1-16.

Boskin, Michael. 1977. “Notes on the Tax Treatment of Human Capital,” *Proceedings of the Treasury Conference on Tax Policy,* U.S. Department of Treasury.

Bovenberg, L. and B. Jacobs. 2005. “Redistribution and Education Subsidies are Siamese Twins,” *Journal of Public Economics*, 89(11-12), 2005–2035.

Bovenberg, L. and B. Jacobs. 2011. “Optimal Taxation of Human Capital and the Earnings Function,” *Journal of Public Economic Theory*, 13, (6), 957-971.

Brody, E. 1994. “Paying back your country through income-contingent student loans,” *San Diego Law Review*, 31, 449-518.

Card, D., 1995a. “Earnings, schooling and ability revisited” in S. Polachek (ed.) *Research in Labor Economics*, Vol. 14, JAI Press, Greenwich, CT, pp. 23-48.

Card, D., 1995b. “Using geographic variation in college proximity to estimate the return to schooling,” in L. N. Christofides, E. K. Grant and R. Swidinsky (eds.), *Aspects of labor market behavior: essays in honour of John Vanderkamp,* University of Toronto Press, Toronto. pp. 201-222.

Caucutt, E, and K. Kumar. 2003. "Higher Education Subsidies and Heterogeneity: A Dynamic Analysis," *Journal of Economic Dynamics and Control*, 27, 1459-1502.
Chamley, C. 1986. "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, 54(3), pp. 607-622.
Chetty, R. 2012. "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply," *Econometrica*, 80(3), 969-1018.
Chetty, R., J. Friedman, N. Hilger, D. Schanzenbach, D. Yagan. 2011. "How Does Your Kindergarden Classroom Affect Your Earnings? Evidence from Project STAR," *Quarterly Journal of Economics* 126(4), 1593-1660.
Cole, H.L., and N. Kocherlakota. 2001. "Efficient allocations with hidden income and hidden storage," *Review of Economic Studies*, 68(3), 523-542.
Cunha, F, J. Heckman and S. Navarro. 2005. "Separating uncertainty from heterogeneity in life cycle earnings," *Oxford Economics Papers*, 57(2), 191.
Cunha, F, and J. Heckman. 2007. "Identifying and Estimating the Distributions of Ex Post and Ex Ante Returns to Schooling," *Labour Economics*, 14(6), 870-893.
Cunha, F, and J. Heckman. 2008. "Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation," *Journal of Human Resources*, 43(4), 738-782.
Cunha, F, J. Heckman, L. Lochner, and D. Masterov. 2006. "Interpreting the Evidence on Life Cycle Skill Formation," in E. Hanushek and F. Welch (eds.), *Handbook of the Economics of Education,* Vol. 1, Chapter 12, 697–812.
Da Costa, C. E., and L. J. Maestri. 2007. "The risk Properties of Human Capital and the Design of Government Policies," *European Economic Review* 51(3), 695-713.
Diamond, P.1998. "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review* 88(1), 83-95.
Diamond, P. and J. Mirrlees. 1971. "Optimal Taxation and Public Production I: Production Efficiency," *American Economic Review* 61(1), 8-27.
Doepke, M., and R. Townsend. 2006. "Dynamic Mechanism Design with Hidden Income and Hidden Actions", *Journal of Economic Theory* 126(1), 235-285.
Dynarski, S. 2003. "Does Aid Matter? Measuring the Effects of Student Aid on College Attendance and Completion," *American Economic Review* 93(1), 279-288.
Farhi, E. and I. Werning 2013. "Insurance and Taxation over the Life Cycle," *Review of Economic Studies* 80, 596-635
Federal Reserve Bank of New York. 2013. *Quarterly Report on Household Debt and Credit,* May 2013
Fernandes, A. and C. Phelan. 2000. "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory* 91(2), 223-247.
Findeisen, S. and D. Sachs 2012. "Education and Optimal Dynamic Taxation: The Role of Income-contingent Student Loans," *Working Paper, University of Konstanz.*
Fryer R. 2011. "Financial Incentives and Student Achievement: Evidence from Randomized Trials," *Quarterly Journal of Economics*, 126, 1755-1798.
Goldin, C. and L. Katz. 2008. "The Race Between Education and Technology," *Harvard University Press, Cambridge MA.*
Golosov, M, A. Tsyvinski and M. Troshkin. 2013. "Redistribution and Social Insurance," *Working paper, Princeton University.*
Golosov, M, A. Tsyvinski and I. Werning. 2006. "New Dynamic Public Finance: A User's Guide," *NBER Macroeconomic Annual*, MIT Press.
Golosov, M, and A. Tsyvinski. 2007. "Optimal Taxation with Endogenous Insurance Markets," *Quarterly Journal of Economics*, 122(2), 487-534.
Grochulski, B, and T. Piskorski. 2010. "Risky Human Capital and Deferred Capital Income Taxation," *Journal of Economic Theory*, 145(3), 908-943.
Hader, S, and K. McGarry. 2012. "Parental Investments in College and Later Cash Transfers," *NBER working paper 18485.*
Heathcote, J., F. Perri and G. Violante. 2010. "Unequal we Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006," *Review of Economic Dynamics,* 13, 15-51.
Heathcote, J., K. Storesletten and G. Violante. 2005. "Two Views of Inequality Over the Life Cycle," *Journal of the European Economic Association,* 3, 765-775.

Heckman, James. 1976a. "Estimates of a Human Capital Production Function Embedded in a Life-Cycle Model of Labor Supply," in N. Terleckyj (ed.), *Household Production and Consumption*, NBER, p. 225-264.
Heckman, James. 1976b. "A Life-Cycle Model of Earnings, Learning, and Consumption," *Journal of Political Economy*, 84(4), S11-S44.
Heckman, J., L. Lochner, and P. Todd. 2005. "Earnings Functions, Rates of Return, and Treatment Effects: The Mincer Equation and Beyond," *NBER Working Paper 11544.*
Hicks, J. 1970. "Elasticity of Substitution Again: Substitutes and Complements," *Oxford Economic Papers,* 25, 289-296.
Hoxby, C. 1998. "Tax Incentives for Higher Education," in J. Poterba (ed.) Tax Policy and the Economy, *MIT Press, Cambridge, Mass,* 49-82.
Huggett, M., G. Ventura and A. Yaron. 2011. "Sources of Lifetime Inequality," *American Economic Review,* 101, 2923-54.
Huggett, M. and G. Kaplan. 2011. "Human Capital Values and Returns: Bounds Implied by Earnings and Asset Returns Data," *Journal of Economic Theory,* 146(3), 897-919.
Jones, L., Manuelli, R. and P. Rossi. 1993. "Optimal taxation in Models of Endogenous Growth," *Journal of Political Economy,* 101(3), 485-517.
Judd, K.L. 1985. "Redistributive taxation in a simple perfect foresight model," *Journal of Public Economics*, 28(1), pp. 59-83.
Kapicka, M. and J. Neira 2013. "Optimal Taxation in a Life-Cycle Economy with Endogenous Human Capital Formation," *Working Paper, University of California, Santa Barbara.*
Kapicka, M. 2013a. "Optimal Mirrleesean Taxation with Unobservable Human Capital Formation," *Working Paper, University of California, Santa Barbara.*
Kapicka, M. 2013b. "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach," *Review of Economic Studies,* 80(3), 1027-1054.
Kaplow, L. 1994. "Human Capital Under and Ideal Income Tax," *Virginia Law Review,* 80(7), 1477-1514.
Kaplow, L. 2006. "On the desirability of commodity taxation even when income taxation is not optimal," *Journal of Public Economics*, 90, 1235-1250.
Karaivanov, A. and R. Townsend 2013. "Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes," *Working Paper, MIT.*
Kocherlakota, N. 2004. "Figuring out the impact of hidden savings on optimal unemployment insurance," *Review of Economic Dynamics*, 7, 541-554.
Kocherlakota, N. 2005. "Zero Expected Wealth taxes: A Mirrlees Approach to Dynamic Optimal Taxation," *Econometrica*, 73(5), 1587-1621.
Kremer, M. 2002. "Should taxes be independent of age?" *mimeo.*
Krueger, D., and A. Ludwig. 2013. "Optimal Capital and Progressive Labor Income Taxation with Endogenous Schooling Decisions and Intergenerational Transfers," *University of Pennsylvania Working Paper.*
Laroque, G. 2005. "Indirect taxation is superfluous under separability and taste homogeneity: A simple proof," *Economic Letters*, 87, 141-144.
Lochner, L. and A. Monge-Naranjo. 2011. "The Nature of Credit Constraints and Human Capital," *American Economic Review*, 101, 2487-2529.
Maldonado, D. 2007. "Education Policies and Optimal Taxation," *International Tax and Public Finance*, 15(2), 131-143.
Milgrom, R. and I. Segal. 2002. "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70(2), 583-601.
Mirrlees, J. 1971. "An Exploration in the Theory of Optimal Income Taxation," *Review of Economic Studies*, 38, 175-208.
Mirrlees, J., S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba 2011. "Tax by Design: The Mirrlees Review," *Oxford University Press.*
Murphy, K., and F. Welch. 1992. "The structure of Wages," *Quarterly Journal of Economics*, 107, 285-326.
Chang Wei, C. 2010. "What is the Price of College?" *Stats in Brief*, National Center for Education Statistics, US Department of Education (NCES 2011-175).
OECD. 2004. "Improving Skills for More and Better Jobs: Does Training Make a Difference?," *OECD Employment Outlook.*
OECD. 2013. "Education at a Glance, 2013," *OECD, Paris.*
Palacios, M. 2004. *Investing in Human Capital: A Capital Markets Approach to Higher Education Funding*, Cambridge: Cambridge University Press.

Pavan, A., I. Segal and J. Toikka. 2013. "Dynamic Mechanism Design: A Myersonian Approach," forthcoming *Econometrica*.
Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies* 68(1), 205-229
Samuelson, P. 1973. "Complementarity – An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory", *Journal of Economic Literature,* 12, (4), 1255-1289.
Scheuer, F. 2013. "Entrepreneurial Taxation with Endogenous Entry," *American Economic Journal: Economic Policy,* forthcoming.
Schultz, T. 1961. "Investment in Human Capital," *American Economic Review,* 51.
Spear, S., and S. Srivastava 1987. "On Repeated Moral Hazard with Discounting", *Review of Economic Studies,* 54(4), 599-617.
Storesletten, K., Telmer, C., and A. Yaron. 2004. "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics* 51, 609-633.
Weinzierl, M. 2011. "The Surprising Power of Age-Dependent Taxes," *Review of Economic Studies,* 78, 1490-1518.
Werning, Ivan 2002 "Optimal Unemployment Insurance with Unobservable Savings," *Working Paper MIT.*
Werning, Ivan 2007a "Nonlinear Capital Taxation," *Working Paper MIT.*
Werning, Ivan 2007b "Optimal Fiscal Policy with Redistribution," Quarterly Journal of Economics, 122, 925-967.
Werning, Ivan 2007c "Pareto Efficient Income Taxation," *Working Paper MIT.*

A Appendix with Proofs

A.1 Observable Human Capital

Notation:

$w_{\theta,t}, w_{s,t}, w_{z,t}$ are the partials of the wage with respect to ability, human capital s and z respectively. $w_{\theta s,t}, w_{\theta z,t}, w_{ss,t}, w_{zs,t}, w_{zz,t}$ are the second order derivatives. $\phi_{l,t}$ and $\phi_{z,t}$ are the partials of the disutility with respect to labor and training respectively, and $\phi_{zz,t}, \phi_{lz,t}, \phi_{ll,t}$ the corresponding second order derivatives. ε_{xyt} is the elasticity of variable x with respect to variable y, $\varepsilon_{xyt} = d\log(x_t)/d\log(y_t)$, hence for instance, $\varepsilon_{w\theta,t}$ is the elasticity of the wage with respect to ability θ. When clear, the history dependence of allocations is omitted, e.g., c_t denotes $c(\theta^t)$.

Link to the Ben-Porath model in discrete time.

In the Ben-Porath model, the general accumulation process for human capital is $z_{t+1} = H(z_t, \alpha_t) + (1-\delta) z_t$ with α_t the fraction of human capital reinvested into human capital acquisition and δ the depreciation rate of human capital. Labor supply is equal to $l_t = 1 - \alpha_t$, the residual time from human capital acquisition. The rental rate of human capital w_t is constant and earnings are $y_t = w_t z_t l_t$. The most general functional form used is $z_{t+1} = \beta_0 \alpha_t^{\beta_1} (\alpha_t z_t)^{\beta_2} + (1-\delta) z_t$, so that $w_t \alpha_t z_t$ is the opportunity cost of human capital acquisition.

In this paper, a variation of this model is used, which is better adapted to a Mirrleesian analysis. The accumulation process is simplified along some dimensions and made more complex along others. First, instead of α_t, the input into human capital is effort or time i_t. i_t plays the same role of α_t but has a disutility cost, potentially nonseparable from the disutility cost of labor $\phi_t(l,i)$, and not just an opportunity cost of time. This is because, first, disutility costs from learning are empirically relevant (see Heckman et al. 2005) and second, with observable human capital, labor would also be observable if we only had: $l_t = 1 - i_t$. Second, the parameters are fixed so that only time, but not the previous stock of human capital matters for the accumulation, and there is no depreciation: $\beta_1 = 1 = \beta_0$ and $\beta_2 = \delta = 0$. This leads to a linear accumulation process $z_{t+1} = z_t + i_t$. The diminishing returns to human capital are instead captured in the more complex rental rate for human capital, i.e., the wage, which is heterogeneous, nonlinear, and stochastic $w_t = w_t(\theta_t, z_t, s_t)$ (and also depends on human capital expenses).

Generalized model a la Ben-Porath: It is possible to generalize the model to feature an accumulation process close to the one in Ben-Porath. The key change is to make the accumulation process depend on the human capital stock. Hence, let $z_t = Z_t(z_{t-1}, i_t) = z_{t-1} + H_t(z_{t-1}, i_t)$ be a more general production function without depreciation (which is without loss of generality). Define the inverse function with respect to i, $I_t(z_t, z_{t-1}) \equiv Z_t^{-1,i}(z_{t-1}, z_t)$. The accumulation process can be time-dependent. Note now that with $\frac{\partial \phi_t}{\partial z_t} = \frac{\partial \phi_t}{\partial i_t}\frac{\partial i_t}{\partial z_t} = \frac{\partial \phi_t}{\partial i_t}\frac{\partial I_t}{\partial z_t}$, the definition for the net bonus t_{zt} is unchanged. The same formula for the optimal bonus on training time as in (32) applies.

Derivations and proofs for Propositions (1) and (3) :

The expenditure function: $\tilde{c}(l, \omega - \beta v, i, \theta)$ defines consumption indirectly as a function of labor l, current period utility ($\tilde{u} = \omega - \beta v$), training, and the current realization of the type (note that conditional on these variables, consumption does not depend on human capital s). Then, $\omega(\theta) =$

$u_t(c(\theta)) - \phi_t(l(\theta), z(\theta) - z_-) + \beta v(\theta)$ becomes redundant as a constraint, and the choice variables are $(l(\theta), s(\theta), z(\theta), \omega(\theta), v(\theta), \Delta(\theta))$. Let the multipliers in program (14) be (in the order of the constraints there) $\mu(\theta)$, λ_-, and γ_-. The problem is solved using the optimal control approach where 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:

$$
\begin{aligned}
&(\tilde{c}(l(\theta), \omega(\theta) - \beta v(\theta), z(\theta) - z_-, \theta) + M_t(s(\theta) - s_-) - w_t(\theta, s(\theta), z(\theta)) l(\theta)) f^t(\theta|\theta_-) \\
&+\frac{1}{R} K(v(\theta), \Delta(\theta), \theta, s(\theta), z(\theta), t+1) f^t(\theta|\theta_-) \\
&+\lambda_- \left[v - \omega(\theta) f^t(\theta|\theta_-)\right] + \gamma_- \left[\Delta - \omega(\theta) \frac{\partial f^t(\theta|\theta_-)}{\partial \theta_-}\right] + \mu(\theta) \left[\frac{w_{\theta,t}}{w_t} l(\theta) \phi_{l,t}(l(\theta), z(\theta) - z_-) + \beta \Delta(\theta)\right]
\end{aligned}
$$

with boundary conditions:

$$
\lim_{\theta \to \bar{\theta}} \mu(\theta) = \lim_{\theta \to \underline{\theta}} \mu(\theta) = 0
$$

Part i) Taking the first order conditions (hereafter, FOC) of the recursive planning problem yields (the variable with respect to which the FOC is taken appears in brackets):

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

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

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

$$
\begin{aligned}
[s(\theta)] \quad : \quad & -M_t'(s(\theta) - s_-) + l(\theta) w_{s,t} + \frac{1}{R} \frac{\partial K(v(\theta), \Delta(\theta), \theta, s(\theta), z(\theta), t+1)}{\partial s(\theta)} \\
= \quad & \frac{\mu(\theta)}{f^t(\theta|\theta_-)} l(\theta) \phi_{l,t}(l(\theta), i(\theta)) \frac{1}{w_t^2} w_{\theta,t} w_{s,t} (\rho_{\theta s,t} - 1)
\end{aligned}
$$

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:

$$
\begin{aligned}
& -M_t'(s(\theta) - s_-) + l(\theta) w_{s,t} + \frac{1}{R} \int M_{t+1}'(s'(\theta') - s(\theta)) f^{t+1}(\theta'|\theta) d\theta' \\
= \quad & \frac{\mu(\theta)}{f^t(\theta|\theta_-)} l(\theta) \phi_{l,t}(l(\theta), i(\theta)) \frac{1}{w_t^2} w_{\theta,t} w_{s,t} (\rho_{\theta s,t} - 1)
\end{aligned}
$$

Use the expression for $w_{s,t} l_t$ from the definition of the human capital wedge τ_{St} in (17) to write the first order condition as as a function of the modified wedge t_{st}:

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

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

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

The law of motion for the co-state $\mu(\theta)$ comes from the first-order condition with respect to the state variable $\omega(\theta)$:

$$[\omega(\theta)] : \left(-\frac{1}{u_t'(c(\theta))} + (\lambda_-) + (\gamma_-)\frac{\partial f^t(\theta|\theta_-)}{\partial\theta_-}\frac{1}{f^t(\theta|\theta_-)}\right) f^t(\theta|\theta_-) = \dot{\mu}(\theta) \tag{51}$$

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

$$\mu(\theta) = \int_{\theta}^{\bar{\theta}} \left(\frac{1}{u_t'(c(\theta))} - (\lambda_-) - (\gamma_-)\frac{\partial f^t(\theta|\theta_-)}{\partial\theta_-}\frac{1}{f^t(\theta|\theta_-)}\right) f^t(\theta|\theta_-) \tag{52}$$

Integrating and using both boundary conditions yields:

$$\lambda_- = \int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{u'(c(\theta))} f^t(\theta|\theta_-)\, d\theta \tag{53}$$

Using the envelope conditions $\frac{\partial K(v(\theta),\Delta(\theta),\theta,s(\theta),z(\theta),t+1)}{\partial v(\theta)} = \lambda(\theta)$ and $\frac{\partial K(v(\theta),\Delta(\theta),\theta,s(\theta),z(\theta),t+1)}{\partial \Delta(\theta)} = -\gamma(\theta)$, the first-order conditions with respect to $v(\theta)$ and $\Delta(\theta)$ respectively lead to:

$$[v(\theta)] : \frac{1}{u'(c)} = \frac{\lambda(\theta)}{R\beta} \tag{54}$$

and

$$[\Delta(\theta)] : -\frac{\gamma(\theta)}{R\beta} = \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \tag{55}$$

Using (53) and (55) in the expression for $\mu(\theta)$ from (52) yields: $\mu\left(\theta^t\right) = \kappa\left(\theta^t\right) + \eta\left(\theta^t\right)$ where

$$\kappa\left(\theta^t\right) = \int_{\theta_t}^{\bar{\theta}} \frac{1}{u'(c(\theta))}\left(1 - u'(c(\theta))\int_{\underline{\theta}}^{\bar{\theta}} \frac{1}{u'(c(m))} f(m|\theta_-)\, dm\right) f^t(\theta|\theta_-)$$

$$\eta\left(\theta^t\right) = -(\gamma_-)\int_{\theta_t}^{\bar{\theta}} \frac{\partial f^t(\theta|\theta_-)}{\partial\theta_-} d\theta = \frac{\mu\left(\theta^{t-1}\right)}{f(\theta_{t-1}|\theta_{t-2})} R\beta \int_{\theta_t}^{\bar{\theta}} \frac{\partial f^t(\theta|\theta_-)}{\partial\theta_-} d\theta$$

where the last equality uses the lag of (55). 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 (1) (respectively, (3)).

Part ii) The proof is immediate by inspection as long as $\mu\left(\theta^t\right) \geq 0\ \forall t,\ \forall \theta^t$, which is now proved.

Lemma 1 *Under assumption (2), $\mu\left(\theta^t\right) \geq 0\ \forall t,\ \forall \theta^t$.*

Proof of Lemma 1:

The proof is close to the one in Golosov, Tsyvinski and Troshkin (2011), for a separable utility function and with human capital. From the envelope condition and the FOC for $v(\theta)$ in (54):

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

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

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

$$\mu(\theta_1)=\int_{\theta_1}^{\bar{\theta}}\left(\frac{1}{u_1'\left(c\left(\tilde{\theta}_1\right)\right)}-\lambda_-\right)f^1\left(\tilde{\theta}_1\right)d\tilde{\theta}_1$$

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

$$\mu(\theta_1)=\int_{\underline{\theta}}^{\theta_1}\left(-\frac{1}{u_1'\left(c\left(\tilde{\theta}_1\right)\right)}+\lambda_-\right)f^1\left(\tilde{\theta}_1\right)d\tilde{\theta}_1$$

Since for $\theta_1\leq\theta_1'$, $\frac{1}{u'(c(\theta_1))}\leq\lambda_-$, we again have $\mu(\theta_1)\geq 0$. Thus, for all θ_1, $\mu(\theta_1)\geq 0$.

By the first-order condition for Δ in (55):

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

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

$$\mu(\theta_2)=\int_{\theta_2}^{\bar{\theta}}\left(\frac{1}{u_2'\left(c\left(\tilde{\theta}_2\right)\right)}-(\lambda_-)-\frac{\partial f\left(\tilde{\theta}_2|\theta_1\right)}{\partial\theta_1}\frac{(\gamma_-)}{f\left(\tilde{\theta}_2|\theta_1\right)}\right)f\left(\tilde{\theta}_2|\theta_1\right)$$

Since by assumption (2) iii), $\frac{g\left(\tilde{\theta}_2|\theta_1\right)}{f\left(\tilde{\theta}_2|\theta_1\right)}$ is increasing in $\tilde{\theta}_2$, and we already showed that $\frac{1}{u'\left(c\left(\tilde{\theta}_2\right)\right)}$ is increasing in $\tilde{\theta}_2$, there is a θ_2' such that

$$\frac{1}{u_2'\left(c\left(\theta_2'\right)\right)}-\frac{\partial f\left(\theta_2'|\theta_1\right)}{\partial\theta_1}\frac{(\gamma_-)}{f\left(\theta_2'|\theta_1\right)}=\lambda_-$$

and such that for $\theta_2\geq\theta_2'$, $\mu\left(\theta^2\right)\geq 0$ (since integrating over non-negative numbers only). Rewriting $\mu\left(\theta^2\right)$ 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 $s^\delta\left(\tilde{\theta}^t\right)=s\left(\tilde{\theta}^t\right)+\delta$ and $y^\delta\left(\tilde{\theta}^t\right)=y\left(\tilde{\theta}^t\right)+dy\left(\tilde{\theta}^t\right)$ such that $\frac{s^\delta\left(\tilde{\theta}^t\right)}{y^\delta\left(\tilde{\theta}^t\right)}=\frac{s\left(\tilde{\theta}^t\right)}{y\left(\tilde{\theta}^t\right)}$, i.e., $dy\left(\tilde{\theta}^t\right)=\delta\frac{y\left(\tilde{\theta}^t\right)}{s\left(\tilde{\theta}^t\right)}$. This perturbation leaves utilities and incentive compatibility constraints unaffected. The change in the resource cost must hence be zero in this neighborhood, and letting $\eta\to 0$, we obtain: $-\frac{y(\theta)}{s(\theta)}+\left(M_t'(e(\theta))-\frac{1}{R}E_t\left(M_{t+1}'\left(e\left(\theta'\right)\right)\right)\right)=0$, which is equivalent to $t_{st}=0$ for the multiplicative wage.

Part iii) If θ is iid, $\gamma_-=0$ and $\eta\left(\theta^t\right)=0$ for all t. In addition, if $u_t'(c_t)=1\ \forall t$, then $\kappa\left(\theta^t\right)=0$ as well.

Part iv) is immediate from the boundary conditions $\mu(\underline{\theta})=\mu\left(\bar{\theta}\right)=0$.

Proof of Corollary (1) :

The general formula for $\rho_{\theta s} \neq 1$ is obtained from the FOC for s_t, together with the definition of the gross wedge in(17), and replacing the multiplier $\mu\left(\theta^t\right)$ by the expression from the optimal tax formula:

$$\begin{aligned}
\tau_{St}\left(\theta^t\right) \quad &= \quad \left(1+\left(1-\rho_{\theta s,t}\right)\frac{\varepsilon_t^c}{1+\varepsilon_t^u}\right)\left(M_t'^d-\tau_{St}\left(\theta^t\right)\right)\left[\frac{\tau_{Lt}\left(\theta^t\right)}{1-\tau_{Lt}\left(\theta^t\right)}\right] \\
&\quad -\frac{1}{R}\left(1-\xi'_{M_t'}\right)E_t\left(M'_{t+1}\left(s_{t+1}\left(\theta^t,\theta_{t+1}\right)-s_t\left(\theta^t\right)\right)\right)\left[\frac{\tau_{Kt}\left(\theta^t\right)}{1-\tau_{Kt}\left(\theta^t\right)}\right]
\end{aligned}$$

When $\rho_{\theta s,t}=1\ \forall t$, the expression becomes $\tau_{St}^*=\tau_{Lt}^*M_t'^d-\frac{\tau_{Kt}^*}{R\left(1-\tau_{Kt}^*\right)}\left(1-\tau_{Lt}^*\right)\left(1-\xi'_{M',t+1}\right)E_t\left(M'_{t+1}\right)$, equivalent to the one in the corollary.

Proof of Proposition (2):

Taking integral of $\dot{\mu}(\theta)$ in equation (51) between the two boundaries, $\bar{\theta}$ and $\underline{\theta}$, and using the boundary conditions $\mu\left(\bar{\theta}\right)=\mu\left(\underline{\theta}\right)=0$, as well as the expression for λ_- from (54), lagged by one period, yields the inverse Euler equation in (24).

Proof of Proposition (4) :

Taking the first order condition of the Hamiltonian in the proof of proposition (1) with respect to z_t yields:

$$\begin{aligned}
[z_t]:&\left(-\frac{\phi_{z,t}}{u'\left(c_t\right)}+w_{z,t}l_t+\int\frac{1}{R}\frac{\phi_{z,t+1}}{u'\left(c_{t+1}\right)}f^{t+1}\left(\theta_{t+1}|\theta_t\right)+\frac{1}{R}E_t\left(\frac{\mu\left(\theta_{t+1}\right)}{f^{t+1}\left(\theta_{t+1}|\theta_t\right)}\frac{w_{\theta,t+1}}{w_{t+1}}l_{t+1}\phi_{lz,t+1}\right)\right) \\
&-\frac{\mu\left(\theta_t\right)}{f^t\left(\theta_t|\theta_{t-1}\right)}\left(l_t\phi_{l,t}\left(l_t,z_t-z_{t-1}\right)\frac{1}{w_t}w_{\theta z,t}-l_t\phi_{l,t}w_{\theta,t}w_{z,t}\frac{1}{w_t^2}+\frac{w_{\theta,t}}{w_t}l_t\phi_{lz,t}\right)=0
\end{aligned}$$

Using the definition of the training time bonus, τ_{Zt} to replace for $w_{z,t}l_t$ yields:

$$-\frac{\phi_{z,t}}{u'\left(c_t\right)}+w_{z,t}l_t+\int\frac{1}{R}\frac{\phi_{z,t+1}}{u'\left(c_{t+1}\right)}f^{t+1}\left(\theta_{t+1}|\theta_t\right)=-\frac{1}{\left(1-\tau_{Lt}\right)}\left(\tau_{Zt}-\tau_{Lt}\left(\frac{\phi_z}{u'\left(c_t\right)}\right)^d+P_{Zt}\right)$$

with P_{Zt} and $\left(\frac{\phi_z}{u'(c_t)}\right)^d$ as defined in the text. The FOC then becomes:

$$\begin{aligned}
[z_t]\quad:\quad t_{zt}=\frac{\left(\tau_{zt}-\tau_{Lt}\left(\frac{\phi_{z,t}}{u'(c_t)}\right)^d+P_{Zt}\right)}{\left(1-\tau_{Lt}\right)\left[\left(\frac{\phi_{z,t}}{u'(c_t)}\right)^d-\tau_{zt}\right]}&=\frac{\mu\left(\theta_t\right)}{f^t\left(\theta_t|\theta_{t-1}\right)}u'\left(c_t\right)\frac{\varepsilon_{w\theta,t}}{\theta_t}\left(1-\rho_{\theta z,t}\right) \\
-\frac{1}{\left[\left(\frac{\phi_{z,t}}{u'(c_t)}\right)^d-\tau_{zt}\right]}&\left(\frac{\mu\left(\theta_t\right)}{f^t\left(\theta_t|\theta_{t-1}\right)}\frac{\varepsilon_{w\theta,t}}{\theta_t}l_t\phi_{lz,t}-\frac{1}{R}E_t\left(\frac{\mu\left(\theta_{t+1}\right)}{f^{t+1}\left(\theta_{t+1}|\theta_t\right)}\frac{\varepsilon_{w\theta,t+1}}{w_{t+1}}l_{t+1}\phi_{lz,t+1}\right)\right)
\end{aligned}$$

We can again replace the multipliers by the labor wedges, and note that

$$\frac{w_t}{w_{z,t}\phi_{l,t}\left(l_t\right)}\phi_{lz,t}=\frac{\varepsilon_{\phi z,t}}{\varepsilon_{wz,t}}\rho_{lz,t}^{\phi}$$

to obtain formula (32).

Proof of Corollary (3) :

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 (25), multiply both sides of the expression by $\pi(\theta) > 0$. Integrating by parts, yields:

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

where the third line uses the expression for λ_- and γ_- from respectively (54) and (55) evaluated at $t-1$. The 4th line uses the optimal wedge from (25) 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:

$$
\begin{aligned}
& E_{t-1}\left(\frac{\tau_{Lt}(\theta^t)}{(1-\tau_{Lt}(\theta^t))} \frac{\theta_t}{\varepsilon_{w\theta,t}} \frac{\varepsilon_t^c}{1+\varepsilon_t^u} \frac{u_{t-1}'(c(\theta^{t-1}))}{u_t'(c(\theta^t))} \frac{\pi(\theta_t)}{\theta_t} \right) \\
= \; & Cov\left(\frac{u_{t-1}'(c(\theta^{t-1}))}{u_t'(c_t(\theta^t))}, \Pi(\theta_t) \right) + \frac{R\beta\tau_{Lt-1}}{(1-\tau_{Lt-1})} \frac{\theta_{t-1}}{\varepsilon_{w\theta,t-1}} \frac{\varepsilon_{t-1}^c}{1+\varepsilon_{t-1}^u} \int \frac{\partial f^t(\theta_t|\theta_{t-1})}{\partial\theta_{t-1}} \Pi(\theta_t) d\theta_t
\end{aligned} \tag{56}
$$

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 (29).

Proof of Corollary (4):

The net wedge on human capital can be rewritten similarly as in the proof of Proposition 3, using the same weighting function $\pi(\theta) = 1$:

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

where the second line uses the expression for λ_- and γ_- from respectively (54) and (55) evaluated at $t-1$. Using the optimal wedge from (20) at time $t-1$ to substitute for the multiplier $\mu(\theta^{t-1})$, the boundary conditions for μ (and the resulting Inverse Euler Equation), and the log AR(1) process for θ_t, formula (30) in the text is obtained.

A.2 Implementation

A sequential reformulation of the recursive allocations:

Any solution to the recursive social planner's problem can be mapped into a solution which depends on the full past history, using a recursive construction. To see this, denote the solutions to the recursive problem at time t, for each realized type θ_t, as a function of all state variables by:

$$\begin{gathered} v_t^*\left(v_{t-1}, \Delta_{t-1}, s_{t-1}, \theta_{t-1}, \theta_t\right), \Delta_t^*\left(v_{t-1}, \Delta_{t-1}, s_{t-1}, \theta_{t-1}, \theta_t\right), \omega_t^*\left(v_{t-1}, \Delta_{t-1}, s_{t-1}, \theta_{t-1}, \theta_t\right), \\ y_t^*\left(v_{t-1}, \Delta_{t-1}, s_{t-1}, \theta_{t-1}, \theta_t\right), s_t^*\left(v_{t-1}, \Delta_{t-1}, s_{t-1}, \theta_{t-1}, \theta_t\right), c_t^*\left(v_{t-1}, \Delta_{t-1}, s_{t-1}, \theta_{t-1}, \theta_t\right) \end{gathered}$$

and the solutions to the planner's sequential problem by $\left\{x_t^*\left(\theta^t\right)\right\} = \left\{y_t^*\left(\theta^t\right), s_t^*\left(\theta^t\right), c_t^*\left(\theta^t\right)\right\}$ The dependence on initial promised utilities in period 1, due to initial heterogeneity, is dropped for notational convenience; they can be just carried as an additional conditioning variable. This allocation gives rise to a sequence of utilities for the agent, generated recursively by:

$$\begin{aligned} \omega_t^*\left(\theta^t\right) &= u_t\left(c_t^*\left(\theta^t\right)\right) - \phi_t\left(\frac{y_t^*\left(\theta^t\right)}{w_t\left(\theta_t, s_t^*\left(\theta^t\right)\right)}\right) + \beta \int \omega_{t+1}^*\left(\theta^t, \theta_{t+1}\right) f^{t+1}\left(\theta_{t+1}|\theta_t\right) d\theta_{t+1} \\ v_t^*\left(\theta^t\right) &= \int \omega_{t+1}^*\left(\theta^t, \theta_{t+1}\right) f^{t+1}\left(\theta_{t+1}|\theta_t\right) d\theta_{t+1} \\ \Delta_t^*\left(\theta^t\right) &= \int \omega_{t+1}^*\left(\theta^t, \theta_{t+1}\right) \frac{\partial}{\partial \theta_t} f^{t+1}\left(\theta_{t+1}|\theta_t\right) d\theta_{t+1} \end{aligned}$$

To initialize the allocations, set $\omega_1^*\left(\theta_1\right) = \omega_1^*\left(v_0, \Delta_0, s_0, \theta_0, \theta_1\right) = \underline{U}\left(\theta_1\right)$ (if there is initial heterogeneity in θ_1, with v_0 arbitrary in that case, $\Delta_0 = 0$), $y_1^*\left(\theta_1\right) = y_1^*\left(v_0, \Delta_0, s_0, \theta_0, \theta_1\right)$, $s_1^*\left(\theta_1\right) = s_1^*\left(v_0, \Delta_0, s_0, \theta_0, \theta_1\right)$, and construct iteratively the full, history dependent allocation for all histories θ^t using the difference equations:

$$\omega_t^*\left(\theta^t\right) = \omega_t^*\left(v_{t-1}^*\left(\theta^{t-1}\right), \Delta_{t-1}^*\left(\theta^{t-1}\right), s_{t-1}^*\left(\theta^{t-1}\right), \theta_{t-1}, \theta_t\right)$$

where v_{t-1}^* and Δ_{t-1}^* are written as functions of ω_{t-1}, i.e.,

$$v_{t-1}^*\left(\theta^{t-1}\right) = \frac{1}{\beta}\left[\omega_{t-1}^*\left(\theta^{t-1}\right) - u_{t-1}\left(c_{t-1}^*\left(\theta^{t-1}\right)\right) + \phi_{t-1}\left(\frac{y_{t-1}^*\left(\theta^{t-1}\right)}{w_{t-1}\left(\theta_{t-1}, s_{t-1}^*\left(\theta^{t-1}\right)\right)}\right)\right]$$

$$\Delta_{t-1}^*\left(\theta^{t-1}\right) = \frac{1}{\beta}\left[\frac{\partial \omega_{t-1}^*\left(\theta^{t-1}\right)}{\partial \theta_{t-1}} - \phi_{t-1}'\left(y_{t-1}^*\left(\theta^{t-1}\right) / w_{t-1}\left(s_{t-1}^*\left(\theta^{t-1}\right), \theta_{t-1}\right)\right) \frac{y_{t-1}^*\left(\theta^{t-1}\right)}{w_{t-1}\left(s_{t-1}^*\left(\theta^{t-1}\right), \theta_{t-1}\right)^2} \frac{\partial w_{t-1}}{\partial \theta_{t-1}}\right]$$

Construct also:

$$\begin{aligned} y_t^*\left(\theta^t\right) &= y_t^*\left(v_{t-1}^*\left(\theta^{t-1}\right), \Delta_{t-1}^*\left(\theta^{t-1}\right), s_{t-1}^*\left(\theta^{t-1}\right), \theta_{t-1}\right) \\ s_t^*\left(\theta^t\right) &= s_t^*\left(v_{t-1}^*\left(\theta^{t-1}\right), \Delta_{t-1}^*\left(\theta^{t-1}\right), s_{t-1}^*\left(\theta^{t-1}\right), \theta_{t-1}\right) \\ c_t^*\left(\theta^t\right) &= c_t^*\left(v_{t-1}^*\left(\theta^{t-1}\right), \Delta_{t-1}^*\left(\theta^{t-1}\right), s_{t-1}^*\left(\theta^{t-1}\right), \theta_{t-1}\right) \end{aligned}$$

Using these sequential optimal policies in the planner's problem we can rewrite the costs as a function of the history, i.e., $K_t\left(\theta^{t-1}\right) = K_t\left(v_{t-1}^*\left(\theta^{t-1}\right), s_{t-1}^*\left(\theta^{t-1}\right), \theta_{t-1}\right)$. or, if agents also have heterogeneous initial wealth levels, $K_t\left(b_0, \theta^{t-1}\right)$. Note that in the planner's problem (i.e., in the direct revelation mechanism) initial wealth holdings are irrelevant since the planner can fully observe and allocate consumption. Since agents can borrow at the same rate as the government, without loss of generality, agents can do all the borrowing and saving on their own. The implicit wealth levels are then defined recursively as: $\frac{1}{R} b_t^*\left(b_0, \theta^t\right) = y_t^*\left(b_0, \theta^t\right) - c_t^*\left(b_0, \theta^t\right) - M\left(e_t^*\left(b_0, \theta^t\right)\right) + b_{t-1}^*\left(b_0, \theta^t\right)$.

General Deductibility Scheme

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

$$
\begin{aligned}
-\frac{\partial T_t}{\partial e_t} &= \sum_{j=1}^{T-t} \beta^{j-1} E_t \left(\frac{u'_{t+j-1}}{u'_t} \frac{\partial T_{t+j-1}}{\partial y_{t+j-1}} \left(M'_{t+j-1} - \frac{1}{R} M'_{t+j} \right) \right) + \beta^{T-t} E_t \left(\frac{u'_T}{u'_t} \left(\frac{\partial T_T}{\partial y_T} \right) \right) \\
&\quad - \sum_{j=1}^{T-t} \beta^j E_t \left(\frac{u'_{t+j}}{u'_t} \left(\left(1 - \xi'_{M',t+j}\right) E_{t+j-1} \left(M'_{t+j}\right) \frac{\partial T_{t+j}}{\partial b_{t+j-1}} - \frac{\partial T_{t+j}}{\partial s_{t+j-1}} \right) \right)
\end{aligned} \tag{57}
$$

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 $(M'_{t+j} - \frac{1}{R} M'_{t+j+1})$, not at the "historic" marginal cost faced by the agent at the time of the purchase M'_t, i.e., a purchase of Δe at time t is deducted as $(M'_{t+j} - \frac{1}{R} M'_{t+j+1})\Delta e$ from y_{t+j} at $t+j$. Otherwise, there would be arbitrage possibilities.[75] 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 $\left(1 - \xi'_{M',t+j}\right) E_{t+j-1} \left(M'_{t+j} \right)$ enters the picture. The deduction is risk adjusted, as witnessed by the insurance factors, $\xi'_{M',t+1} \equiv -Cov\left(\frac{\beta u'_{t+1}}{u'_t} - 1, M'_{t+1} \right) / \left(E_t \left(\frac{\beta u'_{t+1}}{u'_t} - 1 \right) E_t \left(M'_{t+1} \right) \right)$, defined in the text.

As stated in the text, formula (40) can be rewritten in terms of the risk-adjusted, dynamic cost:

$$
-\frac{\partial T_t}{\partial e_t} + \beta E_t \left(\frac{u'_{t+1}}{u'_t} \left(\frac{\partial T_{t+1}}{\partial e_{t+1}} - \frac{\partial T_{t+1}}{\partial s_t} \right) \right) = \frac{\partial T_t}{\partial y_t} M'^d_t - \left(1 - \frac{\partial T_t}{\partial y_t} \right) \frac{1}{R} \left(1 - \xi'_{M',t+1}\right) E_t \left(\beta R \frac{u'_{t+1}}{u'_t} \frac{\partial T_{t+1}}{\partial b_t} \right) E_t \left(M'_{t+1} \right) \tag{58}
$$

Solving this relation forward, yields the analogous to (57) with M'^d_t.

Proof of Proposition (5)

In the first step, we construct the history-independent savings tax, in the spirit of Werning (2011), with added human capital. Consider an incentive compatible allocation expressed as a function of the full history $c_t\left(\theta^t\right)$, $y_t\left(\theta^t\right)$, $e_t\left(\theta^t\right)$,$s_t\left(\theta^t\right)$, and its associated continuation utility $\omega_t\left(\theta^t\right)$, and suppose it is implemented as the outcome of a direct revelation mechanism with no savings. Allow agents to save any desired amount, with the restriction $b_T \geq 0$ (end of period T asset level). Consider a general tax function $T_{Kt}\left(b_t, r^t\right)$ as a function of savings and the history of reports up to period t. Given the report of the agent up to period $t-1$, and the report of the current shock, the planner assigns $c_t\left(r^{t-1}, r_t\right)$, $y_t\left(r^{t-1}, r_t\right)$, $e_t\left(r^{t-1}, r_t\right)$ and the agent chooses savings b_t. Let $V_t\left(b_{t-1}, r^{t-1}, \theta_t\right)$ be the continuation value of an agent with beginning of period savings b_{t-1}, a history of reports r^{t-1}, and a realized shock θ_t. The agent's problem is:

$$
V_t\left(b_{t-1}, r^{t-1}, \theta_t\right) = \max_{r_t, b_t} \tilde{V}_t\left(b_{t-1}, b_t, r^{t-1}, r_t, \theta_t\right)
$$

with $\tilde{V}_t\left(b_{t-1}, b_t, r^{t-1}, r_t, \theta_t\right)$ defined as the value from saving an amount b_t and reporting r_t:

$$
\tilde{V}_t\left(b_{t-1}, b_t, r^{t-1}, r_t, \theta_t\right) = \left\{ \begin{array}{c} u_t\left(c_t\left(r^{t-1}, r_t\right) + b_{t-1} - \left(\frac{1}{R} b_t + T_{Kt}\left(b_t, r^{t-1}, r_t\right)\right)\right) - \phi_t\left(\frac{y\left(r^{t-1}, r_t\right)}{w_t\left(\theta_t, s_t\left(r^{t-1}, r_t\right)\right)} \right) \\ +\beta E\left(V_{t+1}\left(b_t, r^t, \theta_{t+1}\right) | \theta_t\right) \end{array} \right\}
$$

In period $T-1$, define for each type realization θ_{T-1}, current asset level b_{T-2}, history of reports r^{T-1}, and savings levels b_{T-1} a fictitious tax T^{θ}_{Kt} which makes an agent just indifferent between saving

[75] Note that with linear cost, as in the main text, this is just $(1-\beta)$ with $\beta = \frac{1}{R}$ for all $t < T$, and 1 for $t = T$.

b_{T-1} and saving zero.

$$\begin{aligned}\omega_{T-1}\left(\theta^{T-1}\right) &= u_{T-1}\left(c_{T-1}\left(r^{T-1}\right)+b_{T-2}-\frac{1}{R}b_{T-1}-T^{\theta}_{KT-1}\left(b_{T-1},r^{T-1},\theta_{T-1}\right)\right)\\ &\quad -\phi_{T-1}\left(\frac{y_{T-1}\left(r^{T-1}\right)}{w_{T-1}\left(\theta_{T-1},s_{T-1}\left(r^{T-1}\right)\right)}\right)+\beta E\left(V_T\left(b_{T-1},r^{T-1},\theta_T\right)|\theta_{T-1}\right)\end{aligned}$$

Taking the sup over all types θ_{T-1} yields a history-dependent, but type-independent savings tax $T^r_K\left(b_{T-1},r^{T-1}\right)=\sup_{\theta_{T-1}}T^{\theta}_{KT-1}\left(b_{T-1},r^{T-1},\theta_{T-1}\right)$.

By induction, suppose that in period t the agent is faced with a continuation value function $V_{t+1}\left(b_t,r^t,\theta_{t+1}\right)$. Define the tax function for period t as $T_{Kt}\left(b_t,r^t\right)=\sup_{\theta_t}T^{\theta}_{Kt}\left(b_t,r^t,\theta_t\right)$ with $T^{\theta}_{Kt}\left(b_t,r^t,\theta_t\right)$ to equate:

$$\omega_t\left(\theta^t\right)=u_t\left(c_t\left(r^t\right)+b_{t-1}-\frac{1}{R}b_t-T^{\theta}_{Kt}\left(b_t,r^t,\theta_t\right)\right)-\phi_t\left(\frac{y_t\left(r^t\right)}{w_t\left(\theta_t,s_t\left(r^t\right)\right)}\right)+\beta E\left(V_{t+1}\left(b_t,r^t,\theta_{t+1}\right)|\theta_t\right)$$

Work backwards to define the tax functions in this fashion for all periods. The sequence of tax functions $\left\{T_{Kt}\left(b_t,r^t\right)\right\}_{t=1}^{T-1}$ thus defined implement zero savings each period for all sequences of reports. Next, take the supremum over all histories of reports r^t to obtain a history independent tax which implements zero savings.

$$T_K\left(b\right)=\sup_{r^t}T_{Kt}\left(b,r^t\right)$$

In the second step, we construct the loan and repayment schedule that mimics the direct mechanism above. Note that L^t can directly be mapped into a history of human capital and education levels e^t (recall that $s_0=0$), using that $e_t=M_t^{-1}\left(L_t\right)$ Hence, $\left(L^{t-1},y^{t-1}\right)$ and $\left(e^{t-1},y^{t-1}\right)$ will be used interchangeably. First, set implicit finite (but potentially very large) upper and lower limits on asset holdings, $\bar{b}>0$ and $\underline{b}<0$. This can be done either by extending the proposed savings tax so that for $b_t>\bar{b}$ and $b_t<\underline{b}$, $\forall t$, it is confiscatory (e.g., tax away all wealth and imposes a large penalty on borrowing) or by directly setting borrowing and saving limits. Let $B\equiv\left[\underline{b},\bar{b}\right]$.

In each period, all allocations which can arise as outcomes in the optimum are made affordable:

$$D_t\left(L^{t-1},y^{t-1},e^*_t\left(\theta^{t-1},\theta\right),y^*_t\left(\theta^{t-1},\theta\right)\right)+T_Y\left(y^*_t\left(\theta^{t-1},\theta\right)\right) = y^*_t\left(\theta^{t-1},\theta\right)-c^*_t\left(\theta^{t-1},\theta\right) \quad (59)$$

$$L_t\left(e^*_t\left(\theta^{t-1},\theta\right)\right) = M_t\left(e^*_t\left(\theta^{t-1},\theta\right)\right) \quad (60)$$

for all $\left(L^{t-1},y^{t-1}\right)$ such that $\theta^{t-1}\in\Theta^{t-1}\left(\left\{M_1^{-1}\left(L_1\right),...,M_{t-1}^{-1}\left(L_{t-1}\right)\right\},y^{t-1}\right)\neq\emptyset$ and all $\theta\in\Theta$.

To mimic the direct revelation mechanism, we need to exclude pairs (e_t,y_t) which would not be assigned to any type θ_t in the social planner's problem after a history θ^{t-1}, and, consequently, exclude histories $\left(y^{t-1},e^{t-1}\right)$, which do not correspond to any history θ^{t-1}. Call "non-allowed" a choice which is not assigned in the social planner's problem for *any* type θ_t after history θ^{t-1}, i.e., such that $(e_t,y_t)\notin Q^{t-1}_{e,y}\left(\theta^{t-1}\right)$. The repayments at non-allowed levels have to be sufficiently dissuasive to make them strictly dominated by allowed choices. Then the history of reports would exactly be tracked by $r^{t-1}=\theta^{t-1}\in Q^{t-1}_{e,y}\left(y^{t-1},e^{t-1}\right)$, and the agent's problem becomes equivalent to making a report $r_t\in\Theta$ in each period, by choosing a pair (e_t,y_t) designed for some θ_t after θ^{t-1}. We know that in this case, the previously constructed history-independent savings tax would enforce zero savings.

There are several ways to rule out non-allowed allocations, and the goal here is just to provide a possible one, which is to simply set the repayment prohibitively high, so that irrespective of savings, it is never optimal to chose such allocations. For instance, after a history $\theta^{t-1} \in \Theta^{T-1}\left(e^{T-1}, y^{T-1}\right)$ and for any choice $(e_t, y_t) \notin Q^t_{ey}\left(\theta^{t-1}\right)$, set

$$D_t\left(L^{t-1}, y^{t-1}, e_t, y_t\right) + T_Y(y_t) > \left[\bar{b} - \underline{b} + y_t\right]$$

i.e., the repayment plus income tax at least confiscate income and impose an additional penalty such that all wealth is confiscated and agents can never borrow sufficiently to retain positive consumption.[76] This leaves the agent with zero consumption, and will never be chosen. More generally, arbitrarily large repayments can be set. The second and less draconian way is to take the envelope of the repayment schedules which, after each history, for any possible beginning of period wealth and optimal savings choice, would make the agent just indifferent between any non-allowed allocation and his optimal allocation.[77] Whatever the method chosen, once the non-allowed choices are ruled out, each period, after every history, the agent faces a problem equivalent in outcomes to the direct revelation mechanism, i.e., he faces only allocations which are also available to him in the social planner's problem after that history. Accordingly, the savings tax ensures that he will find it optimal not to save. By temporal incentive compatibility, he will chose the allocation designed for him.

Proof of Proposition (6) :

The proof is an extended and modified version of the proof of Proposition in Albanesi and Sleet (2006), adding human capital. With iid shocks, the recursive formulation of the relaxed program no longer requires the states Δ and θ_{t-1}:

$$K(v, s_-, t) = \min_{(c(\theta), l(\theta), \omega(\theta), s(\theta), v(\theta))} \int \left[c(\theta) + M_t\left(s(\theta) - s_-\right) - w\left(\theta, s(\theta)\right) l(\theta) + \frac{1}{R} K\left(v(\theta), s(\theta), t+1\right)\right] f(\theta)\, d\theta$$

subject to:

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

Denote by $K^{-1,s}$ and $K^{-1,v}$ the partial inverse functions of $K(v, s, t)$ with respect to its arguments s and v respectively. Define the set

$$Q^t_{e,y}(b_-, s_-) = \left\{e, y : e = e^*_t(v, s_-, \theta_t), y = y^*_t(v, s_-, \theta_t) \text{ for some } \theta_t \in \Theta, \text{ with } v = K^{-1,v}(b_-, s_-)\right\}$$

to be the set of output levels y_t and education levels e_t which are available to an agent with promised utility v and previous human capital level s_- in the social planner's problem. The value function of

[76] To extend the repayment scheme's domain to histories L^{t-1}, y^{t-1} for which $\Theta^{T-1}\left(e^{T-1}, y^{T-1}\right) = \emptyset$, set for all e_t, y_t after such histories:

$$D_t\left(L^{t-1}, y^{t-1}, e_t, y_t\right) + T_Y(y_t) > \bar{b} - \underline{b} + y_t$$

[77] A more sophisticated implementation, which smooths the repayment schedule to make it differentiable is currently explored by the author. That implementation involves adding wealth as a conditioning variable in the repayment function. See as well the next implementation below.

the agent who starts a period with wealth b_{t-1} and human capital s_{t-1} is denoted by $V_t(b_{t-1}, s_{t-1})$, as in the main text.

In each period t, the agent's problem can be split into two stages, because of the separability between consumption and labor. In stage 1, he chooses labor supply l_t (equivalently, output y_t) and human capital expenses e_t. He pays a tax $T_t(b_{t-1}, s_{t-1}, y_t, e_t)$ and is left with a total resource amount $b_t^m = y_t - T_t(b_{t-1}, s_{t-1}, y_t, e_t) - M(e_t) + b_{t-1}$. In stage 2, he chooses consumption and next-period bond holdings, b_t to maximize $V_t^m(b_t^m, s_{t-1} + e_t)$, the intermediate value function from resource level b_t^m, defined as:

$$\begin{aligned} V_t^m(b_t^m, s_t) &= \max_{c_t, b_t}(u_t(c_t) + \beta V_{t+1}(b_t, s_t)) \\ s.t. &: \quad b_t^m = c_t + \frac{1}{R} b_t \end{aligned}$$

with $V_T^m(b_T^m, s_T) = u_T(b_T^m)$. Denote the market outcomes by $\hat{b}_t(b_t^m, s_t)$ and $\hat{c}_t(b_t^m, s_t)$. In stage 1, the problem of the agent is:

$$\begin{aligned} V_t(b_{t-1}, s_{t-1}) &= \max_{\{y_t(\theta), e_t(\theta), b_t^m(\theta)\}} \int_\Theta (-\phi_t(y_t(\theta)/w_t(\theta, s_{t-1} + e_t(\theta))) + V_t(b_t^m(\theta), s_t(\theta))) f(\theta) d\theta \\ s.t. &: \quad b_t^m = y_t - T_t(b_{t-1}, s_{t-1}, y_t, e_t) - M(e_t) + b_{t-1} \end{aligned}$$

Let the market outcomes be denoted by $\hat{y}_t(b_{t-1}, s_{t-1}, \theta_t)$, $\hat{e}_t(b_{t-1}, s_{t-1}, \theta_t)$, and $\hat{b}_t^m(b_{t-1}, s_{t-1}, \theta_t)$ for each type θ_t.

In each period, the planner solves a two-stage problem as well. In the first stage, he allocates human capital expenses e_t, output y_t, and an intermediate promised utility v_t^m. In the second stage, he allocates consumption c_t, and continuation utility v_t. In stage 2, given an intermediate promised utility v_t^m, and an acquired human capital level $s_{t-1} + e_t = s_t$, the planner solves the program with intermediate continuation cost function $K^m(v_t^m, s_t, t)$:

$$\begin{aligned} K^m(v_t^m, s_t, t) &= \min_{c_t, v_t}\left(c_t + \frac{1}{R} K(v_t, s_t, t+1)\right) \\ s.t. &: \quad u_t(c_t) + \beta v_t = v_t^m \end{aligned}$$

with $K(v_T, s_T, T+1) \equiv 0$. Denote the solutions to this problem by $c_t^*(v_t^m, s_t)$ and $v_t(v_t^m, s_t)$.

In stage 1, the problem of the planner is hence:

$$\begin{aligned} K(v_{t-1}, s_{t-1}, t) &= \min_{\{v_t^m(\theta), e_t(\theta), y_t(\theta)\}} \int_\Theta (M(e_t(\theta)) - y_t(\theta) + K^m(v_t^m(\theta), s_{t-1} + e_t(\theta), t)) f(\theta) d\theta \\ s.t. &: \quad v_t^m(\theta) - \phi_t(y_t(\theta)/w_t(\theta, s_{t-1} + e_t(\theta))) \\ &\geq \quad v_t^m(\theta') - \phi_t(y_t(\theta')/w_t(\theta, s_{t-1} + e_t(\theta'))) \quad \forall \theta, \theta' \\ & \quad \int_\Theta (v_t^m(\theta) - \phi_t(y_t(\theta)/w_t(\theta, s_{t-1} + e_t(\theta)))) f(\theta) d\theta = v_{t-1} \end{aligned}$$

Denote the solutions to the planner problem by $v_t^{m*}(v_{t-1}, s_{t-1}, \theta_t)$, $e_t^*(v_{t-1}, s_{t-1}, \theta_t)$, and $y_t^*(v_{t-1}, s_{t-1}, \theta_t)$.

In stage 2, the problem of an agent who has acquired human capital s_t and who starts with intermediate wealth $b_t^m = K^m(v_t^m, s_t, t)$ is exactly the dual of the planner's problem who has promised utility $v_t^m = (K^m)^{-1,v}(b_t^m, s_t, t)$ for an agent with human capital s_t (where $(K^m)^{-1,v}$ is the partial inverse of K^m with respect to its first argument), so that the consumption choice of the agent and planner will coincide if mapped appropriately, i.e., $\hat{c}_t(K^m(v_t^m, s_t, t), s_t) = c_t^*(v_t^m, s_t)$. Furthermore, $\hat{b}_t(K^m(v_t^m, s_t, t), s_t) = K(v_t(v_t^m, s_t), s_t, t+1)$. To see this, suppose instead that there was another pair $\left(\tilde{c}, \tilde{b}\right) \neq (c^*, K)$ such that $\tilde{c} + \frac{1}{R}\tilde{b} = b_t^m$, but which yields higher utility: $u_t(\tilde{c}) + \beta V_{t+1}\left(\tilde{b}, s_t\right) > v_t^{m*} = (K^m)^{-1,v}(b_t^m, s_t, t)$ Note that the choice of s_t, already made in the previous stage is now fixed. Under the assumption that $V_{t+1}(., s_t)$ is increasing and continuous in its first argument, and given that $u_t(c)$ is increasing and continuous in c, there is also a pair $\left(\tilde{c}', \tilde{b}'\right)$ such that $\tilde{c} \leq c^*$, $\tilde{b} \leq K$ with one or both of these inequalities strict and such that $u_t(\tilde{c}) + \beta V_{t+1}\left(\tilde{b}, s_t\right) = v_t^{m*}$. But then $\left(\tilde{c}', \tilde{b}'\right)$ is better than (c^*, K) in the planner's problem and hence, (c^*, K) could not have been optimal, a contradiction.

For the first stage, consider an agent with initial wealth and human capital levels b_{t-1} and s_{t-1}. First, map the allocations from the social planner's problem to allocations defined on the state space $(b_{t-1}, s_{t-1}, \theta_t)$:

$$\begin{aligned}
y_t^*(b_{t-1}, s_{t-1}, \theta_t) &= y_t^*\left(K^{-1,v}(b_{t-1}, s_{t-1}, t), s_{t-1}, \theta_t\right) \\
e_t^*(b_{t-1}, s_{t-1}, \theta_t) &= e_t^*\left(K^{-1,v}(b_{t-1}, s_{t-1}, t), s_{t-1}, \theta_t\right) \\
b_t^{*m}(b_{t-1}, s_{t-1}, \theta_t) &= K_t^m\left(v_t^{m*}(v_{t-1}, s_{t-1}, \theta_t), s_{t-1} + e_t^*(v_{t-1}, s_{t-1}, \theta_t), t\right)
\end{aligned}$$

Then, set the tax level such that for all $y, e \in Q_{e,y}^t(b_{t-1}, s_{t-1})$,

$$\begin{aligned}
& T_t(b_{t-1}, s_{t-1}, y_t^*(b_{t-1}, s_{t-1}, \theta_t), e_t^*(b_{t-1}, s_{t-1}, \theta_t)) \\
= \; & y_t^*(b_{t-1}, s_{t-1}, \theta_t) - b_t^{*m}(b_{t-1}, s_{t-1}, \theta_t) + b_{t-1} - M_t(e_t^*(b_{t-1}, s_{t-1}, \theta_t))
\end{aligned}$$

(that is, make all allocations consistent with an allocation in the planner's problem just affordable).

To extend the definition of the tax function to the full domain of allocations, even those which would not arise in the social planner's problem, three steps are taken. First, to rule out wealth levels b_{t-1} not observed along the equilibrium path, i.e., such that

$$b_{t-1} \neq K\left(v_{t-1}^*\left(\theta^{t-1}\right), s_{t-1}^*\left(\theta^{t-1}\right), t\right)$$

for any θ^{t-1}, set the borrowing limits to be $\underline{b_{t-1}} = \min_{v,s} K(v, s, t)$ where the min is taken over the possible values of v and s at time t in the planner's program. Second, if $b_{t-1} = K\left(v_{t-1}^*\left(\tilde{\theta}^{t-1}\right), s_{t-1}^*\left(\tilde{\theta}^{t-1}\right), t\right)$ for some $\tilde{\theta}^{t-1}$ but $s_{t-1} \neq s_{t-1}^*\left(\tilde{\theta}^{t-1}\right)$ (that is, the levels of wealth and human capital from the past period are not mutually consistent), set the tax T_t such that, for all e_t, y_t:

$$T_t(b_{t-1}, s_{t-1}, e_t, y_t) = y_t + \max\{b_{t-1}, 0\} - \min\left\{\underline{b_t}, 0\right\}$$

Finally, if the agent is on the equilibrium path with b_{t-1} and s_{t-1}, and letting $v_{t-1} = K^{-1,v}(b_{t-1}, s_{t-1}, t)$, set T_t such that for all pairs $y_t, e_t \notin Q_{e,y}^t(b_{t-1}, s_{t-1})$, and all θ_t

$$\begin{aligned}
& -\phi_t\left(\frac{y_t}{w_t(\theta_t, s_{t-1} + e_t)}\right) + \beta V_t^m(b_{t-1} + y_t - T_t(b_{t-1}, s_{t-1}, y_t, e_t) - M_t(e_t), s_{t-1} + e_t) \\
< \; & -\phi_t\left(\frac{y_t^*(v_{t-1}, s_{t-1}, \theta_t)}{w_t(\theta_t, s_{t-1} + e_t^*(v_{t-1}, s_{t-1}, \theta_t))}\right) + v_t^{m*}(v_{t-1}, s_{t-1}, \theta_t)
\end{aligned}$$

In the first stage, given the dissuasive taxes on choices which never arise in the planner's problem, the agent can either choose the full allocation (y_t and e_t) destined for some type θ_t (that could be his own true type), which will then lead him to choose also the continuation wealth optimal for that same type, or he could choose y_t optimal for some type θ_t^1 but e_t optimal for some type θ_t^2. This will however leave him with lower value given the taxes on the off-equilibrium paths. Hence, if a type deviates, he must deviate to the full allocation of another type. By the temporal incentive compatibility constraint, he will choose not to do so.

A.3 Unobservable Human Capital

Proof of Proposition (7) :

A recursive formulation of the problem with unobservable human capital is:

$$K\left(v,\Delta,\Delta^s,\theta_-,s_-,t\right)=\min\int\left[\begin{array}{c}c^a\left(\theta\right)-w_t\left(\theta,s\left(\theta\right)\right)l\left(\theta\right)\\+\frac{1}{R}K\left(v\left(\theta\right),\Delta\left(\theta\right),\Delta^s\left(\theta\right),\theta,s\left(\theta\right),t+1\right)\end{array}\right]f^t\left(\theta|\theta_-\right)d\theta \tag{61}$$

$$\begin{aligned}
&\dot{\omega}\left(\theta\right)=\frac{w_{\theta,t}}{w_t}l\left(\theta\right)\phi_t'\left(l\left(\theta\right)\right)+\beta\Delta\left(\theta\right)\\
&\omega\left(\theta\right)=u_t\left(c^a\left(\theta\right)-M_t\left(s\left(\theta\right)-s_-\right)\right)-\phi_t\left(l\left(\theta\right)\right)+\beta v\left(\theta\right)\\
&v=\int\omega\left(\theta\right)f^t\left(\theta|\theta_-\right)d\theta\\
&\Delta=\int\omega\left(\theta\right)\frac{\partial f^t\left(\theta|\theta_-\right)}{\partial\theta_-}\left(\theta|\theta_-\right)d\theta\\
&M_t'\left(s\left(\theta\right)-s_-\right)u_t'\left(c^a\left(\theta\right)-M_t\left(s\left(\theta\right)-s_-\right)\right)=\frac{w_{s,t}}{w_t}l\left(\theta\right)\phi_t'\left(l\left(\theta\right)\right)-\beta\Delta^s\left(\theta\right)\\
&\Delta^s=-\int\left(\frac{w_{s,t}}{w_t}l\left(\theta\right)\phi_t'\left(l\left(\theta\right)\right)-\beta\Delta^s\left(\theta\right)\right)f^t\left(\theta|\theta_-\right)d\theta
\end{aligned}$$

where the maximization is over $\left(c^a\left(\theta\right),l\left(\theta\right),s\left(\theta\right),\omega\left(\theta\right),v\left(\theta\right),\Delta\left(\theta\right),\Delta^s\left(\theta\right)\right)$, with $\Delta^s\left(\theta\right)$ as defined in (45) .

Let $\phi\left(l\right)$ be the disutility of labor, $\phi'\left(l\right)$ and $\phi''\left(l\right)$ its first and second order partials. The function $\tilde{c}^a\left(l,\omega-\beta v,s,s_-,\theta\right)$ defines assigned resources as a function of labor l, current period utility (rewritten again as: $\tilde{u}=\omega-\beta v$), current and beginning of period human capital levels s and s_-, and the current realization of the type. Then, with the definition $c\left(\theta\right)=c^a\left(\theta\right)-M_t\left(s\left(\theta\right)-s_-\right)$, the constraint $\omega\left(\theta\right)=u_t\left(c\left(\theta\right)\right)-\phi_t\left(l\left(\theta\right),z\left(\theta\right)-z_-\right)+\beta v\left(\theta\right)$ becomes redundant, and the choice variables are $\left(l\left(\theta\right),s\left(\theta\right),\omega\left(\theta\right),v\left(\theta\right),\Delta\left(\theta\right),\Delta^s\left(\theta\right)\right)$. Let the multipliers on the constraints in program (61) be (in the order of the constraints): $\mu\left(\theta\right)$, λ_-, γ_-, $\gamma^E\left(\theta\right),\gamma_-^S$. The corresponding Hamiltonian is:

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

with the boundary conditions:

$$\lim_{\theta\to\bar{\theta}}\mu\left(\theta\right)=\lim_{\theta\to\underline{\theta}}\mu\left(\theta\right)=0$$

From the FOC of the agents, along the optimum, and all other choice variables of the planner held constant:

$$\frac{d\tilde{c}^a}{dl} = \frac{\phi'(l)}{u'(c)} = w(1-\tau_L), \qquad \frac{d\tilde{c}^a}{ds_t} = M_t'(s_t - s_{t-1}), \qquad \frac{d\tilde{c}^a}{d\omega} = \frac{1}{u_t'(c)}$$

Using the three envelope conditions: $\frac{\partial K(v(\theta),\Delta(\theta),\Delta^s(\theta),\theta,s(\theta),t+1)}{\partial v(\theta)} = \lambda(\theta)$, $\frac{\partial K(v(\theta),\Delta(\theta),\Delta^s(\theta),\theta,s(\theta),t+1)}{\partial \Delta(\theta)} = \gamma(\theta)$, and $\frac{\partial K(v(\theta),\Delta(\theta),\Delta^s(\theta),\theta,s(\theta),t+1)}{\partial \Delta^s(\theta)} = -\gamma^s(\theta)$, the FOCs for $\omega(\theta)$, $v(\theta)$, $\Delta(\theta)$, and $\Delta^s(\theta)$ can be rewritten respectively as:

$$[\omega(\theta)] : -\frac{1}{u_t'(c)} f^t(\theta|\theta_-) + \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} M_t'(e_t) \frac{u_t''(c)}{u_t'(c)} f^t(\theta|\theta_-) + \lambda_- f^t(\theta|\theta_-) + \gamma_- \frac{\partial f^t(\theta|\theta_-)}{\partial \theta_-} = \dot{\mu}(\theta)$$

$$\begin{aligned} [v(\theta)] \quad &: \quad \frac{1}{u_t'(c_t)} = \frac{1}{\beta R}\lambda(\theta) + \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} \frac{u_t''}{u_t'} M_t'(e_t) \qquad (62) \\ [\Delta(\theta)] \quad &: \quad -\frac{\gamma(\theta)}{R\beta} = \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \\ [\Delta^s(\theta)] \quad &: \quad -\beta \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} + \beta\gamma_-^S = \frac{1}{R}\gamma^S(\theta) \end{aligned}$$

Let $\tilde{\gamma}^E(\theta) = \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)}$ and θ_{-2} again denote the shock two periods back. Hence, the multiplier $\mu(\theta)$ solves:

$$\mu(\theta) = \int_\theta^{\bar{\theta}} \left(\frac{1}{u_t'(c)} \left(1 - \tilde{\gamma}^E(\theta_s) M_t'(e(\theta_s)) u_t''(c)\right) - \lambda_- + \frac{R\beta}{f(\theta_-|\theta_{-2})} \mu_- \frac{\partial f^t(\theta_s|\theta_-)}{\partial \theta_-} \frac{1}{f(\theta_s|\theta_-)} \right) f(\theta_s|\theta_-) d\theta_s$$

Using the boundary conditions:

$$\lambda_- = \int_{\underline{\theta}}^{\bar{\theta}} \left(\frac{1}{u_t'(c_t)} \left(1 - \tilde{\gamma}^E(\theta_m) M_t'(e(\theta_m)) u_t''(c_t)\right) f(\theta_m|\theta_-) d\theta_m \right)$$

This equation, together withequation (62), lagged by one period, shows that the standard Inverse Euler no longer holds and is instead as in formula (49). The FOCs for labor and human capital are (with θ' and s' being the next period's type and human capital).

$$[l(\theta)] : \frac{\tau_L(\theta)}{1-\tau_L(\theta)} = \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \frac{w_{\theta,t}}{w_t} u'(c_t) \frac{1+\varepsilon^u}{\varepsilon^c} + u'(c_t) \left(\frac{w_{s,t}}{w_t}\right) \left(\frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} - \gamma_-^S\right) \frac{1+\varepsilon^u}{\varepsilon^c} - \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} M_t' u_t''$$

$$\begin{aligned} [s(\theta)] \quad &: \quad \left(-M_t'(s(\theta) - s_-) + w_{s,t} l(\theta)\right) f^t(\theta|\theta_-) + \frac{1}{R} \int \left(M_{t+1}'\left(s'(\theta') - s(\theta)\right)\right) f^{t+1}(\theta'|\theta) d\theta' f^t(\theta|\theta_-) \\ &\quad + \mu(\theta) \frac{l(\theta)\phi_t'(l(\theta))}{w_t} \frac{w_{\theta,t} w_{s,t}}{w_t} \left(1 - \rho_{\theta s,t}\right) + \left(\gamma^E(\theta) - f^t(\theta|\theta_-)\gamma_-^S\right) l(\theta) \phi'(l(\theta)) \left(\frac{w_{s,t}}{w_t}\right)^2 \left(1 - \rho_{ss,t}\right) \\ &\quad + \gamma^E(\theta) M_t''(s(\theta) - s_-) u_t' - \frac{1}{R} \int \gamma^E(\theta') M_{t+1}''\left(s'(\theta') - s(\theta)\right) u_{t+1}'\left(c(\theta')\right) f^{t+1}(\theta'|\theta) d\theta' f^t(\theta|\theta_-) \\ = &\quad 0 \end{aligned}$$

where $\rho_{ss,t} = \frac{w_{ss,t} w_t}{w_{s,t}}$. Using the definition of the net wedge in (19) with τ_{St} set identically to 0 (because of the agent's Euler equation), the FOC for $s(\theta)$ allows to derive:

$$\left(\frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} - \gamma_-^S\right) = \frac{t_{st} w_t}{u'(c_t) w_{s,t} \left(1 - \rho_{ss,t}\right)} - \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \frac{w_{\theta,t}}{w_{s,t}} \frac{\left(1 - \rho_{\theta s,t}\right)}{\left(1 - \rho_{ss,t}\right)} + \frac{w_t \left(\frac{1}{R} E\left(\tilde{\gamma}_{t+1}^E M_{t+1}'' u_{t+1}'\right) - \tilde{\gamma}_t^E M_t'' u_t'\right)}{u_t' w_{s,t} \left(1 - \rho_{ss,t}\right) w_{s,t} l_t \left(1 - \tau_{Lt}\right)}$$

Using this expression into the FOC for $l(\theta)$ yields the optimal tax formula.

Proof of Proposition (8):

If training time is unobservable, but expenses are observable, another endogenous state variable is introduced, analogous to Δ^s in the unobservable s_t case:

$$\Delta^z(\theta) = -E_t\left(\phi_{z,t+1}\right)$$

As human capital expenses are observable, the agent's consumption is again directly controlled by the planner. Define the expenditure function indirectly as a function of the other choice variables: $\tilde{c}_t(l(\theta), \omega(\theta) - \beta v(\theta), z(\theta), z_-, \theta)$. The recursive program is (with multipliers in brackets after each corresponding constraint):

$$\begin{aligned} & K(v, \Delta, \Delta^z, \theta_-, s_-, z_-, t) \\ = & \min_{\{l,\omega,v,s,z,\Delta,\Delta^z\}} \int \left[\begin{array}{c} \tilde{c}_t(l(\theta), \omega(\theta) - \beta v(\theta), z(\theta), z_-, \theta) + M(s(\theta) - s_-) \\ -w(\theta, s(\theta), z(\theta)) l(\theta) + \frac{1}{R} K(v(\theta), \Delta(\theta), \Delta^z(\theta), \theta, s(\theta), z(\theta), t+1) \end{array} \right] f^t(\theta|\theta_-) d\theta \end{aligned} \tag{63}$$

$$\begin{aligned} \omega(\theta) &= u_t(c(\theta)) - \phi_t(l(\theta), z(\theta) - z_-) + \beta v(\theta) \\ \dot{\omega}(\theta) &= \frac{w_{\theta,t}}{w_t} l(\theta) \phi_{l,t}(l(\theta), z(\theta) - z_-) + \beta \Delta(\theta) \qquad [\mu(\theta)] \\ v &= \int \omega(\theta) f^t(\theta|\theta_-) d\theta \qquad [\lambda_-] \\ \Delta &= \int \omega(\theta) \frac{\partial f^t(\theta|\theta_-)}{\partial \theta_-} d\theta \qquad [\gamma_-] \\ \Delta^z &= -\int \left(\frac{w_{z,t} l(\theta)}{w_t} \phi_{l,t} - \beta \Delta^z(\theta) \right) f^t(\theta|\theta_-) d\theta \qquad [\gamma^Z_-] \\ \phi_{z,t} &= w_{z,t} l(\theta) \frac{1}{w_t} \phi_{l,t} - \beta \Delta^z(\theta) \ldots\ldots [\gamma^E(\theta)] \end{aligned}$$

With the Hamiltonian:

$$\begin{aligned} & \left(\tilde{c}(l(\theta), \omega(\theta) - \beta v(\theta), z(\theta), z_-, \theta) - w_t(\theta, s(\theta), z(\theta)) l(\theta)\right) f^t(\theta|\theta_-) \\ & + \frac{1}{R} K(v(\theta), \Delta(\theta), \Delta^z(\theta), \theta, s(\theta), z(\theta), t+1) f^t(\theta|\theta_-) \\ & + \lambda_- \left[v - \omega(\theta) f^t(\theta|\theta_-)\right] + \gamma_- \left[\Delta - \omega(\theta) \frac{\partial f^t(\theta|\theta_-)}{\partial \theta_-}\right] + \mu(\theta) \left[\frac{w_{\theta,t}}{w_t} l(\theta) \phi_{l,t} + \beta \Delta(\theta)\right] \\ & - \gamma^E(\theta) \left[w_{z,t} l(\theta) \frac{1}{w_t} \phi_{l,t} - \beta \Delta^z(\theta) - \phi_{z,t}\right] - \gamma^Z_- \left[\left(w_{z,t} l(\theta) \frac{1}{w_t} \phi_{l,t} - \beta \Delta^z\left(\theta^t\right)\right) f^t(\theta|\theta_-) + \Delta^z\right] \end{aligned}$$

Taking the FOC with respect to z_t, l_t and s_t (without yet making assumption 4) so as to have the general formulas: (recall that θ' and s', z' denote the subsequent period's choices):

$$
\begin{aligned}
[z(\theta)] \quad : \quad & -\frac{\phi_{z,t}}{u_t'(c(\theta))} + w_{z,t} l(\theta) - \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} \phi_{zz,t} \\
& - \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \left(l(\theta)\phi_{l,t} \frac{1}{w_t} w_{\theta z,t} - l(\theta)\phi_{l,t} w_{\theta,t} w_{z,t} \frac{1}{w_t^2} + \frac{w_{\theta,t}}{w_t} l(\theta)\phi_{lz,t} \right) \\
& + \frac{1}{R} \int \left(\frac{\phi_{z,t+1}}{u_{t+1}'(c(\theta'))} \right) f^{t+1}(\theta'|\theta) + \frac{1}{R} \int \gamma^E(\theta') \phi_{zz,t+1} \\
& + \frac{1}{R} \int \left(\frac{\mu(\theta')}{f^{t+1}(\theta'|\theta)} \frac{w_{\theta,t+1}}{w_{t+1}} l(\theta') \phi_{lz,t+1} \right) f^{t+1}(\theta'|\theta) \\
& - \frac{1}{R} \int \left(\gamma^Z(\theta) + \frac{\gamma^E(\theta')}{f^{t+1}(\theta'|\theta)} \right) \left(w_{z,t+1} l(\theta') \frac{1}{w_{t+1}} \phi_{lz,t+1} \right) f^{t+1}(\theta'|\theta) \\
= \quad & -\left(\gamma_-^Z + \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} \right) \left(w_{zz,t} l(\theta) \frac{1}{w_t} \phi_{l,t} - (w_{z,t})^2 l(\theta) \frac{1}{w_t^2} \phi_{l,t} + w_{z,t} l(\theta) \frac{1}{w_t} \phi_{lz,t} \right)
\end{aligned}
$$

$$
[l(\theta)] : \left(\frac{\phi_{l,t}}{u_t'(c(\theta))} - w_t \right) + \frac{\mu(\theta)}{f^t(\theta|\theta_-)} \frac{w_{\theta,t}}{w_t} \phi_{l,t} \frac{1+\varepsilon_t^u}{\varepsilon_t^c} - \left(\gamma_-^Z + \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} \right) \frac{w_{z,t}}{w_t} \frac{1+\varepsilon_t^u}{\varepsilon_t^c} \phi_{l,t} + \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} \phi_{zl,t} = 0
$$

$$
\begin{aligned}
[s(\theta)] \quad : \quad & M_t'(s(\theta) - s_-) - l(\theta) w_{s,t} - \frac{1}{R} \int M'\left(s'(\theta') - s(\theta)\right) f(\theta'|\theta) \, d\theta = \\
& \frac{\mu(\theta)}{f^t(\theta|\theta_-)} l(\theta) \phi_{l,t} \frac{1}{w_t^2} w_{\theta,t} w_{s,t} (1 - \rho_{\theta s,t}) + \left(\gamma_-^Z(\theta_-) + \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} \right) w_{z,t} w_{s,t} \frac{1}{w_t^2} l(\theta) \phi_{l,t} (\rho_{zs,t} - 1)
\end{aligned}
$$

The FOC with respect to Δ^z yields:

$$
[\Delta^z(\theta)] : \frac{1}{\beta R} \gamma^Z(\theta) = \gamma_-^Z + \frac{\gamma^E(\theta)}{f^t(\theta|\theta_-)} \tag{64}
$$

Suppose now that assumption 4 holds. From the definition of the wedge τ_{Zt} and the net wedge t_{zt}, as well as the formula for the optimal t_{zt}^* in (32), and the relation in (64), the FOC for z_t allows to write:

$$
\left(\gamma_{t-1}^Z + \frac{\gamma_t^E(\theta_t)}{f^t(\theta_t|\theta_{t-1})} \right) = -\frac{[-t_{zt} + t_{zt}^*]}{u'(c_t) \frac{w_{z,t}}{w_t} (\rho_{zz,t} - 1)} + \frac{\left[\frac{\gamma_t^E(\theta_t)}{f^t(\theta_t|\theta_{t-1})} \phi_{zz,t} - \frac{1}{R} \int \gamma_{t+1}^E(\theta_{t+1}) \phi_{zz,t+1} \right]}{w_{z,t} l_t u_t'(c_t)(1 - \tau_{Lt}) \frac{w_{z,t}}{w_t} (\rho_{zz,t} - 1)} \tag{65}
$$

Combining the FOCs for l and s yields:

$$
[l_t] : \left(\frac{\tau_{Lt}}{1 - \tau_{Lt}} - \frac{\tau_{Lt}^*}{1 - \tau_{Lt}^*} \right) \frac{w_t}{w_{z,t}} \frac{\varepsilon_t^c}{1 + \varepsilon_t^u} = -\left(\gamma_{t-1}^Z f^t(\theta_t|\theta_{t-1}) + \gamma_t^E(\theta_t) \right) = \frac{(t_{st} - t_{st}^*)}{u'(c_t)(\rho_{zs,t} - 1)} \frac{w_t}{w_{z,t}}
$$

or, canceling terms, we obtain the expression in proposition 8. For completeness, the full expressions for τ_{Lt} and t_{st} are:

$$
(t_{st} - t_{st}^*) = \frac{(\rho_{zs,t} - 1)}{(\rho_{zz,t} - 1)} (t_{zt}^* - t_{zt}) + \frac{(\rho_{zs,t} - 1)}{(\rho_{zz,t} - 1)} \frac{\left[\frac{\gamma_t^E(\theta_t)}{f^t(\theta_t|\theta_{t-1})} \phi_{zz,t} - \frac{1}{R} \int \gamma_{t+1}^E(\theta_{t+1}) \phi_{zz,t+1} \right]}{w_{z,t} l_t (1 - \tau_{Lt})}
$$

$$
\begin{aligned}
\left[\frac{\tau_{Lt}}{(1 - \tau_{Lt})} - \frac{\tau_{Lt}^*}{(1 - \tau_{Lt}^*)} \right] \quad = \quad & \frac{w_t}{w_{z,t}} \frac{1 + \varepsilon_t^u}{\varepsilon_t^c} \frac{1}{(\rho_{zz,t} - 1)} (t_{zt}^* - t_{zt}) \\
& - \frac{1 + \varepsilon_t^u}{\varepsilon_t^c} \frac{1}{(\rho_{zz,t} - 1)} \frac{\left[\frac{\gamma_t^E(\theta_t)}{f^t(\theta_t|\theta_{t-1})} \phi_{zz,t} - \frac{1}{R} \int \gamma_{t+1}^E(\theta_{t+1}) \phi_{zz,t+1} \right]}{w_{z,t} l_t (1 - \tau_{Lt})}
\end{aligned}
$$