

Education and Optimal Dynamic Taxation: The Role of Income-Contingent Student Loans*

Sebastian Findeisen

University of Mannheim and CEPR

Dominik Sachs

European University Institute

First version: October 21, 2011

This version: March 16, 2016

Abstract

We study the optimal design of integrated education finance and tax systems. The distribution of wages is endogenously determined by the costly education decisions of heterogeneous individuals before labor market entry. Consistent with empirical evidence, this human capital investment decision is risky. We find that an integrated education and tax system in which the government provides education loans to young individuals coupled with income-contingent repayment can always be designed in a Pareto optimal way. We present a simple empirically driven application of the framework to US data in which individuals make a college entry decision. We find the optimal repayment schemes for college loans can be well approximated by a schedule that is linearly increasing in income up to a threshold and constant afterwards. So although the full optimum could lead to complicated non-linear schedules in theory, very simple instruments can replicate it fairly well. The welfare gains from income-contingent repayment are significant.

JEL-classification: H21, H23, I21

Keywords: Optimal taxation, Education, Education Finance, Student Loans

*Contact: findeisen@uni-mannheim.de, dominik.sachs@eui.eu. We are grateful to the editor Wojciech Kopczuk and two anonymous referees for valuable comments and suggestions. We are also in particular grateful to our advisors Friedrich Breyer and Fabrizio Zilibotti for support and valuable comments. We further thank Manuel Amador, Dan Anderberg, Carlos da Costa, Emmanuel Farhi, Mike Golosov, Bas Jacobs, Leo Kaas, Sebastian Koehne, Normann Lorenz, Elena Mattana, Emmanuel Saez, Florian Scheuer, Dirk Schindler, Kjetil Storesletten, Aleh Tsyvinski, Matthew Weinzierl, Iván Werning, Christoph Winter and many seminar audiences for helpful discussions. We thank Stefan Voigt for valuable research assistance. We are grateful to Yale and Stanford for their hospitality, where parts of this paper were written. Dominik Sachs' research was funded by a post-doc fellowship of the Fritz-Thyssen Foundation and the Cologne Graduate School in Management, Economics and Social Sciences.

1 Introduction

How should governments design their higher education finance systems? There exist large differences across countries in the structure of higher education finance. In some countries, such as Denmark, Finland and Sweden, university and college students pay low or no tuition fees and in addition receive grants because of generous public subsidies for higher education. These countries have highly progressive tax systems, which allow to finance these education subsidies. By contrast, in the United Kingdom and the United States, e.g., the burden of educational costs mainly lies on the student and higher education is much less heavily subsidized by public finances. Instead, student loans offered by both the private and the public sector play a big part in financing higher education. From a policy perspective, the choice of an optimal education finance system is intimately linked to the tax system. Both underlie the same basic trade-offs, namely equity concerns in the form of redistribution and insurance against income risk versus efficiency concerns by distorting labor supply and education incentives.

In this paper, we address the optimal design of integrated education finance and tax systems. We build a novel optimal taxation framework in the spirit of Mirrlees (1971) and the vast literature following his footsteps, which allows to study the question from a new angle. In our framework, the distribution of wages is not exogenous but determined by the costly education decisions of individuals before labor market entry. Consistent with what is typically found in empirical studies, this human capital investment decision is risky. To solve the problem, we use an applied mechanism design approach. The benevolent government can observe total income and the education level of individuals, but it has to respect incentive compatibility – first, when individuals decide on education and second, when individuals decide on labor supply. The main novelty of our approach is that in our framework the government is not restricted to the use of predetermined instruments but is free to choose its own instruments, which can condition on education, income and savings. In addition, they are allowed to be fully nonlinear.

We find that an integrated education and tax system in which the government provides education loans to young individuals, coupled with income-contingent repayment rates of these loans after individuals enter the labor market, can effectively deal with all the major trade-offs underlying the education finance and tax problem. In other words, such systems can always be designed such that they are second-best Pareto efficient. This is because income-contingent repayment rates allow the government to *effectively differentiate tax distortions across education groups*, minimizing the efficiency cost of labor supply distortions. At the same time, it can subsidize education by varying the generosity of the loans.¹ Importantly, the government typically will find it optimal that some individuals partially default and never pay back the full value of their loans, while for some individuals the amount of repayment might exceed their loan values because this provides insurance.

¹We do not model credit market imperfection in the form of borrowing constraints. If these are relevant, as is still a debated question in the literature (Carneiro and Heckman, 2005), wide availability of student loans has the additional benefit of lifting these constraints.

We present a simple empirically driven application of the framework to US data in which individuals make a college entry decision. We simulate optimal income taxes and college student loans with income-contingent repayment. The optimal policy simulation provides three important insights. First, we find that the optimal repayment scheme for college loans can be well approximated by a schedule that is linearly increasing in income. So although the full optimum could lead to complicated nonlinear schedules in theory, very simple instruments can replicate it fairly well. Second, for our benchmark parameterization college graduates find it optimal to participate voluntarily in the loan schemes as compared to taking a risk-free loan on the private market. Third, we calculate the welfare gains of moving from a third-best scenario where the government optimally sets the income tax and offers a loan system with non-contingent repayment to the system with contingent repayments. We find welfare gains ranging from about 0.2% to 0.6% of lifetime consumption and we show how these gains vary with risk-aversion.

Several countries like the United Kingdom, Australia and New Zealand currently administer income-contingent college student loans, where repayment is proportional to income.² Very recently in the United States, “the student loan industry was effectively nationalized by provisions of the Health Care and Education Reconciliation Act of 2010” (Brooks 2015, p. 251). Under the new system, student loan programs are directly administered by the Department of Education. The possibilities to opt for income-based repayment have increased since then. Whereas different options exist, they all have in common that repayment is capped at between 10 and 15% of income and they all include loan forgiveness of the remaining debt after 20-25 years (Brooks 2015). Our framework gives these policies a theoretical second-best foundation, based on an applied mechanism design approach to the education finance and taxation problem. Our theoretical considerations suggest that it might be optimal for the government to enforce that very rich individuals pay back more than the capitalized loan value or that repayment might actually be decreasing in income. In the mentioned countries, repayment never exceeds the loan value and repayment schedules are non-decreasing in income. To address these issues, we also consider policy experiments in which we restrict income-contingent repayment not to exceed the actual loan value and to be non-decreasing in income. We find that a large share of the welfare gains from the full optimum can be reaped with these simpler policies and that they are similar to current policies in the U.S.: the marginal repayment rate is 10.5% on average.

More generally, a contribution of this paper is to extend existing studies on taxation and human capital (see the literature review below) by (i) considering ex-ante heterogeneity and uncertainty and (ii) by explicitly looking at education decisions along the extensive margin. The latter is in our view necessary to model the decision to go to college. Certainly, the college decision is not only binary in the real world. Important factors are the quality of college, the major of study, the length of study and learning effort during college – it is a multi-dimensional

²Chapman (2006) provides a survey for practices in those and other countries. Barr (2004) discusses the trade-offs involved in designing these programs. To the best of our knowledge, the first economist to endorse the idea was Milton Friedman (1955). He envisioned repayment amounts to be proportional to income, i.e. a linearly increasing repayment schedule. Something we find as an optimal policy in our simulation for the most part of the income distribution.

decision problem. Modeling education as a binary instead of a continuous variable (as usually done in the literature) is an important complementary comparison case and a necessary step towards more realistic models. Concerning (i), the joint consideration of ex-ante heterogeneity (to have some people going to college and some not) and income risk (to capture the riskiness of educational investment) is crucial to think about the desirability of income-contingent student loans. Having uncertainty in the model is necessary to include the insurance rationale of income-contingent repayment. On top, only the presence of ex-ante heterogeneous individuals with and without a college degree makes it possible to study a realistic loan repayment system, in which income contingency implies that workers with the same income face different effective marginal tax rates.

Relation To Existing Literature. This paper makes a contribution to the literature on optimal income taxation starting from Mirrlees (1971) (see the recent survey of Piketty and Saez (2013)). In Section 3 we discuss how the expression for optimal education-dependent marginal tax rates compares to the seminal optimal tax formulas from Diamond (1998) and Saez (2001) with exogenous human capital.

Bovenberg and Jacobs (2005) and Jacobs and Bovenberg (2011) analyze how endogenous education alters the optimal tax problem and show for which specifications of the earnings function education should be subsidized at a higher or lower rate than the tax distortion. Bohacek and Kapicka (2008) study a dynamic model with certainty and obtain similar results regarding education subsidies. These articles work under certainty whereas we take idiosyncratic human capital risk into account. Importantly, with idiosyncratic education risk, the necessity of education dependent labor wedges and income-contingent loans arises, as intuitively they can be understood as providing an additional source of insurance. As we discuss in Section 2.1, when we review some stylized empirical facts, there is strong evidence that uncertainty about college returns is important and matters for human capital investment decisions.³

Best and Kleven (2013) and Kapicka and Neira (2015), study how human capital acquisition at the working age influence the optimal taxation problem. We focus on a different part of the human capital accumulation process, namely education before labor market entry. Importantly, both papers reasonably assume that tax policies cannot directly condition on human capital acquired while working. In contrast, we allow the government to use information about education before labor market entry in the tax code, as is done in the real world in some countries in the form of student loans with income-contingent repayment. In addition, our focus is on education finance instead of only tax policies.

³One strand of literature has looked at first- versus second-best investment rules of human capital under risk with ex-ante homogenous agents. Da Costa and Maestri (2007) show that human capital should always be encouraged in the second-best optimum. Anderberg (2009) emphasizes that the risk properties of human capital are crucial for the question whether and how education should be distorted relative to a first-best rule. Focusing on linear policy instruments, Anderberg and Andersson (2003) as well as Jacobs et al. (2012) obtain similar results. An early treatment how taxes affect the risk properties of human capital investment is Eaton and Rosen (1980). Grochulski and Piskorski (2010) focus on the implications of unobservable human capital investment for capital taxation in an ex-ante homogeneous agent setting with uncertainty. Kapicka (2006) introduces non-observable endogenous human capital into a dynamic, non-stochastic Mirrlees model where taxes can only be conditioned on current income. He shows that marginal tax rates are lowered due to the education margin.

Stantcheva (2015) studies second-best optimal policies in a rich dynamic-stochastic environment. She considers both, observable and unobservable human capital. Her focus is rather on human capital accumulation over the life-cycle, whereas our focus is on college education.

Working with a two-type model, Gary-Bobo and Trannoy (2015) come to a similar conclusion concerning the income-contingency of loans in a very recent paper. In contrast to their work, we employ a continuous type approach with continuous skill and income distributions in the tradition of the large literature on optimal income taxation as in Mirrlees (1971) and Saez (2001). In particular, we are interested in determining the forces shaping the optimal design of student loan policies both theoretically and numerically, which requires a model with continuous types.

Concerning the implementation of history-dependent allocations, this paper is related to Golosov and Tsyvinski (2006) who consider an environment with absorbing disability shocks and present an implementation in which disability insurance conditions on asset testing. Also in the context of optimal taxation, Scheuer (2014) considers differential taxation of profits and labor income; in our case a comparable logic applies for an endogenous education instead of an occupational choice.

Finally, taking a quantitative approach and working in the Ramsey tradition with simpler but given policy instruments, Krueger and Ludwig (2013) solve for the optimal income tax and education subsidies in a rich macro model.

This paper is organized as follows. Section 2 contains the basics of the model. In Section 3, we investigate dynamic incentive compatibility and describe the major properties of constrained efficient allocations. Decentralized implementations of constrained efficient allocations are provided in Section 4. Simpler policies with history-independent labor wedges (implying loan repayment that does not condition on income) are theoretically discussed in Section 5. We bring our model to US-data and simulate optimal policies in Section 6. Section 7 concludes.

2 The Model

2.1 Structure

We consider a simple life-cycle model, in which individuals first acquire education and work afterwards. Individuals differ in innate ability $\theta_i \in \{\theta_l, \theta_h\}$, which is private information and can be interpreted as a one dimensional aggregate of (non-) cognitive skills, I.Q. and family background.⁴ We refer to the θ_h type as the high type and to θ_l as the low type. Sometimes we also call θ_h the college type and θ_l the high-school type as our quantitative exploration of the model will be to college/high-school context. The share of individuals is given by f_l and f_h , where $f_l + f_h = 1$. In the first period, individuals make a binary education decision $e_i \in \{e_l, e_h\}$, where e_l and e_h reflect resource costs. One can think of e_h as graduating from

⁴In the working paper version of this paper, we show how the results generalize to continuous type space in θ (Findeisen and Sachs 2013).

college and of e_l as entering the labor market directly after high-school. We focus on separating allocations where the high-type is incentivized to take the high education level e_h and the low type is incentivized to take the low education level e_l . The main question we address in this paper is therefore not about the optimal level of education spending but rather how to set optimal incentives for education, labor supply and savings given a realistic modeling of the relevant margins of heterogeneity and uncertainty.

Flow utility during education is denoted by $u^e(c^e)$ with $u_c^e > 0, u_{cc}^e < 0$, where c^e is consumption in the education period. When individuals enter the labor market, they draw their labor market ability a from a continuous conditional cumulative distribution function (*cdf*) $G(a|e_i, \theta_i)$, which depends on *innate* ability θ_i and education e and has bounded support $[\underline{a}, \bar{a}]$, with $\underline{a} \geq 0$. For the working period, we assume that preferences over consumption and leisure are given by the utility function $u^w(c^w, l)$, where labor effort l is equal to $\frac{y}{a}$, so that gross income is $y = a \times l$. We assume that $u^w(\cdot, \cdot)$ obeys the Spence-Mirrlees condition.

Expected lifetime utility of an individual of type θ_i is given by

$$\beta_i^e u^e(c_i^e) + \beta_i^w \int_{\underline{a}}^{\bar{a}} u^w\left(c_i^w(a), \frac{y_i(a)}{a}\right) dG(a|\theta_i, e_i) \quad \forall i = l, h,$$

where β_i^e and β_i^w reflect discounting and the different period length of the education and working period. Thus, the education decision also determines the amount time needed to complete the education stage. For example, in the simulations graduating from college will take 5 years as in the US data, whereas high-school graduates start working directly. For the theory, assume that the education period lasts T_i^e years and the working period T_i^w years. Then β_i^e and β_i^w can be thought of as $\beta_i^e = \sum_{t=1}^{T_i^e} \beta^{t-1}$ and $\beta_i^w = \sum_{t=T_i^e+1}^{T_i^e+T_i^w} \beta^{t-1}$, where β is the yearly discount factor. The rate of transformation for goods (the implicit interest rate) between the education and working period is R_i . We set $R_i = \frac{\beta_i^e}{\beta_i^w}$, which is the same standard assumption as setting $\beta R = 1$ in a model with equal period length. Essentially, our set-up corresponds to a two-period model, where the lengths of the periods are allowed to differ.

We capture many empirical regularities with this specification of the model. First, assuming $G(a|\theta_i, e_i)$ to be non-degenerate, our model captures the important fact of uncertainty in the labor market and risky educational investment. See e.g. Cunha and Heckman (2008) or Chen (2008) for recent contributions.

Second, we allow this *cdf* to be a function of innate ability θ_i and thereby capture the fact that inequality in earnings is – to a certain extent – also determined by innate ability. Taber (2001) and Hendricks and Schoellman (2014) suggest that much of the rise in the college premium may be attributed to a rise in the demand for unobserved skills, which are predetermined and independent of education. Indirect evidence for the importance of unobserved skills comes from the strong persistence of within-education-group inequality (Acemoglu and Autor 2011).

Third, the *cdf* G being a function of e_i captures the returns to education. Importantly, for most of our results, we do not impose a certain assumption on the pattern of these returns.

Fourth, as long as $G(a|\theta_h, e_h) - G(a|\theta_h, e_l) \neq G(a|\theta_l, e_h) - G(a|\theta_l, e_l)$, returns to educational investment differ in innate ability θ_i . E.g., Carneiro and Heckman (2005) document that the returns can differ by as much as 19% points across individuals for one year of college.

2.2 Informational Asymmetries and Incentive Compatibility

We cast the problem as a sequential mechanism – agents report an initial type θ_i in the education period and, after uncertainty has materialized, report their productivity a in the working period. The planner assigns initial consumption levels c_i^e and education levels e_i to individuals with innate ability $\theta_i \forall i = l, h$. Moreover, with each report there comes a sequence of utility promises for the next period $\{v^w(\theta_i, a)\}_{a \in [\underline{a}, \bar{a}]}$. In the second period, the screening takes place over consumption levels $c_i^w(a)$ and labor supply $y_i(a)$. All these quantities define an allocation in the economy. Further, note that this implies that education, consumption and income are assumed to be observable. Assuming consumption to be observable implies that we either assume that there are no private markets for savings or that there are (potentially imperfect) capital markets and the planner can observe the amount of savings. Dynamic incentive compatibility is ensured backwards, so we start analyzing the problem from the second period.

2.2.1 Working Period Incentive Compatibility

By the revelation principle, we can restrict attention to direct mechanisms. Suppose that in the first period agents have made truthful reports $r_\theta(\theta_i) = \theta_i \forall i = l, h$, albeit this is not necessary and just simplifies the exposition.⁵ Conditions for this to be true are given in the next subsection. Conditional on this report, the second period incentive constraint must be met for any history of types (θ_i, a) and reporting strategy $r_a(a)$:

$$u^w \left(c_i^w(a), \frac{y_i(a)}{a} \right) \geq u^w \left(c_i^w(r_a(a)), \frac{y_i(r_a(a))}{a} \right) \quad \forall i = l, h, \forall a, r_a(a) \in [\underline{a}, \bar{a}].$$

Like in a standard Mirrlees problem, we assume that preferences satisfy single-crossing for given first-period reports. For global incentive compatibility it is, hence, necessary and sufficient that all local envelope conditions hold:

$$\frac{\partial v^w(\theta_i, a)}{\partial a} = -u_l^w \left(c_i^w(a), \frac{y_i(a)}{a} \right) \frac{y_i(a)}{a^2} \quad \forall i = l, h, \forall a \in [\underline{a}, \bar{a}] \quad (1)$$

and the usual monotonicity condition, stating that $y_i(a)$ is non-decreasing in ability levels a , is satisfied:

$$\frac{dy_i(a)}{da} \geq 0 \quad \forall i = l, h, \forall a \in [\underline{a}, \bar{a}]. \quad (2)$$

⁵The reason is that in the second period the utility is a function of $a, r_a(a)$ and $r_\theta(\theta_i)$ but not of θ_i .

2.2.2 Education Period Incentive Compatibility

In the education period, an agent takes into account the effect of her report about θ_i on future utility. Education period incentive compatibility is ensured only if both of the following incentive constraints hold:

$$\begin{aligned} & \beta_l^e u^e(c_l^e) + \beta_l^w \int_{\underline{a}}^{\bar{a}} u^w\left(c_l^w(a), \frac{y_l(a)}{a}\right) dG(a|\theta_l, e_l) \\ & \geq \beta_h^e u^e(c_h^e) + \beta_h^w \int_{\underline{a}}^{\bar{a}} u^w\left(c_h^w(a), \frac{y_h(a)}{a}\right) dG(a|\theta_l, e_h) \end{aligned} \quad (3)$$

and

$$\begin{aligned} & \beta_h^e u^e(c_h^e) + \beta_h^w \int_{\underline{a}}^{\bar{a}} u^w\left(c_h^w(a), \frac{y_h(a)}{a}\right) dG(a|\theta_h, e_h) \\ & \geq \beta_l^e u^e(c_l^e) + \beta_l^w \int_{\underline{a}}^{\bar{a}} u^w\left(c_l^w(a), \frac{y_l(a)}{a}\right) dG(a|\theta_h, e_l). \end{aligned} \quad (4)$$

Which of these constraints is binding in equilibrium will depend on set of Pareto weights assigned to the different types θ_l and θ_h . For social welfare functions which are commonly used in the literature such as the Utilitarian or Rawlsian, the relevant constraint is usually (4). In Appendix A.1, we present a set of sufficient conditions that guarantee that constraint (3) is always fulfilled if (4) is fulfilled so it does not have to be included in the planning problem. Finally, note that mimicking the other type here implies having the other type's discount factor as the same amount of time used for a specific education level is independent of the type.

3 Constrained Pareto Optimal Allocations

In this section, we characterize constrained Pareto-efficient allocations, where “constrained” refers to the asymmetric information problem that the government faces and which results in incentive compatibility constraints as discussed in Section 2.2. In Subsection 3.1, we set up the problem of the government. In Subsections 3.2-3.4, we analyze Pareto-optimal distortions of the labor supply, the savings and the education decision.

3.1 The Planning Problem

To characterize the whole Pareto frontier, we assign Pareto weights \tilde{f}_h and $\tilde{f}_l = 1 - \tilde{f}_h$ to the different types. Any distribution of these weights corresponds to one point on the Pareto frontier. We assume that the planner discounts the future at the same rate as individuals. The planning problem therefore reads as

$$\begin{aligned} \max \quad & \tilde{f}_l \times \left[\beta_l^e u^e(c_l^e) + \beta_l^w \int_{\underline{a}}^{\bar{a}} u^w \left(c_l^w(a), \frac{y_l(a)}{a} \right) dG(a|\theta_l, e_l) \right] \\ & + \tilde{f}_h \times \left[\beta_h^e u^e(c_h^e) + \beta_h^w \int_{\underline{a}}^{\bar{a}} u^w \left(c_h^w(a), \frac{y_h(a)}{a} \right) dG(a|\theta_h, e_h) \right] \end{aligned}$$

subject to the resource constraint:

$$\begin{aligned} f_l \times & \left[-\beta_l^e (c_l^e + e_l) + \beta_l^w \int_{\underline{a}}^{\bar{a}} (y_l(a) - c_l^w(a)) dG(a|\theta_l, e_l) \right] \\ + f_h \times & \left[-\beta_h^e (c_h^e + e_h) + \beta_h^w \int_{\underline{a}}^{\bar{a}} (y_h(a) - c_h^w(a)) dG(a|\theta_h, e_h) \right] \geq 0 \end{aligned}$$

and the incentive constraints (1), (3) and (4).

Consistent with standard practice in screening problems with a continuous type space, our strategy for solving the second-best problem is to work with a relaxed problem with only restriction (1) being imposed for the working period and then check ex-post in our numerical applications whether incentive compatibility is fulfilled by checking (2). Indeed in our numerical explorations we find that incentive compatibility is always satisfied and therefore the first-order approach is valid for the primitives we consider.

λ denotes the multiplier on the resource constraint and η_h and η_l the multipliers of the first-period incentive compatibility conditions. To keep the exposition simple, we focus on the case where only (4) is binding, so that $\eta_h > 0$ and $\eta_l = 0$.

3.2 Labor Wedge

The optimal labor wedge is history dependent, so in addition to the working period skill level a , it depends on θ_i . The labor wedge is positive (negative) if an individual works less (more) than she would at the intervention-free market price (which is her productivity level a). Formally, it reads as:

$$\tau_i^y(a) = 1 - \frac{u_l^w \left(c_i^w(a), \frac{y_i(a)}{a} \right) \frac{1}{a}}{u_c^w \left(c_i^w(a), \frac{y_i(a)}{a} \right)} \quad \forall a, i.$$

The following proposition characterizes the optimal labor wedge.⁶

⁶Golosov et al. (2013) provide formulas for dynamic optimal labor wedges with exogenous human capital, connecting them to empirical observables in the spirit of the contributions of Diamond (1998) and Saez (2001) for the static Mirrlees model.

Proposition 1. *At any constrained Pareto optimum where $\eta_h > 0$, labor wedges at income level y for type θ_i with $i = l, h$ satisfy:*

$$\frac{\tau_i^y(a)}{1 - \tau_i^y(a)} = \frac{1 + \varepsilon_i^u(a)}{\varepsilon_i^c(a) a g(a|\theta_i, e_i)} \int_a^{\bar{a}} \exp\left(\int_a^x \left(1 - \frac{\varepsilon_i^u(s)}{\varepsilon_i^c(s)}\right) \frac{y_i'(s)}{y_i(s)} ds\right) \{\mathcal{A}_i(x) + \mathcal{B}_i(x)\} dx,$$

where $\mathcal{A}_i(x) = g(x|\theta_i, e_i) \left(1 - \frac{u_c^w(c_i(x), \frac{y_i(x)}{x})}{u_c^e(c_i^e)}\right)$, $\mathcal{B}_l(x) = u_c^w\left(c_l(x), \frac{y_l(x)}{x}\right) \eta_h / \lambda \{g(a|\theta_h, e_l) - g(a|\theta_l, e_l)\}$, $\mathcal{B}_h(x) = 0$ and $\varepsilon_i^u(a)$ ($\varepsilon_i^c(a)$) is the uncompensated (compensated) labor supply elasticity of type (θ_i, a) and the optimal values for the Lagrangian multipliers are given by

$$\eta_h = \frac{\tilde{f}_l u_c^e(c_l^e) - \tilde{f}_h u_c^e(c_h^e)}{\frac{u_c^e(c_h^e)}{f_h} + \frac{u_c^e(c_l^e)}{f_l}} \text{ and } \lambda = \frac{1}{\left(f_l \frac{1}{u_c^e(c_l^e)} + f_h \frac{1}{u_c^e(c_h^e)}\right)}.$$

Proof. See Appendix A.2.1. □

First, consider the labor wedge of the high type θ_h . Since $\mathcal{B}_h(x) = 0$, the formula shows close resemblance with the standard Mirrlees formula (Saez 2001). Optimal effective marginal tax rates on labor income are decreasing in the compensated elasticity and larger for higher values of risk aversion. Additionally, the shape of these effective marginal tax rates crucially depends on the distribution of skills for college graduates $g(a|\theta_h, e_h)$. The only difference to the static formula is the marginal utility of consumption during the education period that shows up in $\mathcal{A}_h(a)$; it replaces the Lagrangian multiplier of the resource constraint as compared to the static formula.

For the low type, we can clearly see how the government has to take into account incentive compatibility in the education period when designing labor wedges for the working period. This is captured by the term $\mathcal{B}_l(x)$. This term is proportional to the value of the Lagrangian multiplier on the incentive constraint η_h . The Lagrangian multiplier on the incentive constraint η_h is positive and larger the larger the inequality in consumption between the high and low type and the larger the Pareto weight \tilde{f}_l (and therefore the lower \tilde{f}_h). In other words, the stronger the desire of the planner to redistribute from the high type to the low type at the margin, the higher the value of relaxing the incentive constraint and the larger is η_h .⁷ The role of this additional term for self-selection is particularly intuitive for the case where preferences between consumption and labor are separable:

Corollary 1. *Assume that preferences satisfy $u^w(c^w, \frac{y}{a}) = u(c^w) - \Psi(\frac{y}{a})$ where $\Psi'(\cdot), \Psi''(\cdot), u'(\cdot) > 0$ and $u''(\cdot) < 0$. At any constrained Pareto optimum where $\eta_h > 0$, labor wedges at income level y for the low type θ_l satisfy:*

$$\frac{\tau_l^y(a)}{1 - \tau_l^y(a)} = \frac{1 + \varepsilon_l^u(a)}{\varepsilon_l^c(a)} \frac{u_c^w(c_l^w(a), \frac{y_l(a)}{a})}{a g(a|\theta_l, e_l)} [\mathcal{A}_l(a) + \mathcal{B}_l(a)],$$

⁷In the polar case where the incentive constraint is binding from low to high, in contrast, this additional distortion shows up for the high type and is zero for the low type.

where

$$\mathcal{A}_l(a) = g(a|\theta_l, e_l) \left[\int_a^{\bar{a}} \frac{1}{u_c^w \left(c_l^w(a^*), \frac{y_l(a^*)}{a^*} \right)} dG(a^*|\theta_l, e_l) \right. \\ \left. - \frac{1 - G(a|\theta_l, e_l)}{G(a|\theta_l, e_l)} \int_{\underline{a}}^a \frac{1}{u_c^w \left(c_l^w(a^*), \frac{y_l(a^*)}{a^*} \right)} dG(a^*|\theta_l, e_l) \right]$$

$$\mathcal{B}_l(a) = \frac{\eta_h}{\lambda f_l} [G(a|\theta_l, e_l) - G(a|\theta_h, e_l)].$$

Proof. See Appendix A.2.2. □

Notice the term $\mathcal{B}_l(a)$ is proportional to the difference of the *cdf* of the low type and the counterfactual distribution of the high type if she deviates and mimics the low type with educational level e_l . The larger this difference (i.e. the stronger the difference between the high type and the low type), the higher marginal tax rates are for the low type. Intuitively, if the low type is very unlikely to be of higher ability than some level a^* , but the high type is very likely – even if she chooses e_l – then having high labor distortions for the low type with $a \geq a^*$ is not very costly in terms of distorting labor supply of the low type but very efficient in deterring the high type from mimicking. Term $\mathcal{A}_l(a)$ again captures the standard forces from the static model. We relate it to Saez (2001) in Appendix A.2.2.

Finally, we also consider a special case that has received a lot of attention in the literature: preferences without income effects on labor supply.

Corollary 2. *Assume that preferences satisfy $u^w(c^w, \frac{y}{a}) = u(c^w - \Psi(\frac{y}{a}))$ where $\Psi'(\cdot), \Psi''(\cdot), u'(\cdot) > 0$ and $u''(\cdot) < 0$.*

At any constrained Pareto optimum where $\eta_h > 0$, labor wedges at income level y for the high type θ_i with $i = l, h$ satisfy:

$$\frac{\tau_i^y(a)}{1 - \tau_i^y(a)} = \frac{1 + \varepsilon_i^c(a)}{\varepsilon_i^c(a) a g(a|\theta_i, e_i)} \int_a^{\bar{a}} \{ \mathcal{A}_i(x) + \mathcal{B}_i(x) \} dx.$$

where all terms are defined as in Proposition 1.

We omit a proof because it follows directly from Proposition 1. In line with previous findings of the literature, optimal tax formulas become significantly simpler in the absence of income effects. This corollary will also serve as an important benchmark for history-independent labor wedge policies that we will study in Section 5 and for which we assume away income effects for tractability.

3.3 Savings Wedge

We now look at optimal distortions of an individual's Euler equation between the education and the working period

$$\tau_i^s = 1 - \frac{u_c^e(c_i^e)}{\int_{\underline{a}}^{\bar{a}} u_c^w\left(c_i^w(a), \frac{y_i(a)}{a}\right) g(a|\theta_i, e_i) da} \quad \forall i.$$

$\tau_i^s > (<) 0$ implies a downward (upward) distortion of savings.

The following proposition transparently displays how our findings about the savings distortion connect to the recent literature.

Proposition 2. *In any Pareto-optimal allocation, the intertemporal allocation of consumption is governed by*

$$\frac{1}{u_c^e(c_i^e)} = \int_{\underline{a}}^{\bar{a}} \frac{1}{u_c^w\left(c_i^w(a), \frac{y_i(a)}{a}\right)} dG(a|\theta_i, e_i) - \int_{\underline{a}}^{\bar{a}} \frac{\tau_i^y(a)}{1 - \tau_i^y(a)} \frac{y_i(a)}{a} \frac{\varepsilon_i^c(a)}{1 + \varepsilon_i^u(a)} \frac{u_{c,l}^w\left(c_i^w(a), \frac{y_i(a)}{a}\right)}{u_c^{w^2}\left(c_i^w(a), \frac{y_i(a)}{a}\right)}.$$

Proof. See Appendix A.2.3 □

In case of separable preferences (i.e. $u_{c,l}^w = 0$), we obtain the famous inverse Euler equation (Diamond and Mirrlees, 1978; Rogerson, 1985; Golosov et al. 2003), which implies a positive savings wedge and holds in very general dynamic, stochastic, private information settings as the latter authors have shown. One intuition is that a small positive savings distortion has only a second-order effect on individual utility, but a first-order effect on incentive constraints. For the case of linear policy instruments, Jacobs and Schindler (2012) have derived a related result that capital should be taxed to boost labor supply, which in turn has a positive fiscal externality through labor income taxes.

In case of non-separable preferences, i.e. $u_{c,l}^w \neq 0$, the inverse Euler equation is augmented by an additional term whose sign depends on the sign of $u_{c,l}^w$. If consumption and labor are Pareto-Edgeworth substitutes ($u_{c,l}^w < 0$) – which implies that more consumption increases the disutility of labor –, the savings wedge is larger than according to the inverse Euler logic. The intuition is simple: if $u_{c,l}^w < 0$, lower savings boost labor supply (put differently, it relaxes incentive constraints in the working period). In case of Pareto-Edgeworth complements this effect works the other way around and therefore the sign of the savings wedge is ambiguous in this case.

We now provide an alternative characterization of the savings distortion, which gives a direct expression for the savings wedge and which also highlights the self-selection problem in the first period:

Proposition 3. *In any Pareto-optimal allocation with $\eta_h > 0$ and $\eta_l = 0$, savings wedges are given by*

$$\frac{\tau_h^s}{1 - \tau_h^s} = \int_a^{\bar{a}} \frac{\tau_h^y(a)}{1 - \tau_h^y(a)} \frac{a \varepsilon_h^c(a)}{1 + \varepsilon_h^u(a)} \left(\frac{u_{cl}^w \left(c_h(a), \frac{y_h(a)}{a} \right)}{u_c^w \left(c_h(a), \frac{y_h(a)}{a} \right)} \frac{y_h'(a)}{a} - \frac{u_{cc}^w \left(c_h(a), \frac{y_h(a)}{a} \right)}{u_c^w \left(c_h(a), \frac{y_h(a)}{a} \right)} c_h'(a) \right) da$$

$$\begin{aligned} \frac{\tau_l^s}{1 - \tau_l^s} &= \int_a^{\bar{a}} \frac{\tau_l^y(a)}{1 - \tau_l^y(a)} \frac{a \varepsilon_l^c(a)}{1 + \varepsilon_l^u(a)} \left(\frac{u_{cl}^w \left(c_l(a), \frac{y_l(a)}{a} \right)}{u_c^w \left(c_l(a), \frac{y_l(a)}{a} \right)} \frac{y_l'(a)}{a} - \frac{u_{cc}^w \left(c_l(a), \frac{y_l(a)}{a} \right)}{u_c^w \left(c_l(a), \frac{y_l(a)}{a} \right)} c_l'(a) \right) da \\ &\quad + \frac{\eta_h}{\lambda f_l} \int_a^{\bar{a}} u_c^w \left(c_l^w(a), \frac{y_l(a)}{a} \right) \{g(a|\theta_h, e_l) - g(a|\theta_l, e_l)\} da. \end{aligned}$$

Proof. See Appendix A.2.4 □

First consider the savings wedge of the high type. It captures the impact of savings on future labor supply. The second term in brackets $-\frac{u_{cc}^w}{u_c^w} c_h'(a)$ captures the income effect on future labor supply, a force for a positive savings wedge. The first term in brackets $\frac{u_{cl}^w}{u_c^w} \frac{y_h'(a)}{a}$ captures the second effect on labor supply coming from the non-separability of the utility function – as discussed above, if consumption and labor are Pareto-Edgeworth substitutes (complements) this effect works in favor of a positive (negative) savings wedge.

For the low type, both of these effects also show up. In addition, however, there is also a force to distort the savings decision to relax the self-selection problem in the education period. This is a force towards a negative savings wedge, i.e. an implicit savings subsidy. What determines this? A high type that mimicks a low type in the education period always would like to save less than a truth-telling low type. The reason is a higher expected income in the working period. Subsidizing savings therefore hurts the high-type mimicker more than the low-type truth-teller.⁸ This intuition can be related to the results from Jacobs and Bovenberg (2010), who find a motive for positive capital taxation to boost education investment incentives. In our discrete model, the savings wedge on the low type relaxes incentive constraints a the margin.

Finally, we look at the case, where there are no income effects on labor supply, in which case things simplify a lot.

Corollary 3. *If preferences satisfy $u^w(c^w, \frac{y}{a}) = u(c^w - \Psi(\frac{y}{a}))$ where $\Psi'(\cdot), \Psi''(\cdot), u'(\cdot) > 0$ and $u''(\cdot) < 0$, then in any Pareto optimal allocation with $\eta_h > 0$ and $\eta_l = 0$, we have $\tau_h^s = 0$ and*

$$\frac{\tau_l^s}{1 - \tau_l^s} = \frac{\eta_h}{f_l \lambda} \int_a^{\bar{a}} u' \left(c_l^w(a) - \Psi \left(\frac{y_l(a)}{a} \right) \right) (g(a|\theta_l, e_l) - g(a|\theta_h, e_l)) da.$$

⁸Note that there is no such effect for the high type which follows a “no distortion at the top” logic. On other points of the Pareto frontier, where the incentive constraint of the low type is binding, the result would turn: there would be no such additional effect for the low type, however, the savings decision of the high type would be taxed at the margin to relax the incentive constraint of the low type.

Proof. In case of no income effects, we have $u_{cl}^w = -u''\Psi'$ and $u_{cc}^w = u''$, which gives $\left(\frac{u_{cl}^w}{u_c^w} \frac{y_i'(a)}{a} - \frac{u_{cc}^w}{u_c^w} c_h'(a)\right) = \frac{u''}{u'} \left(\Psi' \frac{y'(a)}{a} - c'(a)\right) = 0$, where the last equality follows from working-period incentive compatibility. Therefore, Corollary 3 directly follows from Proposition 3. \square

Intuitively, the amount of wealth in the working period now has no impact on the labor supply decision. The only reason to distort the savings margin is therefore to relax the incentive constraint in the first period. Consequentially, the savings decision of the high type is undistorted and that of low type is subsidized.

3.4 Education Distortion

The characterization of education distortions resembles the logic behind the extensive margin labor supply model (Diamond (1980), Saez (2002)), where individuals can choose to work or not. In particular, one can define an education tax similar to the binary labor supply case, where a participation tax can be defined (Choné and Laroque (2011)). In the labor supply model, a participation tax is defined as the increase in resources due to labor market participation minus the increase in consumption of the respective individual due to labor market participation.

In our case, the former – i.e. the impact on resources due to type θ_h choosing e_h over e_l – is given by:⁹

$$\Delta R = \beta_h^w \int_{\underline{a}}^{\bar{a}} y_h(a) dG(a|\theta_h, e_h) - \beta_l^w \int_{\underline{a}}^{\bar{a}} y_l(a) dG(a|\theta_h, e_l) - (\beta_h^e e_h - \beta_l^e e_l).$$

First, resources change by the expected discounted change in income – taking into account that more education takes more time. Second, resources decrease since more education is costly.

To obtain the education tax ΔT , one has to subtract the value of the change in consumption:

$$\Delta T = \Delta R - \beta_h^w \int_{\underline{a}}^{\bar{a}} c_h(a) dG(a|\theta_h, e_h) - \beta_h^e c_h^e + \beta_l^w \int_{\underline{a}}^{\bar{a}} c_l(a) dG(a|\theta_h, e_l) + \beta_l^e c_l^e.$$

This concept of this education tax ΔT measures the public gains from education and captures three different things: (i) while on the labor market, the income distribution of better educated individuals differs – we label this the return effect. (ii) Better educated individuals spend less time on the labor market – we label this the time effect. (iii) Even holding time on the labor market and the distribution of wages fixed, the contribution to public funds differs because the allocation conditions on education – we label this the policy effect.

$$\Delta T = \Delta T_{return} + \Delta T_{time} + \Delta T_{policy},$$

where

⁹We could define everything equivalently for the low type as well. Only subscripts h would be l instead. But given that we focus on separating allocations – i.e. where the high type always obtains the high education level and the low type obtains the low education level – we focus on the high type in this section.

$$\Delta T_{return} = \beta_l^w \int_{\underline{a}}^{\bar{a}} (y_l(a) - c_l^w(a)) (g(a|\theta_h, e_h) - g(a|\theta_h, e_l)) da$$

and

$$\Delta T_{time} = (\beta_l^e - \beta_h^e) (c_l^e + e_l) + (\beta_l^e - \beta_h^e) \int_{\underline{a}}^{\bar{a}} (c_l^w(a) - y_l(a)) dG(a|\theta_h, e_l)$$

and

$$\Delta T_{policy} = \beta_h^e (c_l^e + e_l - (c_h^e + e_h)) + \beta_h^w \int_{\underline{a}}^{\bar{a}} ((y_h(a) - c_h^w(a)) - (y_l(a) - c_l^w(a))) dG(a|\theta_h, e_h).$$

ΔT_{return} captures the contribution to public funds due to higher earnings through better education: holding the difference between gross income and consumption constant $y_l(a) - c_l^w(a)$, the distribution ameliorates from $G(a|\theta_h, e_l)$ to $G(a|\theta_h, e_h)$.

ΔT_{time} captures the impact on public funds through the time investment of education. Holding the distribution of a and the allocation variables constant, it captures how much individuals contribute to public funds solely through this time channel.

ΔT_{policy} captures that part that better educated individuals contribute to public funds that is not due to returns and time. Holding the distribution over a (i.e assuming zero returns) and time on the labor market constant, contributions to public funds still differ. First, because the difference between gross income and consumption conditional on a varies with education. Second, because of the differences in consumption during education and the differential resource costs of education.

The decomposition might seem a bit abstract at this point, in particular the policy term ΔT_{policy} since we did not explicitly define policies that implement the desired allocations. In Section 4, we show how ΔT_{return} , ΔT_{time} and ΔT_{policy} can be re-expressed in terms of taxes, subsidies and loan policies. In Section 6, we quantitatively evaluate the different parts of ΔT .

3.5 Special Cases: Model Without Human Capital and IID Shocks

We conclude the analysis of second-best optimal allocations by looking at two special cases of the general model, which convey additional intuition on the workings of the model. To this end, we will place two restrictions on the distribution function $G(a|\theta_i, e_i)$.

Model Without Human Capital. Our framework nests a model without any human capital or education choices. The distribution function is then given by $G(a|\theta_i)$.¹⁰ How would this affect optimal labor and savings wedges? It turns out that the formulas from Propositions 1 and 3 would still be valid, with only $g(a|\theta_i, e_i)$ being replaced by $g(a|\theta_i)$. This implies that the efficient allocation would still depend on the innate type θ_i in this case. The intuition is that having agents reveal their true type θ_i in the first period is still useful for the planner as it brings additional information in the second period, when the planner has to solve the incentive

¹⁰One could call this special case also exogenous unobservable human capital.

problem for the revelation of a . In an optimal taxation language, the planner could still use θ_i as a ‘tag’.¹¹ Note that in our general model with education choices the implementation of this history dependence in the labor wedge becomes particularly appealing. This is because the revelation of θ_i comes with the education choice, and on this education choice the government can condition policies. We show this in Section 4, where we discuss the implementation of incentive compatible allocations.

As said above, the labor wedges are only altered by the fact that all conditional distribution functions only depend on innate ability θ_i . How does the absence of human capital affect optimal labor wedges? The main intuition here comes from the incentive terms $\mathcal{B}_l(a)$. Again, it is particularly transparent for separable preferences (compare Corollary 1), where the term now reads as $\frac{\eta_h}{\lambda_{f_h}} [G(a|\theta_l) - G(a|\theta_h)]$. This effect to relax the incentive constraint of the high type is stronger because the ability distribution of the high type would not change by mimicking the low type as there is no additional change in the education level. In other words, a high type that mimics the low type is “more different” compared to the low type. Therefore, deterring the high type from mimicking by setting high marginal tax rates for low types with high shock realizations becomes more effective in the model without human capital.

IID Shocks. Next we look at the special case where $G(a|\theta_i, e_i)$ is such that ability a is independent of the innate skill level θ_i , i.e. $G(a|\theta_i, e_i) = G(a|e_i)$. This naturally implies that innate ability does not influence labor market outcomes directly and also that the returns to education are homogenous across different θ_i . Essentially, this brings the model to a to an ex-ante homogenous agents framework with risk. What happens to the optimal labor and savings wedges? For the labor wedge it directly follows that the formula would be as in Proposition 1 and Corollaries 1 and 2 with $\mathcal{B}_i(x)$ set to zero. The formula would then be equivalent to the seminal optimal taxation formula (Saez 2001). Intuitively with a representative agent only, labor taxation has a pure insurance role, leading to the static formula with a Utilitarian planner. Additionally, the allocation is no longer history dependent, since the original θ_i is meaningless, and labor wedges can be directly seen as marginal tax rates.

For savings wedges, the terms $\frac{\eta_h}{\lambda_{f_h}} \int_a^{\bar{a}} u'_l(a) \{g(a|\theta_h, e_l) - g(a|\theta_l, e_l)\}$ for the low type in Proposition 3 and Corollary 3 disappear. The intuition here is that the planner does not need to exploit differential incentives to save across different θ_i types anymore, as everybody faces the same distribution. From Corollary 3 it becomes clear that savings wedges are then zero in the absence of income effects: without the Inverse Euler logic as well as no heterogeneity in savings motives, the planner does not have any incentives to distort savings, even in the presence of risk. Finally, education distortions can still be defined as in Section 3.4, with the small adjustment that the labor wedge is not education dependent.

¹¹This is a result also known from the dynamic pricing literature; Courty and Li (2000) present a model where a consumer gets an early signal about her taste valuation in a later period and the monopolist prices goods according to that signal.

4 Implementation

So far we only considered a direct mechanism, in which individuals make reports about their realized type and the planner assigns bundles of consumption, labor supply and education as functions of the reports. The focus in the characterization of the optimal allocation was on wedges or *implicit* price distortions of the allocation. In this section, we explore two decentralized implementations of constrained Pareto optima. We start with a direct approach, in spirit close to the mechanism, where education-dependent labor wedges are mapped into education-dependent taxes. We then move to our main implementation and one of the central results of the article, by showing how this history-dependence of the wedges can also be achieved by income-contingent repayment of loans.

4.1 Implementation One: Education-Dependent Taxes

In the first implementation that we consider, the government offers a set of student grants to the agents in the education period: $\mathcal{G}(e) : \{e_l, e_h\} \rightarrow \mathbb{R}$. In the working period, there is a tax function, which does not only condition on earnings but also on education acquired, i.e. a history-dependent labor income tax schedule $\mathcal{T}(y, e_i) : \mathbb{R}_+ \rightarrow \mathbb{R} \forall i = l, h$. Finally, individuals face a savings tax schedule $\mathcal{T}^s(s) : \mathbb{R} \rightarrow \mathbb{R}$. Given these policy instruments, the individual problem for type $i = h, l$ reads as:

$$\max_{e, c^e, s} \beta^e u^e(c^e) + \int_{\underline{a}}^{\bar{a}} \beta^w v^I(a, e, s) dG(a|\theta_i, e) \text{ s.t. } \mathcal{G}(e) \geq e + c^e + s,$$

where

$$v^I(a, e, s) = \max_{c^w, y} u^w\left(c^w, \frac{y}{a}\right) \text{ s.t. } y - \mathcal{T}(y, e) \geq c^w + \frac{\beta_i^e}{\beta_i^w} s - \mathcal{T}^s(s),$$

where $R_i = \frac{\beta_i^e}{\beta_i^w}$ is set to the return on savings – see Section 2. The following proposition states that for any desired incentive-compatible allocation, there indeed exists a combination of these instruments that implements this desired allocation.

Proposition 4. *Any incentive-compatible allocation can be implemented by a grant schedule $\mathcal{G}(e) : \{e_l, e_h\} \rightarrow \mathbb{R}$, an education-dependent labor income tax schedule $\mathcal{T}(y, e_i) : \mathbb{R}_+ \rightarrow \mathbb{R} \forall i = l, h$ and a savings tax schedule $\mathcal{T}^s(s) : \mathbb{R} \rightarrow \mathbb{R}$.*

Proof. See Appendix A.3.1 □

It is a standard result from the *NDPF*-literature that history-dependent labor income tax functions can implement the desired labor supply allocation. Werning (2011) has shown that history-independent savings taxes can do the job of implementing the desired intertemporal allocation of consumption.

There exist various combinations of the three policy instruments that can implement the desired allocation. In the appendix we are more explicit about that. For the main body, we focus on the most simple implementation; an implementation with zero savings. In this case,

grants are chosen such that they cover education costs and consumption in the education period: $\mathcal{G}(e_i) = e_i + c_i^e$. Labor income taxes are chosen such that $c_i^w(a) = y_i(a) - \mathcal{T}(y_i(a), e_i)$. Finally, the savings tax function $\mathcal{T}^s(s)$ is chosen prohibitively high such that both agents choose $s = 0$. Note again, that this is just one possible implementation and that there also exist implementations with non-zero savings choices, see Appendix A.3.1 for more details.

Education Tax. For the considered policy instruments, it is intuitive to decompose the education tax ΔT . The return effect measures how much more tax revenue the government collects solely – i.e. holding the tax function and time on the labor market constant – because of returns:

$$\Delta T_{return} = \beta_l^w \int_{\underline{a}}^{\bar{a}} \mathcal{T}(y_l(a), e_l) (g(a|\theta_h, e_h) - g(a|\theta_h, e_l)) da.$$

The impact due to a different time – i.e. holding the tax function and the distribution constant – on the labor market is given by:

$$\Delta T_{time} = -\mathcal{G}(e_l)(\beta_h^e - \beta_l^e) + (\beta_h^w - \beta_l^w) \int_{\underline{a}}^{\bar{a}} (\mathcal{T}(y_l(a), e_l)) dG(a|\theta_h, e_h).$$

The impact on public funds solely – i.e. holding the distribution of a and time on the labor market constant – through policies then given by:

$$\Delta T_{policy} = \beta_h^e (\mathcal{G}(e_l) - \mathcal{G}(e_h)) - \beta_h^w \int_{\underline{a}}^{\bar{a}} (\mathcal{T}(y_l(a), e_l) - \mathcal{T}(y_h(a), e_h)) dG(a|\theta_h, e_h).$$

First, individuals receive a different amount of grants \mathcal{G} because of going to college. Second, taxes are education dependent and therefore, for a given realization of a , tax payments differ.

4.2 Implementation Two: Income-Contingent Loans

The previous implementation required that people with the same income but different levels of education pay different taxes. In reality people might perceive this as a violation of horizontal equity concerns, which could hinder the political feasibility of such policies. In this light we now present a potentially more appealing alternative implementation with only one labor income tax schedule and a repayment scheme of the education grant.¹²

Technically, this can be seen as a simple reinterpretation of the previous implementation – we take the tax system of the θ_l -type as the *common* labor income tax schedule and introduce an income-contingent repayment schedule, which conditions on the size of the loan.¹³ Repayment can potentially both exceed the loan value or be below it. The latter can be considered as a

¹²Diamond and Saez (2011) argue that practical policy prescription from optimal tax models should not go against commonly held normative views (horizontal equity for example) and limit complexity to a reasonable degree. The second implementation seems in line with these recommendations.

¹³Related implementations are of course possible. For example one, where the tax function of the θ_h -type is the labor income tax schedule in place. The extreme case would just be that income taxes do not exist and all schedules that were interpreted as history-dependent labor income schedules in implementation 1 can now

partial default. Together both instruments are sufficient to replicate the optimal labor wedges. Formally we summarize this in the following corollary:

Corollary 4. *As opposed to the policy instruments in Proposition 4, an incentive-compatible allocation can also be implemented by a (compulsory) loan schedule $\mathcal{L}(e)$, a loan repayment schedule $\mathcal{R}(y, \mathcal{L})$, an income tax $\mathcal{T}(y)$ and a savings tax $\mathcal{T}^s(s)$ where*

- $\mathcal{L}(e_i) = \mathcal{G}(e_i)$
- $\mathcal{T}(y) = \mathcal{T}(y, e_l)$
- $\mathcal{R}(y, \mathcal{L}(e_l)) = 0$
- $\mathcal{R}(y, \mathcal{L}(e_h)) = \mathcal{T}(y, e_h) - \mathcal{T}(y)$
- $\mathcal{T}^s(s)$ is defined as in Appendix A.3.1.

Proof. The budget constraints of the individuals are unchanged as compared to Proposition 4. Therefore the corollary is a direct consequence Proposition 4. \square

Starting from Proposition 4, we have first reinterpreted grants as loans. Second, we have set the tax schedule for the low type as the tax schedule for everybody. Given that the high type then faces the ‘wrong tax schedule’, we set the repayment schedule for the high type as the difference between the ‘right tax schedule’ ($\mathcal{T}(y, e_h)$) and the ‘wrong tax schedule’ ($\mathcal{T}(y) = \mathcal{T}(y, e_l)$). Finally, we have set loan repayment for the low type to zero given that she faces the ‘right tax schedule’.

As a pure theoretical result, this corollary does not necessarily make a case for income-contingent loans as they are understood in the real-world policy debate. Depending on the allocation, repayment might be decreasing in income $\mathcal{R}_{\mathcal{L}}(\mathcal{L}(e_h), \cdot) < 0$. Whether repayment is increasing or decreasing depends on whether labor wedges are higher for θ_l -types or θ_h -types for given levels of income y . In general, repayment might be also negative $\mathcal{R}(\mathcal{L}(e_h), \cdot) < 0$, or – to the contrary – be much higher than what a normal market repayment of the loan would be.

All this can potentially contrast the typical idea of income-contingent repayment, which is (i) non-negative, (ii) weakly increasing in income and typically (iii) bounded from above. Whether income-contingent repayment according to Corollary 4 fulfills these criteria depends on the properties of the considered allocation. In Section 6, we quantify our model with US-data and arrive at the interesting result that repayment schedules according to Corollary 4 mostly fulfill (i), (ii) and (iii). In particular, we find very strong support for (ii): individuals with a college degree should typically face higher marginal distortions on their labor supply decision.

be interpreted as repayment schedules. Taking the labor income tax schedule of the low type, however, is more natural in our view. Especially in our application of the theory in Section 6.

Education Tax. In the case of an income-contingent loan implementation, the return effect reads as:

$$\Delta T_{return} = \beta_l^w \int_{\underline{a}}^{\bar{a}} \mathcal{T}(y_l(a)) (g(a|\theta_h, e_h) - g(a|\theta_h, e_l)) da.$$

The time effect reads as

$$\Delta T_{time} = -\mathcal{L}(e_l)(\beta_h^e - \beta_l^e) + (\beta_h^w - \beta_l^w) \int_{\underline{a}}^{\bar{a}} (\mathcal{T}(y_l(a), e_l)) dG(a|\theta_h, e_h).$$

The policy effect is now given by

$$\Delta T_{policy} = \beta_h^e (\mathcal{L}(e_l) - \mathcal{L}(e_h)) + \beta_h^w \int_{\underline{a}}^{\bar{a}} (\mathcal{R}(y_h(a), \mathcal{L}(e_h)) + T(y_h(a)) - T(y_l(a))) dG(a|\theta_h, e_h).$$

First, individuals receive a different loan during the education period. Second, repayment in the working period differs (with repayment of the low type set to zero).

Other Implementations. Besides education-dependent taxes and income-contingent loans, other policy instruments that implement education-dependent labor wedges could do the job. Another implementation could, e.g., be downstream or delayed tuitions fees, where graduates pay tuition fees after their studies and dependent on their income. Further, human capital risk insurance contracts where graduates pay some premium before labor market uncertainty realizes and then receive (potentially negative) payments once uncertainty has materialized. We choose to focus on income-contingent loans because such policy instruments are already used in the real world. Our intention is not to claim that they are superior over other potential implementations.

5 Comparison Case: Flat Reypament

The previous section has revealed that relatively sophisticated policy instruments are needed to implement constrained efficient allocations as characterized in Section 3. A natural question is how much better these constrained efficient allocations perform in terms of welfare than allocations that can be implemented through simpler policies.

It is not obvious what the right ‘simple’ comparison is. Often the literature has analyzed how much welfare is lost from imposing history-independence and linearity on wedges (Farhi and Werning 2013, Golosov, Troshkin and Tsyvinski 2015, Stantcheva 2015). For the purpose of quantifying the welfare gains from income-contingent student loans, imposing linearity on wedges would be a too strong restriction. Such an approach would not only restrict the loan repayment but also the labor income tax and education subsidies and would therefore overestimate the welfare gains from income-contingency of the loan repayment.

We therefore want to allow the labor income tax and the education subsidies to be nonlinear. For the respective allocation, this implies that we restrict labor wedges to be history independent

but allow for nonlinearity. Formally, the only additional restriction is that individuals with the same income should face the same labor wedge. To characterize such allocations, we apply a first-order approach for history-independent policies that we elaborate in greater detail in Findeisen and Sachs (2015).

5.1 The Planning Problem

The planning problem is almost equivalent to that of the second-best problem in Section 3.1. However, a restriction on preferences has to be made: the absence of income effects on labor supply, i.e. $u^w(c^w, \frac{y}{a}) = u(c^w - \Psi(\frac{y}{a}))$ where $\Psi'(\cdot), \Psi''(\cdot), u'(\cdot) > 0$ and $u''(\cdot) < 0$. In this case, the restriction that individuals with the same income face the same labor wedge is equivalent to

$$y_h(a) = y_l(a) \tag{5}$$

because only the shock a matters for the labor supply decision – the level of wealth (which typically varies between the high and the low type) does not play a role.

Proposition 5. *If preferences satisfy $u^w(c^w, \frac{y}{a}) = u(c^w - \Psi(\frac{y}{a}))$ where $\Psi'(\cdot), \Psi''(\cdot), u'(\cdot) > 0$ and $u''(\cdot) < 0$, then any incentive-compatible allocation that satisfies (5), can be implemented with an education-independent tax function $\mathcal{T}(y)$, a loan schedule $\mathcal{L}(e)$ with fixed repayment rates $\mathcal{R}(\mathcal{L})$ and a savings tax $\mathcal{T}^s(s)$.*

Proof. See Appendix A.4.1. □

As in the case with history-dependent labor wedges, there are some degrees of freedom. In line with above, we focus on an implementation with zero private savings. As in Proposition 4, we normalize the labor income tax schedule such that the low type does not have to repay anything of the loan.

5.2 Labor Wedges

What determines the level of optimal labor wedges if they are constrained to depend on income only? Given that we had to restrict preferences such that there are no income effects on labor supply, the history-dependent benchmark are the formulas in Corollary 2. As is formally stated in the following proposition, optimal history-independent labor wedges are governed by the same forces; the forces are just averaged across education levels. In Sections 6.3, we quantitatively compare optimal history-independent with history-dependent labor wedges.

Proposition 6. *Assume that preferences satisfy $u^w(c^w, \frac{y}{a}) = u(c^w - \Psi(\frac{y}{a}))$ where $\Psi'(\cdot), \Psi''(\cdot), u'(\cdot) > 0$ and $u''(\cdot) < 0$. At any constrained Pareto optimum where $\eta_h > 0$, optimal history-independent labor wedges at income level $y(a)$ satisfy:*

$$\frac{\tau^y(a)}{1 - \tau^y(a)} = \frac{1 + \varepsilon^c(a)}{\varepsilon^c(a)a(g(a|\theta_h, e_h) + g(a|\theta_l, e_l))} \int_a^{\bar{a}} \{\mathcal{A}_h(x) + \mathcal{A}_l(x) + \mathcal{B}_l(x)\} dx$$

where terms \mathcal{A}_l , \mathcal{A}_h and \mathcal{B}_l are defined as in Proposition 1 and $\varepsilon^c(a)$ is the compensated elasticity for individuals with wage a .

Proof. See Appendix A.4.2. □

5.3 Education Subsidies and Savings Wedges

As we show in Appendix A.4.2, savings wedges are still as in Corollary 3. In the absence of income effects, there is no role for boosting labor supply also if labor wedges are history independent. Therefore the high type's decision is undistorted whereas the low type faces a marginal savings subsidy (borrowing tax). The latter serves as an effective tool to relax the incentive constraint of the high type, see the discussion in Section 3.3.

For education policies in this case, the education tax ΔT can now be decomposed into:

$$\Delta T_{return} = \beta_l^w \int_{\underline{a}}^{\bar{a}} \mathcal{T}(y(a)) (g(a|\theta_h, e_h) - g(a|\theta_h, e_l)) da$$

and

$$\Delta T_{time} = -\mathcal{L}(e_l)(\beta_h^e - \beta_l^e) + (\beta_h^w - \beta_l^w) \int_{\underline{a}}^{\bar{a}} (\mathcal{T}(y_l(a), e_l)) dG(a|\theta_h, e_h)$$

and

$$\Delta T_{policy} = \beta_h^e (\mathcal{L}(e_l) - \mathcal{L}(e_h)) + \beta_h^w \mathcal{R}(\mathcal{L}(e_h)).$$

In particular the policy effect differs from the previous cases: education policies do *not* condition on income since we impose (5). Quantitatively as one will see in the next section, the optimal value of ΔT_{policy} will be very different for the case with income-contingent loans (or education-dependent labor taxes) compared to the case with standard loans and standard labor taxes.

6 A Quantitative Exploration

Building on our theoretical findings above, the purpose of this section is to ask whether optimal dynamic tax theory can make a case for income-contingent student loans. In the previous sections we have laid out the theoretical foundations to answer this questions. We have elaborated the properties of second-best efficient allocations in Section 3. Although these allocations are solely constrained by informational asymmetries, we have shown that the second-best can indeed be implemented with policy instruments that come close to real world policies: non-linear income taxes combined with a student loan system with income-contingent repayment. Whether this is indeed a case for repayment that increases in income is not clear, however. If optimal labor wedges for highly educated individuals are lower than for less educated individuals, this would imply repayment that is *decreasing* in income – a property that is not in line with real world income-contingent loan practices where repayment is *increasing* in income.

The first purpose of this section is therefore to ask whether a quantitative exploration of the model indeed makes a case for loan repayment that increases in income. A robust result

is that it indeed makes a case for repayment that increases in income – at least for low and intermediate income levels. For high incomes it might be slightly decreasing.

The second purpose of this section is then to quantify the welfare gains from income-contingency of loan repayment. We therefore compare the second-best with history-independent policies as discussed in Section 5. We find that income-contingency of repayment yields significant welfare gains.

It might be considered as an undesirable property that the second-best is partly characterized by locally decreasing repayment. Further, the second best implies that some income levels repay more than the loan value which might also go against real world policy practices. We therefore ask how much welfare is lost if repayment schedules are restricted such that repayment (i) is non-decreasing in income and (ii) never exceeds the loan value. We find that such "real-world repayment schedules" can yield a large part of the welfare gains from the second best.

6.1 Parametrization

In our quantitative analysis, individuals start life as high school graduates and then live for 47 years (age 18-65). If individuals do not go to college, they spent 47 years on the labor market. If individuals go to college, they spend 4 years in college and 43 on the labor market. Formally, we set $\beta_l^e = 0$ and $\beta_l^w = \sum_{t=1}^{47} \beta^{t-1}$ for those who don't go to college and $\beta_h^e = \sum_{t=1}^4 \beta^{t-1}$ and $\beta_h^w = \sum_{t=5}^{47} \beta^{t-1}$ for those who do go to college, where we set $\beta = \frac{1}{1.04}$. We maximize a Utilitarian social welfare function so we set \tilde{f}_i to the population mass f_i .

To get the ability distributions $g(a|e_i, \theta_j)$ with $i, j = l, h$, we take the factual and the estimated counterfactual earnings distributions for high-school graduates from Cunha and Heckman (2007), based on white males.¹⁴ The factual lifetime income distributions for high-school and college graduates come directly from the National Longitudinal Survey of Youth 1979 (NLSY79), covering cohorts born between 1957-64. For the counterfactual distributions, i.e. the distribution of outcomes if a high-school type goes to college and vice-versa, we rely on the structural estimates from their paper. They use the Armed Services Vocational Aptitude Battery (ASVAB) test taken by participants of the NLSY79 to have a measure of skills around high-school graduation, corresponding to θ in our model. Given this information, Cunha and Heckman (2007) estimate a structural model of selecting into education levels.

As in any survey data, top incomes are underrepresented. We therefore append Pareto tails at earnings of \$88,000. In the benchmark scenario we chose a Pareto parameter of 1.5 for all distributions (Atkinson et al. 2011; Diamond and Saez, 2011). To the best of our knowledge, there exists no systematic evidence on the conditional distributions of top incomes for college graduates and non-graduates separately. In Section 6.5, we show how results change under

¹⁴Precisely, we use the estimates from Figures 1 and 2 from their paper. We used the software GetData Graph Digitizer to read out the data from the graphs. Since Cunha and Heckman (2007) consider the present value of lifetime earnings (18-65), we take a 47 years annuity with the same present value. Since the data in these graphs are not smoothed, we apply a standard Kernel smoother.




Figure 1: Skill Distributions

the assumption that the tail is fatter for college graduates. Finally, we smooth the resulting distribution again to overcome the kink from the appended tail. Given a (linear) approximation of the real world tax system we calibrate the implied skill distributions as input for the model from the optimality conditions of the agents. The resulting calibrated skill distributions are illustrated in Figure 1. We choose a tax rate of 25%, matching the patterns about average marginal tax rates documented in Guner, Kaygusuz, and Ventura (2014). The results are robust to using a more accurate approximation of marginal tax rates. Lastly, we assume there is an atom of workers equal to five percent for each distribution reflecting unemployment or disability as in Mankiw et al. (2009).

The share of high school and college types are set to 64.19% and 35.81%, respectively, as reported in Cunha and Heckman (2008). Following Gallipoli et al. (2011), we set the annual monetary cost of college education to \$11,100. The yearly interest rate is set to 4% and the yearly discount factor β to 1/1.04. We work with a CRRA specification and focus on the case with no income effects so that:

$$U(c, y, a) = \frac{\left(c - \frac{(y/a)^\sigma}{\sigma}\right)^{1-\rho}}{1-\rho},$$

with $\sigma = 3$, implying a constant labor supply elasticity of 0.5 and set $\rho = 2$. In unreported simulations, we also varied the values for σ and ρ ; the main results do not change.

Performance of the model. Our calibrated model yields a college wage premium measured as the percentage difference of average college to high school earnings of 57% for identical Pareto tails and 61% if the tails are education-dependent.¹⁵ This is consistent with US data where the

¹⁵The calibration implies about \$57,000 in average earnings for college graduates and about 37,000\$ for high-school graduates.

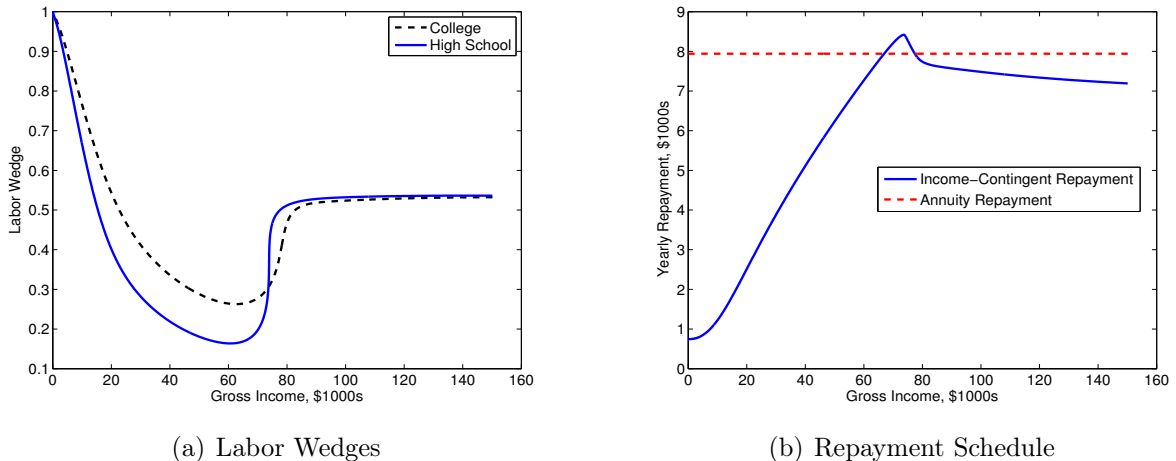


Figure 2: Utilitarian Optimum

college earnings premium rose from 40% in the 1980's to about 80% in the early 2000's (Lee, Lee, and Shin (2015)).

The model also matches measures of within-group, here defined as within-education group, inequality. For high-school graduates, the variance of log earnings conditional on working is 0.16 in the calibration and around 0.18 in US data from the early 2000's (Lemieux 2006a). The distribution of college-graduates features more dispersion in the model with a variance of 0.28, the data counterpart to this number is 0.29 (Lemieux (2006a)). The overall income distribution of both education groups has variance of the log of 0.23, also close to the number of 0.21 reported by Lemieux (2006a).

6.2 Second-Best Optimal Policies

Optimal Labor Wedges: Figure 2(a) displays the optimal labor wedges as a function of yearly income up to \$160,000. Both schedules follow a U-shaped pattern, reflecting a result from the static Mirrlees problem (Diamond 1998, Saez 2001). The intuition for the pattern is simple: for very low incomes, marginal distortions are high for two reasons: first, distorting their labor supply is relatively harmless since they are rather unproductive. Second, the inverse hazard rate $\frac{1-G(a|\cdot,\cdot)}{g(a|\cdot,\cdot)}$ is rather high. Note that $1 - G(a|\cdot,\cdot)$ is proportional to the additional revenue generated by the (implicit or explicit) marginal tax rate and $g(a|\cdot,\cdot)$ is the mass of individuals whose labor supply is distorted. For intermediate incomes the density $g(a|\cdot,\cdot)$ strongly increases making distortions more and more harmful, leading to a decrease in optimal distortions. Finally, due to the properties of the Pareto distribution, the ratio $\frac{1-G(a|\cdot,\cdot)}{ag(a|\cdot,\cdot)}$ converges to a constant and as a consequence the labor wedges start to converge. The impact of the $\mathcal{B}(\theta, a)$ -term of Corollary 2 on the marginal tax rates of the high-school type is quantitatively very small, which is why we do not illustrate the impact.

Looking at Figure 1, one can see in which way tax distortions are tailored to the different income distributions. At every point of the skill support before the Pareto tail kicks in, college

labor distortions generate much bigger mechanical revenue effects for the government. In the top income tails, the wedges converge to almost the same top tax rate (Saez, 2001), with a very small difference caused by the education incentive force $\mathcal{B}(\theta, a)$, which we discussed in the theoretical section of this paper, that leads to slightly higher top tax rates for high school types to increase the attractiveness of going to college.¹⁶

Repayment Schedule. We now build on the implementation result from Section 4.2 and illustrate optimal income-contingent repayment schedules. The (common) labor income tax schedule is determined by the high-school labor wedges. The overall loan (in net present value) that students take on is \$143,745, which covers annual consumption during college of \$26,977 and coverage of the annual fees of \$11,100. Figure 2(b) shows the yearly repayment of college debt as a function of income. The slope of the repayment schedule is given by the difference in the labor wedges as we outlined in Section 4.2. As the optimal labor wedge for college graduates lies above the high-school wedge, repayment is increasing in income up to incomes of US-\$80,000. Repayments for college graduates start at about US-\$1,000. Remarkably, in this income region, the repayment schedule of loans is almost linear with a slope of roughly 0.1, because the difference in the labor wedges is almost constant apart from very low incomes. Afterwards, there is a very small decreasing range and the repayment schedule flattens out as the top labor wedges converge. In sum, optimal repayments can be very well approximated by an intercept of US-\$1,000, a US-\$1,000 increase in repayment for every US-\$10,000 earned up to earnings of US-\$70,000 and no additional repayments for incomes above that threshold. So although we did not place any restrictions on the shape of the repayment schedule, linearity comes very close to the second-best optimum.

The red dotted horizontal line shows the yearly repayment that would occur if individuals chose a standard loan (with a yearly interest rate of 4%) where the repayment is not contingent on income and they repay the same amount every year. As can be seen, only some individuals pay back more than in the income-contingent case, but for most income levels partial default is optimal. The expected interest on repayment is therefore only 2.2%. The maximal ex-post rate of return is 4.4% and the minimal ex-post rate of return -5.1%.¹⁷

As discussed in the implementation section, we assume the college loan system to be mandatory. We check if this is a restrictive assumption by allowing college graduates to opt out and instead take a loan with a yearly interest rate of 4% to finance tuition and early consumption. We find that given the choice, individuals would opt into the loan system with income contingent repayment rates. This is also true for an interest rate of 3%. However, this is arguably a strict test of the assumption since it is not clear whether individuals would be able to borrow

¹⁶Some of these results are related to the simulations of Luttmer and Zeckhauser (2008) who consider a static setting where going to college is purely a signal and not an investment; hence counterfactual and factual distributions are equal.

¹⁷This set of results is sensitive to the interest rate, however. For 3%, e.g., more individuals would pay back more than the loan value in the income-contingent case. For 5%, nobody would pay back more.

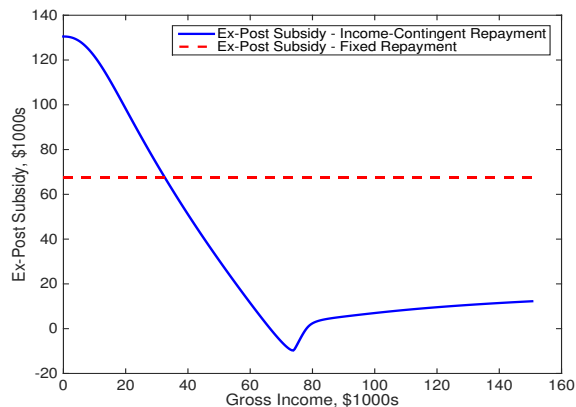


Figure 3: Ex-Post College Subsidies

up to their desired amount and might face a substantial risk premium on their interest rate if they borrow in the private market.

Subsidization of College. We now turn to the question by how much college education is subsidized. First, on net, college is *taxed* by \$33,727, i.e. $\Delta T = \$33,727$. Thus, \$33,727 of the monetary gains from college education are reaped by the government. If policies were not conditional on education and individuals would be on the labor market for the same time, the government would reap $\Delta T_{return} = \$105,068$. However, individuals that go to college are on the labor market for less years which costs the government \$18,784, i.e. $\Delta T_{time} = -\$18,784$.

Finally, policies do condition on education. What is the effect of this on public funds? The absolute subsidy is \$52,557, i.e. $\Delta T_{policy} = -\$52,557$. Thus, in expectation individuals are subsidized by this amount for going to college. Given the income contingency of loan repayment, the subsidy varies with the amount of income. Thus, it is an ex-post subsidy. The blue bold line in Figure 3 illustrates how this ex-post subsidy varies with income and nicely illustrates the power of income-contingent repayment: if individuals end up poor after college, receive a very high education subsidy. This subsidy then decreases in income and the pattern how the subsidy evolves with income mirrors the repayment schedule of income-contingent loans.

Finally, savings distortions are zero in our application. This is a direct consequence of Corrolary 3. As we assume an education period of length zero for the high school type, there is no transition from an education to a working period, where the planner would find it optimal to distort savings for the high school type. For the college type, we get a no distortion at the top result, i.e. $\eta_h = 0$ for the high type.

6.3 The Welfare Gains From Income-Contingent Repayment

We now aim at quantifying what the potential welfare effects of income-contingency are and therefore compare the second best to history-independent policies as described in Section 5. Figure 4 shows the optimal education independent labor income tax in this case. In line with




Figure 4: Optimal Education-Independent Taxes

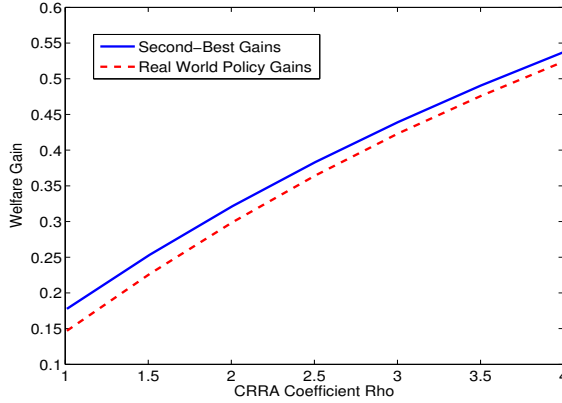


Figure 5: Welfare Gains

the discussion around Proposition 5, optimal education-independent marginal tax rates lie between their education-dependent counterparts from the second-best optimum.

In case of standard loans with fixed repayment, the education subsidy does not depend on realized income. The flat subsidy here is \$67,557, see also Figure 3 to compare it to the income-contingent ex-post subsidy in the second best. The subsidy is higher in the absence of income-contingent loans. Intuitively, if the government cannot tailor the subsidy to different income realizations – and therefore to different levels of marginal utility of consumption – it is more expensive to incentivize education.

We next calculate the welfare gains from income-contingent repayment schemes. In Figure 5 the blue bold line presents the consumption equivalent welfare gains as the CRRA parameter ρ varies from 1 to 4.¹⁸ The welfare gains are increasing in risk-aversion which underscores the

¹⁸Following common practices in the literature, we asked by how much percent consumption of each individual in every period would have to be increased (for the case with a flat repayment) such that welfare is as high as with income-contingent repayment.




Figure 6: Real World Adjustment

role of the loans as an insurance device. For a CRRA coefficient of two, the gains are about 0.32%.¹⁹

6.4 Real World Policies: Cap on Repayment and Non-Decreasing Repayment

There might be two limitations to the full second-best optimum which could reduce its real world appeal. First, for some (small) range of the income distribution, repayment for college graduates actually exceeds the loan value, as becomes obvious from Figure 2(b). Second, for high earners the repayment schedule actually decreases in income. These properties are likely to go against commonly held normative views, when it comes to the actual implementation of an income-contingent loan system. Indeed, actual income-contingent repayment systems in the United States, United Kingdom or Australia are never decreasing and cap repayment at the loans values. To deal with these concerns, we calculate an allocation which can be implemented with a repayment schedule respecting these constraints – i.e. it is never decreasing and capped at the loan value. In this scenario, effective marginal tax rates for college graduates are adjusted so that they are equal to the marginal tax rates for high school graduates as soon as repayment reaches the capitalized loan value. These modified policies still respect incentive compatibility and budget feasibility, of course.²⁰ Figures 6(a) and 6(b) show the resulting labor wedges and the repayment schedule.

¹⁹The welfare gains are evenly distributed in the benchmark case ($\rho = 2$), implying that both the college and the high school type achieve a utility gain of 0.32% of consumptions equivalents. For lower values of ρ , a larger share of the gain is reaped by the high-school graduates, for higher values of ρ the result is reversed.

²⁰More technically, we first adjusted the lump sum element of the common labor income tax schedule such that the government budget constraint holds. In case, the resulting allocation is not incentive compatible, we adjusted the lump sum elements of the labor income tax and the repayment schedule such that the government budget constraint holds and the incentive constraint of the college type binds.

By construction, this repayment system is, of course, inferior in welfare terms to the optimal repayment schedule. Figure 5 illustrates that welfare gains are not much smaller as compared to the full second best. For $\rho = 2$, roughly 78% of the welfare gain from optimal income-contingent repayment be reaped with the restricted repayment.²¹

It is also interesting to note that this repayment schedule comes very close to current U.S. policies. The average marginal repayment rate in the increasing region is 10.5% and the maximal marginal rate is 15.8%. In the U.S., three income-contingent repayment options were recently introduced: “Income-Based Repayment”, “Pay as You Earn” and “Revised Pay as You Earn”. These options differ in some details, however, they have in common that repayment is capped at between 10% and 15% of income and remaining debt is (partly) forgiven after 20-25 years (Brooks 2015). Further these programs take into account family size. When making these comparisons, one of course has to take into account that our optimal repayment schedule applies for an optimal income tax schedule and not for the current tax schedule. Further, due to the simplicity of our model, we naturally neglect issues of how repayment should evolve over the life cycle. We leave a more detailed elaboration of current policies from an optimal policy perspective for future research.

6.5 Robustness: Differing Top Income Tails

We now test if and how a different assumption on top incomes across income distributions changes the results. We focus on the case, where the college income distribution has a thicker tail than the high school income distribution. For college graduates, we choose a Pareto parameter of 1.28. For high-school graduates we choose a Pareto parameter of 3.²² These values lie within the range of what has been typically found in empirical studies covering many countries and time periods (Atkinson et al, 2011). If we aggregate the two distributions to the aggregate income distribution, we find that the resulting tail for top incomes resembles a Pareto tail with a parameter not far away from 1.5.²³

Second-Best Optimal Policies. Figures 7(a) and 7(b) display the corresponding schedules for labor wedges and the repayment schedule. The college labor wedge now lies above the high school labor wedge everywhere, leading to a strictly increasing repayment schedule. The implicit top tax rate for college graduates is higher than for high-school graduates, driven by the differences in the Pareto parameter. Interestingly, again a simple linear approximation of

²¹In case of an interest rate of 3%, 68% of the welfare gain can be reaped with simpler policies. For an interest rate of 5%, second-best optimal income-contingent repayment would actually never exceed the loan value.

²²The top tails are not dependent on innate type θ but are just determined by the education level. In an earlier working paper version (Findeisen and Sachs, 2013), we also explore the case in which the tails are determined by innate type θ instead. The results are very similar.

²³The sum of two Pareto distributions tends to behave like a Pareto distribution, where the heavier tail distribution seems to dominate (Ramsay, 2006). This implies that, in the tails, the resulting aggregate distribution is very close to the college distribution.

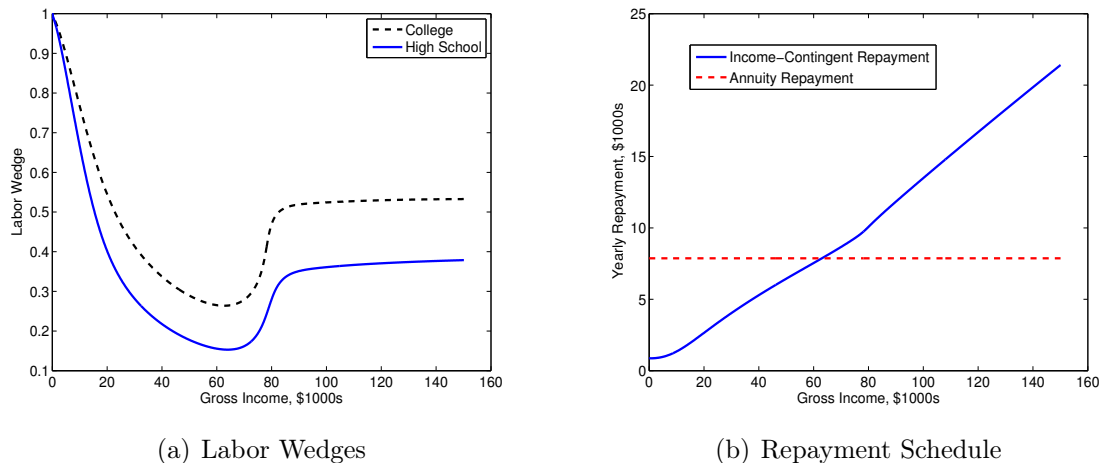


Figure 7: Utilitarian Optimum With Thick College Tails

the repayment schedule with a linear slope of about 11% could almost perfectly implement the second-best optimum. Repayment of college graduates now exceeds the annuity loan value by a much more significant amount and for much bigger fraction of the population.²⁴ We check again if a college type would prefer not to choose the income-contingent loan in this case and find that the loans indeed have to be compulsory. However, as we show next, one can again construct slightly different policies which respect a cap on repayment. These yield a large share of the welfare gain and do not require the loans to be compulsory.

Real World Policies: Cap on Repayment. As in Section 6.4, we now adjust the second-best optimum towards policies that satisfy the same two mentioned real-world restrictions. The adjustment we make is slightly different this time. In Section 6.4, we lowered the labor wedges of the college types such that they equal the optimal ones for the high school types above all income levels, where the second-best repayment starts to exceed the loan value. Here, we do the opposite and increase the labor wedges of the high school types such that they are equal to the college labor wedges. The reason for this is that optimal history-independent wedges (see Figure 9) are closer to the college wedges for high incomes, which is driven by the fatter college top income tail “dominating” the top income tail for the high school types, see footnote 23. The new adjusted policies respect incentive compatibility and budget feasibility. In order to avoid bunching because of a discrete upward jump in marginal tax rates, we smooth out the increase over an interval of roughly US-\$5,000. The resulting labor wedges and repayment are illustrated in Figures 8(a) and 8(b).

The Welfare Gains From Income-Contingent Repayment As in Section 6.3, we now calculate the welfare gains over student loans without income-contingent repayment. Due to

²⁴The results on ex-post education subsidies then of course mirror this result. College graduates with incomes above \$62,000 actually do not receive an ex-post subsidy, but pay an ex-post tax.

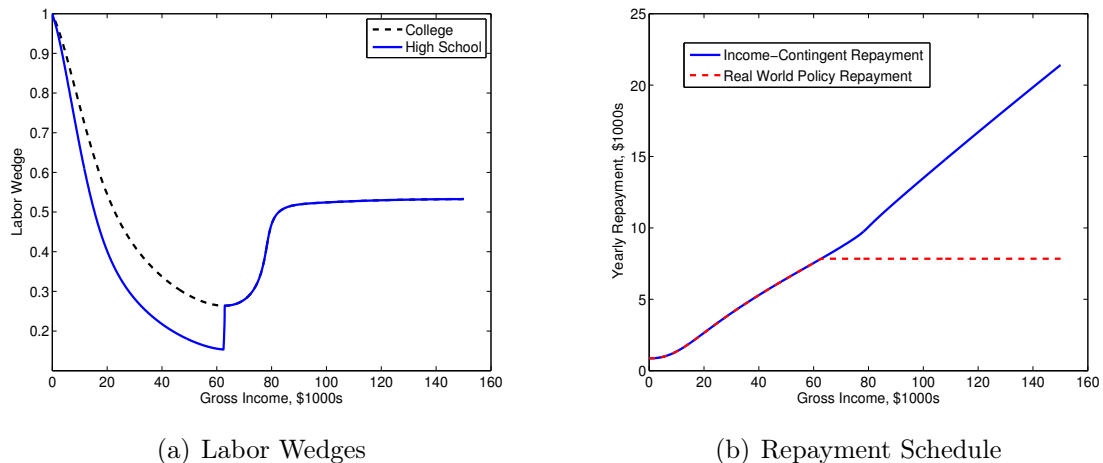


Figure 8: Real World Optimum With Thick College Tails

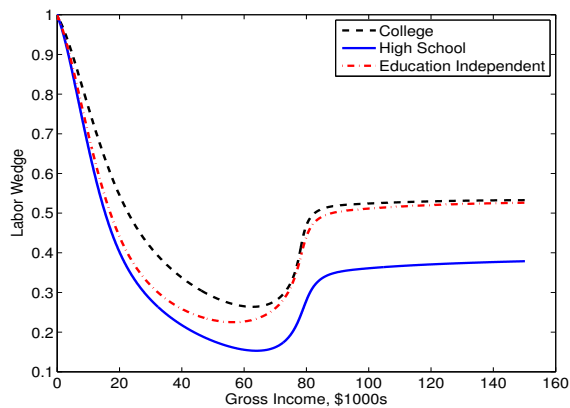


Figure 9: Optimal Education-Independent Taxes

the differing top income tails, the college and high school wedges are more distinct from each other (see Figure 9) than in the benchmark case. This yields to welfare gains (see Figure 10) that are slightly higher. They are 0.36% of lifetime consumption for a CRRA coefficient of 2. Again, the adjusted system respecting a cap can yield a large part of those gains: in fact, they lead to a gain of 0.33%, which is almost 92% of the welfare gain. For an interest rate of 3% (5%) the latter value is 75% (95 %).

7 Conclusion

This paper has studied the implications of endogenous education decisions before labor market entry on Pareto optimal tax policies in a dynamic environment with heterogeneous agents and uncertainty. An attractive way to decentralize Pareto optimal allocations is to have the government support students to finance consumption and tuition during education. During their working life students pay back these loans, contingent on income and loan size. We therefore

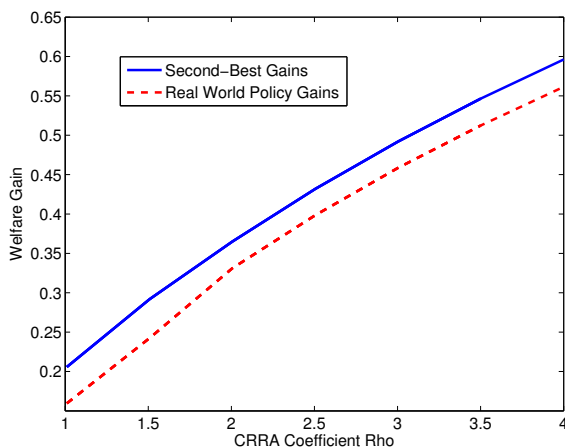


Figure 10: Welfare Gains and Risk-Aversion

make a second-best argument in favor of student loans with income-contingent repayment rates and, in addition, provide guidance for the optimal design of such repayment schedules.

We have abstracted from several aspects that can be tackled in future work. First, we have abstracted from initial wealth heterogeneity. In an environment where individuals differ concerning the income and wealth of their parents, typically the question arises to what extent optimal education policies should depend on parents' income and wealth. Second, due to our assumption that all labor market risk is realized directly after labor market entry, some aspects concerning the optimal timing of repayment were naturally disregarded. Relatedly, we did not consider human capital accumulation after labor market entry like on-the-job training. Third, we assumed full commitment to all policies from the government side. Relaxing these assumptions might be a fruitful area for future research.²⁵

A Appendix

A.1 Incentive Compatibility

We look at the case where $\frac{\beta_l^w}{\beta_h^w} \leq 1$, so people with lower education levels enter the labor market earlier.

Lemma 1. *Suppose there is first-order stochastic dominance in innate abilities $G(a|\theta_h, \cdot) \leq G(a|\theta_l, \cdot)$ and increasing differences $|G(a|\theta_h, e') - G(a|\theta_h, e'')| \geq |G(a|\theta_l, e') - G(a|\theta_l, e'')|$ for $e' > e''$, conditions (1), (2) hold, and (4) holds with equality, and we have:*

(i) $y_h(a) - y_l(a) \geq 0$ and $c_h(a) - c_l(a) \geq 0 \forall a$

(ii) $u_{cl} \leq 0$,

then the considered allocation is incentive compatible.

²⁵In Findeisen and Sachs (2016), we study education and tax policies without commitment, however, in a simpler environment without uncertainty. Our main finding is that education subsidies are more (less) progressive because of the lack of commitment if tax instruments are linear (only constrained by informational asymmetries).

Note that when $u_{cl} \geq 0$, than (i) $y_h(a) - y_l(a) \leq 0$ and $c_h(a) - c_l(a) \leq 0 \forall a$ is sufficient.

This lemma implies that instead of directly ex-post verifying whether period two incentive compatibility is satisfied in an allocation, one can alternatively check these two simple monotonicity conditions; if they are fulfilled, then the allocation is incentive compatible. The lemma relies on two plausible empirical conditions, namely that a higher innate skill level leads to a better distribution of outcomes in a first-order stochastic dominance sense and that a higher innate skill level implies higher returns to education. Both conditions are fulfilled in our calibrated economy in Section 6.

Adding (4), which holds with equality by assumption, to (3) gives:

$$\begin{aligned} & \beta_l^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_l, a) dG(a|\theta_l, e_l) - \beta_h^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_h, a) dG(a|\theta_l, e_h) + \beta_h^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_h, a) dG(a|\theta_h, e_h) \\ & - \beta_l^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_l, a) dG(a|\theta_h, e_l) \geq 0. \end{aligned}$$

where we have used the fact the allocation is incentive compatible at the working stage so that we can use the value functions $v(\theta_i, a)$ directly. Adding and subtracting $\beta_l^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_h, a) dG(a|\theta_h, e_l)$ and $\beta_l^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_h, a) dG(a|\theta_l, e_l)$ gives:

$$\begin{aligned} & \beta_l^w \int_{\underline{a}}^{\bar{a}} [v^w(\theta_h, a) - v^w(\theta_l, a)] (g(a|\theta_h, e_l) - g(a|\theta_l, e_l)) da \\ & + \beta_h^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_h, a) (g(a|\theta_h, e_h) - \frac{\beta_l^w}{\beta_h^w} g(a|\theta_h, e_l) - (g(a|\theta_l, e_h) - \frac{\beta_l^w}{\beta_h^w} g(a|\theta_l, e_l))) da \geq 0. \end{aligned}$$

Integration by parts yields:

$$\begin{aligned} & - \beta_l^w \int_{\underline{a}}^{\bar{a}} \left[\frac{\partial(v^w(\theta_h, a) - v^w(\theta_l, a))}{\partial a} \right] (G(a|\theta_h, e_l) - G(a|\theta_l, e_l)) da \\ & - \beta_h^w \int_{\underline{a}}^{\bar{a}} \left(\frac{\partial v^w(\theta_h, a)}{\partial a} \left(G(a|\theta_h, e_h) - \frac{\beta_l^w}{\beta_h^w} G(a|\theta_h, e_l) \right) - \left(G(a|\theta_l, e_h) - \frac{\beta_l^w}{\beta_h^w} G(a|\theta_l, e_l) \right) \right) da > 0. \end{aligned}$$

Consider the first line. The second term in the first integral is negative at every point under FOSD in education. Remember that $\frac{\partial v^w(\theta_i, a)}{\partial a} = -u_l^w \left(c_w(\theta_i, a), \frac{y(\theta_i, a)}{a} \right) \frac{y(\theta_i, a)}{a^2}$. So the first term in the first integral is positive everywhere whenever:

$$y_h(a) - y_l(a) \geq 0, c_h(a) - c_l(a) \geq 0, u_{cl} \leq 0,$$

as stated in Lemma 1. As a special case, under separable preferences (i.e. $u_{cl} = 0$), $y_h(a) - y_l(a) \geq 0$ is enough.

The second line of the condition is positive whenever the returns to education are increasing in the innate type. In other words whenever the high type profits more from a higher education level, captured by:

$$(G(a|\theta_h, e(\theta_h)) - G(a|\theta_h, e(\theta_l))) - (G(a|\theta_l, e(\theta_h)) - G(a|\theta_l, e(\theta_l))) < 0. \quad (6)$$

To see this we have to show that

$$\left(G(a|\theta_h, e_h) - \frac{\beta_l^w}{\beta_h^w} G(a|\theta_h, e_l) \right) - \left(G(a|\theta_l, e_h) - \frac{\beta_l^w}{\beta_h^w} G(a|\theta_l, e_l) \right) < 0. \quad (7)$$

Equation 6 is a special case for which $\frac{\beta_l^w}{\beta_h^w} = 1$, so when education takes the same amount of time for education levels. In general one expects $\frac{\beta_l^w}{\beta_h^w} \leq 1$, because lower education levels come with earlier labor market entry. Equation 7 is decreasing when $\frac{\beta_l^w}{\beta_h^w}$ gets smaller because of first-order condition dominance in skills. So for all $\frac{\beta_l^w}{\beta_h^w}$ between 0 and 1, equation 7 is negative, which completes the proof.

A.2 Optimal Wedges

After integrating by parts and using the transversality conditions $\mu(\theta, \underline{a}) = \mu(\theta, \bar{a}) = 0 \quad \forall \theta$, the Lagrangian for the social planner's problem reads as

$$\begin{aligned} \mathcal{L} = & \sum_{i=l,h} \beta_i^e u^e(c_i^e) \tilde{f}_i + \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_i, a) dG(a|\theta_i, e_i) \tilde{f}_i \\ & + \lambda \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} y_i(a) dG(a|\theta_i, e_i) f_i \\ & - \lambda \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} \gamma(v^w(\theta_i, a), y_i(a)/a) dG(a|\theta_i, e_i) f_i - \lambda \sum_{i=l,h} \beta_i^e (c_i^e + e_i) f_i \\ & - \sum_{i=l,h} \int_{\underline{a}}^{\bar{a}} \left[\mu'_i(a) v^w(\theta_i, a) + \mu_i(a) u_l \left\{ \gamma \left(v^w(\theta_i, a), \frac{y_i(a)}{a} \right), \frac{y_i(a)}{a} \right\} \cdot \frac{y_i(a)}{a^2} \right] da \\ & + \sum_{i=l,h; j=l,h; j \neq i} \eta_i \left[\beta_i^e u^e(c_i^e) + \beta_i^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_i, a) dG(a|\theta_i, e_i) da \right. \\ & \left. - \beta_j^e u^e(c_j^e) - \beta_j^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_j, a) dG(a|\theta_j, e_j) da \right] \end{aligned}$$

where $\gamma(u^w, l)$ is the inverse of $u^w(\cdot, l)$. The first-order conditions are:

$$\frac{\partial \mathcal{L}}{\partial c_h^e} = u_c^e(c_h^e) (\tilde{f}_h + \eta_h) - \lambda f_h = 0 \quad (8)$$

$$\frac{\partial \mathcal{L}}{\partial c_l^e} = u_c^e(c_l^e)(\tilde{f}_l - \eta_h) - \lambda f_l = 0 \quad (9)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial v^w(\theta_h, a)} &= \left(\tilde{f}_h + \eta_h \right) g(a|\theta_h, e_h) - \frac{\lambda g(a|\theta_h, e_h) f_h}{u_c^w \left(c_h^w(a), \frac{y_h(a)}{a} \right)} \\ &\quad - \mu_h(a) \frac{u_{c,l}^w y_h(a)}{u_c^w a^2 \beta_h^w} - \frac{\mu'_h(a)}{\beta_h^w} = 0 \end{aligned} \quad (10)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial v^w(\theta_l, a)} &= \tilde{f}_l g(a|\theta_l, e_l) - \eta_h g(a|\theta_h, e_l) - \frac{\lambda g(a|\theta_l, e_l) f_l}{u_c^w \left(c_l^w(a), \frac{y_l(a)}{a} \right)} \\ &\quad - \mu_l(a) \frac{u_{c,l}^w y_l(a)}{u_c^w a^2 \beta_l^w} - \frac{\mu'_l(a)}{\beta_l^w} = 0 \end{aligned} \quad (11)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial y_i(a)} &= \lambda g(a|\theta_i, e_i) f_i - \frac{\mu_i(a)}{\beta_i^w} \left[\frac{u_{cl}^w u_l^w y_i(a)}{u_c^w a^3} + u_{ll}^w \frac{y_i(a)}{a^3} + u_l^w \frac{1}{a^2} \right] \\ &\quad + \lambda g(a|\theta_i, e_i) f_i \frac{u_l^w}{u_c^w} \frac{1}{a} = 0. \end{aligned} \quad (12)$$

A.2.1 Proof of Proposition 1

Rewriting (12):

$$\begin{aligned} &\lambda g(a|\theta_i, e_i) f_i \left[1 + \frac{u_l^w}{a u_c^w} \right] \\ &- \frac{1}{\beta_i^w} \mu_i(a) \left[\frac{u_{cl}^w u_l^w y_i(a)}{u_c^w a^3} + u_{ll}^w \frac{y_i(a)}{a^3} + u_l^w \frac{1}{a^2} \right] = 0. \end{aligned}$$

Dividing by $\frac{-u_l^w}{a u_c^w}$ and $\lambda g(a|e, \theta_i) f_i$ and using the definition of the labor wedge, i.e. $u_c^w(1 - \tau^y) = -u_l^w \frac{1}{a}$ yields

$$\frac{\tau_i^y(a)}{1 - \tau_i^y(a)} = \frac{1}{\beta_i^w} \frac{u_c^w \mu_i(a)}{\lambda g(a|\theta_i, e_i) f_i a} \frac{\left[u_{cl}^w \frac{u_l^w y_i(a)}{a} + u_{ll}^w \frac{y_i(a)}{a} + u_l^w \right]}{u_l^w},$$

which can be written as

$$\frac{\tau_i^y(a)}{1 - \tau_i^y(a)} = \frac{1}{\beta_i^w} \cdot \frac{u_c^w \mu_i(a)}{\lambda g(a|\theta_i, e_i) f_i a} \frac{1 + \varepsilon_u(\theta_i, a)}{\varepsilon_c(\theta_i, a)},$$

where $\frac{[u_{cl}^w \frac{u_l^w}{u_c^w} \frac{y_i(a)}{a} + u_{ll}^w \frac{y_i(a)}{a} + u_l^w]}{u_l} = \frac{1 + \varepsilon_u(\theta_i, a)}{\varepsilon_c(\theta_i, a)}$ can be shown by some algebra, see Saez (2001, p.227). In particular, with the isoelastic specification used in the computations $\frac{(y/a)^\sigma}{\sigma}$ one can verify that this term is equal to $\frac{1}{\sigma}$.

Inserting (8) into (10) and solving for the differential equation yields:

$$\frac{\mu_h(a)}{\beta^w} = \int_a^{\bar{a}} \exp\left(-\int_a^x \frac{u_{c,l}^w y_h(a)}{u_c^w a^2} ds\right) \frac{1}{u_c^w(x)} \left(1 - \frac{u_c^w(x)}{u_c^e(c_h^e)}\right) \lambda f_h g(x|\theta_h, e_h) dx$$

yielding:

$$\frac{\tau_h^y(a)}{1 - \tau_h^y(a)} = \frac{1 + \varepsilon_h^u(a)}{\varepsilon_h^c(a)} \frac{1}{ag(a|\theta_h, e_h)} \int_a^{\bar{a}} \exp\left(-\int_a^x \frac{u_{c,l}^w y_h(a)}{u_c^w a^2} ds\right) \frac{u_c^w(a)}{u_c^w(x)} \left(1 - \frac{u_c^w(x)}{u_c^e(c_h^e)}\right) dG(x|\theta_h, e_h).$$

For the low type, we get by similar steps:

$$\begin{aligned} \frac{\mu_l(a)}{\beta^w} = & \int_a^{\bar{a}} \exp\left(-\int_a^x \frac{u_{c,l}^w y_h(a)}{u_c^w a^2} ds\right) \frac{1}{u_c^w(x)} \left(1 - \frac{u_c^w(x)}{u_c^e(c_h^e)}\right) \\ & + \frac{\eta_h}{\lambda} u_c^w(x) \{g(a|\theta_h, e_l) - g(a|\theta_l, e_l)\} \lambda f_h g(x|\theta_h, e_h) dx \end{aligned}$$

Using the same arguments as in Saez (2001, p. 227), one can show that this formula can be written as the one in Proposition 1.

A.2.2 Proof of Corollary 1

If preferences are separable of the form $u(c) - \Psi(l)$, where Ψ are the convex utility costs of labor and we further assume that $u(\cdot) = u^e(\cdot)$, the Lagrangian for the social planner's problem reads as

$$\begin{aligned}
\mathcal{L} = & \sum_{i=l,h} \beta_i^e u'(c_i^e) \tilde{f}_i + \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_i, a) dG(a|\theta_i, e_i) \tilde{f}_i \\
& + \lambda \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} y_i(a) dG(a|\theta_i, e_i) f_i \\
& - \lambda \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} u^{-1} [v^w(\theta_i, a) + \Psi(y_i(a)/a)] dG(a|\theta_i, e_i) f_i \\
& - \lambda \sum_{i=l,h} \beta_i^e (c_i^e + e_i) f_i \\
& - \sum_{i=l,h} \int_{\underline{a}}^{\bar{a}} \left(\mu_i'(a) v^w(\theta_i, a) + \mu_i(a) \Psi' \left(\frac{y_i(a)}{a} \right) \frac{y_i(a)}{a^2} \right) da \\
& + \sum_{i=l,h;j=l,h;j \neq i} \eta_j \left[\beta_i^e u^e(c_i^e) + \beta_i^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_i, a) dG(a|\theta_i, e_i) da \right. \\
& \left. - \beta_j^e u^e(c_j^e) - \beta_j^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_j, a) dG(a|\theta_j, e_j) da \right].
\end{aligned}$$

With first-order conditions:

$$\frac{\partial \mathcal{L}}{\partial c_h^e} = u'(c_h^e)(\tilde{f}_h + \eta_h) - \lambda f_h = 0 \quad (13)$$

$$\frac{\partial \mathcal{L}}{\partial c_l^e} = u'(c_l^e)(\tilde{f}_l - \eta_h) - \lambda f_l = 0 \quad (14)$$

$$\frac{\partial \mathcal{L}}{\partial v^w(\theta_h, a)} = (\tilde{f}_h + \eta_h) g(a|\theta_h, e_h) - \lambda \frac{1}{u'(c_h^w(a))} g(a|\theta_h, e_h) f_h - \frac{\mu_h'(a)}{\beta_h^w} = 0 \quad (15)$$

$$\frac{\partial \mathcal{L}}{\partial v^w(\theta_l, a)} = \tilde{f}_l g(a|\theta_l, e_l) - \eta_h g(a|\theta_h, e_l) - \lambda \frac{1}{u'(c_l^w(a))} g(a|\theta_l, e_l) f_l - \frac{\mu_l'(a)}{\beta_l^w} = 0 \quad (16)$$

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial y_i(a)} = & \lambda g(a|\theta_i, e_i) f_i - \frac{\mu_i(a)}{\beta_i^w} \left[\Psi'' \left(\frac{y_i(a)}{a} \right) \frac{y_i(a)}{a^3} + \frac{1}{a^2} \Psi' \left(\frac{y_i(a)}{a} \right) \right] \\
& - \lambda g(a|\theta_i, e_i) f_i \frac{\Psi' \left(\frac{y_i(a)}{a} \right)}{a u'(c_i^w(a))} = 0,
\end{aligned} \quad (17)$$

Rewriting (17):

$$\lambda g(a|\theta_i, e_i) f(\theta) \left[1 - \frac{\Psi' \left(\frac{y_i(a)}{a} \right)}{a u'(c_i^w(a))} \right]$$

$$-\frac{1}{\beta_i^w} \mu_i(a) \left[\Psi'' \left(\frac{y_i(a)}{a} \right) \frac{y_i(a)}{a^3} + \frac{1}{a^2} \Psi' \left(\frac{y_i(a)}{a} \right) \right] = 0.$$

Dividing by $\frac{\Psi'}{au'}$ and $\lambda g(a|e, \theta_i) f_i$ and using the definition of the labor wedge, i.e. $u'(1-\tau^y) = \Psi' \frac{1}{a}$ yields

$$\frac{\tau_i^y(a)}{1 - \tau_i^y(a)} = \frac{1}{\beta_i^w} \frac{\mu_i(a)}{\lambda g(a|\theta_i, e_i) f_i a} \left[\frac{\Psi'' \frac{y}{a^2} + \Psi' \frac{1}{a}}{\frac{\Psi'}{au'}} \right],$$

which can be written as

$$\frac{\tau_i^y(a)}{1 - \tau_i^y(a)} = \frac{1}{\beta_i^w} \cdot \frac{u' \mu_i(a)}{\lambda g(a|\theta_i, e_i) f_i a} \frac{1 + \varepsilon_u(\theta_i, a)}{\varepsilon_c(\theta_i, a)},$$

where $\frac{\Psi'' \frac{y}{a^2} + \Psi' \frac{1}{a}}{\frac{\Psi'}{au'}} = \frac{1 + \varepsilon_u^u(a)}{\varepsilon_c^c(a)}$ can be shown by simple algebra, see Saez (2001, p.227).

The multiplier $\mu_h(a)$ can be obtained using (15) and (13):

$$\frac{\mu_h(a)}{\beta^w} = \frac{\lambda f_h}{u'(c_h^e)} g(a|\theta_h, e_h) - \lambda f_h \int_{\underline{a}}^a \frac{1}{u'(c_h^w(a^*))} dG(a^*|\theta_h, e_h).$$

The multiplier $\mu_l(a)$ can be obtained using (16) and (14):

$$\begin{aligned} \frac{\mu_l(a)}{\beta^w} &= \frac{\lambda f_l}{u'(c_l^e)} g(a|\theta_l, e_l) - \lambda f_l \int_{\underline{a}}^a \frac{1}{u'(c_l^w(a^*))} dG(a^*|\theta_l, e_l) \\ &\quad + \eta_h (g(a|\theta_l, e_l) - g(a|\theta_h, e_h)) \end{aligned}$$

yielding:

$$\frac{\tau_l^y(a)}{1 - \tau_l^y(a)} = \frac{1 + \varepsilon_l^u(a)}{\varepsilon_l^c(a)} \frac{u'(c_l^w(a))}{ag(a|\theta_l, e_l)} [\mathcal{A}_l(a) + \mathcal{B}_l(a)]$$

where

$$\mathcal{A}_l(a) = \frac{G(a|\theta_l, e_l)}{u'(c_l^e)} - \int_{\underline{a}}^a \frac{1}{u'(c_l^w(a^*))} dG(a^*|\theta_l, e_l)$$

$$\mathcal{B}_l(a) = \frac{\eta_h}{\lambda f_l} [g(a|\theta_l, e_l) - G(a|\theta_h, e_h)].$$

Using the inverse Euler equation, the term $\mathcal{A}_i(a)$ can be written as in the proposition:

$$\begin{aligned}
& \frac{G(a|\theta_i, e_i)}{u'(c_i^e)} - \int_{\underline{a}}^a \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) \\
&= G(a|\theta_i, e_i) \int_{\underline{a}}^{\bar{a}} \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) - \int_{\underline{a}}^a \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) \\
&= G(a|\theta_i, e_i) \int_{\underline{a}}^a \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) + G(a|\theta_i, e_i) \int_a^{\bar{a}} \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) \\
&\quad - \int_{\underline{a}}^a \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) \\
&= G(a|\theta_i, e_i) \int_a^{\bar{a}} \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) - (1 - G(a|\theta_i, e_i)) \int_{\underline{a}}^a \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i).
\end{aligned}$$

Relation to the formula of Saez (2001)

The insurance part of the labor wedge can be expressed as in Saez (2001), for our case with separable preferences. This relation applies if agents do not differ ex-ante. Using the inverse Euler equation, we obtain

$$\begin{aligned}
\mathcal{A}_i(a) &= \int_{\underline{a}}^{\bar{a}} \frac{G(a|\theta_i, e_i)}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) - \int_{\underline{a}}^a \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) \\
&= \int_{\underline{a}}^{\bar{a}} \frac{G(a|\theta_i, e_i)}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) - \int_{\underline{a}}^{\bar{a}} \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) \\
&\quad + \int_a^{\bar{a}} \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) \\
&= \int_a^{\bar{a}} \frac{1}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i) - \int_{\underline{a}}^{\bar{a}} \frac{1 - G(a|\theta_i, e_i)}{u'(c_i^w(a^*))} dG(a^*|\theta_i, e_i)
\end{aligned}$$

where the second equality follows from the transversality condition. This term can be expressed as in Saez (2001) as shown by Mankiw, Weinzierl and Yagan (2009) in their online appendix.

A.2.3 Proof of Proposition 2

Solving (8) for $\tilde{f}_h + \eta$, inserting into (10) and integrating over all a yields the result. For the low type the proof works almost equivalent.

A.2.4 Proof of Proposition 3

Define $\hat{\mu}_i(a) = \mu_i(a)u_c^w(c_i^w, \frac{y_i(a)}{a})$. This implies

$$\hat{\mu}'_i(a) = \mu'_i(a)u_c^w\left(c_i^w, \frac{y_i(a)}{a}\right) + \mu_i(a)\left(u_{cc}^w c'_h(a) + u_{cl}^w\left(\frac{y'_i(a)}{a} - \frac{y_i(a)}{a^2}\right)\right). \quad (18)$$

Inserting $\mu'_h(a)u_c^w(c_h^w, \frac{y_h(a)}{a})$ as implicitly above and inserting into (10) yields:

$$\begin{aligned} & u_c^w \left(c_h^w, \frac{y_h(a)}{a} \right) \left(\tilde{f}_h + \eta_h \right) g(a|\theta_h, e_h) - \lambda g(a|\theta_h, e_h) f_h \\ & - \frac{\hat{\mu}'_h(a)}{\beta^w} + \mu_h(a) \left(u_{cc}^w c'_h(a) + u_{cl}^w \frac{y'_h(a)}{a} \right) = 0. \end{aligned}$$

Inserting $\tilde{f}_h + \eta_h$ as implicitly defined by (8) and integrating yields:²⁶

$$\frac{\tau_h^s}{1 - \tau_h^s} = \int_{\underline{a}}^{\bar{a}} \frac{\mu_h(a)}{\lambda f_h} \left(u_{cl}^w \frac{y'_i(a)}{a} - u_{cc}^w c'_h(a) \right).$$

For the low type, first one has to add and subtract $\eta_h g(a|e_l, \theta_l)$ in (11). Then inserting $\mu'_l(a)u_c^w(c_l^w, \frac{y_l(a)}{a})$ as implicitly defined in (18) and inserting $\tilde{f}_l - \eta_h$ as implicitly defined in (9) into (11), and finally integrating and rearranging yields:

$$\frac{\tau_l^s}{1 - \tau_l^s} = \int_{\underline{a}}^{\bar{a}} \frac{\mu_h(a)}{\lambda f_h} \left(u_{cl}^w \frac{y'_i(a)}{a} - u_{cc}^w c'_h(a) \right) + \frac{\eta_h}{\lambda f_h} \int_{\underline{a}}^{\bar{a}} u_c^w \left(c_l^w(a), \frac{y_l(a)}{a} \right) \{g(a|\theta_h, e_l) - g(a|e_l, \theta_l)\} da.$$

A.3 Implementation

A.3.1 Proof of Proposition 4

Starting from a direct mechanism we show that optimal allocations can indeed be implemented with the policy instruments as defined in Proposition 4.

Step 1: Introduce savings

This step of the proof closely follows Werning (2011). Note that the implementation takes care of double deviations – where agents misreport and save too much – by making sure savings deviations are appropriately punished for any reporting strategy. Imagine the desired incentive-compatible allocation is implemented by a direct mechanism. From that point on, assume that individuals could freely save s at rate R . Let r_θ denote the report about θ_i for $i = l, h$. Given a savings tax schedule $T^s(s, r_\theta) : \mathbb{R} \rightarrow \mathbb{R}$, the budget constraints read as

$$\tilde{c}^e(r_\theta) + s = c^e(r_\theta)$$

$$\tilde{c}^w(r_\theta, r_a) = c^w(r_\theta, r_a) + Rs - T^s(s, r_\theta).$$

²⁶For that step it is useful to recall that $\frac{\tau_h^s}{1 - \tau_h^s} = \frac{\int_{\underline{a}}^{\bar{a}} u_c^w dG(a|\theta_h, e_h)}{u_c^e} - 1$.

Define the optimal report r_a about a , for a given report r_θ about θ_i , a given savings tax schedule $T^s(s, r_1)$ and a given level of savings s :

$$r_a^*(a, r_\theta, s, T^s) = \arg \max_{r_a} u \left[c^w(r_\theta, r_a) + Rs - T^s(s, r_\theta) - \psi \left(\frac{y(r_\theta, r_a)}{a} \right) \right].$$

Then the optimal report in period one, for a given level of savings s and a given savings tax schedule $T^s(s, r_\theta)$, is defined by

$$r_\theta^*(\theta_i, s, T^s(r_\theta, s)) = \arg \max_{r_\theta} \beta_i^e u^e(c^e(r_\theta) - s) + \beta_i^w \int_{\underline{a}}^{\bar{a}} u \left[c^w(r_\theta, r_a^*) + R_i s - T^s(s, r_\theta) - \psi \left(\frac{y(r_\theta, r_a^*)}{a} \right) \right] dG(a|\theta_i, e(r_\theta)).$$

Then define a hypothetical tax schedule $T^*(r_\theta, s, \theta)$ for each θ implicitly, where $V(\theta)$ is the (lifetime) value function of a truth teller of type θ_i

$$V(\theta_i) = V(\theta_i, s, r_\theta^*, T^*(r_\theta, s, \theta)) \forall s.$$

This hypothetical tax schedule would make individuals of type θ_i indifferent between truth telling and the optimal joint deviation for any s . It is hypothetical since it does not only depend on the report r_θ , which is observable but also on the unobservable type θ_i . However, we know that for both levels θ_i such a tax schedule exists. Therefore taking the upper envelope over these functions yields a savings tax function $\hat{T}(s, r_\theta)$ that also implements zero savings and is feasible since it does not condition on θ_i :

$$\hat{T}(s, r_\theta) = \sup_{\theta_i} T^*(s, r_\theta, \theta_i).$$

Lemma 2. *An incentive-compatible allocation can be implemented by a direct mechanism extended by a savings choice and history-dependent savings tax schedules $\hat{T}(s, r_\theta)$.*

In a last step, we make the savings tax independent of the report r_θ . Therefore, we simply take the upper envelope of $\hat{T}(s, r_\theta)$:

$$T^s(s) = \sup_{r_\theta} \hat{T}(s, r_\theta).$$

Lemma 3. *An incentive-compatible allocation can be implemented by a direct mechanism extended by a savings choice and a savings tax schedules $T^s(s)$.*

Step 2: Introduction of an education-dependent tax schedule.

Starting from an incentive-compatible allocation we have established that a savings tax function exists which implements zero savings in Step 1. We next decentralize the labor-leisure decision.

By the same arguments as in the standard Mirrlees model, it follows that this extended direct mechanism can also implement the desired incentive-compatible allocation.

In period two, individuals have already chosen the education level e intended for them by the planner, as the direct mechanism was the starting point. After drawing a , instead of directly revealing their type, individuals of type θ_i face an income tax schedule $T(\cdot, e_i) : \mathbb{R}_+ \rightarrow \mathbb{R}$ that is defined by

$$T(y_i(a), e_i) = y_i(a) - c_i^w(a) \quad \forall a. \quad (19)$$

It is a standard application of the taxation principle that this tax schedule satisfies

$$(y_i(a), c_i^w(a)) \in \arg \max_{y, c^w} u^w \left(c^w, \frac{y}{a} \right) \quad s.t. \quad c \leq y - T(y, e),$$

and therefore implements the desired allocation.

Step 3: Complete Decentralization – allow for educational investment

Finally in the last step we decentralize education decisions. The goal of the planner is to implement the right e choices, using a grant function $\mathcal{G}(e_i) : \{e_l, e_h\} \rightarrow \mathbb{R}$ and condition taxes such that holding income fixed: $T(y, \cdot) : \{e_l, e_h\} \rightarrow \mathbb{R}$. One way to implement the desired allocation is to set $\mathcal{G}(e_i) = e_i + c_i^e$ for $i = h, l$ and $T(y, e_i)$ for $i = h, l$ as defined in (19).

Indeterminacy of the Implementation and the Role of Zero-Savings. The just presented implementation is indeterminate for two reasons: first of all, other policy instruments (e.g. income-contingent student loans, see Section 4.2) can implement the desired incentive-compatible allocation. Second, even considering education grants, savings taxes and education-dependent income taxes, there are various combinations of these instruments that can implement the desired allocations. For example, it is easy to show that – for any s – $\mathcal{G}^*(e_i) = \mathcal{G}(e_i) + s \quad \forall i = l, h$, $T^*(\cdot, e_i) = T(\cdot, e_i) + s$ and $T^{s*}(x) = T^s(x - s)$ also implement the desired allocation, however, do not imply zero savings but equilibrium savings of s .

A.4 History-Independent Policies

A.4.1 Implementation

The proof is very similar to that in Appendix A.3.1. The difference in step 1 is that y is constrained only to be a function of the second period report (or report), hence $y(r_a)$ ($y(a)$) instead of $y(r_\theta, r_a)$ ($y_i(a)$).

The difference in step 2 is that the tax schedule is just defined differently and history-independent: $\mathcal{T}(y(a)) = y(a) - c_l(a)$. Nevertheless, the standard insights of the taxation principle apply.

For step 3, the goal of the planner is to implement the right e choices, using a loan schedule function $\mathcal{L}(e) : \mathbb{R}_+ \rightarrow \mathbb{R}$ with repayment rates $\mathcal{R}(e) : \mathbb{R}_+ \rightarrow \mathbb{R}$. The simplest way to implement

the desired allocation is to set $\mathcal{L}(e_i) = e_i + c_i^e$ for $i = h, l$, $\mathcal{R}(e_l) = 0$ and $\mathcal{R}(e_h) = y(a) - \mathcal{T}(y(a)) - c_h^w(a)$.

A.4.2 Wedges

$$\begin{aligned}
\mathcal{L} &= \sum_{i=l,h} \beta_i^e u(c_i^e) \tilde{f}_i + \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_i, a) dG(a|\theta_i, e_i) \tilde{f}_i \\
&+ \lambda \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} y(a) dG(a|\theta_i, e_i) f_i \\
&- \lambda \sum_{i=l,h} \beta_i^w \int_{\underline{a}}^{\bar{a}} (u^{-1}(v^w(\theta_i, a)) + \Psi(y(a)/a)) dG(a|\theta_i, e_i) f_i - \lambda \sum_{i=l,h} \beta_i^e (c_i^e + e_i) f_i \\
&- \sum_{i=l,h} \int_{\underline{a}}^{\bar{a}} \left[\mu'_i(a) v^w(\theta_i, a) + \mu_i(a) u' (u^{-1}(v_i^w(a)) \Psi' \left(\frac{y(a)}{a} \right) \frac{y(a)}{a^2} \right] da \\
&+ \sum_{i=l,h; j=l,h; j \neq i} \eta_i \left[\beta_i^e u(c_i^e) + \beta_i^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_i, a) dG(a|\theta_i, e_i) da \right. \\
&\left. - \beta_j^e u(c_j^e) - \beta_j^w \int_{\underline{a}}^{\bar{a}} v^w(\theta_j, a) dG(a|e_j, \theta_i) da \right].
\end{aligned}$$

The first-order conditions read as:

$$\frac{\partial \mathcal{L}}{\partial c_h^e} = u'(c_h^e)(\tilde{f}_h + \eta_h) - \lambda f_h = 0 \quad (20)$$

$$\frac{\partial \mathcal{L}}{\partial c_l^e} = u'(c_l^e)(\tilde{f}_l - \eta_h) - \lambda f_l = 0 \quad (21)$$

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial v^w(\theta_h, a)} &= (\tilde{f}_h + \eta_h) g(a|\theta_h, e_h) - \frac{\lambda g(a|\theta_h, e_h) f_h}{u'(c_h^w(a) - \Psi(\frac{y_h(a)}{a}))} \\
&- \mu_h(a) u'' \Psi' \frac{y(a)}{a^2 \beta_h^w} - \frac{\mu'_h(a)}{\beta_h^w} = 0
\end{aligned} \quad (22)$$

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial v^w(\theta_l, a)} &= \tilde{f}_l g(a|\theta_l, e_l) - \eta_h g(a|\theta_h, e_l) - \frac{\lambda g(a|\theta_l, e_l) f_l}{u'(c_l^w(a) - \Psi(\frac{y_l(a)}{a}))} \\
&- \mu_l(a) u'' \Psi' \frac{y(a)}{a^2 \beta_l^w} - \frac{\mu'_l(a)}{\beta_l^w} = 0
\end{aligned} \quad (23)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial y(a)} = \sum_{i=h,l} \left(\lambda g(a|\theta_i, e_i) f_i - \frac{\mu_i(a)}{\beta_i^w} u' \left(c_i^w(a) - \Psi \left(\frac{y_i(a)}{a} \right) \right) \left[\Psi' \frac{1}{a^2} + \Psi'' \frac{y(a)}{a^3} \right] \right. \\ \left. + \lambda g(a|\theta_i, e_i) f_i \Psi' \frac{1}{a} \right) = 0, \end{aligned} \quad (24)$$

Some algebra reveals that (24) implies:

$$\frac{\tau^y(a)}{1 - \tau^y(a)} = \frac{1 + \varepsilon^c(a) \sum_{i=h,l} \mu_i(a) u' \left(c_i^w(a) - \Psi \left(\frac{y_i(a)}{a} \right) \right)}{\varepsilon^c(a) \lambda a \sum_{i=l,h} f_i g(a|\theta_i, e_i)}. \quad (25)$$

Define $\hat{\mu}_i(a) = \mu_i(a) u'(\theta_i, a)$. This yields:

$$\hat{\mu}_i'(a) = \mu_i'(a) u' \left(c_i^w(a) - \Psi \left(\frac{y_i(a)}{a} \right) \right) + \mu_i(a) u'' \left(c_i^w(a) - \Psi \left(\frac{y_i(a)}{a} \right) \right) \left(c_i'(a) - \Psi' \left(\frac{y'(a)}{a} - \frac{y(a)}{a^2} \right) \right).$$

Inserting this and (20) and (21) into (22) and (23) yields:

$$\mu_h(a) u' \left(c_h^w(a) - \Psi \left(\frac{y_h(a)}{a} \right) \right) = \int_a^{\bar{a}} \left(1 - \frac{u_c^w(x)}{u_c^e(c_h^e)} \right) dG(x|\theta_h, e_h)$$

and

$$\begin{aligned} \mu_l(a) u' \left(c_l^w(a) - \Psi \left(\frac{y_l(a)}{a} \right) \right) = \int_a^{\bar{a}} \left(g(x|\theta_l, e_l) \left(1 - \frac{u_c^w(x)}{u_c^e(c_l^e)} \right) \right. \\ \left. + u' \left(c_l^w(x) - \Psi \left(\frac{y_l(x)}{x} \right) \right) \eta_h \{ g(x|\theta_h, e_l) - g(x|\theta_l, e_l) \} \right) dx. \end{aligned}$$

Inserting this into (25) yields the formula in Proposition 6.

Savings Wedge: The steps are basically equivalent as in Appendix A.2.4 since the respective first-order conditions are unchanged. In case of no income effects, we then have $u_{cl}^w = -u''\Psi'$ and $u_{cc}^w = u''$, which gives $\left(\frac{u_{cl}^w y_i'(a)}{u_c^w} - \frac{u_{cc}^w}{u_c^w} c_h'(a) \right) = \frac{u''}{u'} \left(\Psi' \frac{y'(a)}{a} - c'(a) \right) = 0$, where the last equality follows from working-period incentive compatibility, just as in the proof of Corollary 3.

References

- ACEMOGLU, D., AND D. AUTOR (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics Vol. 4*, ed. by O. Ashenfelter, and D. Card, pp. 1043–1171. Elsevier.
- ANDERBERG, D. (2009): “Optimal Policy and the Risk Properties of Human Capital Recon-

- sidered,” *Journal of Public Economics*, 93(9-10), 1017–1026.
- ANDERBERG, D., AND F. ANDERSSON (2003): “Investments in Human Capital, Wage Uncertainty, and Public Policy,” *Journal of Public Economics*, 87(7-8), 1521 – 1537.
- ATKINSON, A. B., T. PIKETTY, AND E. SAEZ (2011): “Top Incomes in the Long Run of History,” *Journal of Economic Literature*, 49(1), 3–71.
- BARR, N. (2004): “Higher education funding,” *Oxford review of economic policy*, 20(2), 264–283.
- BEST, M., AND H. KLEVEN (2013): “Optimal Income Taxation with Career Effects of Work Effort,” *Working Paper, London School of Economics*.
- BOHACEK, R., AND M. KAPICKA (2008): “Optimal Human Capital Policies,” *Journal of Monetary Economics*, 55(1), 1–16.
- BOVENBERG, L., AND B. JACOBS (2005): “Redistribution and Education Subsidies are Siamese Twins,” *Journal of Public Economics*, 89(11-12), 2005–2035.
- BROOKS, J. R. (2015): “Income-Driven Repayment and the Public Financing of Higher Education,” *Georgetown Law Journal*, 104.
- CARNEIRO, P., AND J. HECKMAN (2005): “Human Capital Policies,” in *Inequality in America: What Role for Human Capital Policies*, pp. 77–240. The MIT Press.
- CHAPMAN, B. (2006): “Income contingent loans for higher education: International reforms,” *Handbook of the Economics of Education*, 2, 1435–1503.
- CHEN, S. (2008): “Estimating the Variance of Wages in the Presence of Selection and Unobserved Heterogeneity,” *Review of Economics and Statistics*, 90(2), 275–289.
- CHONÉ, P., AND G. LAROQUE (2011): “Optimal Taxation in the Extensive Model,” *Journal of Economic Theory*, 146(2), 425–453.
- COURTY, P., AND H. LI (2000): “Sequential Screening,” *The Review of Economic Studies*, 67(4), 697–717.
- 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.
- (2008): “The Evolution of Labor Earnings Risk in the US Economy,” *Working Paper, University of Chicago*.
- 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. (1980): “Income Taxation with Fixed Hours of Work,” *Journal of Public Economics*, 13(1), 101–110.
- DIAMOND, P. A. (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. A., AND J. A. MIRRLEES (1978): “A model of Social Insurance with Variable Retirement,” *Journal of Public Economics*, 10(3), 295–336.
- DIAMOND, P. A., AND E. SAEZ (2011): “The Case for a Progressive Tax: From Basic Research to Policy Recommendations,” *Journal of Economic Perspectives*, 25(4), 165–90.
- EATON, J., AND H. S. ROSEN (1980): “Taxation, human capital, and uncertainty,” *The American Economic Review*, 70(4), 705–715.
- FARHI, E., AND I. WERNING (2013): “Insurance and taxation over the life cycle,” *The Review of Economic Studies*, 80(2), 596–635.
- FINDEISEN, S., AND D. SACHS (2012): “Education and Optimal Dynamic Taxation: The Role of Income-Contingent Student Loans,” *Working Paper No. 40, Department of Economics, University of Zurich*.
- (2015): “Insurance and Redistribution with Simple Tax Instruments,” *Working Paper*.
- (2016): “Education Policies and Taxation without Commitment,” *Working Paper*.
- FRIEDMAN, M. (1955): “The role of government in education,” *Rutgers University Press*.
- GALLIPOLI, G., C. MEGHIR, AND G. VIOLANTE (2011): “Equilibrium effects of education policies: A quantitative evaluation,” *Working Paper*.
- GARY-BOBO, R., AND A. TRANNOY (2015): “Optimal Student Loans and Graduate Tax under Moral Hazard and Adverse Selection,” *The RAND Journal of Economics*, 46(3), 546–576.
- GOLOSOV, M., N. KOCHERLAKOTA, AND A. TSYVINSKI (2003): “Optimal Indirect and Capital Taxation,” *Review of Economic Studies*, 70(3), 569–587.
- GOLOSOV, M., M. TROSHKIN, AND A. TSYVINSKI (2016): “Redistribution and Social Insurance,” *American Economic Review*, 106(2), 359–86.
- GOLOSOV, M., AND A. TSYVINSKI (2006): “Designing optimal disability insurance: A case for asset testing,” *Journal of Political Economy*, 114(2), 257–279.
- GROCHULSKI, B., AND T. PISKORSKI (2010): “Risky Human Capital and Deferred Capital Income Taxation,” *Journal of Economic Theory*, 145(3), 908–943.
- GUNER, N., R. KAYGUSUZ, AND G. VENTURA (2014): “Income Taxation of U.S. Households: Facts and Parametric Estimates,” *Review of Economic Dynamics*, 17(4), 559–581.

- HENDRICKS, L., AND T. SCHOELLMAN (2014): “Student Abilities During the Expansion of US Education,” *Journal of Monetary Economics*, 63, 19–36.
- HUGGETT, M., G. VENTURA, AND A. YARON (2011): “Sources of Lifetime Inequality,” *American Economic Review*, 101(7), 2923–2954.
- JACOBS, B., AND A. L. BOVENBERG (2010): “Human capital and optimal positive taxation of capital income,” *International Tax and Public Finance*, 17(5), 451–478.
- JACOBS, B., AND A. L. BOVENBERG (2011): “Optimal Taxation of Human Capital and the Earnings Function,” *Journal of Public Economic Theory*, 13(6), 957–971.
- JACOBS, B., D. SCHINDLER, AND H. YANG (2012): “Optimal Taxation of Risky Human Capital,” *Scandinavian Journal of Economics*, 114(3), 908–931.
- KAPICKA, M. (2006): “Optimal Income Taxation with Human Capital Accumulation and Limited Record Keeping,” *Review of Economic Dynamics*, 9(4), 612–639.
- KAPICKA, M., AND J. NEIRA (2015): “Optimal Taxation with Risky Human Capital,” *Working paper*.
- KRUEGER, D., AND A. LUDWIG (2013): “Optimal Progressive Labor Income Taxation and Education Subsidies when Education Decisions and Intergenerational Transfers are Endogenous,” *American Economic Review Papers & Proceedings*, 103(3), 496–501.
- LEE, D., S. Y. LEE, AND Y. SHIN (2015): “The Option Value of Human Capital,” *NBER Working Paper No. 21724*.
- LEMIEUX, T. (2006a): “Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill?,” *The American Economic Review*, pp. 461–498.
- (2006b): “Postsecondary Education and Increasing Wage Inequality,” *American Economic Review*, 96(2), 195–199.
- LUTTMER, E., AND R. ZECKHAUSER (2008): “Schedule selection by agents: from price plans to tax tables,” *NBER Working Paper No. 13808*.
- MANKIW, N., M. WEINZIERL, AND D. YAGAN (2009): “Optimal Taxation in Theory and Practice,” *Journal of Economic Perspectives*, 23(4), 147–174.
- MIRRLEES, J. A. (1971): “An Exploration in the Theory of Optimum Income Taxation,” *The Review of Economic Studies*, 38(2), 175–208.
- PIKETTY, T., AND E. SAEZ (2013): “Optimal labor income taxation,” *Handbook of Public Economics*, Vol 5.

- RAMSAY, C. M. (2006): “The distribution of sums of certain iid Pareto variates,” *Communications in Statistics- Theory and Methods*, 35(3), 395–405.
- SAEZ, E. (2001): “Using Elasticities to Derive Optimal Income Tax Rates,” *Review of Economic Studies*, 68(1), 205–229.
- SAEZ, E. (2002): “Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses,” *Quarterly Journal of Economics*, 117(3), 1039–1073.
- SCHEUER, F. (2014): “Entrepreneurial Taxation with Endogenous Entry,” *American Economic Journal: Economic Policy*, 6(2), 126–63.
- STANTCHEVA, S. (2015): “Optimal Taxation and Human Capital Policies Over the Lifecycle,” *Working Paper, Harvard University*.
- STIGLITZ, J. E. (1982): “Self-selection and Pareto efficient taxation,” *Journal of Public Economics*, 17(2), 213–240.
- STORESLETTEN, K., C. TELMER, AND A. YARON (2004): “Consumption and risk sharing over the life cycle,” *Journal of Monetary Economics*, 51(3), 609–633.
- TABER, C. (2001): “The Rising College Premium in the Eighties: Return to College or Return to Unobserved Ability?,” *Review of Economic Studies*, 68(3), 665–691.
- WERNING, I. (2007): “Optimal Fiscal Policy with Redistribution,” *The Quarterly Journal of Economics*, 122(3), 925–967.