Optimal Taxation with Risky Human Capital and Retirement Savings*

Radoslaw Paluszynski[†] Pei Cheng Yu[‡]

This version:
February 14, 2020

Abstract

We study optimal tax policies with human capital investment and retirement savings for present-biased agents. Agents are heterogeneous in their innate ability and make risky education investments which determines their labor productivity. We demonstrate that the optimal distortions vary with education status. In particular, the optimal policy encourages human capital investment with savings incentives. Our implementation uses income-contingent student loans and existing retirement policies, augmented by a new tax instrument that subsidizes retirement savings for college graduates. The instrument mimics the latest policy proposals by allowing employers to offer 401(k) matching contributions proportional to student loans repayment.

Keywords: Present bias, Human capital, Retirement, Sequential screening

*The authors would like to thank Gonzalo Castex, Taha Choukhmane, German Cubas, Mikhail Golosov, David Rahman, Satoshi Tanaka, Kei Mu Yi and Hsin-Jung Yu.

[†]University of Houston (e-mail: *rpaluszynski@uh.edu*)
[‡]University of New South Wales (e-mail: *pei-cheng.yu@unsw.edu.au*)

1 Introduction

The average cost of higher education in the US has been growing nearly eight times faster than median household income over the last two decades. Due to the lack of insurance against labor market uncertainties, this rise in college costs can reduce investment in higher education. At the same time, policymakers have been concerned about the seemingly insufficient amount of private retirement savings. Raising the welfare of retirees with more generous social security benefits would require imposing distortionary taxes, which makes policies that increase private savings preferable. Though human capital investment and retirement savings are usually treated as separate policy issues, this paper argues that they are interdependent when people are present-biased.

There are currently multiple policy proposals in the US that suggest making retirement savings contingent on student loan repayment, which establishes a link between retirement and education policies. These proposals are based on a pathbreaking IRS ruling in 2018 that allowed a company to make contributions to the retirement plans of employees who are paying off their student debt even if they do not make any 401(k) contributions.[1] In essence, individuals automatically save for retirement while repaying student loans. Other private employers have since offered similar benefits. Following these developments, the Retirement Parity for Student Loans Act and the Retirement Security and Savings Act were introduced in Congress. Both bills allow employer 401(k) matching based on student loan payments. Despite the enthusiasm of policymakers, the benefit of connecting student loan repayment to retirement savings is not apparent. This paper provides a theoretical guidance in linking education with retirement savings—two seemingly unrelated areas of government policy.

We study a Mirrlees life-cycle model with present-biased agents. We focus on present-biased agents to capture the self-control problem documented in recent empirical studies on the underinvestment in education (Cadena and Keys, 2015) and insufficient retirement savings (Angeletos et al., 2001; Laibson et al., 2017). In our framework, agents initially differ in their innate ability which could either be high or low. Based on their innate ability, agents choose their level of education: college or high school. Afterwards, they work before they retire. The likelihood of having higher productivity when working increases with innate ability and education status. Both innate ability and productivity are the agents' private information, so the government sequentially screens the agents and designs policies conditioned on the observed education investment and income. Crucially, the government separates agents so that high innate ability agents go to college while low innate ability agents do not. The government also attempts to paternalistically offset the present bias.

[1] It was revealed that the company involved in the ruling was Abbott Laboratories, a health care company.

Our theoretical framework characterizes the optimal interdependence between retirement savings and education investment. When agents are saving for retirement, college graduates are rewarded with a retirement savings subsidy. Intuitively, present-biased agents who are deciding on their education investment want to prevent their future-selves from under-saving for retirement. Therefore, the optimal retirement savings policy incentivizes human capital investment by providing college graduates with a savings vehicle that offsets their present bias. For non-college graduates, the retirement savings subsidy is lower and it is monotonically decreasing in income. Since high innate ability agents are more likely to earn higher income, a savings subsidy for non-college graduates that decreases with income discourages high innate ability agents from entering the workforce without a college degree.

We also show that the usual inverse Euler equation for time-consistent agents does not hold. When agents invest in higher education, the inverse marginal utility of consumption is *strictly higher* than the working period's expected inverse marginal utility. This implies that consumption is more front-loaded for college graduates compared to the time-consistent case. As a result, the optimal student loans are more generous to gratify the present-biased agents and encourage the ones with high innate ability to invest in college.

We also derive the optimal labor wedge for our environment. Though the theoretical characterization differs from those obtained with time-consistent agents, we show quantitatively that the optimal labor wedge with present-biased agents is very close to the one for time-consistent agents. More precisely, two new opposing forces are introduced when agents are present-biased. First, present-biased agents are less sensitive to future incentives. As a result, the labor distortion decreases to mitigate the present-biased agents' tendency to undervalue future incentives. Second, by helping agents commit to saving for retirement, the government is distorting intertemporal decisions from the present-biased agents' perspective. Therefore, the introduction of commitment increases labor distortions. In the quantitative analysis, we show that these two economic forces approximately offset each other.

We consider two ways of decentralizing the constrained efficient allocations. For one of them, we introduce retirement accounts with income and education contingent savings subsidies. We also consider one where non-college graduates rely mainly on social security benefits during retirement, while college graduates are supplemented with deposits worth a fraction of their student loan repayments in their retirement savings accounts. This latter implementation is inspired by the recent IRS ruling and policy proposals that treat student loans repayment as salary reduction contribution to retirement accounts. For both schemes, the government provides individuals with income-contingent student loans.

We bring our model to the US data by calibrating the structural parameters and by approximating the current tax system to infer realistic distributions of skills among high school and college graduates.[2] We show that our theoretical predictions are quantitatively significant. The optimal tax schedules involve extensive use of the intertemporal wedge during working life which, crucially, differs across income and education groups. College graduates are offered savings subsidies to smooth their consumption over the life cycle, which ex-ante incentivizes them to choose college education. The difference in savings subsidies between the two education groups declines with income, because the utility is close to linear at high levels of income which results in low gains from consumption smoothing. Finally, we show that the welfare gains from our optimal tax are potentially significant, exceeding 1% of lifetime consumption relative to the world with optimal policies dedicated to time-consistent agents.

1.1 Related Literature

This paper contributes to the literature on optimal human capital policies. Bovenberg and Jacobs (2005) studies optimal education and income policies in an environment where schooling increases productivity. However, human capital investment in their environment is riskless. This is contrary to empirical studies that find returns to human capital investments to be risky (Cunha and Heckman, 2007). This paper captures the risky returns to education by modeling productivity as a random draw from a distribution determined by human capital. There are other papers that have studied how risk from human capital investments affects the design of optimal policy. Anderberg (2009) finds that how human capital affects the degree of wage risk matters for optimal policy. Grochulski and Piskorski (2010) focuses on the optimal capital taxation in an environment where agents share the same innate ability and human capital investment is unobservable. Craig (2019) studies a setting where employers observe informative but imperfect signals to infer the human capital investment of ex-ante heterogeneous workers. In contrast, our paper focuses on how initial differences in innate ability affects the design of policies when investment in education is observable.

Several papers have also examined the optimal policy for human capital acquisition over the working age. Bohacek and Kapicka (2008) and Kapicka (2015) study the optimal tax policy when human capital investment is deterministic while the agent works. Stantcheva (2017) studies an environment where agents make monetary investments in each period to build up their stock of human capital. Makris and Pavan (2019) examine the learning-by-doing aspect of human capital accumulation, so human capital is acquired stochastically as a by-product from labor effort. Kapicka and Neira (2019) considers risky but unobservable human capital investment, so tax policies are not conditional upon this investment. In

[2] Our calibration uses an extended definition of college that includes Master's, Doctoral and Professional degrees.

contrast, our work focuses on human capital acquired before agents enter the labor force.

Gary-Bobo and Trannoy (2015) and Findeisen and Sachs (2016) consider environments most similar to ours. They examine optimal education and income tax policies in a setting where agents differ in initial ability and make risky investments in education before they enter the labor market. Our paper models initial ability and the risk from human capital investment in a similar fashion to their paper. However, we consider present-biased agents. In our setting, we demonstrate the importance of linking retirement policies to education outcomes. For time-consistent agents, as in their papers, introducing distortions in retirement savings does not increase investments in human capital.

Our paper contributes to the literature on Mirrlees taxation when agents have behavioral biases. Farhi and Gabaix (2019) use sparse maximization (Gabaix, 2014) to study optimal taxation of behavioral agents in a static setting. Lockwood (2018) studies optimal income taxation with present-biased agents where wages depend on past work effort. He shows how present bias has a potentially large effect on the optimal marginal income tax rate. We find an economic force similar to the one in Lockwood (2018). However, we also discover a novel opposing force that quantitatively negates the effect of present bias on the optimal labor wedge. There are also papers that focus on the design of retirement savings policies for time-inconsistent agents in a Mirrlees setting. Moser and de Souza e Silva (2019) consider a multi-dimensional screening environment, where agents are heterogeneous in present bias and productivity. Yu (2019a) examines a multi-dimensional screening environment, where the agents' degree of present bias, productivity and level of sophistication are hidden.

The rest of the paper is organized as follows. Section 2 presents the dynamic life-cycle model and Section 3 characterizes the optimal savings and labor wedges. In Section 4, we calibrate the model and present the quantitative results and welfare analysis. Section 5 demonstrates two policies that decentralizes the optimum. Section 6 extends the theory to include differences in present bias and also considers the case with naïve agents.

2 Model

We consider a life-cycle model with three periods: $t = 0, 1, 2$. At $t = 0$, agents learn their innate ability $\gamma \in \{H, L\}$ with $H > L$, and proceed to choose their education investment $e \in \{e_L, e_H\}$ where $e_H \geq e_L$. We will refer to agents with innate ability γ as γ-agents. The share of γ-agents is $\pi_\gamma \in (0, 1)$ with $\pi_H + \pi_L = 1$. The level of education investment e represents the binary decision of whether to invest in higher-education (invest e_H) or not (invest e_L). Human capital depends on both γ and e, which we denote as $\kappa(e, \gamma)$. We assume κ is strictly increasing in both arguments. Intuitively, this captures the fact that education helps raise human capital, and also how H-agents are more effective in human capital accumulation than L-agents. The government observes e while κ and γ are the agents' private information. We refer to γ as the *ex-ante* private information.

At $t = 1$, agents enter the labor market, and privately learn their productivity $\theta \in \Theta = [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}_+$. Productivity is drawn from a differentiable distribution with c.d.f. $F(\theta|\kappa)$, which depends on human capital κ and is ranked according to first order stochastic dominance: if $\kappa > \kappa'$, $F(\theta|\kappa) < F(\theta|\kappa'), \forall\theta \in \Theta$. Also, let $f(\theta|\kappa)$ denote the p.d.f. and assume $f(\theta|\kappa)$ is bounded away from zero for any θ and κ. This models the riskiness of human capital investment, where agents with higher human capital are more likely to be productive. An agent with productivity θ who provides work effort l produces output $y = \theta l$. The government observes output y, but not productivity θ nor labor supply l. We refer to θ as the *ex-post* private information. Finally, at $t = 2$, agents retire and consume their savings.

To model present bias, we adopt the quasi-hyperbolic discounting model (Laibson, 1997). Let $\beta < 1$ denote the short-run discount factor, which represents the degree of present bias. Let δ denote the long-run discount factor. Agents of type (γ, θ) have the following utility at $t = 1:$

$$U_1(c_1, c_2, y; \gamma, \theta, e) = u(c_1) - h\left(\frac{y}{\theta}\right) + \beta\delta_2 u(c_2).$$

The flow utilities u and h are defined for consumption $c_t \geq 0$ and output $y \geq 0$, respectively. Utility from consumption u is twice differentiable, strictly increasing, and strictly concave: $u', -u'' > 0$. Disutility from labor $h(l)$ is twice differentiable, strictly increasing, and strictly convex: $h', h'' > 0$, with $h(0) = 0$. Agents with innate ability γ have the following utility at $t = 0:$

$$U_0(\{c_t\}, e, y; \gamma) = \delta_0(e)\, u(c_0) + \beta\delta_1(e) \int_\Theta \left[u(c_1) - h\left(\frac{y}{\theta}\right) + \delta_2 u(c_2)\right] f(\theta|\kappa(e,\gamma))\, d\theta.$$

The length of each period is different, so the long-run discount factor δ_t is determined by the annual discount factor and the number of years in that period. Furthermore, the length of the schooling period ($t = 0$) is different across education groups. Therefore, the long-run discount factors δ_0 and δ_1 are functions of e to reflect how the number of years in school affects the length of $t = 0$. We assume that all agents work the same number of years in $t = 1$, so δ_2 is constant across education groups. Under our specification, the flow utility and allocations are in annual terms. For example, (c_1, y) is the annual consumption-output bundle during the working period ($t = 1$). More details are provided in Section 4.

Crucially, since $\beta < 1$, present-biased agents discount the immediate future more than the distant future. As a result, agents in $t = 0$ dislike the fact that their future-selves in $t = 1$ under-save for retirement.

2.1 Planning Problem

To characterize the constrained efficient allocation, we use the direct mechanism—agents report their private information to the government. This theoretical characterization is useful because, in later sections, we will use it as a blueprint to decentralize the optimum as a competitive equilibrium. The government designs

$$P = \left\{ c_0(\gamma), \left[c_1(\gamma,\theta), c_2(\gamma,\theta), y(\gamma,\theta) \right]_{\theta\in\Theta} \right\}_{\gamma\in\{H,L\}}.$$

Since agents privately learn their innate ability γ and productivity θ sequentially, we require P to be incentive compatible for each period (Myerson, 1986). Let $U_1(\theta';\gamma,\theta)$ denote the utility of a type (γ,θ) agent who reported γ truthfully and reports $\theta' \in \Theta$ in $t=1$. The ex-post incentive compatibility constraints ensure the agents report θ truthfully: $\forall \theta,\theta' \in \Theta$,

$$U_1(\gamma,\theta) \equiv U_1(\theta;\gamma,\theta) \geq U_1(\theta';\gamma,\theta). \tag{1}$$

Let the utility in $t=0$ of γ-agents who reported innate ability γ' be denoted as

$$U_0(\gamma';\gamma) = \delta_0(e_{\gamma'})\, u(c_0(\gamma')) + \beta\delta_1(e_{\gamma'}) \int_\Theta \left[U_1(\gamma',\theta) + (1-\beta)\delta_2 u(c_2(\gamma',\theta)) \right] dF(\theta|\kappa_{\gamma',\gamma}),$$

where $\kappa_{\gamma',\gamma} = \kappa(e_{\gamma'},\gamma)$ and let $\kappa_{\gamma,\gamma} = \kappa_\gamma$. Then, the ex-ante incentive compatibility constraints ensure that the agents report γ truthfully at $t=0$: for any innate ability γ,γ',

$$U_0(\gamma) \equiv U_0(\gamma;\gamma) \geq U_0(\gamma';\gamma). \tag{2}$$

The government is paternalistic in that it treats present bias as an error and attempts to correct it. The basis for this is because $\beta \neq 1$ reflects a self-control problem that agents disapprove of in every other period (O'Donoghue and Rabin, 1999). The government attempts to increase investment in education and raise retirement savings by maximizing the sum of long-run utilities:

$$\sum_\gamma \pi_\gamma \left\{ \delta_0(e_\gamma)\, u(c_0(\gamma)) + \delta_1(e_\gamma) \int_\Theta \left[u(c_1(\gamma,\theta)) - h\left(\frac{y(\gamma,\theta)}{\theta}\right) + \delta_2 u(c_2(\gamma,\theta)) \right] f(\theta|\kappa_\gamma)\, d\theta \right\}$$

subject to the ex-post incentive constraints (1), the ex-ante incentive constraints (2) and the resource constraint

$$\sum_\gamma \pi_\gamma \left\{ \frac{-c_0(\gamma) - e_\gamma}{R_0(e_\gamma)} + \frac{1}{R_1(e_\gamma)} \int_\Theta \left[y(\gamma,\theta) - c_1(\gamma,\theta) - \frac{1}{R_2} c_2(\gamma,\theta) \right] f(\theta|\kappa_\gamma)\, d\theta \right\} \geq 0,$$

where R_t denotes the gross rate of return. We will assume that $\delta_t R_t = 1$.

It is worth emphasizing that, apart from the inherent investment risk, education is costly for two additional reasons. First, it is costly in terms of resources. Second, it is costly in terms of time, because receiving education delays entry into the labor market. Due to these reasons, present-biased agents are less willing to invest in human capital.

2.2 Characterizing Incentive Compatibility

Here, we derive a lemma that simplifies ex-post incentive compatibility and discuss the difficulties in theoretically characterizing ex-ante incentive compatibility. The following lemma characterizes the set of policies that are ex-post incentive compatible. Its proof is standard so it is omitted.

Lemma 1 *For any γ, P is ex-post incentive compatible if and only if (i.) $y(\gamma, \theta)$ is non-decreasing in θ, and (ii.) $U_1(\gamma, \theta)$ is absolutely continuous in θ, with* $\frac{\partial U_1(\gamma,\theta)}{\partial \theta} = \frac{y(\gamma,\theta)}{\theta^2} h'\left(\frac{y(\gamma,\theta)}{\theta}\right)$.

There are three main difficulties in characterizing ex-ante incentive compatibility for time-inconsistent agents. First, local ex-ante incentive compatibility does not necessarily imply global ex-ante incentive compatibility when agents are time-inconsistent (Halac and Yared, 2014; Galperti, 2015; Yu, 2019b). This paper simplifies the problem by examining the case with two levels of innate ability.

The second difficulty lies in the direction of the relevant deviation at $t = 0$. Usually, the relevant deviation is downwards when agents are time consistent. Findeisen and Sachs (2016) showed that part of the sufficient condition for this to be true requires output $y(\gamma, \theta)$ to be weakly increasing with innate ability γ. However, Yu (2019b) showed that the optimal allocations are usually non-monotonic with respect to ex-ante information. The non-monotonicity helps relax the ex-ante incentive constraints when agents are time-inconsistent. Therefore, it is unclear in which direction the ex-ante incentive constraints binds. For our theoretical analysis, we focus on the case where only the downward ex-ante incentive compatibility constraint—the incentive constraint for H-agents—binds. Then, in our quantitative analysis, we verify that the downward ex-ante incentive constraint is indeed the relevant constraint.

Finally, it is unclear whether the government should screen ex-ante private information in the first place. Given the parameters of the model, it is difficult to theoretically determine who should go to college: everyone, no one, or only the H-agents. For our theoretical analysis, we will focus on the case where only the H-agents go to college, and verify that this is indeed optimal given the cost of college in our quantitative analysis.

2.3 Wedges

To understand how present-bias and informational frictions affect the optimal tax policy, the paper will focus on characterizing the optimal intertemporal and labor wedges.

The intertemporal wedge in $t = 0$ for innate ability γ is

$$\tau_0^k(\gamma) = 1 - \frac{u'\left(c_0\left(\gamma\right)\right)}{\mathbb{E}_\theta\left[u'\left(c_1\left(\gamma,\theta\right)\right)|\gamma\right]},$$

and the intertemporal wedge in $t = 1$ for type (γ, θ) is

$$\tau_1^k(\gamma,\theta) = 1 - \frac{u'\left(c_1\left(\gamma,\theta\right)\right)}{u'\left(c_2\left(\gamma,\theta\right)\right)}.$$

When $\tau^k \neq 0$, then consumption is not smoothed across time. More specifically, if $\tau_t^k > 0$, then savings is restricted in t. Similarly, if $\tau_t^k < 0$, then savings is distorted upwards in t.

The labor wedge in $t = 1$ for type (γ, θ) is

$$\tau^w(\gamma,\theta) = 1 - \frac{h'\left(\frac{y(\gamma,\theta)}{\theta}\right)}{\theta u'\left(c_1\left(\gamma,\theta\right)\right)}.$$

Since agents' equilibrium wage is equal to their productivity θ in a competitive labor market, if $\tau^w \neq 0$, then agents are not working at the efficient level. In particular, if $\tau^w > 0$, then there is an under-supply of labor given the market wage. On the other hand, if $\tau^w < 0$, then agents are over-supplying labor given the market wage.

3 Theoretical Results

In this section, we derive the optimal intertemporal and labor wedges, which are crucial for determining the optimal tax policies.

3.1 Intertemporal Wedges

The following proposition provides the inverse Euler equations for present-biased agents.

Proposition 1 *The constrained efficient allocation satisfies (i.) the inverse Euler equation in aggregate:*

$$\sum_\gamma \frac{\pi_\gamma}{u'\left(c_0\left(\gamma\right)\right)} = \sum_\gamma \pi_\gamma \mathbb{E}_\theta\left(\frac{1}{u'\left(c_1\left(\gamma,\theta\right)\right)}\middle|\gamma\right), \tag{3}$$

(ii.) for any $\gamma \in \{H, L\}$,

$$\mathbb{E}_\theta \left(\frac{1}{u'(c_1(\gamma,\theta))} \middle| \gamma \right) = \mathbb{E}_\theta \left(\frac{1}{u'(c_2(\gamma,\theta))} \middle| \gamma \right), \tag{4}$$

and (iii.) for any $\theta \in \Theta$,

$$\frac{1}{\beta u'(c_2(H,\theta))} = \frac{1}{u'(c_1(H,\theta))} + \left(\frac{1-\beta}{\beta}\right)\left(\frac{\pi_H + \beta\mu}{\pi_H + \mu}\right)\frac{1}{u'(c_0(H))} \tag{5}$$

$$\frac{1}{\beta u'(c_2(L,\theta))} = \frac{1}{u'(c_1(L,\theta))} + \left(\frac{1-\beta}{\beta}\right)\left(\frac{\pi_L - \beta\mu\left(\frac{f(\theta|\kappa_{L,H})}{f(\theta|\kappa_L)}\right)}{\pi_L - \mu}\right)\frac{1}{u'(c_0(L))}, \tag{6}$$

where $\mu = [u'(c_0(L)) - u'(c_0(H))]\left[\frac{u'(c_0(L))}{\pi_L} + \frac{u'(c_0(H))}{\pi_H}\right]^{-1}$.

Proposition 1 follows from considering variations around any incentive compatible allocation that preserve incentive compatibility. The optimal allocation minimizes the resources expended, which satisfies (3) and (4).

To understand Proposition 1, first consider the standard inverse Euler equation at $t = 0$ for time-consistent agents:

$$\frac{1}{u'(c_0(\gamma))} = \mathbb{E}_\theta \left(\frac{1}{u'(c_1(\gamma,\theta))} \middle| \gamma \right) \text{ for any } \gamma.$$

By Jensen's inequality, the inverse Euler equation for time-consistent agents implies $u'(c_0(\gamma)) < \mathbb{E}_\theta[u'(c_1(\gamma,\theta))]$ for any γ. Due to informational constraints, the transfer of consumption from $t = 0$ to $t = 1$ for time-consistent agents is restricted regardless of their ex-ante private information.[3] By restricting savings, the government can induce effort in $t = 1$ at a lower cost, which relaxes the ex-post incentive constraint. Therefore, the intertemporal wedge τ_0^k is strictly positive for any innate ability γ.

With present-biased agents, intertemporal distortions have an additional effect on welfare. The government can relax the ex-ante incentive compatibility constraint by increasing the intertemporal wedge for H-agents while decreasing the wedge for L-agents at $t = 0$. To see this, if we take expectation of (5) and (6) with respect to θ, then by (4) we can derive the following inverse Euler inequalities:

$$\frac{1}{u'(c_0(H))} > \mathbb{E}_\theta \left(\frac{1}{u'(c_1(H,\theta))} \middle| H \right)$$

[3]See Golosov et al. (2003) for more on the inverse Euler equation for time-consistent agents.

and

$$\frac{1}{u'(c_0(L))} < \mathbb{E}_\theta\left(\frac{1}{u'(c_1(L,\theta))}\middle| L\right).$$

Comparing it with the standard inverse Euler equation, the consumption for H-agents is even more front-loaded while the consumption is relatively back-loaded for L-agents.[4] In other words, the government exacerbates the intertemporal distortion for H-agents to satisfy their temptation, encouraging them to accumulate human capital. Furthermore, the less front-loaded consumption path for L-agents helps discourage downward deviations. As a result, the best the government can do is to choose consumption such that the inverse marginal utility is equalized in aggregate, which is implied by (3).

Next, note that the optimal intertemporal decision at $t = 1$ for time-consistent agents satisfies the standard Euler equation:

$$u'(c_1(\gamma,\theta)) = u'(c_2(\gamma,\theta)) \text{ for any } \gamma,\theta.$$

This implies that it is optimal for time-consistent agents to smooth consumption between work and retirement periods. This is because there is no additional uncertainty beyond $t = 1$, so the government does not face informational constraints in the future. As a result, there is no need to distort the intertemporal margin at $t = 1$.

From the government's perspective, present-biased agents save too little for their retirement. In essence, if present-biased agents were allowed to save freely, then $u'(c_1(\gamma,\theta)) = \beta u'(c_2(\gamma,\theta))$, so the intertemporal wedge would be $\tau_1^k(\gamma,\theta) = 1-\beta > 0$ for any type (γ,θ). In contrast, by (5) and (6), it is optimal for the intertemporal wedge τ_1^k to depend on reported innate ability and productivity. This is because, with present-biased agents, an intertemporal wedge in $t = 1$ that depends on the reported innate ability relaxes the ex-ante incentive constraint. To see why, recall that agents at $t = 0$ are concerned that they will not save enough at $t = 1$. Thus, they have demand for a commitment device that incentivizes them to save more. By introducing a wedge on retirement savings that depends on past reports, the government is able to influence the incentives for education attainment.

More specifically, by (5), we immediately notice that

$$\frac{u'(c_1(H,\theta))}{u'(c_2(H,\theta))} > \beta.$$

[4] Grochulski and Piskorski (2010) found that the inverse marginal utility of consumption is a strict supermartingale when agents are time consistent and ex-ante identical. In their paper, human capital investments are unobservable, so under investing in education is complementary to shirking in future periods. Hence, in addition to the usual distortion to deter over-saving, the optimal policy makes the intertemporal distortion worse at the education stage to deter under-investing in education. If education investment was observable in their environment, like ours, then the intertemporal distortion disappears.

In other words, the government helps the H-agents save for retirement. The government essentially rewards H-agents for going to college with a commitment device that helps them smooth consumption across work and retirement. This commitment device helps substitute part of the information rent to H-agents.

By (6), for L-agents, the marginal rate of intertemporal substitution is

$$\frac{u'\left(c_1\left(L,\theta\right)\right)}{u'\left(c_2\left(L,\theta\right)\right)}\begin{cases} > \beta & \text{if } \pi_L > \beta\mu\frac{f\left(\theta|\kappa_{L,H}\right)}{f\left(\theta|\kappa_L\right)} \\ = \beta & \text{if } \pi_L = \beta\mu\frac{f\left(\theta|\kappa_{L,H}\right)}{f\left(\theta|\kappa_L\right)} \\ < \beta & \text{if } \pi_L < \beta\mu\frac{f\left(\theta|\kappa_{L,H}\right)}{f\left(\theta|\kappa_L\right)} \end{cases}.$$

Notice that the retirement savings for L-agents depend on the likelihood ratio $\frac{f\left(\theta|\kappa_{L,H}\right)}{f\left(\theta|\kappa_L\right)}$. If $\frac{f\left(\theta|\kappa_{L,H}\right)}{f\left(\theta|\kappa_L\right)}$ is relatively large, meaning that the observed productivity is likely to have come from an agent with high innate ability, then it is optimal to distort the retirement savings such that the present bias is exacerbated. The government uses this additional intertemporal distortion to deter the H-agents from under-investing in education. It is also a cost effective method since L-agents are unlikely to have that level of productivity. On the other hand, if $\frac{f\left(\theta|\kappa_{L,H}\right)}{f\left(\theta|\kappa_L\right)}$ is relatively small, meaning that the observed productivity is unlikely to have come from a H-agent, then the government helps offset the present bias.

Assumption 1 *f satisfies the monotone likelihood ratio property: $\frac{f(\theta|\kappa)}{f(\theta|\kappa')}$ is increasing in θ for any $\kappa > \kappa'$.*

Assumption 1 implies that higher productivity θ is more likely to come from higher accumulated human capital κ. When Assumption 1 holds, the optimal intertemporal wedge for L-agents increases with productivity. This implies that the government helps the L-agents who are less productive with their retirement savings, while the retirement savings of L-agents who are highly productive are restricted. This is because Assumption 1 implies that H-agents who do not invest in higher education are more likely than L-agents to be productive. As a result, the government exacerbates the present bias of low-educated and productive agents to relax the ex-ante incentive constraint and induce H-agents to increase education attainment. We assume Assumption 1 holds for the rest of the paper.[5]

The intertemporal distortion of τ_1^k described above provides us with the theoretical basis for linking retirement savings to education investent. Hence, it is worth noting that the intertemporal wedge τ_1^k is distorted because γ is the private information of present-biased

[5] Assumption 1 implies first order stochastic dominance.

agents. If the government can observe γ, then the standard Euler equation would hold: $u'(c_1(\gamma,\theta)) = u'(c_2(\gamma,\theta))$ for any γ and θ. As was mentioned before, if agents were time-consistent instead, then the standard Euler equation for $t = 1$ also holds. As a result, the interdependence between retirement savings and education investment is solely used to relax the ex-ante incentive compatibility constraint of present-biased agents.

Also, if productivity θ is a deterministic function of human capital κ, so human capital investment is riskless, then there will also be no intertemporal distortions. In fact, Yu (2019a) showed that the government can implement the full-information efficient optimum through the use of off-path policies when there is no dynamic private information.

3.2 Labor Wedge

The dynamic incentive problem and the agents' present bias also affects the labor wedge. To separate the economic forces that determine the optimal labor distortions, we define

$$A_\gamma(\theta) = \frac{1 - F(\theta|\kappa_\gamma)}{\theta f(\theta|\kappa_\gamma)},$$

$$B_\gamma(\theta) = 1 + \frac{\frac{y(\gamma,\theta)}{\theta} h''\left(\frac{y(\gamma,\theta)}{\theta}\right)}{h'\left(\frac{y(\gamma,\theta)}{\theta}\right)},$$

$$C_\gamma(\theta) = \int_\theta^{\bar\theta} \frac{u'(c_1(\gamma,\theta))}{u'(c_1(\gamma,x))}\left[1 - \frac{u'(c_1(\gamma,x))}{\phi}\right] \frac{f(x|\kappa_\gamma)}{1 - F(\theta|\kappa_\gamma)} dx,$$

$$D_\gamma(\theta) = u'(c_1(\gamma,\theta))\left[\frac{1}{u'(c_0(\gamma))} - \frac{1}{\phi}\right],$$

$$E_\gamma(\theta) = \left[\frac{u'(c_1(\gamma,\theta))}{\beta u'(c_2(\gamma,\theta))} - 1\right] - \left(\frac{1-\beta}{\beta}\right)\frac{u'(c_1(\gamma,\theta))}{\phi},$$

where $\phi > 0$ is the shadow price on the resource constraint.

Proposition 2 *The labor wedge for any $\theta \in \Theta$ satisfies*

$$\frac{\tau^w(H,\theta)}{1 - \tau^w(H,\theta)} = A_H(\theta) B_H(\theta) \left[C_H(\theta) - D_H(\theta) + E_H(\theta)\right], \tag{7}$$

$$\frac{\tau^w(L,\theta)}{1 - \tau^w(L,\theta)} = A_L(\theta) B_L(\theta) \left[C_L(\theta) - \left(\frac{1 - F(\theta|\kappa_{L,H})}{1 - F(\theta|\kappa_L)}\right) D_L(\theta) + \frac{g(\theta|\kappa_L)}{g(\theta|\kappa_{L,H})} E_L(\theta)\right], \tag{8}$$

where $g(\theta|\kappa) = \frac{f(\theta|\kappa)}{1-F(\theta|\kappa)}$ *and* $\frac{1}{\phi} = \mathbb{E}_\gamma\left[\mathbb{E}_\theta\left(\frac{1}{u'(c_1(\gamma,\theta))}\middle|\gamma\right)\right]$.

Proposition 2 presents the optimal labor wedge for present-biased agents in a sequential screening environment. Following Golosov et al. (2016), we decompose the economic forces into three distinct components: intratemporal, intertemporal and present-bias components. The intratemporal component summarizes the trade-off between production efficiency and insurance against productivity differences. The intertemporal component captures how labor distortions affect the education decision in the previous period. Unique to our paper, the present-bias component encompasses the effects of time inconsistency on the optimal labor distortions. We rewrite (7) and (8) to pinpoint each component:

$$\frac{\tau^w(H,\theta)}{1-\tau^w(H,\theta)} = \underbrace{A_H(\theta)\,B_H(\theta)\,C_H(\theta)}_{\text{intratemporal component}} - \underbrace{A_H(\theta)\,B_H(\theta)\,D_H(\theta)}_{\text{intertemporal component}} + \underbrace{A_H(\theta)\,B_H(\theta)\,E_H(\theta)}_{\text{present-bias component}},$$

$$\frac{\tau^w(L,\theta)}{1-\tau^w(L,\theta)} = \underbrace{A_L(\theta)\,B_L(\theta)\,C_L(\theta)}_{\text{intratemporal component}} - \underbrace{\left(\frac{1-F(\theta|\kappa_{L,H})}{1-F(\theta|\kappa_L)}\right)A_L(\theta)\,B_L(\theta)\,D_L(\theta)}_{\text{intertemporal component}} + \underbrace{\frac{g(\theta|\kappa_L)}{g(\theta|\kappa_{L,H})}A_L(\theta)\,B_L(\theta)\,E_L(\theta)}_{\text{present-bias component}}.$$

All components are affected by $A_\gamma(\theta)$ and $B_\gamma(\theta)$. To understand these terms, first note that by introducing a labor wedge for type (γ,θ) agents, their labor supply changes according to their Frisch elasticity of labor supply, which is $B_\gamma(\theta)$. Furthermore, an increase in the labor distortion for agents of type (γ,θ) decreases their total output in proportion to $\theta f(\theta|\kappa)$. Meanwhile, the incentive constraints for higher productivity agents of mass $1-F(\theta|\kappa)$ are relaxed. This trade-off is captured by $A_\gamma(\theta)$.

Without dynamic information, the optimal labor wedge is determined by the intratemporal component, which summarizes the economic forces in static models, such as Diamond (1998) and Saez (2001). In addition to $A_\gamma(\theta)$ and $B_\gamma(\theta)$, the intratemporal component also consists of $C_\gamma(\theta)$, which captures the strength of the government's insurance motive against the productivity shock. In static Mirrlees, the inverse marginal utility is the cost of a marginal increase in utility in consumption terms, so the cost of a marginal increase in average utility in $t=1$ is $\frac{1}{\phi}$. Hence, if the cost of increasing average utility is small relative to the cost of increasing the utility of (γ,x) agents ($\frac{1}{\phi}<\frac{1}{u'(c_1(\gamma,x))}$), then $C_\gamma(\theta)$ is positive. This is because the benefits of increasing the labor wedge of type (γ,θ) agents to relax the ex-post incentive constraints of higher productivity agents ($x\geq\theta$ types) outweighs the cost. Furthermore, the degree of labor distortion increases with consumption inequality, which is represented by $\frac{u'(c_1(\gamma,\theta))}{u'(c_1(\gamma,x))}$.

When there is dynamic information and agents are time-consistent, then the labor wedge is shaped by both the intratemporal and intertemporal components. This is similar to the labor distortions in Findeisen and Sachs (2016). The intertemporal component contains the term $D_\gamma(\theta)$ and is augmented by $\frac{1-F\left(\theta|\kappa_{L,H}\right)}{1-F(\theta|\kappa_L)}$ for L-agents. Notice that $D_\gamma(\theta)$ can be rewritten as $\frac{[u'(c_0(\gamma))]^{-1}-\phi^{-1}}{[u'(c_1(\gamma,\theta))]^{-1}}$. Therefore, by Proposition 1, we have $D_H(\theta) > 0$ and $D_L(\theta) < 0$. This implies that the government can encourage investment in education through promising a smaller labor wedge $\tau^w(H,\theta)$ rather than raising $c_0(H)$. Similarly, it increases the labor wedge of L-agents to discourage H-agents from working without a college degree. To that end, the government also exploits the fact that H-agents who mimicked L-agents are more likely to have higher productivity than actual L-agents, which is captured by $\frac{1-F\left(\theta|\kappa_{L,H}\right)}{1-F(\theta|\kappa_L)}$. This shows how the optimal labor distortion for non-college grads leverages the difference in productivity distribution between actual L-agents and H-agents who eschewed college.

When agents are present-biased, the present-bias component highlights the additional force that influences the labor wedge. The present-bias component is comprised of two potentially off-setting forces: the disagreement and myopic components,

$$E_\gamma(\theta) = \underbrace{\left[\frac{u'(c_1(\gamma,\theta))}{\beta u'(c_2(\gamma,\theta))} - 1\right]}_{\text{disagreement component}} - \underbrace{\left(\frac{1-\beta}{\beta}\right)\frac{u'(c_1(\gamma,\theta))}{\phi}}_{\text{myopic component}}.$$

The disagreement component characterizes how the labor distortion is affected by the government's policy of increasing retirement savings. The myopic component captures the fact that present-biased agents do not fully internalize the returns from working. Observe that when $\beta = 1$, the intertemporal wedge at $t = 1$ is zero so $E_\gamma(\theta) = 0$ for any γ and θ. Furthermore, the present-bias component for L-agents is amplified by $\frac{g(\theta|\kappa_L)}{g\left(\theta|\kappa_{L,H}\right)}$, which is greater than 1 by Assumption 1.[6] Again, this shows how the difference in productivity distribution for agents with varying innate abilities is used to relax the ex-ante incentive constraint.

To understand the disagreement component, recall from Proposition 1, where we showed how the H-agents are offered a commitment device in $t = 1$ to relax the ex-ante incentive constraints. It also provides commitment to L-agents conditional on their realized productivity. However, from the present-biased agents' perspective in $t = 1$, the optimal consumption path satisfies $u'(c_1) = \beta u'(c_2)$. Hence, the commitment the government provides is perceived as an intertemporal distortion by agents in $t = 1$, which distorts their labor supply by decreasing their incentives to work. This is reflected in the disagreement component which increases the labor wedge when the government provides more commitment. As a result,

[6] Assumption 1 also implies hazard rate dominance: $g(\theta|\kappa) \leq g(\theta|\kappa')$ for all θ when $\kappa > \kappa'$.

there is a trade-off between the provision of work incentives and commitment. This suggests that the consumption of more productive agents would be more front-loaded, because the gains from higher production efficiency outweighs the cost of lower retirement savings.

To understand the myopic component, note that labor supply is incentivized with more consumption in both $t = 1$ and $t = 2$. However, present-biased agents underestimate the benefits of work because they discount retirement consumption more heavily. As a result, the myopic component shows how the optimal labor wedge is decreased to correct for the tendency of present-biased agents to undervalue the rewards for effort. Furthermore, due to insurance motives, the myopic component decreases the labor wedge more for agents with lower working-period consumption c_1. So, less productive agents receive more help in internalizing the benefits of working. The optimal labor wedge in Lockwood (2018) also contains an economic force similar to the myopic component. However, Lockwood (2018) does not characterize the optimal labor wedge with dynamic consumptiom, so the disagreement component is absent in his characterization.

Lastly, notice that the myopic component always decreases the labor distortion, while the disagreement component increases it unless the government does not provide commitment: $\frac{u'(c_1)}{u'(c_2)} < \beta$. As a result, the present-bias component is shaped by two opposing economic forces. For H-agents, we can show that the disagreement component dominates the myopic component for all levels of productivity. Meanwhile, for L-agents, the myopic component dominates the disagreement component. To see this, notice that by (5) and (6) we can rewrite the present-bias component as

$$E_H(\theta) = \left(\frac{(1-\beta)\,\mu}{\pi_H}\right) \frac{u'(c_1(H,\theta))}{\phi},$$

$$E_L(\theta) = -\left(\frac{(1-\beta)\,\mu \frac{f(\theta|\kappa_{L,H})}{f(\theta|\kappa_L)}}{\pi_L}\right) \frac{u'(c_1(L,\theta))}{\phi}.$$

Beyond the government's paternalistic motives, providing H-agents with commitment also encourages them to enroll in college, so the disagreement component dominates for H-agents. At the same time, the value of providing L-agents with commitment is smaller to the government, so the myopic component dominates.

Since all of the components are composed of endogenous variables, we rely on our quantitative exercise to characterize the labor wedge. Appendix C quantifies the decomposition of the optimal labor wedge. As we show there, the shape of the labor wedge is mostly determined by our assumptions on the distribution of skills and the size of the intratemporal component. By contrast, the present-bias component plays a minor role quantitatively.

4 Quantitative Analysis

In this section, we quantify the model by imposing specific functional forms and calibrating their parameters. Then, we measure the quantitative significance of the theoretical results presented in Section 3, as well as the welfare gains under the optimal tax system.

Table 1 presents the calibrated parameter values. We select the standard functional forms for the utility of consumption $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ and the disutility of labor $h(\ell) = \frac{\ell^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}$, and we assume they are constant over time. The risk aversion and the Frisch elasticity of labor supply are then set to standard values of 2 and 0.5, respectively. The short- and long-run discount factors are calibrated based on the values proposed by Nakajima (2012). The short-term discount factor is 0.7, in the ballpark of the empirical estimates of Laibson et al. (2017). The long-run discount factors are based on the annual rate of 0.9852 and compounded to take into account the relative length of different periods. In the subsequent analysis we will also make comparisons with a variant of our model for time-consistent agents (i.e. $\beta = 1$). In that case, following Nakajima (2012), we recalibrate the effective discount factors based on the annual rate of 0.9698. The purpose of such a recalibration is to separate the effect of time-inconsistency in agents' behavior from their effectively increased impatience.

In our calibrated model, we expand the definition of high school and college graduates by admitting a wide range of real-world education outcomes. We associate the former with all individuals who hold an Associate's degree or less. The share of such low types in the 2015 Current Population Survey is 0.68. We associate the latter with all individuals who hold a Bachelor's, Master's, Professional or Doctoral degree. We assume that $t = 0$ begins at age 18 and lasts 5.12 years for the high types (reflecting a weighted average across all degree durations), or 0 years otherwise (hence, $\delta_0(e_L) = 0$). Agents work for 45 years and then retire and live for 20 years in retirement. The annual cost of higher education is calculated to be \$15,700. Section B.2 in the Appendix discusses more details of our calibration.

In order to calibrate the distributions of skills for agents of different innate ability and education, we create a separate model which we refer to as the "current policies" world. This model is described in detail in Appendix B. We take this model to the data, solve for optimal behavior and simulate large population of agents from each of the four groups: (i.) factual high school graduates, (ii.) high school graduates, had they gone to college (high school counterfactual), (iii.) factual college graduates, and (iv.) college graduates, had they not gone to college (college counterfactual). We select the parameters governing the distributions of skills such that the simulated distribution of lifetime earnings for each group matches the one reported by Cunha and Heckman (2007). In particular, this study uses a variation of the Roy model to infer counterfactual distributions of earnings for both high

Table 1: Parameter values in the model

Symbol	**Meaning**	**Value**
$\pi_0(L)$	Share of low type	0.68
$\pi_0(H)$	Share of high type	0.32
σ	Risk aversion	2
η	Frisch elasticity	0.5
e_H	Cost of higher education	1.57
Discount factors: present bias		
β	Short-term discount factor	0.7
$\delta_0(e_L)$	High school period 0 long-term discount factor	0.00
$\delta_1(e_L)$	High school period 1 long-term discount factor	1.00
$\delta_0(e_H)$	College period 0 long-term discount factor	0.15
$\delta_1(e_H)$	College period 1 long-term discount factor	0.93
δ_2	Retirement discount factor	0.27
Discount factors: time-consistent benchmark		
$\delta_0(e_L)$	High school period 0 long-term discount factor	0.00
$\delta_1(e_L)$	High school period 1 long-term discount factor	1.00
$\delta_0(e_H)$	College period 0 long-term discount factor	0.19
$\delta_1(e_H)$	College period 1 long-term discount factor	0.85
δ_2	Retirement discount factor	0.15

school and college graduates had they made the opposite education decision. Also, to correct for the under-representation of high-end earnings in the data, we add an upper Pareto-tail to each distribution such that the upper 10% of the mass is distributed according to a shape parameter of 1.5, as in Saez (2001). Figure 1 presents the four distributions backed out as a result of this procedure.

In what follows, we discuss our quantitative results. We begin with Table 2 which shows the optimal intertemporal wedge in $t = 0$. In line with the hallmark dynamic Mirrlees result, the government finds it optimal to tax savings in $t = 0$ in order to induce higher labor effort from agents in the next period. Notice also that the optimal wedge amounts are in the ballpark of the model with time-consistent agents, which is a result of our calibration that holds the effective discount factor constant across the two models. Importantly though, the wedge for present-biased agents is slightly higher, raising the consumption of college students and providing additional incentive to make the college investment.[7]

[7] The intertemporal wedge for L-agents $\tau_0^k(L)$ is not shown since for our quantitative exercise, we assumed L-agents do not have a student period ($\delta_0(e_L) = 0$).

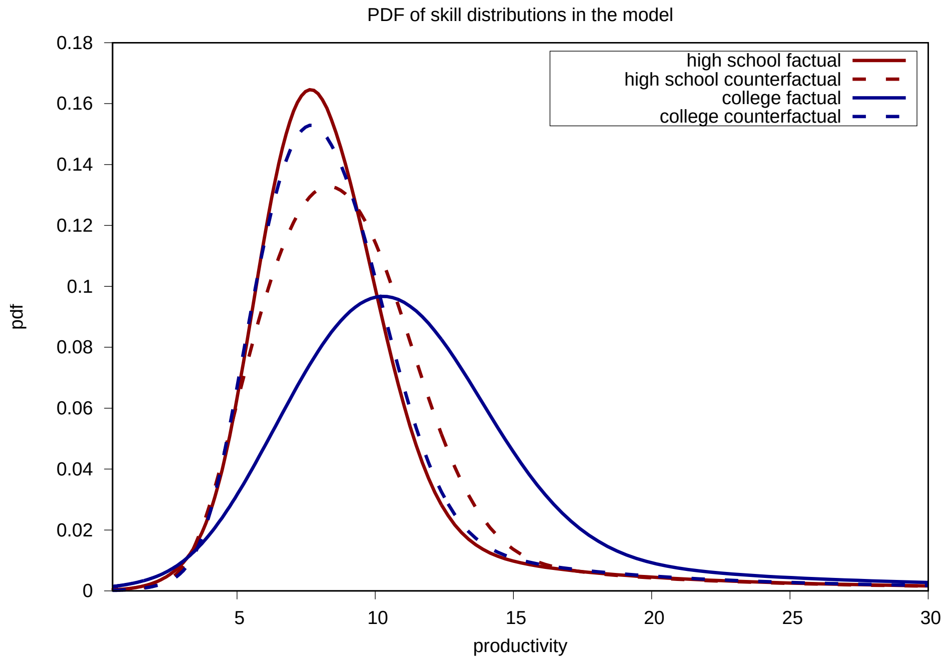

Figure 1: Calibrated distributions of skills for the four groups of agents

Table 2: Intertemporal wedge in period zero: present-bias vs. time-consistent case

	Present-biased	Time-consistent
$\tau_0^k(H)$	0.3244	0.2805

Figure 2 shows the optimal intertemporal wedges in $t = 1$ and conveys a key quantitative result. The wedges are negative for a wide interval of low incomes, and are always smaller than $1 - \beta$, which implies that the government offers savings subsidies.[8] This is an expected outcome in a model with paternalistic policies and agents who suffer from present bias. More importantly, the wedges are significantly different for the two education groups. College graduates enjoy a much higher subsidy than high school graduates at all income levels, with the difference eventually disappearing for higher incomes. The government does so in order to provide them with incentives to invest in college education ex-ante. Without such incentives, H-agents worry that additional education will not deliver a sufficient increase in their welfare, because their own present bias will prevent them from smoothing their higher working-age income across the life cycle. By contrast, notice that in the variant of our model with time-consistent agents, the optimal intertemporal wedge in the working-age period is equal to zero for both education groups. This is because time-consistent agents are able to raise retirement savings on their own.

[8]The theoretical result where the government decreases savings for sufficiently high θ is not quantitatively significant since the distributions $f(\theta|\kappa)$ and $f(\theta|\kappa_{L,H})$ are similar.

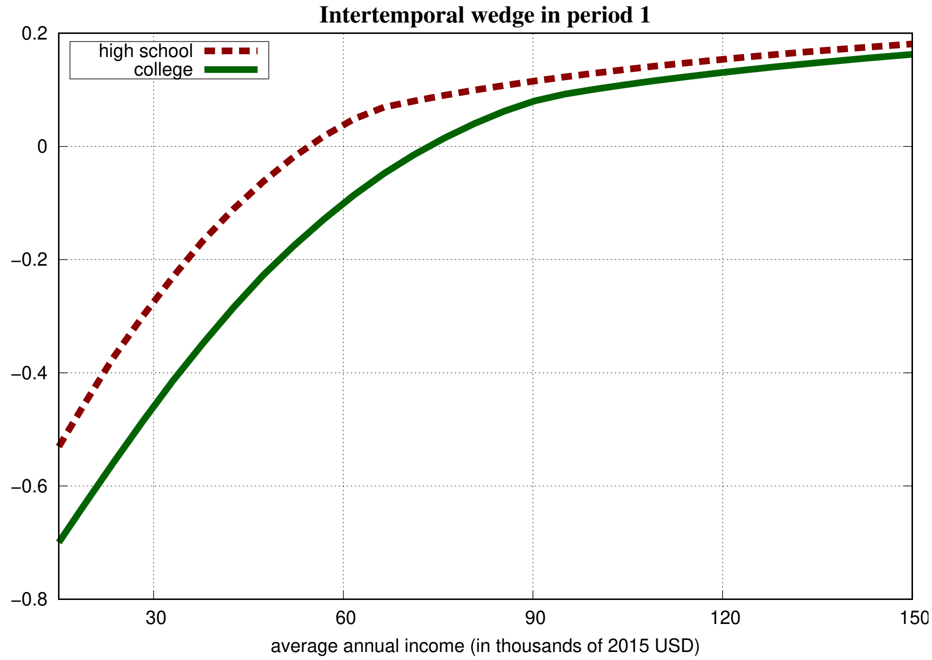

Figure 2: Intertemporal wedge in the model with present-biased agents

Figure 3 presents the optimal labor wedges for both education groups according to the two variants of our model: with present-biased agents or with time-consistent agents. The optimal labor wedges follow a U-shaped pattern and converge to a constant for top income levels, which is standard in Mirrlees taxation with Pareto-tailed productivity distributions (Diamond, 1998; Saez, 2001). It is important to notice that optimal wedges mostly decline with income and are significantly different for the two education groups. This resembles the main result of Findeisen and Sachs (2016) and stems from the fact that H-agents must be offered a separate income tax schedule to provide them with incentives to optimally choose to go to college. Notice that the differences in optimal wedges between the present-biased and time-consistent settings are very small, and arise predominantly at lowest incomes. This implies that the presence of present-biased agents does not alter the normative prescriptions in terms of the design of income tax schedules that the literature has established so far. Appendix C reinforces this point by showing that the present-bias component of the optimal labor wedge, as introduced in Section 3.2, is in general small quantitatively and declines monotonically with income. This is in contrast to a setting without dynamic consumption, where the myopic component is not mitigated by the disagreement component and marginal income tax rates could become negative (Lockwood, 2018).

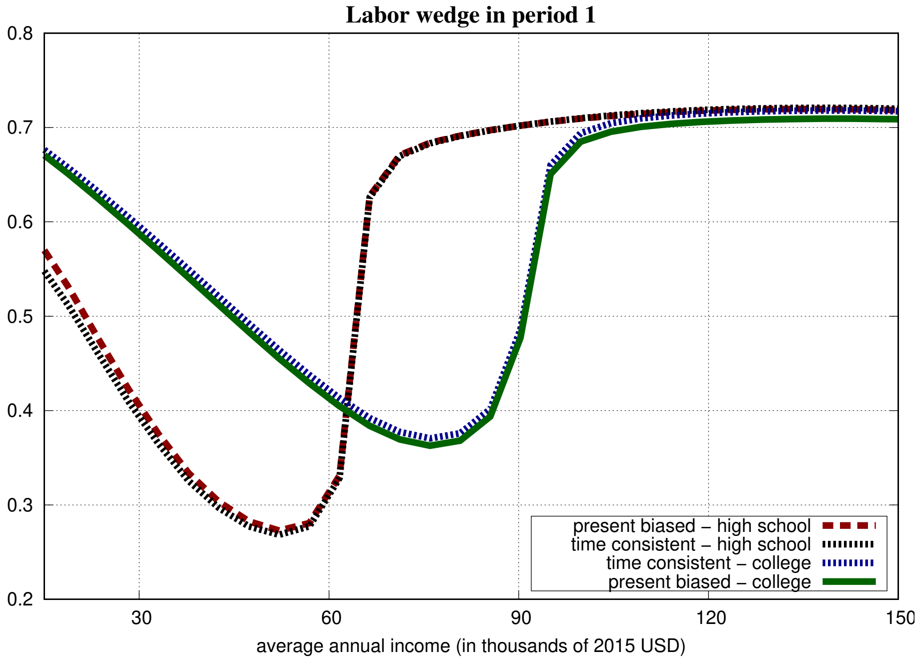

Figure 3: Labor wedge in the model with present-biased agents

4.1 Welfare Gains from Optimal Policies

We now turn our attention to the calculation of potential welfare gains arising from the optimal policies. We will compare our optimum to three separate benchmarks: optimal policies for time-consistent agents, as well as optimal policies with present-biased agents when either labor or intertemporal wedges are restricted to be education-independent.

4.1.1 Welfare Gains Relative to Optimal Time-Consistent Policies

As the first benchmark, we use the optimal policies dedicated to time-consistent agents, for whom $\beta = 1$. We consider two possible policy implementations for time-consistent agents.[9] The first one, called the laissez-faire implementation, leaves the agents alone in their retirement savings decision in period $t = 1$. Because the policy is dedicated for time-consistent agents, the government is confident that agents will smooth consumption in line with their time preferences. This is not the case for present-biased agents though, and we expect our optimal policies to bring about significant welfare gains relative to this benchmark.

In order to isolate the effect of education-dependent savings incentives from mere subsidization of retirement savings, we also consider a second implementation for time-consistent agents which features mandatory savings. In this world, agents are forced to smooth their

[9]The details of these implementations are presented in Appendix D.

consumption between working-life and retirement in line with the Euler equation. It does not make a difference for time-consistent agents who would have made the same choice anyway. On the other hand, the government helps present-biased agents save for retirement under this implementation, but without taking advantage of the education-dependent intertemporal wedge.

Table 3: Welfare gains over optimal policies for time-consistent agents

	Mandatory savings	Laissez-faire
% increase in lifetime consumption	0.96	2.19

Table 3 presents the welfare gains under our baseline parametrization relative to the two time-consistent benchmarks. Consistent with prior expectations, the gains over laissez-faire policies are the highest and amount to 2.19% of lifetime consumption. The gains come from increased retirement savings and also from improved production efficiency. On the other hand, the gains relative to mandatory savings is 0.96% of lifetime consumption, which is considerably lower than laissez-faire policies, but still significant. Since mandatory savings policy already forces agents to smooth their consumption, this implies the welfare gains of the optimal education-dependent policies largely come from more efficient production.

4.1.2 Welfare Gains from Education-Dependent Wedges

We now turn our attention to the benchmark of optimal policies with present-biased agents, but where some policies are restricted to be education-independent. Education-independent policies are policies that are conditioned only on observed income y. We will calculate the consumption-equivalent welfare gain from moving to the tax system where both wedges depend on education attainment, relative to either of the two cases: (i.) education-independent labor income tax and education-dependent savings subsidy and (ii.) education-independent savings subsidy but education-dependent labor income tax.

Specifically, for case (i.) we solve the government's problem under an additional constraint that, for any $\hat{\theta}$ and $\tilde{\theta}$ such that $y(H,\tilde{\theta}) = y(L,\hat{\theta})$ we have

$$\frac{h'\left(\frac{y(L,\hat{\theta})}{\hat{\theta}}\right)}{\hat{\theta}u'\left(c_1(L,\hat{\theta})\right)} = \frac{h'\left(\frac{y(H,\tilde{\theta})}{\tilde{\theta}}\right)}{\tilde{\theta}u'\left(c_1(H,\tilde{\theta})\right)} \tag{9}$$

In essence, regardless of education, agents with the same reported income face the same labor tax. For case (ii.), on the other hand, we require the agents with the same reported income, $y(H,\tilde{\theta}) = y(L,\hat{\theta})$, face the same savings subsidy

$$\frac{u'\left(c_1(L,\hat{\theta})\right)}{u'\left(c_2(L,\hat{\theta})\right)} = \frac{u'\left(c_1(H,\tilde{\theta})\right)}{u'\left(c_2(H,\tilde{\theta})\right)} \tag{10}$$

Solving for optimal policies under constraints (9) and (10) is non-trivial because these restrictions are contingent on allocations (declared income) rather than the underlying state variables (productivity). We overcome this challenge by designing a computational algorithm, described in Appendix E, which allows us to make the constraints conditional on allocations. Figures 4(a) and 4(b) present the optimal wedges obtained under the two restrictions, along with the education-dependent benchmark.

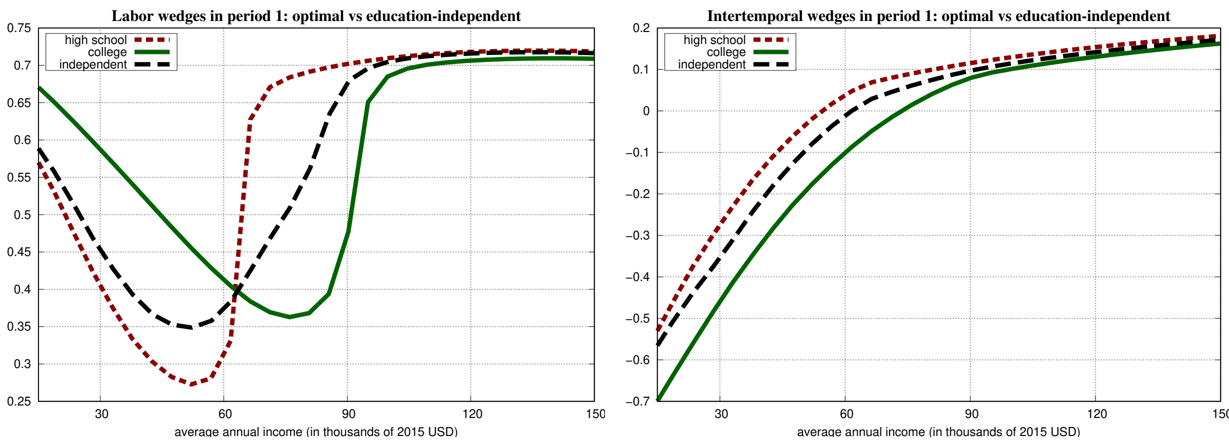

(a) Education-independent labor wedge (b) Education-independent intertemporal wedge

Figure 4: Optimal education-dependent and independent wedges

Table 4: Welfare gains over optimal education-independent policies

Education-independent:	**labor wedge**	**intertemporal wedge**
% increase in lifetime consumption	0.51	0.02

Table 4 shows welfare gains measured as a corresponding percentage increase in lifetime consumption that would result from moving from a tax system in which one of the wedges is education-independent to the optimal tax system (where all taxes depend on educational attainment). First, if labor income taxes are allowed to be conditioned on education, the resulting welfare gain is equivalent to a 0.51% increase in lifetime consumption. Second, if savings subsidies are allowed to depend on education, the corresponding gain in lifetime consumption would be 0.02%.

The welfare implications of optimal education-dependent income taxes are in the ballpark of the numbers reported by previous literature, for example Findeisen and Sachs (2016). The novel component in this paper, education-dependent savings subsidies, are shown to have much smaller welfare gains. This result is consistent with the broad literature in macroeconomics which has found that consumption smoothing yields relatively small welfare gains, given the standard parameter values. Most notably, Lucas (1987) showed that the gain from eliminating *all* post-war business cycle fluctuations in the US would be equivalent to a 0.05% increase in average consumption.

4.2 Testing Policies Without Screening

As a final step in our quantitative analysis of the model, we test whether the optimal Mirrleesian screening approach is indeed preferable quantitatively to simpler alternatives. In particular, we calculate the government's value derived under the policy that all agents get higher education, as well as one where no agents do.[10]

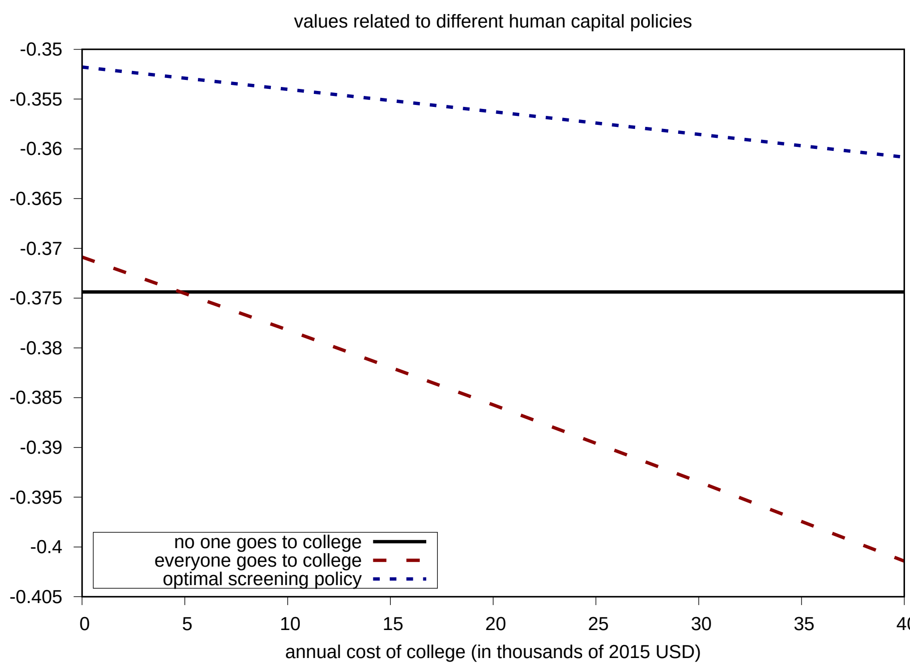

Figure 5: Comparing optimal policies to alternatives without screening

Figure 5 presents the values associated with these alternative policies, along with the optimal screening one. The values are presented as function of the annual monetary cost of higher education, ranging from zero up to 40,000 USD (the actual calibrated cost, as Table

[10]In evaluating these policies, we use the counterfactual distributions of skills presented in Figure 1, as well as counterfactual values for the discount factors $\delta_0(e)$ and $\delta_1(e)$

1 shows, is 15,700 USD). It can immediately be noticed that the optimal Mirrleesian policy dominates both alternatives at all cost values, including zero. This is due to the fact that going to college and beyond entails a significant opportunity cost of time spent in education. It is also worth noticing that for realistic levels of the calibrated cost, sending no one to college dominates the alternative of sending everyone to college.

5 Implementation

After characterizing the wedges, we are ready to discuss the implications of our findings on the design of student loans, income taxes and retirement policies. In particular, this section highlights how to decentralize policies where retirment savings is contingent on education investment, which is the main innovation of the paper.[11]

For education policies, we consider a decentralization with student loans and income-contingent repayment plans. Agents can take out a loan amount of $L(e)$, which is a function of the education investment. After agents enter the work force, the loan repayment depends on realized income. We abstract from parental financial assistance, so students solely rely on student loans in $t = 0$.

For retirement savings, we consider two ways of implementing the optimum. First, we present an implementation where the subsidy for retirement savings is both income and education contingent. Finally, we examine an implementation with social security and a retirement savings account where student loan repayments are also considered as contributions to the account. The latter captures the spirit of the recently proposed bills in the US Congress, the Retirement Parity for Student Loans Act and the Retirement Security and Savings Act, which intend to qualify student loan repayments for employer matching.

Before presenting the decentralized economy, it is important to note that we are departing from the direct revelation mechanism in which agents report their type (γ, θ). Instead, for our implementation, policies are based on the observed education investment e, income y, and savings. To do this, we first need to show that the optimal consumption from the direct revelation mechanism $\{c_0(\gamma), c_1(\gamma, \theta), c_2(\gamma, \theta)\}_{\gamma, \theta \in \Theta}$ can be expressed as a function of income y and education e. It is immediate that, by separating the agents according to their innate ability, the optimal allocations can be rewritten as a function of education instead of reported innate ability: $c_0(\gamma) = c_0(e_\gamma)$ and $c_t(\gamma, \theta) = c_t(e_\gamma, \theta)$. The next lemma shows that reported productivity can be replaced with income, so the government can implement the optimum using policies that depend on income and education.

[11] We present the three period formulation of the problem here. Details on the life-cycle model are presented in the appendix.

Lemma 2 *For any $e \in \{e_L, e_H\}$, the optimal consumption $c_1(e,\theta)$ and $c_2(e,\theta)$ are functions of $y(e,\theta)$: $c_t(e,\theta) = c_t(y(e,\theta))$ for any $t \geq 1$.*

5.1 Education-Contingent Retirement Savings Subsidy

For our first implementation, agents are offered a student loan $L(e)$ in $t = 0$. They are required to make income contingent repayments of $[1 - \tau^e(e,y)] L(e)$ in $t = 1$, where the subsidy $\tau^e(e,y)$ is a function of education expenses and income. In $t = 1$, agents also face an income tax $T(y)$ independent of education. Most importantly, agents can save s_2 in a retirement account at $t = 1$, where the savings are subsidized at a rate $\tau^s(e,y)$ which is a function of income and education investment. Furthermore, retirement savings s_2 come from after-tax funds, so the income and education dependent retirement savings account is similar to a Roth 401(k). Finally, in each period, agents can save via the risk-free bond b, which are taxed with a history-independent bond savings tax $T^k(b)$.[12]

Given the proposed policies, at $t = 1$, agents with education level e and productivity θ solve the following:

$$\max_{c_1, y, c_2, s_2, b_2} u(c_1) - h\left(\frac{y}{\theta}\right) + \beta\delta_2 u(c_2)$$

subject to

$$c_1 + s_2 + b_2 + \tilde{R}_1(e)(1 - \tau^e(e,y)) L(e) = y - T(y) + \tilde{R}_1(e) b_1 - T^k(b_2),$$

$$c_2 = R_2(1 + \tau^s(e,y)) s_2 + R_2 b_2,$$

where $\tilde{R}_1(e) = \frac{R_1(e)}{R_0(e)}$ is the gross interest rate normalized by the difference between the period lengths of $t = 0$ and $t = 1$. For example, $\tilde{R}(e_L) = 0$ since we assumed $\delta_0(e_L) = 0$ for our quantitative analysis. Let $\{c_1^*(e,\theta), y^*(e,\theta), c_2^*(e,\theta)\}$ denote the solution to the agents' problem at $t = 1$ for any $\theta \in \Theta$ and $e \in \{e_L, e_H\}$. Also, let $U_1(e,\theta)$ denote the value function for the agents' problem at $t = 1$. The agents' problem with innate ability γ at $t = 0$ is

$$\max_{c_0, e, b_1} \delta_0(e) u(c_0) + \beta\delta_1(e) \int_{\underline{\theta}}^{\overline{\theta}} [U_1(e,\theta) + (1-\beta)\delta_2 u(c_2^*(e,\theta))] f(\theta|\kappa(e,\gamma)) d\theta$$

subject to

$$c_0 + e + b_1 = L(e) - T^k(b_1) \text{ and } e \in \{e_L, e_H\}.$$

Let $P^{ss} = \left\{[L(e), \tau^e(e,y)], \tau^s(e,y), \left[T(y), T^k(b)\right]\right\}$. The following proposition states

[12] The bond savings tax helps the government deter agents from over-saving while simultaneously under-supplying labor (Werning, 2011).

that the optimum can be decentralized with an income-contingent student loans policy $(L(e), \tau^e(e,y))$ combined with an income and education dependent retirement subsidy $\tau^s(e,y)$ and tax policy $\left(T(y), T^k(b)\right)$.

Proposition 3 *The optimum can be implemented with* P^{ss}.

Figure 6 presents the optimal student loan repayment and retirement savings subsidies for the two education groups as function of income. Panel 6(a) shows that for the H-agents with income below 60,000 in present value, the repayment subsidy starts at over 80% and decreases with income. Once the Pareto tail kicks in, the trend reverts and the optimal subsidy increases before settling at around 120%. Panel 6(b) shows the savings subsidy schedules. Notice that the optimal subsidies closely mirror the intertemporal wedges τ_1^k depicted in Figure 2, where lower income levels receive more subsidies and the subsidy for college graduates is higher for virtually all income levels.

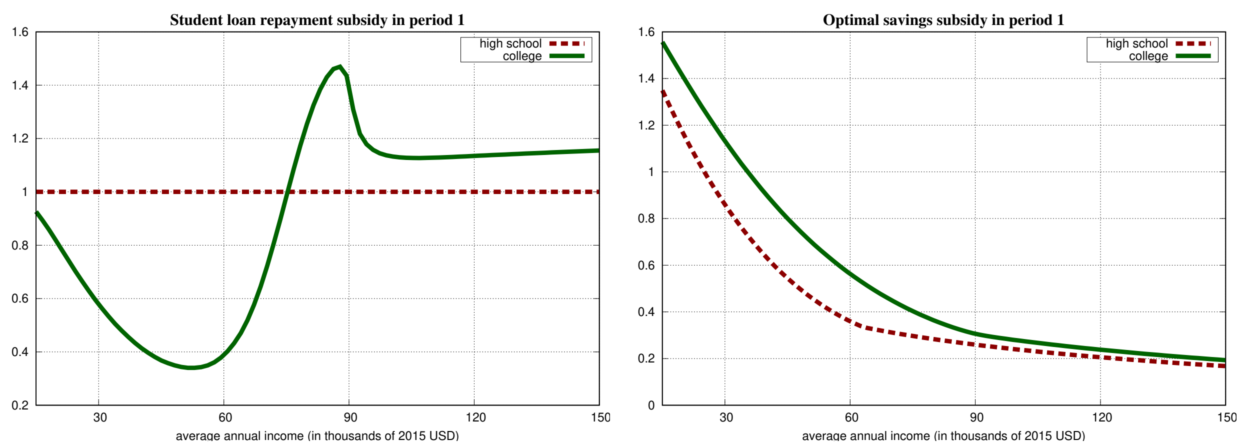

(a) Optimal student loan subsidy $\tau^e(e,y)$

(b) Optimal savings subsidy $\tau^s(e,y)$

Figure 6: Optimal education-contingent subsidies

It is worth pointing out that the optimal student loans subsidy is determined by the labor wedges. We set the income tax to match the labor wedge for L-agents, while the income contingent student loans subsidies coupled with the marginal income tax rate replicates the optimal labor wedge for H-agents. Since the optimal labor wedge for H-agents is larger with the difference growing until income $60,000$, the student loans subsidy is decreasing up to that amount. Beyond $60,000$, the difference in the labor wedges decreases initially and with the optimal labor wedge for L-agents eventually rising above the labor wedge of H-agents, which causes the student loans subsidy to increase. Also, since we appended the Pareto-tail to the top 10% of the productivity distribution for each education group, the U-shaped dip in the labor wedge for H-agents is much higher than the one for L-agents. As a result, the labor wedge for L-agents is much larger than the labor wedge for H-agents with incomes between $70,000$ and $90,000$. This drives the significant increase in student loans subsidy beyond $60,000$.

5.2 Student Loan Payment as Contribution to Retirement Savings

In this section, we consider an implementation with social security benefits and retirement savings accounts that are linked to student loan repayments. The advantage of this decentralization is that it adopts the main features of existing retirement policies. Furthermore, it demonstrates how the aforementioned retirement bills proposed in the US Congress could be used to implement the optimum.

For this implementation, similar to the current system, all agents receive an income-contingent social security benefit $a\left(y\right)$ upon retirement. The retirement savings account is defined by the contribution matching rate $\alpha \in [0,1]$ and a contribution limit $\bar{s}$. Retirement account contributions come from pre-tax income and are only lump-sum taxed T^{ra} upon withdrawal, so the retirement savings account is similar to a traditional 401(k). Furthermore, similar to current retirement savings accounts, matched contributions are not subject to the contribution limit $\bar{s}$.

The novelty of this implementation is that the amount of student loan repaid $r\left(e,y\right)$ is considered a contribution, so employers can further contribute $\alpha r\left(e,y\right)$ into the account. Let $\phi\left(s_2,r\right)$ denote the amount of assets in the retirement savings account as a function of the deposit s_2 and the student loan repayment r, so we have $\phi\left(s_2,r\right) = \left(1+\alpha\right)s_2 + \alpha r$. Similar to the previous implementation, $L\left(e\right)$ denotes the student loan taken out in $t=0$, and $T\left(y\right)$ denotes the income tax in $t=1$. Here, the student loan repayment is tax deductible and reduces income tax by $g\left(r\right)$.

Given the proposed policies, at $t=1$, agents with education investment e and productivity θ solve

$$\max_{c_1,y,c_2,s_2,b_2} u\left(c_1\right) - h\left(\frac{y}{\theta}\right) + \beta\delta_2 u\left(c_2\right)$$

subject to

$$c_1 + s_2 + b_2 + r\left(e,y\right) = y - T\left(y - s_2\right) + g\left(r\left(e,y\right)\right) + \tilde{R}_1\left(e\right)b_1 - T^k\left(b_2\right),$$

$$c_2 = a\left(y\right) + R_2\phi\left(s_2, r\left(e,y\right)\right) + R_2 b_2 - \mathbf{1}_{\phi>0}T^{ra},$$

$$0 \le s_2 \le \bar{s},$$

where $\mathbf{1}_{\phi>0}$ is an indicator function with $\mathbf{1}_{\phi>0} = 1$ if and only if there are assets in the account, otherwise $\mathbf{1}_{\phi>0} = 0$. Adopting the notation introduced in the previous implementation, the agents' problem with innate ability γ at $t = 0$ is

$$\max_{c_0,e,b_1} \delta_0(e)\, u(c_0) + \beta\delta_1(e) \int_{\underline{\theta}}^{\overline{\theta}} \left[U_1(e,\theta) + (1-\beta)\,\delta_2 u\left(c_2^*(e,\theta)\right)\right] f\left(\theta|\kappa(e,\gamma)\right) d\theta$$

subject to

$$c_0 + e + b_1 = L(e) - T^k(b_1) \text{ and } e \in \{e_L, e_H\}.$$

Let $P^{ra} = \left\{ [L(e), r(e,y)], a(y), [\alpha, \bar{s}], \left[T(y), T^k(b), T^{ra}, g(r)\right] \right\}$ denote the policy instruments for the proposed implementation. The following proposition shows that it is possible to decentralize the optimum using P^{ra}.

Proposition 4 *The optimum can be implemented through P^{ra} where student loan repayments are considered contributions to the retirement savings account.*

Figure 7 presents the student loan repayment schedule in our second implementation. The solid green line shows the face value of the repayment which starts high and then decreases initially, allowing college-graduates to accumulate additional (and decreasing in income) contributions in their 401(k) plans. In order to maintain the optimal income-contingency in loan repayments, these agents are offered the tax deduction, which makes their repayment schedule increase in income, as represented by the dashed red line. Notice that once the Pareto tail for high school graduates kicks in, the two trends revert.

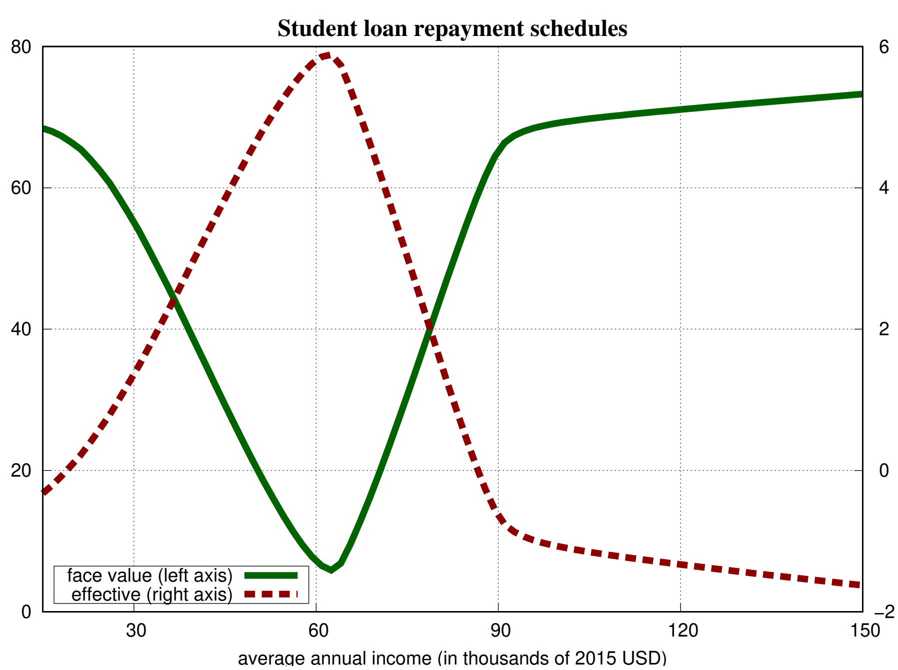

Figure 7: Optimal student loan repayment schedule $r(e, y)$

The effective student loan repayment schedule, which is determined by the difference between the repayment schedule and the tax deduction, is largely shaped by the optimal labor distortion as explained in the previous implementation. What is significant is that this decentralization uses student loan repayments as a retirement savings vehicle for college-educated agents. More specifically, our implementation constructs the social security benefits to match the optimal retirement consumption of high school graduates. The student loan repayments at its face value are designed to supplement the social security benefits so college-educated agents consume the optimum during retirement.

At the heart of the implementation is the idea that college graduates can save for retirement while paying off their student loans. The contribution matching rate α that arises in our proposed implementation amounts to 2.00%, which is in the ballpark of the actual rate used by the IRS ruling from May 2018.

6 Discussion

6.1 Heterogeneous Present Bias

We extend our results to an environment with heterogeneous present bias by assuming that agents with innate ability γ have present bias β_γ, where $1 \geq \beta_H > \beta_L$. The perfect correlation between innate ability and the degree of present bias allows us to bypass the multi-dimensional screening problem, simplifying the analysis.[13] Proposition 5 characterizes the distortions and shows that the link between retirement savings and education investment persists when the degree of present bias is heterogeneous.

Proposition 5 *The constrained efficient allocation with heterogeneous present bias satisfies*

i. the inverse Euler equations (3), (4) and for any $\theta \in \Theta$,

$$\frac{1}{\beta_H u'\left(c_2\left(H,\theta\right)\right)} = \frac{1}{u'\left(c_1\left(H,\theta\right)\right)} + \left(\frac{1-\beta_H}{\beta_H}\right)\left(\frac{\pi_H + \beta_H\mu}{\pi_H + \mu}\right)\frac{1}{u'\left(c_0\left(H\right)\right)},$$

$$\frac{1}{\beta_L u'\left(c_2\left(L,\theta\right)\right)} = \frac{1}{u'\left(c_1\left(L,\theta\right)\right)} + \left(\frac{1-\beta_L}{\beta_L}\right)\left(\frac{\pi_L - \beta_H\mu\left(\frac{f\left(\theta|\kappa_{L,H}\right)}{f\left(\theta|\kappa_L\right)}\right)}{\pi_L - \mu}\right)\frac{1}{u'\left(c_0\left(L\right)\right)},$$

where $\mu = \left[u'\left(c_0\left(L\right)\right) - u'\left(c_0\left(H\right)\right)\right]\left[\frac{u'(c_0(L))}{\pi_L} + \frac{u'(c_0(H))}{\pi_H}\right]^{-1}$.

[13] This setup is related to Golosov et al. (2013). They consider an environment with time-consistent agents where productivity is perfectly correlated with the long-run discount factor.

ii. the labor wedge for H-agents satisfies (7) and for L-agents:

$$\frac{\tau^w(L,\theta)}{1-\tau^w(L,\theta)} = A_L(\theta)\,B_L(\theta)$$
$$\times\left[C_L(\theta) - \left(\frac{1-F(\theta|\kappa_{L,H})}{1-F(\theta|\kappa_L)}\right)D_L(\theta) + \frac{\frac{\beta_L}{1-\beta_L}g(\theta|\kappa_L)}{\frac{\beta_H}{1-\beta_H}g(\theta|\kappa_{L,H})}E_L(\theta)\right],$$

where $E_\gamma(\theta) = \left[\frac{u'(c_1(\gamma,\theta))}{\beta_\gamma u'(c_2(\gamma,\theta))} - 1\right] - \left(\frac{1-\beta_\gamma}{\beta_\gamma}\right)\frac{u'(c_1(\gamma,\theta))}{\phi}$ *and* $\frac{1}{\phi} = \mathbb{E}_\gamma\left[\mathbb{E}_\theta\left(\frac{1}{u'(c_1(\gamma,\theta))}\middle|\gamma\right)\right]$.

Though the economic forces determining the wedges for H-agents remain unchanged, Proposition 5 shows us how the optimal policy leverages the difference in β for the L-agents' wedges. For the intertemporal wedge $\tau_1^k(L,\theta)$, recall that the optimal policy recommends front-loading consumption for high-income L-agents. Here, this front-loading could be more perverse. It takes advantage of the fact that H-agents value retirement consumption more than L-agents, so a restriction on retirement savings further deters downward deviations by H-agents. This logic is similar to the key finding in Golosov et al. (2013) which shows that discouraging the consumption of a good preferred by high types among low types raises welfare. The labor wedge for L-agents $\tau^w(L,\theta)$ also differs from the case with homogeneous β. Recall that, in the previous sections, the present-bias component $E_L(\theta)$ for L-agents is enhanced by the differences in the factual and counterfactual distributions to deter H-agents from mimicking. Here, the labor distortion for L-agents coming from the present-bias component is weakened. This is because H-agents are less tempted to mimic L-agents due to the larger intertemporal distortion, which relieves the labor distortions stemming from present bias.

A special case is when H-agents are time consistent while only L-agents are present-biased ($\beta_H = 1 > \beta_L$). From Proposition 5, the H-agents' wedges share the same properties as the wedges for time-consistent agents. Also, the present-bias component $E_L(\theta)$ no longer influences the labor wedge of L-agents. Instead, the optimal policy takes advantage of present-biased L-agents entirely through the intertemporal distortion in retirement savings $\tau_1^k(L,\theta)$, which is worsened with time-consistent H-agents. This implies that even though linking student loan payments to retirement savings is not essential for time-consistent college graduates, education-dependent savings policies are still optimal. We believe this case is a theoretical curiosity, since empirical studies have demonstrated pervasive present-biased behavior among college students (Ariely and Wertenbroch, 2002; Steel, 2007).

6.2 Non-Sophistication

The paper has thus far assumed that the agents are sophisticated—fully aware of their present bias. Sophisticated agents have a demand for commitment to prevent their future-selves from under-saving. The optimal policy in this paper takes advantage of this demand by assisting college graduates with their retirement savings to incentivize them to go to college in the first place. We may also want to investigate the optimal education and retirement savings policies for non-sophisticated agents.

For non-sophisticated agents, the government can use *off-path policies* to take advantage of their incorrect beliefs. Following Yu (2019a), the government can introduce a menu of savings options in $t = 1$. One of the options in the menu will be selected by the agents on the equilibrium path while the other option is a decoy, the off-path policy. The decoy option features a relatively backloaded consumption path—high retirement consumption but lower working period consumption—compared to the on-path option. At $t = 0$, the non-sophisticated agents underestimate their present bias and thus overestimate the value of retirement consumption to their future-selves. As a result, they mispredict that they will select the decoy option in $t = 1$. In reality, their future-selves prefer the more front-loaded on-path option instead. Therefore, the government can exploit this incorrect belief by promising college graduates with high retirement benefits—which never needs to be implemented on the equilibrium path—to induce investment in higher education. In other words, the inclusion of a decoy option in the menu can relax the ex-ante incentive constraint. In fact, Yu (2019a) showed that if the consumption utility is unbounded above and below, then the ex-ante incentive constraints can be fully relaxed. More details are provided in Appendix F.

Off-path policies are powerful, but the optimal policy should still feature the interdependence between retirement savings and education investment. This is due to two reasons. First, the economy is most likely populated by agents with heterogeneous levels of sophistication. A menu with decoy options would not be able to fool sophisticated agents, so it is optimal for the government to rely on the present paper's policies for relatively more sophisticated agents. Future work should explore the optimal combination of these two policies. Second, governments may object to the use of off-path policies to mislead agents due to moral or reputational reasons. In this case, it is optimal to implement the education-dependent retirement savings even for non-sophisticated agents. As long as agents have some demand for commitment, albeit lower than what is optimal, the government can still take advantage of this demand by making retirement savings contingent on education investment. However, this interdependence disappears when agents are naïve—fully unaware of their present bias. This is because naïve agents believe their future-selves to be time-consistent, so this paper's retirement policies would not encourage them to increase investment in education.

7 Conclusion

This paper formulates the optimal education and retirement policies in a dynamic Mirrlees model with present-biased agents. A novel contribution of this paper is to show that the optimal retirement savings policy depends on education. More specifically, we show how linking student loan repayments to retirement savings along with some qualitative changes to existing policies can implement the optimum. We estimate the welfare gains from these policies to be significant. Also, the inverse Euler equation does not hold with present-biased agents, but the labor wedge is quantitatively similar to the case with time-consistent agents.

References

Anderberg, Dan, "Optimal Policy and the Risk Properties of Human Capital Reconsidered," *Journal of Public Economics*, 2009, *93*, 1017–1026.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 2001, *15* (3), 47–68.

Ariely, Dan and Klaus Wertenbroch, "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," *Psychological Science*, 2002, *13* (3), 219–224.

Bohacek, Radim and Marek Kapicka, "Optimal Human Capital Policies," *Journal of Monetary Economics*, 2008, *55*, 1–16.

Bovenberg, A. Lans and Bas Jacobs, "Redistribution and Education Subsidies are Siamese Twins," *Journal of Public Economics*, 2005, *89*, 2005–2035.

Cadena, Brian and Benjamin Keys, "Human Capital and the Lifetime Costs of Impatience," *American Economic Journal: Economic Policy*, 2015, *7* (3), 126–153.

Craig, Ashley, "Optimal Income Taxation with Spillovers from Employer Learning," *Working Paper*, 2019.

Cunha, Flavio and James Heckman, "Identifying and Estimating the Distributions of Ex Post and Ex Ante Returns to Schooling," *Labour Economics*, 2007, *14* (6), 870–893.

Diamond, Peter, "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, 1998, *96* (3), 694–719.

Farhi, Emmanuel and Xavier Gabaix, "Optimal Taxation with Behavioral Agents," *American Economic Review*, 2019, *Forthcoming.*

Findeisen, Sebastian and Dominik Sachs, "Education and Optimal Dynamic Taxation: The Role of Income-Contingent Student Loans," *Journal of Public Economics*, 2016, *138*, 1–21.

Gabaix, Xavier, "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of*

Economics, 2014, *129* (4), 1661–1710.

Galperti, Simone, "Commitment, Flexibility, and Optimal Screening of Time Inconsistency," *Econometrica*, 2015, *83* (4), 1425–1465.

Gary-Bobo, Robert and Alain Trannoy, ""Optimal Student Loans and Graduate Tax under Moral Hazard and Adverse Selection," *The RAND Journal of Economics*, 2015, *46* (3), 546–576.

Golosov, Mikhail, Maxim Troshkin, Aleh Tsyvinski, and Matthew Weinzierl, "Preference Heterogeneity and Optimal Capital Income Taxation," *Journal of Public Economics*, 2013, *97*, 160–175.

_ , _ , and _ , "Redistribution and Social Insurance," *American Economic Review*, 2016, *106* (2), 359–386.

_ , Narayana Kocherlakota, and Aleh Tsyvinski, "Optimal Indirect and Capital Taxation," *Review of Economic Studies*, 2003, *70* (3), 569–587.

Grochulski, Borys and Tomasz Piskorski, "Risky Human Capital and Deferred Capital Income Taxation," *Journal of Economic Theory*, 2010, *145*, 908–943.

Halac, Marina and Pierre Yared, "Fiscal Rules and Discretion Under Persistent Shocks," *Econometrica*, 2014, *82* (5), 1557–1614.

Heathcote, Jonathan and Hitoshi Tsujiyama, "Optimal Income Taxation: Mirrlees Meets Ramsey," *Working Paper*, 2017.

_ , Kjetil Storesletten, and Giovanni Violante, "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, 2017, *132* (4), 1693–1754.

Kapicka, Marek, "Optimal Mirrleesean Taxation in a Ben-Porath Economy," *American Economic Journal: Macroeconomics*, 2015, *7* (2), 219–248.

_ and Julian Neira, "Optimal Taxation with Risky Human Capital," *American Economic Journal: Macroeconomics*, 2019, *11* (4), 271–309.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

_ , Peter Maxted, Andrea Repetto, and Jeremy Tobacman, "Estimating Discount Functions with Consumption Choices over the Lifecycle," *Working Paper*, 2017.

Lockwood, Benjamin, "Optimal Income Taxation with Present Bias," *Working Paper*, 2018.

Lucas, Robert E. Jr., "Models of business cycles," *New York: Basil Blackwell*, 1987.

Makris, Miltiadis and Alessandro Pavan, "Taxation Under Learning-by-Doing," *Working Paper*, 2019.

Meara, Ellen, Seth Richards, and David Cutler, "The Gap Gets Bigger: Changes in Mortality and Life expectancy by Education, 1981−2000," *Health Affairs*, 2008, *27* (2), 350–360.

Moser, Christian and Pedro Olea de Souza e Silva, "Paternalism vs Redistribution: Designing Retirement Savings Policies with Behavioral Agents," *Working Paper*, 2019.

Myerson, Roger, "Multistage Games with Communication," *Econometrica*, 1986, *54* (2).

Nakajima, Makoto, "Rising indebtedness and temptation: A welfare analysis," *Quantitative Economics*, 2012, *3*, 257–288.

Nigai, Sergey, "A Tale of Two Tails: Productivity Distribution and the Gains from Trade," *Journal of International Economics*, 2017, *104*, 44–62.

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 1999, *89* (1), 103–124.

_ **and** _ , "Choice and Procrastination," *Quarterly Journal of Economics*, 2001, *116*, 121–160.

Saez, Emmanuel, "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 2001, *68*, 205–229.

Stantcheva, Stefanie, "Optimal Taxation and Human Capital Policies over the Life Cycle," *Journal of Political Economy*, 2017, *125* (6), 1931–1990.

Steel, Piers, "The Nature of Procrastination: A Meta-Analytic and Theoretical Review of Quintessential Self-Regulatory Failure," *Psychological Bulletin*, 2007, *133* (1), 65–94.

Werning, Ivan, "Nonlinear Capital Taxation," *Working Paper*, 2011.

Yu, Pei Cheng, "Optimal Retirement Policies with Present-Biased Agents," *Working Paper*, 2019.

_ , "Seemingly Exploitative Contracts," *Working Paper*, 2019.

Appendices (for online publication)

A Derivation of the Theoretical Results

A.1 The Optimization Problem

Given Lemma 1, the relaxed optimal tax problem is

$$\max_{P} \sum_{\gamma} \pi_{\gamma} \left[\delta_0 \left(e_{\gamma}\right) u \left(c_0 \left(\gamma\right)\right) + \delta_1 \left(e_{\gamma}\right) \int_{\underline{\theta}}^{\overline{\theta}} \left[U_1 \left(\gamma, \theta\right) + \left(1 - \beta\right) \delta_2 u \left(c_2 \left(\gamma, \theta\right)\right)\right] f \left(\theta | \kappa_{\gamma}\right) d\theta \right]$$

subject to

$$U_1 \left(\gamma, \theta\right) = u \left(c_1 \left(\gamma, \theta\right)\right) - h \left(\frac{y \left(\gamma, \theta\right)}{\theta}\right) + \beta \delta_2 u \left(c_2 \left(\gamma, \theta\right)\right), \tag{11}$$

$$\frac{\partial U_1 \left(\gamma, \theta\right)}{\partial \theta} = \frac{y \left(\gamma, \theta\right)}{\theta^2} h' \left(\frac{y \left(\gamma, \theta\right)}{\theta}\right), \tag{12}$$

$$\begin{aligned} &\delta_0 \left(e_H\right) u \left(c_0 \left(H\right)\right) + \beta \delta_1 \left(e_H\right) \int_{\underline{\theta}}^{\overline{\theta}} \left[U_1 \left(H, \theta\right) + \left(1 - \beta\right) \delta_2 u \left(c_2 \left(H, \theta\right)\right)\right] f \left(\theta | \kappa_H\right) d\theta \\ &\qquad \geq \delta_0 \left(e_L\right) u \left(c_0 \left(L\right)\right) + \beta \delta_1 \left(e_L\right) \int_{\underline{\theta}}^{\overline{\theta}} \left[U_1 \left(L, \theta\right) + \left(1 - \beta\right) \delta_2 u \left(c_2 \left(L, \theta\right)\right)\right] f \left(\theta | \kappa_{L,H}\right) d\theta, \end{aligned}$$

and the resource constraint. As is standard, we ignore the monotonicity constraint—$y \left(\gamma, \theta\right)$ is non-decreasing in θ—and check it later. Also, we assume that the ex-ante incentive constraint for H-agents binds and show that the incentive constraint for L-agents holds.

Let $\left(\lambda_{\gamma} \left(\theta\right), \xi_{\gamma} \left(\theta\right), \mu, \phi\right)$ be the multipliers on (11), (12), ex-ante incentive compatibility and resource constraint respectively. Using standard Hamiltonian techniques, we derive the following necessary conditions for optimality

$$\left(1 + \frac{\mu}{\pi_H}\right) u' \left(c_0 \left(H\right)\right) = \left(1 - \frac{\mu}{\pi_L}\right) u' \left(c_0 \left(L\right)\right) = \phi,$$

$$\left(\pi_H + \beta \mu\right) \delta_1 \left(e_H\right) f \left(\theta | \kappa_H\right) - \xi_H' \left(\theta\right) = \lambda_H \left(\theta\right),$$

$$\left[\pi_L - \beta \mu \left(\frac{f \left(\theta | \kappa_{L,H}\right)}{f \left(\theta | \kappa_L\right)}\right)\right] \delta_1 \left(e_L\right) f \left(\theta | \kappa_L\right) - \xi_L' \left(\theta\right) = \lambda_L \left(\theta\right),$$

$$\left(1 - \beta\right) \left(\pi_H + \beta \mu\right) \delta_1 \left(e_H\right) f \left(\theta | \kappa_H\right) + \beta \lambda_H \left(\theta\right) = \frac{\phi \pi_H \delta_1 \left(e_H\right) f \left(\theta | \kappa_H\right)}{u' \left(c_2 \left(H, \theta\right)\right)},$$

$$\left(1 - \beta\right) \left[\pi_L - \beta \mu \left(\frac{f \left(\theta | \kappa_{L,H}\right)}{f \left(\theta | \kappa_L\right)}\right)\right] \delta_1 \left(e_L\right) f \left(\theta | \kappa_L\right) + \beta \lambda_L \left(\theta\right) = \frac{\phi \pi_L \delta_1 \left(e_L\right) f \left(\theta | \kappa_L\right)}{u' \left(c_2 \left(L, \theta\right)\right)},$$

and for all γ, the boundary conditions hold: $\xi_\gamma\left(\underline{\theta}\right)=\xi_\gamma\left(\overline{\theta}\right)=0$, and

$$\lambda_\gamma\left(\theta\right)u'\left(c_1\left(\gamma,\theta\right)\right)=\phi\pi_\gamma\delta_1\left(e_\gamma\right)f\left(\theta|\kappa_\gamma\right),$$

$$\lambda_\gamma\left(\theta\right)\frac{1}{\theta}h'\left(\frac{y\left(\gamma,\theta\right)}{\theta}\right)+\xi_\gamma\left(\theta\right)\left[\frac{1}{\theta^2}h'\left(\frac{y\left(\gamma,\theta\right)}{\theta}\right)+\frac{y\left(\gamma,\theta\right)}{\theta^3}h''\left(\frac{y\left(\gamma,\theta\right)}{\theta}\right)\right]=\phi\pi_\gamma\delta_1\left(e_\gamma\right)f\left(\theta|\kappa_\gamma\right).$$

Below, we show that the theoretical results follow from these conditions.

A.2 Proofs

Proof of Proposition 1: Conditions (5) and (6) and μ follow from the first order conditions. The inverse Euler equations (3) and (4) are derived using the perturbation argument.

Let $P=\left\{c_0\left(\gamma\right),\left[c_t\left(\gamma,\theta\right),y\left(\gamma,\theta\right)\right]_{t>0,\theta\in\Theta}\right\}_\gamma$ be the allocation that solves the constrained efficient planning problem. We first derive (4) by considering a small increase in $c_2\left(\gamma,\theta\right)$ across θ for a fixed γ. That is, for all θ, define $u\left(\tilde{c}_2\left(\gamma,\theta\right)\right)=u\left(c_2\left(\gamma,\theta\right)\right)+\Delta$ for some small Δ. We simultaneously decrease $c_1\left(\gamma,\theta\right)$ for all θ such that $u\left(\tilde{c}_1\left(\gamma,\theta\right)\right)=u\left(c_1\left(\gamma,\theta\right)\right)-\delta_2\Delta$. Such perturbations do not affect the objective function, the ex-ante incentive compatibility and the ex-post incentive compatibility. It only affects the resource constraint. Note that the perturbation must be the same for all θ or else it may violate ex-post incentive compatibility, which is not the case if $\beta=1$. If P is optimal, then it must be that $\Delta=0$ minimizes the resource used, i.e.,

$$0=\arg\min_{\Delta}\int_\Theta\left[-u^{-1}\left[u\left(c_1\left(\gamma,\theta\right)\right)-\delta_2\Delta\right]-\frac{1}{R_2}u^{-1}\left[u\left(c_2\left(\gamma,\theta\right)\right)+\Delta\right]\right]f\left(\theta|\kappa_\gamma\right)d\theta.$$

Evaluating the first order condition of this problem at $\Delta=0$ yields (4).

Similarly, to derive (3), we consider a small decrease in $c_1\left(\gamma,\theta\right)$ for all θ and γ such that $u\left(\tilde{c}_1\left(\gamma,\theta\right)\right)=u\left(c_1\left(\gamma,\theta\right)\right)-\frac{1}{\delta_1(e_\gamma)}\Delta$ for some small Δ. We simultaneously increase $c_0\left(\gamma\right)$ for all γ such that $u\left(\tilde{c}_0\left(\gamma\right)\right)=u\left(c_0\left(\gamma\right)\right)+\frac{1}{\delta_0(e_\gamma)}\Delta$. Since it is perturbed for all θ, the ex-post incentive compatibility constraint is not affected. Also, notice that the ex-ante incentive compatibility constraint and objective function are not affected, but the resource constraint changes. Crucially, the perturbation must be the same for all γ or it may violate ex-ante incentive compatibility, which is not the case if $\beta=1$. If P is optimal, then $\Delta=0$ solves,

$$\min_{\Delta}\sum_\gamma\pi_\gamma\left\{\frac{-u^{-1}\left[u\left(c_0\left(\gamma\right)\right)+\frac{\Delta}{\delta_0(e_\gamma)}\right]}{R_0\left(e_\gamma\right)}-\frac{1}{R_1\left(e_\gamma\right)}\int_\Theta u^{-1}\left[u\left(c_1\left(\gamma,\theta\right)\right)-\frac{\Delta}{\delta_1\left(e_\gamma\right)}\right]f\left(\theta|\kappa_\gamma\right)d\theta\right\}.$$

Evaluating the first order condition of this problem at $\Delta=0$ yields (3). ■

Proof of Proposition 2: From the first order conditions, we have

$$\xi_H(\theta) = \int_\theta^{\bar{\theta}} \left[\lambda_H(x) - (\pi_H + \beta\mu)\,\delta_1(e_H)\, f(x|\kappa_H)\right] dx,$$

$$\xi_L(\theta) = \int_\theta^{\bar{\theta}} \left[\lambda_L(x) - \left[\pi_L f(x|\kappa_L) - \beta\mu f(x|\kappa_{L,H})\right]\delta_1(e_L)\right] dx.$$

First, we derive (7). Since $\lambda_\gamma(\theta) = \frac{\phi\pi_\gamma\delta_1(e_\gamma)f(\theta|\kappa_\gamma)}{u'(c_1(\gamma,\theta))}$, we rewrite the first order condition on $y(H,\theta)$ as

$$\phi\pi_H\delta_1(e_H)\, f(\theta|\kappa_H)\left[1 - \frac{\frac{1}{\theta}h'\left(\frac{y(H,\theta)}{\theta}\right)}{u'(c_1(H,\theta))}\right]$$
$$= \left[\frac{1}{\theta^2}h'\left(\frac{y(H,\theta)}{\theta}\right) + \frac{y(H,\theta)}{\theta^3}h''\left(\frac{y(H,\theta)}{\theta}\right)\right]\int_\theta^{\bar{\theta}} \left[\lambda_H(x) - (\pi_H + \beta\mu)\,\delta_1(e_H)\, f(x|\kappa_H)\right] dx.$$

Let $A_\gamma(\theta) = \frac{1-F(\theta|\kappa_\gamma)}{\theta f(\theta|\kappa_\gamma)}$ and $B_\gamma(\theta) = 1 + \frac{\frac{y(\gamma,\theta)}{\theta}h''\left(\frac{y(\gamma,\theta)}{\theta}\right)}{h'\left(\frac{y(\gamma,\theta)}{\theta}\right)}$, then dividing both sides by $\frac{1}{\theta}h'\left(\frac{y(H,\theta)}{\theta}\right)\phi\pi_H\delta_1(e_H)\, f(\theta|\kappa_H)$ yields

$$\frac{1}{\frac{1}{\theta}h'\left(\frac{y(H,\theta)}{\theta}\right)} - \frac{1}{u'(c_1(H,\theta))}$$
$$= A_H(\theta)\, B_H(\theta)\int_\theta^{\bar{\theta}}\left[\frac{\lambda_H(x)}{\phi\pi_H\delta_1(e_H)\, f(x|\kappa_H)} - \frac{\pi_H+\beta\mu}{\phi\pi_H}\right]\frac{f(x|\kappa_H)}{1-F(\theta|\kappa_H)}dx.$$

By definition $\frac{1}{\theta}h'\left(\frac{y(\gamma,\theta)}{\theta}\right) = (1-\tau^w(\gamma,\theta))\, u'(c_1(\gamma,\theta))$ and from the first order condition, $\frac{\lambda_\gamma(x)}{\phi\pi_\gamma\delta_1(e_\gamma)f(x|\kappa_\gamma)} = \frac{1}{u'(c_1(\gamma,x))}$, so we have

$$\frac{1}{u'(c_1(H,\theta))}\left(\frac{\tau^w(H,\theta)}{1-\tau^w(H,\theta)}\right) = A_H(\theta)\, B_H(\theta)\left[\int_\theta^{\bar{\theta}}\frac{1}{u'(c_1(H,x))}\frac{f(x|\kappa_H)}{1-F(\theta|\kappa_H)}dx - \frac{\pi_H+\beta\mu}{\phi\pi_H}\right].$$

Observe that $\frac{\beta\mu}{\phi\pi_H} = \frac{\mu}{\phi\pi_H} - \frac{(1-\beta)\mu}{\phi\pi_H}$, then by the first order conditions, we can substitute in $\frac{\mu}{\phi\pi_H} = \frac{1}{u'(c_0(H))} - \frac{1}{\phi}$ and $\frac{(1-\beta)\mu}{\phi\pi_H} = \frac{1}{\beta u'(c_2(H,\theta))} - \frac{1}{u'(c_1(H,\theta))} - \left(\frac{1-\beta}{\beta}\right)\frac{1}{\phi}$. Define $C_\gamma(\theta) = \int_\theta^{\bar{\theta}}\frac{u'(c_1(\gamma,\theta))}{u'(c_1(\gamma,x))}\left[1 - \frac{u'(c_1(\gamma,x))}{\phi}\right]\frac{f(x|\kappa_\gamma)}{1-F(\theta|\kappa_\gamma)}dx$, $D_\gamma(\theta) = u'(c_1(\gamma,\theta))\left[\frac{1}{u'(c_0(\gamma))} - \frac{1}{\phi}\right]$, and $E_\gamma(\theta) = \left[\frac{u'(c_1(\gamma,\theta))}{\beta u'(c_2(\gamma,\theta))} - 1\right] - \left(\frac{1-\beta}{\beta}\right)\frac{u'(c_1(\gamma,\theta))}{\phi}$ then multiplying both sides by $u'(c_1(H,\theta))$ to yield (7).

Using a similar process as above, we have the following expression for $\gamma = L$

$$\frac{\tau^w(L,\theta)}{1-\tau^w(L,\theta)} = A_L(\theta)\,B_L(\theta)\left[C_L(\theta) - \int_{\theta}^{\overline{\theta}}\left(\frac{\pi_L f(x|\kappa_L) - \beta\mu f(x|\kappa_{L,H})}{\phi\pi_L}\right)\frac{u'(c_1(L,\theta))}{1-F(\theta|\kappa_L)}dx\right],$$

and integrating gives us (8). Furthermore, from the first order condition for c_0, we have $\phi = \left[\frac{\pi_H}{u'(c_0(H))} + \frac{\pi_L}{u'(c_0(L))}\right]^{-1}$, combining it with (3) yields $\phi = \left\{\mathbb{E}_\gamma\left[\mathbb{E}_\theta\left(\frac{1}{u'(c_1(\gamma,\theta))}\middle|\gamma\right)\right]\right\}^{-1}$. ■

Proof of Lemma 2: For a fixed γ, suppose there exists $\tilde{\theta}$ and $\hat{\theta}$ such that $y\left(\gamma,\tilde{\theta}\right) = y\left(\gamma,\hat{\theta}\right)$. Let $\Phi(\gamma,\theta) = u(c_1(\gamma,\theta)) + \beta\delta_2 u(c_2(\gamma,\theta))$. There are two cases to consider. First, suppose $\Phi\left(\gamma,\tilde{\theta}\right) \neq \Phi\left(\gamma,\hat{\theta}\right)$, then clearly the allocations are not incentive compatible. Next, suppose $\Phi\left(\gamma,\tilde{\theta}\right) = \Phi\left(\gamma,\hat{\theta}\right)$, and without loss of generality $c_1\left(\gamma,\tilde{\theta}\right) > c_1\left(\gamma,\hat{\theta}\right)$ and $c_2\left(\gamma,\tilde{\theta}\right) < c_2\left(\gamma,\hat{\theta}\right)$. Let $\tilde{\pi}$ and $\hat{\pi}$ denote the measure of $\left(\gamma,\tilde{\theta}\right)$ and $\left(\gamma,\hat{\theta}\right)$ agents. Let $\bar{u}_t = \frac{1}{\tilde{\pi}+\hat{\pi}}\left[\tilde{\pi}u\left(c_t\left(\gamma,\tilde{\theta}\right)\right) + \hat{\pi}u\left(c_t\left(\gamma,\hat{\theta}\right)\right)\right]$. By assigning these agents the average utility, the total welfare is unchanged and incentive compatibility is preserved. However, since u is strictly concave, the consumption level that gives $\bar{u}_1$ and $\bar{u}_2$ relaxes the resource constraint. This means that it is not optimal for $c_1\left(\gamma,\tilde{\theta}\right) > c_1\left(\gamma,\hat{\theta}\right)$ and $c_2\left(\gamma,\tilde{\theta}\right) < c_2\left(\gamma,\hat{\theta}\right)$ with $\Phi\left(\gamma,\tilde{\theta}\right) = \Phi\left(\gamma,\hat{\theta}\right)$. In other words, the consumption paths are equivalent for agents of the same level of income. ■

Proof of Proposition 3: First, following Werning (2011), we construct bond savings tax $T^k(b)$ such that agents never purchase bonds. To see how, consider the government assigning the optimal allocation from the direct revelation mechanism given past and current reports, while agents are allowed to purchase any desired amount of bonds. Define a fictitious tax $T_1^k(b_2,\tilde{r},\theta)$ paid in $t=1$ for each productivity realization θ, current bond level b_1, past report $\tilde{r}_\gamma$, current report $\tilde{r}_\theta$, and bond savings b_2, where $\tilde{r} = (\tilde{r}_\gamma, \tilde{r}_\theta)$. The tax $T_1^k(b_2,\tilde{r},\theta)$ is set such that

$$u\left(c_1(\tilde{r}) + \tilde{R}_1(e(\tilde{r}_\gamma))\,b_1 - b_2 - T_1^k(b_2,\tilde{r},\theta)\right) - h\left(\frac{y(\tilde{r})}{\theta}\right) + \beta\delta_2 u(c_2(\tilde{r}) + R_2 b_2)$$
$$= u(c_1(\gamma,\theta)) - h\left(\frac{y(\gamma,\theta)}{\theta}\right) + \beta\delta_2 u(c_2(\gamma,\theta)).$$

Next, by taking the supremum over all $\theta \in \Theta$, we obtain a bond savings tax $T_1^k(b_2,\tilde{r}) = \sup_{\theta\in\Theta} T_1^k(b_2,\tilde{r},\theta)$ that is independent of productivity. Before we derive the bond savings tax in $t=0$, let

$$V\left(b_1,\tilde{r}_\gamma,\theta\right)=u\left(c_1\left(\tilde{r}_\gamma,\hat{r}_\theta\right)+\tilde{R}_1\left(e\left(\tilde{r}_\gamma\right)\right)b_1-\hat{b}_2-T_1^k\left(\hat{b}_2,\tilde{r}_\gamma,\hat{r}_\theta\right)\right)\\ -h\left(\frac{y\left(\tilde{r}_\gamma,\hat{r}_\theta\right)}{\theta}\right)+\delta_2 u\left(c_2\left(\tilde{r}_\gamma,\hat{r}_\theta\right)+R_2\hat{b}_2\right),$$

where

$$\left(\hat{r}_\theta,\hat{b}_2\right)\in\arg\max_{\tilde{r}_\theta,b_2}\left\{u\left(c_1\left(\tilde{r}\right)+\tilde{R}_1\left(e\left(\tilde{r}_\gamma\right)\right)b_1-b_2-T_1^k\left(b_2,\tilde{r}\right)\right)\right.\\ \left.-h\left(\frac{y\left(\tilde{r}\right)}{\theta}\right)+\beta\delta_2 u\left(c_2\left(\tilde{r}\right)+R_2 b_2\right)\right\}.$$

Next, define $T_0^k\left(b_1,\tilde{r}_\gamma\right)=\sup_{\gamma\in\{H,L\}}T_0^k\left(b_1,\tilde{r}_\gamma,\gamma\right)$ with $T_0^k\left(b_1,\tilde{r}_\gamma,\gamma\right)$ chosen such that

$$\delta_0\left(e_\gamma\right)u\left(c_0\left(\tilde{r}_\gamma\right)-b_1-T_0^k\left(b_2,\tilde{r}_\gamma,\gamma\right)\right)+\beta\delta_1\left(e\left(\tilde{r}_\gamma\right)\right)\mathbb{E}\left[V\left(b_1,\tilde{r}_\gamma,\theta\right)|\gamma\right]=\delta_0\left(e_\gamma\right)u\left(c_0\left(\gamma\right)\right)\\ +\beta\delta_1\left(e_\gamma\right)\int_{\underline{\theta}}^{\overline{\theta}}\left[u\left(c_1\left(\gamma,\theta\right)\right)-h\left(\frac{y\left(\gamma,\theta\right)}{\theta}\right)+\delta_2 u\left(c_2\left(\gamma,\theta\right)\right)\right]dF\left(\theta|\kappa\left(e_\gamma,\gamma\right)\right).$$

Finally, by taking the supremum over all reports, we obtain a bond savings tax $T^k\left(b\right)=\sup_{\tilde{r}^t}T_t^k\left(b,\tilde{r}^t\right)$, where $\tilde{r}^1=\tilde{r}$ and $\tilde{r}^0=\tilde{r}_\gamma$, that only depends on bond purchases. With $T^k\left(b\right)$, agents do not purchase bonds while misreporting in equilibrium.

Next, we construct the other policy instruments. By Lemma 2, we can define the optimal consumption derived from the direct mechanism as $\left(c_0\left(e\right),c_1\left(e,y\right),c_2\left(e,y\right)\right)$. First, we construct the student loans and its income-contingent repayment schedule along with the income tax. Let the loan amount be defined as

$$L\left(e\right)=\begin{cases}c_0\left(e\right)+e & \text{if } e\in\{e_L,e_H\}\\ 0 & \text{otherwise}\end{cases},$$

and the income-contingent repayment subsidy is $\tau^e\left(e_L,y\right)=1$ and

$$\tau^e\left(e_H,y\right)=1+\frac{1}{\tilde{R}_1\left(e_H\right)L\left(e_H\right)}\left[c_1\left(e_H,y\right)-c_1\left(e_L,y\right)+\frac{c_2\left(e_H,y\right)}{R_2\left(1+\tau^s\left(e_H,y\right)\right)}-\frac{c_2\left(e_L,y\right)}{R_2\left(1+\tau^s\left(e_L,y\right)\right)}\right].$$

Let $y\left(\gamma,\theta\right)$ be the optimal output of type $\left(\gamma,\theta\right)$ agents in a direct revelation mechanism and define $Y=\{y|y=y\left(\gamma,\theta\right) \text{ with } \gamma\in\{L,H\} \text{ and } \theta\in\Theta\}$ to be the admissible set of income. The income tax is

$$T(y)=\begin{cases} y-c_1(e_L,y)-\frac{c_2(e_L,y)}{R_2(1+\tau^s(e_L,y))} & \text{if } y\in Y \\ y & \text{if } y\notin Y \end{cases}.$$

Next, we define the income and education contingent retirement savings subsidy as

$$1+\tau^s(e,y)=\begin{cases} \frac{u'(c_1(e,y))}{\beta u'(c_2(e,y))} & \text{if } e\in\{e_L,e_H\} \\ 0 & \text{otherwise} \end{cases}.$$

Finally, we check that the policy instruments implement the optimum. First, notice that all agents choose $e\in\{e_L,e_H\}$, otherwise they will not have any retirement consumption. Similarly, due to the income tax, all agents produce output $y\in Y$. Next, for any $e\in\{e_L,e_H\}$ and $y\in Y$, agents at $t=1$ choose consumption to satisfy

$$\frac{u'(c_1)}{\beta u'(c_2)}=1+\tau^s(e,y) \text{ and } c_1+\frac{c_2}{R_2(1+\tau^s(e,y))}=c_1(e,y)+\frac{c_2(e,y)}{R_2(1+\tau^s(e,y))}.$$

Clearly, agents optimally choose $c_1=c_1(e,y)$ and $c_2=c_2(e,y)$. Also, by the taxation principle, agents with productivity θ choose $y=y(e,\theta)$. For the final step, notice that given $L(e)$, agents with innate ability γ optimally choose education level e_γ. ■

Proof of Proposition 4: By Lemma 2, we can define the optimal consumption derived from the direct mechanism as $(c_0(e),c_1(e,y),c_2(e,y))$. Similarly, we focus on an implementation where agents do not purchase bonds due to the bond savings tax $T^k(b)$, which is constructed in the proof of Proposition 3.

For the policy instruments, we focus on an implementation where none of the agents save in the retirement savings account, so $s_2=0$. Agents with education e_L rely on social security for retirement consumption while agents with education e_H depend on social security benefits plus student loan repayment contributions in the retirement account. Let $y(\gamma,\theta)$ be the optimal output of type (γ,θ) agents in a direct revelation mechanism and define $Y=\{y|y=y(\gamma,\theta) \text{ with } \gamma\in\{L,H\} \text{ and } \theta\in\Theta\}$ to be the set of admissible income. First, we construct the matching rate α to be

$$1+\alpha=\inf_{y\in Y,e\in\{e_L,e_H\}}\frac{u'(c_1(e,y))}{\beta u'(c_2(e,y))}.$$

Next, we construct the social security benefit $a(y)=c_2(e_L,y)$. We set the income tax to be

$T(y - s_2) = y - s_2 - c_1(e_L, y)$, and the tax deduction from student loan repayment is

$$g(r(e_H, y)) = r(e_H, y) - [c_1(e_L, y) - c_1(e_H, y)] \text{ and } g(0) = 0.$$

Finally, we construct the student loans and its income-contingent repayment schedule along with the tax on retirement savings account. Let the loan amount be defined as

$$L(e) = \begin{cases} c_0(e) + e & \text{if } e \in \{e_L, e_H\} \\ 0 & \text{otherwise} \end{cases},$$

and the income-contingent repayment schedule is $r(e_L, y) = 0$ and

$$r(e_H, y) = \frac{1}{\alpha R_2}\left[c_2(e_H, y) - c_2(e_L, y) + T^{ra}\right].$$

We choose T^{ra} such that $r(e_H, y)$ and $g(R_1 r)$ are weakly positive. Let $T^{ra}(y)$ be a fictitious tax schedule defined as

$$T^{ra}(y) = \max\{0, c_2(e_L, y) - c_2(e_H, y), c_2(e_L, y) - c_2(e_H, y) + \alpha R_2[c_1(e_L, y) - c_1(e_H, y)]\}.$$

Observe that given $T^{ra}(y)$, both the repayment schedule and the tax deduction are weakly positive for any income. Lastly, by taking the supremum over all income, we obtain an income-independent lump-sum tax:

$$T^{ra} = \sup_{y \in Y} T^{ra}(y).$$

For our last step, we check that the policy instruments implement the optimum. First, notice that all agents would choose $e \in \{e_L, e_H\}$, otherwise $c_0 = 0$. Next, due to the low matching rate, all agents choose $s_2 = 0$. As a result, given the taxes and social security benefit, agents who invested e_L consume $c_1 = c_1(e_L, y)$ and $c_2 = c_2(e_L, y)$. Next, for agents who invested e_H, given the taxes, $c_1 = y - T(y) + g(r(e_H, y)) - r(e_H, y)$ and $c_2 = a(y) + \alpha R_2 r(e_H, y) - T^{ra}$, so they optimally choose $c_1 = c_1(e_H, y)$ and $c_2 = c_2(e_H, y)$. Also, by the taxation principle, agents with productivity θ choose $y = y(e, \theta)$. Finally, notice that given $L(e)$, agents with innate ability γ optimally choose education level e_γ. ■

Proof of Proposition 5: With heterogeneous β, the government's problem remains the same except (11) is now

$$U_1(\gamma,\theta) = u(c_1(\gamma,\theta)) - h\left(\frac{y(\gamma,\theta)}{\theta}\right) + \beta_\gamma \delta_2 u(c_2(\gamma,\theta)),$$

for all γ and the ex-ante incentive constraint is

$$\delta_0(e_H)u(c_0(H)) + \beta_H\delta_1(e_H)\int_{\underline{\theta}}^{\overline{\theta}}\left[U_1(H,\theta) + (1-\beta_H)\delta_2 u(c_2(H,\theta))\right] f(\theta|\kappa_H)\,d\theta$$
$$\geq \delta_0(e_L)u(c_0(L)) + \beta_H\delta_1(e_L)\int_{\underline{\theta}}^{\overline{\theta}}\left[U_1(L,\theta) + (1-\beta_L)\delta_2 u(c_2(L,\theta))\right] f(\theta|\kappa_{L,H})\,d\theta.$$

The results follow from the procedures outlined in the proofs for Proposition 1 and Proposition 2. ■

B Approximating Current Policies

To approximate current income taxes in the United States, we follow Heathcote et al. (2017) and assume an income tax function $T(y) = y - \lambda y^{1-\tau}$. College students have access to low-interest federal loans. In $t = 0$, agents take out loans with gross interest $\underline{R} < R_1(e_H)$ for loans below $\overline{L}$. We model this as if agents who chose e_H receive a lump-sum transfer of $T(e_H) = \underline{R}\overline{L}$ when they start working. Agents who choose e_L do not receive this transfer.

Upon retirement, agents receive social security benefits, which are income-dependent. The regulation below has been translated to fit the context of our model. To derive an agent's social security benefits, first calculate the agent's average indexed monthly earnings (AIME) which is defined as $AIME = \frac{y}{12}$ for annual income y. In practice, the social security administration takes 35 of the highest annual incomes from the 45 years of the agent's work life and calculate the average monthly earnings. Next, based on 2015 social security regulations, the agent's monthly benefit $a(AIME)$ is determined by the following replacement rates and bend points:

$$a(AIME) = \begin{cases} 0.9 \times AIME & \text{if } AIME \leq 826 \\ 743.4 + 0.32 \times (AIME - 826) & \text{if } 826 < AIME \leq 4,980 \\ 2,072.68 + 0.15 \times (AIME - 4,980) & \text{if } 4,980 < AIME \leq 9,875 \\ 2,806.93 & \text{if } AIME > 9,875 \end{cases}.$$

This immediately implies that the agent receives $A(y) = 12 \times a\,(AIME)$ every year in social security benefits.

Using the 2015 regulations, agents are subject to a flat social security tax $T_s\,(y)$, which is defined as

$$T_s\,(y) = \begin{cases} 0.124 \times y & \text{if } y \leq 118,500 \\ 14,694 & \text{if } y > 118,500 \end{cases}.$$

The tax is capped at an annual income of $118,500$. Furthermore, the social security benefits are distributed from the social security tax.

We assume that agents accumulate retirement savings in a 401(k) account and a regular savings account which pays a gross interest of R_2. Let s_2 denote savings in a 401(k) account and b_2 in the regular savings account. Contributions to the 401(k) account are capped at an annual amount of $18,000$. We also assume an employer matching rate of 50%. Contributions to defined contribution plans, such as 401(k), are pre-tax. This means that income tax payments are deferred upon withdrawal when retiring. However, social security tax is not deferred. Since contributions to 401(k) are matched, agents would first save in their 401(k) accounts until the cap binds, before saving in their regular accounts.

B.1 Deriving Allocations for Current Policies

To determine the allocation of present-biased agents under the current policy, we adopt subgame perfect Nash equilibrium as our solution concept.

B.1.1 The Working Period Problem

By backward induction, agents with productivity θ who took out a loan of b_1 in $t = 0$ and invested e in education solve the following problem:

$$\max u\,(c_1) - h\,(l) + \beta\delta_2 u\,(c_2)$$

subject to

$$c_1 + b_2 + s_2 = \theta l - T\,(\theta l - s_2) - T_s\,(\theta l) - \frac{R_1\,(e)}{R_0\,(e)} b_1 + \mathbf{1}_{e=e_H} T\,(e_H)\,,$$

$$c_2 = 1.5 R_2 s_2 + R_2 b_2 + A\,(\theta l) - T\,(1.5 R_2 s_2)\,,$$

$$s_2 \leq \bar{c},$$

where $\overline{c}$ is the upper-bound on contributions to the 401(k) account and $\mathbf{1}_{e=e_H} = 1$ only if agents invested e_H, or else $\mathbf{1}_{e=e_H} = 0$. Let $\overline{\chi}_t(\theta)$ denote the multiplier on the period t budget constraint for agents who invested e_H, and $\underline{\chi}_t(\theta)$ be the multiplier for low-educated agents.

Using Only 401(k): When agents only use 401(k), then it means that agents choose to save $s_2 < \overline{c}$.

We first look at agents who invested e_H. The first order conditions for consumption and savings s_2 are

$$u'(c_1) = \overline{\chi}_1(\theta), \;\; \beta\delta_2 u'(c_2) = \overline{\chi}_2(\theta) \;\text{ and } \overline{\chi}_1(\theta) = \overline{\chi}_2(\theta)\, 1.5R_2 \left(\frac{\theta l - s_2}{1.5R_2 s_2}\right)^{\tau}.$$

This provides us with the following Euler equation:

$$u'(c_1) = 1.5\beta \left(\frac{\theta l - s_2}{1.5R_2 s_2}\right)^{\tau} u'(c_2).$$

For labor supply, we have four different income regions to consider:

$$h'(l) = \begin{cases} \overline{\chi}_2(\theta)\left\{1.5R_2\left(\frac{\theta l - s_2}{1.5R_2 s_2}\right)^{\tau} B(\theta, y, s_2) + 0.9\theta\right\} & \text{if } y \leq 9,912 \\ \overline{\chi}_2(\theta)\left\{1.5R_2\left(\frac{\theta l - s_2}{1.5R_2 s_2}\right)^{\tau} B(\theta, y, s_2) + 0.32\theta\right\} & \text{if } 9,912 < y \leq 59,760 \\ \overline{\chi}_2(\theta)\left\{1.5R_2\left(\frac{\theta l - s_2}{1.5R_2 s_2}\right)^{\tau} B(\theta, y, s_2) + 0.15\theta\right\} & \text{if } 59,760 < y \leq 118,500 \\ \overline{\chi}_2(\theta)\, 1.5R_2\left(\frac{1}{1.5R_2 s_2}\right)^{\tau} \theta\lambda(1-\tau) & \text{if } y > 118,500 \end{cases},$$

where $B(\theta, y, s_2) = \theta\lambda(1-\tau)(y - s_2)^{-\tau} - 0.124\theta$ As for agents who invested e_L, the first order conditions are the same except for replacing $\overline{\chi}_t(\theta)$ with $\underline{\chi}_t(\theta)$.

Using Both 401(k) and Savings: When agents start saving in the regular savings account—$b_2 > 0$, then it means that $s_2 = \overline{c}$.

We first analyze the case where agents invested e_H in $t = 0$. Suppose the agent has saved $s_2 = \overline{c}$, then the agent can only continue to save with the standard savings account. We can rewrite the sequential budget constraint into its present value terms:

$$c_1 + \frac{c_2 - \lambda(1.5R_2\overline{c})^{1-\tau} - A(\theta l)}{R_2} = \lambda(\theta l - \overline{c})^{1-\tau} - T_s(\theta l) - \frac{R_1(e_H)}{R_0(e_H)} b_1 + T(e_H).$$

Let $\overline{\chi}(\theta)$ denote the multiplier on the present-valued budget constraint. The first order conditions on consumption are

$$u'(c_1) = \overline{\chi}(\theta) \text{ and } \beta u'(c_2) = \overline{\chi}(\theta).$$

The first order condition for labor is

$$h'(l) = \begin{cases} \overline{\chi}(\theta)\left[\theta\lambda(1-\tau)(\theta l - \overline{c})^{-\tau} - 0.124\theta + \frac{0.9}{R_2}\theta\right] & \text{if } y \leq 9,912 \\ \overline{\chi}(\theta)\left[\theta\lambda(1-\tau)(\theta l - \overline{c})^{-\tau} - 0.124\theta + \frac{0.32}{R_2}\theta\right] & \text{if } 9,912 < y \leq 59,760 \\ \overline{\chi}(\theta)\left[\theta\lambda(1-\tau)(\theta l - \overline{c})^{-\tau} - 0.124\theta + \frac{0.15}{R_2}\theta\right] & \text{if } 59,760 < y \leq 118,500 \\ \overline{\chi}(\theta)\,\theta\lambda(1-\tau)(\theta l - \overline{c})^{-\tau} & \text{if } y > 118,500 \end{cases},$$

We can derive a similar set of first order conditions for agents who obtained education level e_L.

B.1.2 The Schooling Period Problem

Let $(\check{c}_1(e,\theta), \check{y}(e,\theta), \check{c}_2(e,\theta))$ denote the solution to the problem in Section B.1.1, which is the optimal consumption path and output agents choose in $t = 1$ given education e and producticity θ. Agents with innate ability γ solve the following problem:

$$\max_{c_0,e,b_1} \delta_0(e)\,u(c_0) + \beta\delta_1(e)\int_{\underline{\theta}}^{\overline{\theta}}\left[u(\check{c}_1(e,\theta)) - h\left(\frac{\check{y}(e,\theta)}{\theta}\right) + \delta_2 u(\check{c}_2(e,\theta))\right] f(\theta|\kappa(e,\gamma))\,d\theta$$

subject to

$$c_0 + e = b_1 \text{ and } e \in \{e_L, e_H\}.$$

In essence, agents take out a yearly loan of b_1 to pay for their schooling and consumption in $t = 1$.

B.2 Calibration

In this section we calibrate the model to resemble the "real world" as closely as possible. The goal is to back out the distribution of productivities across different education groups. To this extent, we first pick a number of parameters externally and summarize them in Table 5. Then, we calibrate the distributions of skills internally to match the evidence on lifetime earning provided by Cunha and Heckman (2007).

The values of risk aversion and Frisch elasticity of labor are standard and set to 2 and

Table 5: Parameter values in the model

Symbol	Meaning	Value	Source
γ	Risk aversion	2	Standard values
η	Frisch elasticity	0.5	
τ	Tax progressivity	0.161	Heathcote and Tsujiyama (2017)
λ	Taxation level	0.839	
$\bar{c}$	401(k) contribution limit	1.80	Approximated from data
e_H	Cost of college	1.57	
r_m	Commercial interest on student loans	0.1	
r_g	Government interest on student loans	0.05	
$\bar{b}_g$	Cap on government-subsidized student loans	5.75	
β	Short-term discount factor	0.7	Based on Nakajima (2012)
$\delta_0(e_L)$	High school period 0 long-term discount factor	0.00	
$\delta_1(e_L)$	High school period 1 long-term discount factor	1.00	
$\delta_0(e_H)$	College period 0 long-term discount factor	0.15	
$\delta_1(e_H)$	College period 1 long-term discount factor	0.93	
δ_2	Retirement discount factor	0.270	
Time-consistent benchmark ($\beta = 1$)			Based on Nakajima (2012)
$\delta_0(e_L)$	High school period 0 long-term discount factor	0.00	
$\delta_1(e_L)$	High school period 1 long-term discount factor	1.00	
$\delta_0(e_H)$	College period 0 long-term discount factor	0.19	
$\delta_1(e_H)$	College period 1 long-term discount factor	0.85	
δ_2	Retirement discount factor	0.154	

Note: All monetary parameters are denominated in 10,000 of 2015 US dollars.

0.5, respectively. Next, we discuss the calibration of the current tax system. The parameters of the income tax function τ and λ are borrowed from Heathcote and Tsujiyama (2017) and apply to income level normalized by average income in the economy.[14] The upper bound for 401(k) contributions $\bar{c}$ is set to \$18,000 and reflects the present value of lifetime contribution limits based on the limit in 2015. As for the financing of student loans, we assume for simplicity that the annual interest rates an agent may obtained through private market and through a government-subsidized scheme are 10% and 5%, respectively. The amount of subsidized loan is capped at \$57,500, in line with the regulations for Stafford loans in the US. We further assume that an agent takes ten years to repay the student loans.

[14]We calculate average income directly using the factual distributions of lifetime income from Cunha and Heckman (2007) and the shares of high school and college graduates (and beyond) of 0.68 and 0.32, respectively, from the CPS. The average lifetime income amounts to \$1,570,900.

The annual cost of higher education e_H is assumed to be \$15,700, which is calculated for 2015 based on average tuition costs of private and public colleges plus different types of graduate degrees[15] as well as relative enrollment data for both types of college.[16] Table 6 presents a breakdown of different higher education outcomes, along with average costs and durations, which we use to calculate this parameter.

Table 6: Breakdown of higher education outcomes

Degree type	*% of population*	*Duration*	*Annual cost*
Associate's and less	67.7	0	0
Bachelor's only	20.3	4	15,396
Master's	8.0	6	16,140
Professional	1.9	8	27,210
Doctoral	2.1	10	6,158
Total	100	5.12	15,695

Note: distribution of educational attainment is from CPS 2015. The durations and annual costs are cumulative. The data on costs of various higher degrees are taken from NCES, Digest of Education Statistics and expressed in 2015 dollars. We ignore the cost and duration of Associate's degrees as those are often combined with jobs.

In calibrating the short- and long-term discount factors we primarily follow Nakajima (2012) who uses a general equilibrium model with present-biased agents and targets a capital-output ratio of 3. We adopt his assumed value of the short-term discount factor of 0.7 which places in the midrange of estimates found by Laibson et al. (2017). The annual long-run discount factor is $\delta_{annual} = 0.9852$ following Nakajima (2012) which we in turn use to calculate effective discount rates across the three periods in our model. These effective discount rates also reflect the relative lengths of the periods, which may differ across agents of different education groups. Because high school graduates start working right away, they never actually experience the education period 0; hence their parameter $\delta_0(e_L)$ is zero and $\delta_1(e_L)$ is one. On the other hand, college graduates spend 5.12 years in period 0, which reflects the average duration of undergraduate and graduate studies in the US (Table 6 presents a detailed breakdown), and then another 45 years in period 1. This yields $\delta_0(e_H) = \frac{1-\delta_{annual}^{5.12}}{1-\delta_{annual}^{45}} = 0.15$ and $\delta_1(e_H) = \frac{\delta_{annual}^{5.12}-\delta_{annual}^{50.12}}{1-\delta_{annual}^{45}} = 0.93$. We assume that both education types spend 45 years working and 20 years in retirement. This yields a common retirement period discount factor of $\delta_2 = \frac{\delta_{annual}^{45}-\delta_{annual}^{65}}{1-\delta_{annual}^{45}} = 0.27$.[17]

[15] Source: College Board, Annual Survey of Colleges and NCES, Digest of Education Statistics

[16] Source: NCES, Digest of Education Statistics, and Current Population Survey 2015

[17] Because the college type first spends four years on education before they start to work, we assume that they also retire four years later, at age 67, and live four years longer. This is consistent with a significant

For our analysis in the main body of the paper we also use the benchmark of time-consistent agents, i.e. the world where $\beta = 1$. For reference, we present here the analogous derivations of the effective long-run discount factors for that case. Once again following Nakajima (2012) we assume an annual discount rate $\delta_{annual} = 0.9698$. Then, with the same reasoning we assume $\delta_0(e_L) = 0$ and $\delta_1(e_L) = 1$ for high school graduates, compared to $\delta_0(e_H) = 0.19$ and $\delta_1(e_H) = 0.85$ for college graduates. The discount factor for retirement amounts to $\delta_2 = 0.15$.

Having established the external parameters, we turn to the parameters governing the distribution of skills which are set through solving and simulating the model. For each of the four groups of agents: (i.) factual high school graduates, (ii.) high school graduates, had they gone to college, (iii.) factual college graduates, and (iv.) college graduates, had they not gone to college, we observe the empirical distributions of lifetime earnings reported by Cunha and Heckman (2007). Roughly speaking, these distributions are obtained by estimating a Roy-type model on combined NLSY and PSID data and generating counterfactuals for both education groups. As it is commonly known, panel surveys such as these tend to under-represent the upper tail of the earnings distribution. For this reason, similar to Findeisen and Sachs (2016), we add an upper Pareto-tail with the shape parameter of 1.5 (Saez (2001)). For each distribution, we select an income threshold at which we attach the Pareto tail such that the upper 10% of the mass is distributed according to it. We pick the scale parameter such that the (smoothed out) PDF of the empirical distribution of earnings from Cunha and Heckman (2007) intersects at the threshold with the Pareto PDF Table 7 summarizes the parameters of the Pareto tail that we add to each of the empirical distribution of lifetime earnings.

Table 7: Adding a Pareto tail to lifetime income distributions

	HS fact.	*HS counter.*	*COL fact.*	*COL counter.*
Threshold	203.7	257.6	288.4	212.8
Scale parameter	62.9	80.2	89.4	63.8

Note: The thresholds refer to present value of lifetime earnings and are expressed in $10,000s of 2015 dollars. Thresholds are selected in each case such that 10% of total mass is distributed according to Pareto distribution with the shape parameter of 1.5.

To capture the earnings distribution with a fat upper tail in our model, we assume that agents' skills θ follow a mixture of two distributions, a normal distribution and a two-piece distribution (lognormal-Pareto) as described in Nigai (2017). The probability density

body of research which shows college graduates live longer than non-college graduates (Meara et al., 2008).

function of our mixture is then given by

$$f(\theta) = p \times \left[\frac{1}{2\pi\sigma_1} \exp\left\{ \frac{(\theta-\mu_1)^2}{2\sigma_1^2} \right\} \right] \tag{13}$$
$$+(1-p) \times \begin{cases} \frac{\rho}{\Phi\left(\alpha s(\alpha,\rho)\right)} \frac{1}{\sqrt{2\pi} s(\alpha,\rho)\theta} \exp\left\{ -\frac{1}{2}\left(\alpha s(\alpha,\rho) - \frac{\log(\theta^T)-\log(\theta)}{s(\alpha,\rho)}\right) \right\}, & \text{if } \theta \in (0, \theta^T) \\ (1-\rho)\frac{\alpha(\theta^T)^\alpha}{\theta^{\alpha+1}}, \quad \text{if } \theta \in [\theta^T, \infty) \end{cases}$$

In equation (13), μ and σ are the mean and standard deviation of the normal distribution, and p is the probability of drawing it. The two-piece lognormal-Pareto distribution comes with a shape parameter α, which we fix at 1.5, and two scale parameters, ρ and θ^T. Intuitively, θ^T is the threshold value at which the standard lognormal distribution turns into Pareto, while $\rho \in (0, 1)$ represents the fraction of total mass that is distributed according to lognormal. We have hence 5 parameters to pin down for each of the four groups of agents, $(\mu, \sigma, \rho, \theta^T, p)$, in order to replicate the empirical distributions of earning provided by Cunha and Heckman (2007) and augmented with the Pareto tail. To do so, we solve for the optimal policy functions in each of the four cases and simulate random draws for 100,000 agents. We use a global optimization algorithm to minimize the distance between the simulated CDF of lifetime earnings and the targeted one. Table 8 shows all components of our mixture density defined in (13) matter quantitatively and altogether result in a good fit for model-derived distributions of earnings in each group. Figures 8-9 depict the CDF of lifetime earnings in the model and their empirical targets across the four groups of agents. Notice that all estimations result in an excellent fit to the data, with an exception of High School counterfactual. However, this distribution does not affect the model solution in any way and it is only necessary to verify that the low type indeed prefers to reveal truthfully.

Table 8: Parameters of productivity distributions

Symbol	Meaning	Value			
		HS fact.	*HS counter.*	*COL fact.*	*COL counter.*
μ	Mean of normal	7.91	8.99	10.34	8.25
σ	St. dev. of normal	2.15	2.44	3.48	2.15
θ^T	Threshold for Pareto	8.36	8.45	18.24	8.10
ρ	Fraction of lognormal	0.39	0.48	0.65	0.43
p	Probability of normal	0.64	0.60	0.66	0.56

Note: Productivities drawn from these distributions are in annual terms.

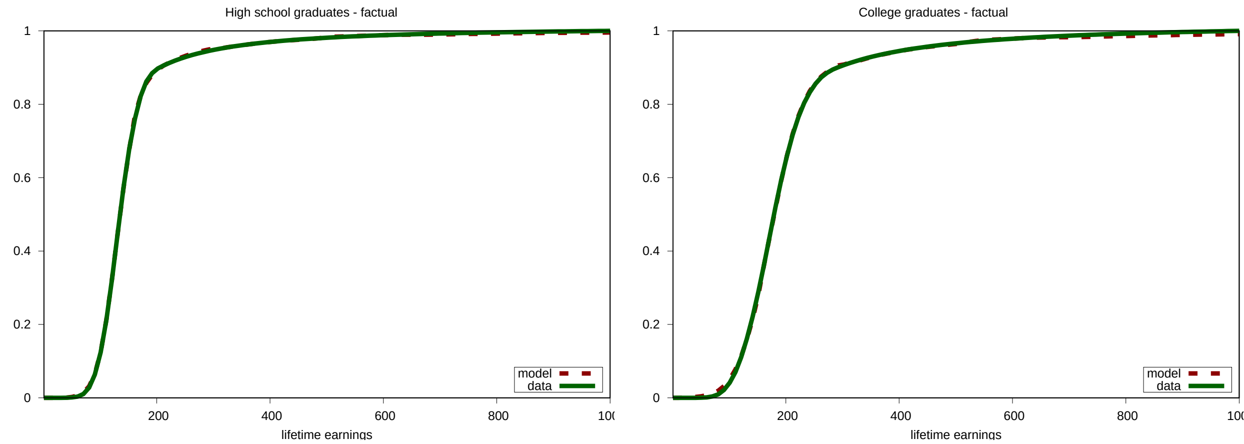

Figure 8: Distributions of lifetime income - factual

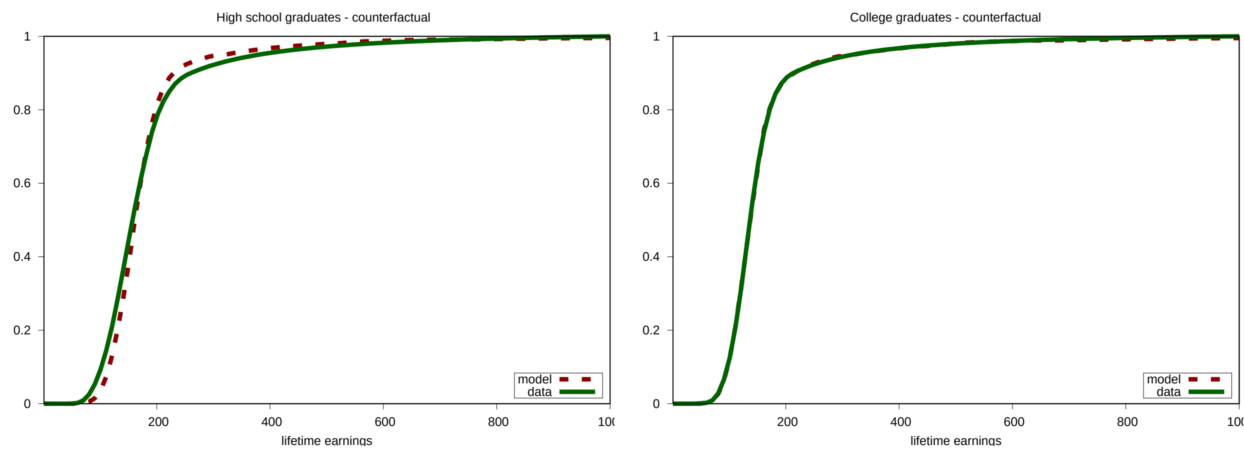

Figure 9: Distributions of lifetime income - counterfactual

C Decomposition of the Labor Wedge

In this section, we quantify the decomposition of the labor wedge introduced in Section 3.2. Figure 10 presents the numerical approximation of the labor wedge components A,C,D and E as function of income for the high innate ability type.[18] The A component depends on the inverse hazard rate of the distribution of θ and declines at first, before increasing and converging to a constant due to the presence of a Pareto tail. By contrast, the intratemporal component C increases and then converges, resulting in the overall convergence of the labor

[18]We ignore the B components because, given the functional forms we impose, it reduces to a constant. We also omit the decomposition for the low innate ability type because in our calibrated model the period-zero consumption of L-agents is not pinned down. As a result, the intertemporal component is not well-defined.

wedge at the top of the distribution. The offsetting role that comes from the intertemporal component D is much smaller in size and decreases monotonically.

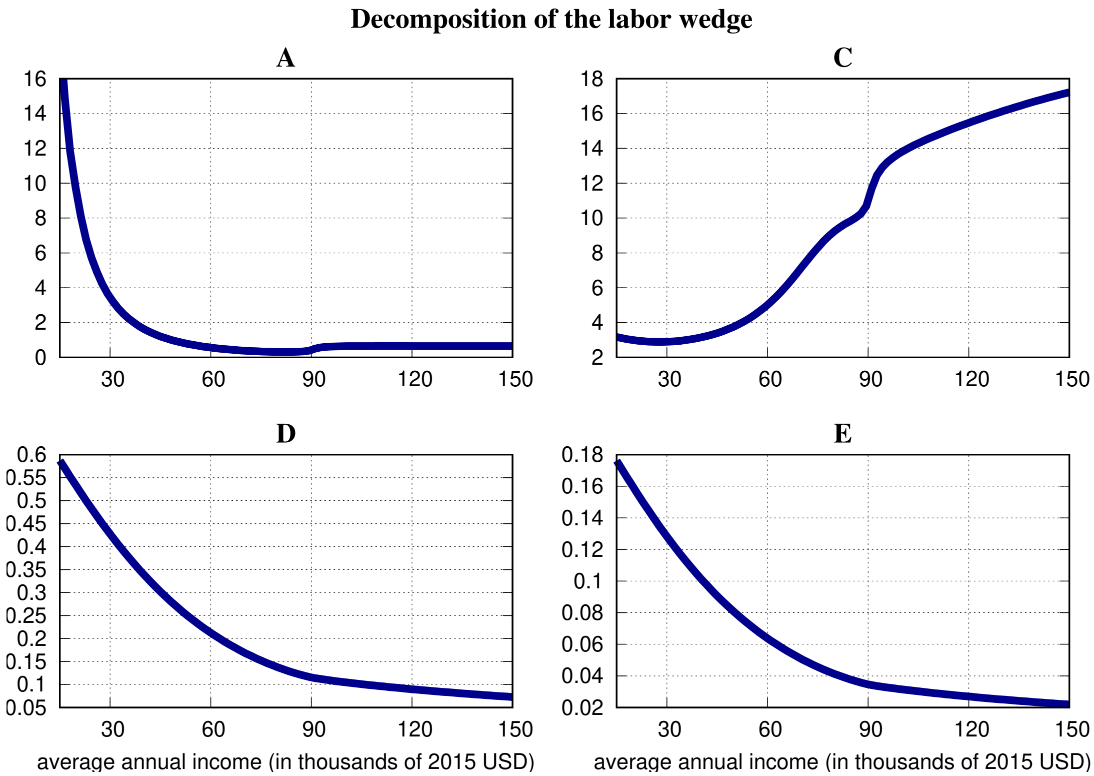

Figure 10: Labor wedge components for the high innate ability type

The novel aspect of our paper is the introduction of E, the present bias component. As can be noticed, this components declines monotonically but its magnitude is also very small compared to components D, or especially C. Consequently, the labor wedge is generally not much affected by the present bias, and any difference shows up most prominently at the lowest levels of income, as evident in Figure 3.

To get a better understanding of this result, we can decompose the present bias element further into the *disagreement component* and the *myopic component*, as explained in Section 3.2. Figure 11 plots these two components on a single graph for the high innate ability type. Notice immediately that they offset each other for the most part and so the difference between them, the present bias component of the labor wedge, is quantitatively very small. This is similar for L-agents, which is plotted in Figure 12.

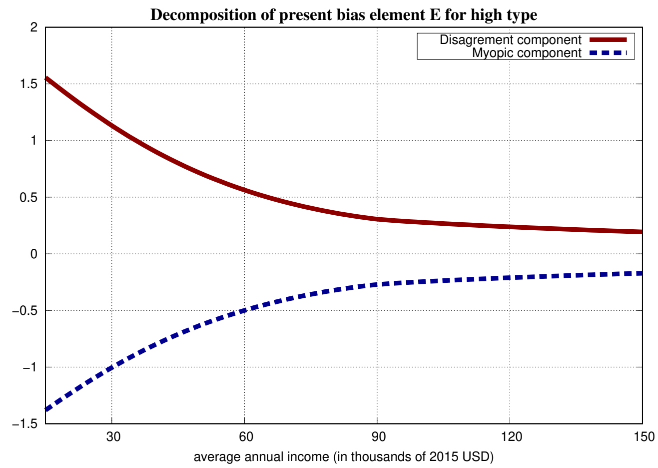

Figure 11: Decomposition of the present bias element of the labor wedge for H-agents

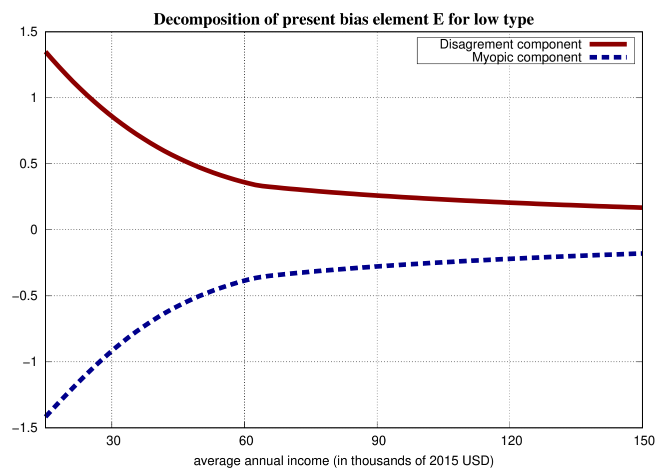

Figure 12: Decomposition of the present bias element of the labor wedge for L-agents

D Time-Consistent Benchmarks

In this section, we present details of the implementation of the time-consistent optimal policies, which we use as benchmark for welfare calculations in Section 4.1. We consider two ways to implement the optimal allocation for time-consistent agents: mandatory retirement savings and laissez faire retirement savings. The two different implementations lead to different measures of welfare improvement.

First, we characterize the optimal allocations for time-consistent agents in a direct mechanism. Let $\left\{\tilde{c}_0(\gamma), [\tilde{c}_t(\gamma,\theta), \tilde{y}(\gamma,\theta)]_{t>0,\theta\in\Theta}\right\}$ be the optimal allocation for time-consistent agents. The optimal allocation for time-consistent agents satisfies the following:

$$u'(\tilde{c}_1(\gamma,\theta)) = u'(\tilde{c}_2(\gamma,\theta)).$$

This implies that $\tilde{c}_1(\gamma,\theta) = \tilde{c}_2(\gamma,\theta) = \tilde{c}(\gamma,\theta)$. For $t=0$, the government implements $\tilde{c}_0(\gamma)$ by providing agents a student loan of

$$L(e_H) = \tilde{c}_0(H) + e_H \text{ and } L(e_L) = \tilde{c}_0(L) + e_L.$$

Next, we proceed to consider two different methods to decentralize the optimal allocations in $t=1$ and $t=2$.

D.1 Mandatory Savings

Consider a mandatory minimum savings rule that forces agents to smooth consumption: $\tilde{c}_1(\gamma,\theta) = \tilde{c}_2(\gamma,\theta)$. For time-consistent agents, the policy implements the optimum. However, for present-biased agents, the minimum savings rule is not incentive compatible.

To see how the minimum savings rule changes the behavior of present-biased agents, we first analyze how agents would change their reports of θ. Since for our quantitative exercise, $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ and $h\left(\frac{y}{\theta}\right) = \frac{1}{1+\frac{1}{\eta}}\left(\frac{y}{\theta}\right)^{1+\frac{1}{\eta}}$. Then, for a given report of innate ability $\hat{\gamma}$ and the time-consistent allocations, present-biased agents choose a report $\hat{\theta}$ to maximize the utility at $t=0$. In essence, a θ-agent solves

$$\max_{\hat{\theta}} u\left(\tilde{c}_1\left(\hat{\gamma},\hat{\theta}\right)\right) - h\left(\frac{\tilde{y}\left(\hat{\gamma},\hat{\theta}\right)}{\theta}\right) + \beta\delta_2 u\left(\tilde{c}_2\left(\hat{\gamma},\hat{\theta}\right)\right).$$

From the argument above and the assumptions on the utility function, the problem can be rewritten as

$$\max_{\hat{\theta}} u\left(\tilde{c}\left(\hat{\gamma},\hat{\theta}\right)\right) - \frac{1}{1+\frac{1}{\eta}}\left(\frac{\tilde{y}\left(\hat{\gamma},\hat{\theta}\right)}{(1+\beta\delta_2)^{\frac{1}{1+\frac{1}{\eta}}}\theta}\right)^{1+\frac{1}{\eta}}.$$

We know that when $\beta=1$, the solution to the problem above is $\hat{\theta}=\theta$, because the mechanism satisfies incentive compatibility for time-consistent agents by assumption. Thus, we can transform the problem into the following alternative problem:

$$\max_{\hat{\theta}} u\left(\tilde{c}\left(\hat{\gamma},\hat{\theta}\right)\right)-\frac{1}{1+\frac{1}{\eta}}\left(\frac{\tilde{y}\left(\hat{\gamma},\hat{\theta}\right)}{\alpha\left(1+\delta_2\right)^{\frac{1}{1+\frac{1}{\eta}}}\theta}\right)^{1+\frac{1}{\eta}},$$

where $\alpha=\left(\frac{1+\beta\delta_2}{1+\delta_2}\right)^{\frac{1}{1+\frac{1}{\eta}}}$. Immediately, we can see that agents optimally report $\hat{\theta}=\alpha\theta$, because the problem is similar to a time-consistent agent with productivity $\alpha\theta$. As a result, the present-biased agents with productivity θ do not report truthfully and instead report

$$\left(\frac{1+\beta\delta_2}{1+\delta_2}\right)^{\frac{1}{1+\frac{1}{\eta}}}\theta.$$

This result is intuitive, because the reward for working is spread evenly between the two periods with mandatory savings. Since present-biased agents put less weight on retirement consumption, the mandatory savings policy provides less incentives for them to work. Their optimal strategy is to under-report their productivity to work less.

Finally, in $t=0$, agents know that they will report $\left(\frac{1+\beta\delta_2}{1+\delta_2}\right)^{\frac{1}{1+\frac{1}{\eta}}}\theta$ in $t=1$. As a result, given the optimal time-consistent allocation, H-agents solve the following:

$$\max\left\{u\left(\tilde{c}_0\left(H\right)\right)+\beta\frac{\delta_1\left(e_H\right)}{\delta_0\left(e_H\right)}\int_{\underline{\theta}}^{\overline{\theta}}\left[u\left(\tilde{c}_1\left(H,\hat{\theta}\right)\right)-h\left(\frac{\tilde{y}\left(H,\hat{\theta}\right)}{\theta}\right)+\delta_2 u\left(\tilde{c}_2\left(H,\hat{\theta}\right)\right)\right]dF\left(\theta|\kappa_H\right),\right.$$
$$\left.\frac{\delta_0\left(e_L\right)}{\delta_0\left(e_H\right)}u\left(\tilde{c}_0\left(L\right)\right)+\beta\frac{\delta_1\left(e_L\right)}{\delta_0\left(e_H\right)}\int_{\underline{\theta}}^{\overline{\theta}}\left[u\left(\tilde{c}_1\left(L,\hat{\theta}\right)\right)-h\left(\frac{\tilde{y}\left(L,\hat{\theta}\right)}{\theta}\right)+\delta_2 u\left(\tilde{c}_2\left(L,\hat{\theta}\right)\right)\right]dF\left(\theta|\kappa_{L,H}\right)\right\},$$

where $\hat{\theta}=\left(\frac{1+\beta\delta_2}{1+\delta_2}\right)^{\frac{1}{1+\frac{1}{\eta}}}\theta$.

D.2 Laissez Faire Savings

Another way to implement the optimum is for the government to allow agents to save freely for retirement. This is because with time-consistent agents, it is not necessary for the government to introduce any additional incentives for retirement savings. Hence, to implement the optimal allocation for time-consistent agents, the government only needs to introduce appropriate income taxes at $t=1$ and student loans in $t=0$. However, under laissez faire savings, present-biased agents do not smooth consumption and it is also not incentive compatible.

To find out how present-biased agents behave, we first derive the income tax $\tilde{T}(y)$ that implements the optimum for time-consistent agents. At $t=1$, time-consistent agents solves the following:

$$\max_{c_1,c_2,y} u(c_1) - h\left(\frac{y}{\theta}\right) + \delta_2 u(c_2)$$

subject to

$$c_1 + s_2 = y - \tilde{T}(y) \text{ and } c_2 = R_2 s_2.$$

Let $\tilde{Y}$ be the set of optimal income for TC agents:

$$\tilde{Y} = \{y | y = \tilde{y}(\gamma,\theta), \forall \gamma \in \{L,H\}, \theta \in \Theta\}.$$

By Lemma 2, we can rewrite the allocations in terms of income: $\tilde{c}_t(\tilde{y}(\gamma,\theta)) = \tilde{c}(\gamma,\theta)$. As a result, we can define the following income tax, which implements the optimum for time-consistent agents:

$$\tilde{T}(y) = \begin{cases} y & \text{if } y \notin \tilde{Y} \\ y - \tilde{c}(y) - \frac{1}{R_2}\tilde{c}(y) & \text{if } y \in \tilde{Y}. \end{cases}.$$

For simplicity, we assume that if the government observes an off-path income level that it didn't expect, it usurps all of the output and leaves the agent without any consumption.

Next, we outline how present-biased agents behave under laissez faire savings. Given laissez faire savings and the income tax above, present-biased agents solve the following at $t=1$,

$$\max_{c_1,c_2,y} u(c_1) - h\left(\frac{y}{\theta}\right) + \beta\delta_2 u(c_2)$$

subject to

$$c_1 + s_2 = y - \tilde{T}(y) \text{ and } c_2 = R_2 s_2.$$

We can rewrite the problem as

$$\max_{c_1,c_2,y} u(c_1) - h\left(\frac{y}{\theta}\right) + \beta\delta_2 u(c_2)$$

subject to

$$c_1 + \frac{1}{R_2}c_2 = \left(1 + \frac{1}{R_2}\right)\tilde{c}(y) \text{ and } y \in \tilde{Y}.$$

It is clear that agents never choose $y \notin \tilde{Y}$, because all of the output would be confiscated. As a result, for any given $y \in \tilde{Y}$, present-biased agents choose consumption $(\hat{c}_1(y), \hat{c}_2(y))$

to satisfy

$$u'\left(\hat{c}_1\left(y\right)\right)=\beta u'\left(\hat{c}_2\left(y\right)\right)$$

and

$$\hat{c}_1\left(y\right)+\frac{1}{R_2}\hat{c}_2\left(y\right)=\tilde{c}\left(y\right)+\frac{1}{R_2}\tilde{c}\left(y\right).$$

It is obvious that agents do not choose the optimum for time-consistent agents. Hence, there will be intertemporal inefficiencies. In addition to the intertemporal inefficiencies, the agents might also choose suboptimal output. Agents choose y by solving the following:

$$\max_{y\in\tilde{Y}} u\left(\hat{c}_1\left(y\right)\right)-h\left(\frac{y}{\theta}\right)+\beta\delta_2 u\left(\hat{c}_2\left(y\right)\right),$$

where $\left(\hat{c}_1\left(y\right),\hat{c}_2\left(y\right)\right)$ is defined above.

After solving for the optimal allocations for $t=1,2$, we can solve for the agent's education choices in $t=0$. The process is the same as the one for mandatory savings.

D.3 Welfare Comparisons

To evaluate the welfare improvement of the paper's proposed policies, we measure the change of moving from mandatory savings or laissez faire savings to the policies introduced in Section 5.

However, this welfare evaluation is not straightforward. We need to guarantee the allocations chosen by present-biased agents under mandatory savings or laissez faire savings are feasible. This is because, from the analysis above, output of present-biased agents is further distorted under policies designed for TC agents. Therefore, the government budget constraint does not hold with present-biased agents under mandatory savings or laissez faire savings.

To facilitate the welfare comparison, we introduce an external government expenditure $G>0$ in the time-consistent setup, so that the resource constraint becomes

$$\sum_{\gamma}\pi_{\gamma}\left\{-\frac{\tilde{c}_0\left(\gamma\right)}{R_0\left(e_{\gamma}\right)}-e_{\gamma}+\frac{1}{R_1\left(e_{\gamma}\right)}\int_{\Theta}\left[\tilde{y}\left(\gamma,\theta\right)-\tilde{c}_1\left(\gamma,\theta\right)-\frac{1}{R_2}\tilde{c}_2\left(\gamma,\theta\right)\right]f\left(\theta|\kappa_{\gamma}\right)d\theta\right\}\geq G.$$

We interpret G as an emergency fund the government uses to supplement the agents' consumption when total output is lower than expected. Hence, we require G to be sufficiently large so that the allocations chosen by the present-biased agents,

$\left\{ \hat{c}_0(\gamma), [\hat{c}_t(\gamma,\theta), \hat{y}(\gamma,\theta)]_{t>0,\theta\in\Theta} \right\}$, are feasible:

$$\sum_{\gamma} \pi_\gamma \left\{ -\frac{\hat{c}_0(\gamma)}{R_0(e_\gamma)} - e_\gamma + \frac{1}{R_1(e_\gamma)} \int_\Theta \left[\hat{y}(\gamma,\theta) - \hat{c}_1(\gamma,\theta) - \frac{1}{R_2} \hat{c}_2(\gamma,\theta) \right] f(\theta|\kappa_\gamma)\, d\theta \right\} \geq 0. \quad (14)$$

D.4 Quantitative Implementation

In our quantitative exercise, we design a fixed-point algorithm to find the value of G such that the resource constraint in (14) binds. The algorithm can be summarized as follows:

1. Start with an initial value for government spending G_0.
2. Solve for the optimal allocations with time-consistent agents.
3. Use the allocations, implemented either through mandatory savings or laissez-faire arrangement, to solve for the best response of present-biased agents. Calculate the resulting gap in the resource constraint which stems from present-biased agents under-reporting their productivity type. Denote the gap G_1.
4. Check if $|G_0 - G_1| < \varepsilon$, where ε is a tolerance criterion. If yes, we have found a fixed point. If not, update G_0 and go back to step 1.

Table 9 summarizes the fixed-point amount of government spending G under both implementations which balances the resource constraint under present-biased agents.

Table 9: Fixed-point amount of government spending that balances the resource constraint

Mandatory savings	**Laissez-faire**
10.19	11.27

E Solving for Optimal Education-Independent Policies

In this section, we describe the computational algorithm used to solve for the benchmark case of optimal allocations conditional on the labor or the intertemporal wedge being education-independent. The challenge lies in the fact that while the allocations that we solve for are a function of productivity θ, the education-independence constraint on wedges is imposed for each observable income $y(\gamma,\theta)$ which itself is an allocation.

To overcome this challenge we adopt the following approach:

1. Create an exogenous grid for income consisting of N_y points. This grid is separate from the original grid for productivity θ on which the allocations are defined.

2. Consider a generic set of allocations $\{c_1(\gamma,\theta), c_2(\gamma,\theta), y(\gamma,\theta)\}_{\gamma\in\{H,L\}}$. For each point on the exogenous income grid defined in step 1, $\hat{y}_i$, find $\hat{\theta}_i$ and $\tilde{\theta}_i$ such that:

$$y(L,\hat{\theta}_i) = \hat{y}_i = y(H,\tilde{\theta}_i)$$

   We use linear interpolation to evaluate income at off-grid productivity values.

3. Solve for the optimal allocations under an additional set of N_y constraints, each one defined by a point on the exogenous grid for income from step 1, as follows.

   - For the case of education-independent labor wedge:

$$\frac{h'\left(\frac{y(L,\hat{\theta}_i)}{\hat{\theta}_i}\right)}{\hat{\theta}_i u'\left(c_1(L,\hat{\theta}_i)\right)} = \frac{h'\left(\frac{y(H,\tilde{\theta}_i)}{\tilde{\theta}_i}\right)}{\tilde{\theta}_i u'\left(c_1(H,\tilde{\theta}_i)\right)}$$

   - For the case of education-independent intertemporal wedge:

$$\frac{u'\left(c_1(L,\hat{\theta}_i)\right)}{u'\left(c_2(L,\hat{\theta}_i)\right)} = \frac{u'\left(c_1(H,\tilde{\theta}_i)\right)}{u'\left(c_2(H,\tilde{\theta}_i)\right)}$$

   Once again, we use linear interpolation to evaluate consumption in both period at off-grid values of productivity.

F Off-Path Policies for Non-Sophisticated Agents

The paper has focused on sophisticated present-biased agents. Sophisticated agents fully anticipate the behavior of their future-selves, so they have a demand for commitment. On the other hand, non-sophisticated agents underestimate the severity of their bias and tend to demand too little commitment. We explore the implications of non-sophistication on the design of optimal policy in this section.

To model non-sophistication, we follow O'Donoghue and Rabin (2001). Agents at $t = 0$ perceive their present bias in $t = 1$ to be $\hat{\beta} \in [\beta, 1]$. Let $W_1\left(c_1, c_2, y; \gamma, \theta, e, \hat{\beta}\right)$ denote the non-sophisticated agents' perceived utility in $t = 1$:

$$W_1\left(c_1, c_2, y; \gamma, \theta, e, \hat{\beta}\right) = u\left(c_1\right) - h\left(\frac{y}{\theta}\right) + \hat{\beta}\delta_2 u\left(c_2\right).$$

If $\hat{\beta} = \beta$, agents are *sophisticated* and fully aware of the bias. If $\hat{\beta} = 1$, agent are *fully naïve* and believe their future-selves to be time-consistent. *Partially naïve* agents know they are present-biased, $\hat{\beta} < 1$, but they underestimate its severity, $\hat{\beta} > \beta$. For this extension, we assume all agents are non-sophisticated and have heterogeneous and unobservable sophistication distributed within support $\left[\underline{\hat{\beta}}, 1\right]$, where $\underline{\hat{\beta}} \in (\beta, 1]$.

Yu (2019a) showed that it is optimal for the government to take advantage of the misspecified beliefs of present-biased agents through the preference arbitrage mechanism (PAM). PAM features off-path allocation used to exploit the incorrect beliefs, which are referred to as the *imaginary allocations* and denoted as (c^I, y^I). The allocation the government implements on-path is called the *real allocations* denoted as (c, y). We assume that u is unbounded below and above ($u(\mathbb{R}_+) = \mathbb{R}$). For H-agents, the government designs the menu

$$\hat{P}_H = \left\{c_0(H), \left[c_1^I, y^I, c_2^I\right], \left[c_1(H,\theta), y(H,\theta), c_2(H,\theta)\right]_{\theta\in\Theta}\right\}.$$

At $t = 1$, H-agents choose between imaginary and real allocations. The consumption path of the imaginary allocations is backloaded ($c_2^I > c_1^I$), while the consumption path of the real allocation is relatively less back-loaded. It is designed this way so that at $t = 0$, the agents mistakenly believe their future-selves will choose the imaginary allocations. However, they end-up selecting the real allocations instead. Since the ex-ante incentive constraints are non-binding for L-agents, the government does not need to design imaginary allocations for them, so $\hat{P}_L = \left\{c_0(L), \left[c_1(L,\theta), y(L,\theta), c_2(L,\theta)\right]_{\theta\in\Theta}\right\}$. Similar to Yu (2019a), it is not necessary to design imaginary allocations tailored for each level of sophistication. It is possible to find a single set of imaginary allocations such that it implements the same real allocations for agents of any sophistication.

Lemma 3 *For non-sophisticated present-biased agents, the ex-ante incentive compatibility constraint is non-binding at the optimum.*

Proof From Yu (2019a), we first choose the imaginary allocations such that it satisfies the preference arbitrage constraint: for any θ,

$$u\left(c_1^I\right) - h\left(\frac{y^I}{\theta}\right) + \hat{\beta}\delta_2 u\left(c_2^I\right) \geq \max_{\hat{\theta}}\left\{u\left(c_1\left(H,\hat{\theta}\right)\right) - h\left(\frac{y\left(H,\hat{\theta}\right)}{\theta}\right) + \hat{\beta}\delta_2 u\left(c_2\left(H,\hat{\theta}\right)\right)\right\}.$$

In essence, in $t = 0$, agents believe their future-selves would choose the imaginary allocation over the real allocation. Notice that the real allocations may not be incentive compatible under the erroneous belief. Next, the imaginary allocations have to satisfy the executability constraints to makes sure that agents eventually choose the real allocation: for any θ,

$$U_1(H,\theta) = u(c_1(H,\theta)) - h\left(\frac{y(H,\theta)}{\theta}\right) + \beta\delta_2 u(c_2(H,\theta)) \geq u\left(c_1^I\right) - h\left(\frac{y^I}{\theta}\right) + \beta\delta_2 u\left(c_2^I\right).$$

Thus, the ex-ante incentive compatibility constraint is

$$\delta_0(e_H)\, u(c_0(H)) + \beta\delta_1(e_H)\int_{\underline{\theta}}^{\overline{\theta}}\left[u\left(c_1^I\right) - h\left(\frac{y^I}{\theta}\right) + \delta_2 u\left(c_2^I\right).\right] f(\theta|\kappa_H)\, d\theta \geq U_0(L;H).$$

Next, we show how the imaginary allocations can be designed such that the ex-ante incentive constraint is non-binding for all sophistication levels. Denote $I(\theta) = u\left(c_1^I\right) - h\left(\frac{y^I}{\theta}\right)$ and set $I(\theta) = U_1(H,\theta) - \beta\delta_2 u\left(c_2^I\right)$ so the executability constraints are binding. Hence, the preference arbitrage constraints can be expressed as $u\left(c_2^I\right) \geq J\left(\theta,\hat{\beta}\right)$, for all $\theta \in \Theta$ where $J\left(\theta,\hat{\beta}\right) = \frac{\max_{\hat{\theta}}\left\{u\left(c_1\left(H,\hat{\theta}\right)\right) - h\left(\frac{y(H,\hat{\theta})}{\theta}\right) + \hat{\beta}\delta_2 u\left(c_2\left(H,\hat{\theta}\right)\right)\right\} - U_1(H,\theta)}{\left(\hat{\beta}-\beta\right)\delta_2}$. Since the preference arbitrage constraints need to hold for all productivity realizations and sophistication, it is clear that c_2^I is chosen to satisfy $u\left(c_2^I\right) \geq \max_{\theta,\hat{\beta}} J\left(\theta,\hat{\beta}\right)$. Similarly, the ex-ante incentive constraint can be rewritten as $u\left(c_2^I\right) \geq K$, where $K = \frac{1}{(1-\beta)\delta_2}\left\{\frac{U_0(L;H) - \delta_0(e_H)u(c_0(H))}{\beta\delta_1(e_H)} - \int_{\underline{\theta}}^{\overline{\theta}} U_1(H,\theta)\, dF(\theta|\kappa_H)\right\}$. Since u is unbounded above, for any real allocation, the imaginary retirement consumption c_2^I can be chosen to satisfy $u\left(c_2^I\right) \geq \max\left\{\max_{\theta,\hat{\beta}} J\left(\theta,\hat{\beta}\right), K\right\}$. Also, since u is unbounded below, it is possible to adjust c_1^I for any given y^I so that $I(\theta) = U_1(H,\theta) - \beta\delta_2 u\left(c_2^I\right)$. As a result, it is always possible to find a single set of imaginary allocations for all levels of sophistication such that the ex-ante incentive constraints are non-binding for any allocation implemented on the equilibrium path. ■

To understand Lemma 3, note that non-sophisticated agents at $t = 0$ overestimate the value of retirement consumption to their future-selves. PAM takes advantage of incorrect beliefs by encouraging education investment through an increased imaginary retirement consumption c_2^I, which H-agents believe they will choose in $t = 1$. However, their future-selves forsake it for more immediate gratification—the relatively less back-loaded real allocations. The following proposition describes the optimal wedges for non-sophisticated agents.

Proposition 6 *The constrained efficient allocation for non-sophisticated agents satisfies*

i. full insurance in $t = 0$: $c_0(H) = c_0(L)$,

ii. the inverse Euler equations: for any γ, $\frac{1}{u'(c_0(\gamma))} = \mathbb{E}_\theta\left(\frac{1}{u'(c_1(\gamma,\theta))}\right) = \mathbb{E}_\theta\left(\frac{1}{u'(c_2(\gamma,\theta))}\right)$, *and, for any* θ, $\frac{1}{\beta u'(c_2(\gamma,\theta))} = \frac{1}{u'(c_1(\gamma,\theta))} + \left(\frac{1-\beta}{\beta}\right)\frac{1}{u'(c_0(\gamma))}$, *iii. the labor wedge for any γ and θ satisfies* $\frac{\tau^{w}(\gamma,\theta)}{1-\tau^{w}(\gamma,\theta)} = A_{\gamma}(\theta) B_{\gamma}(\theta) C_{\gamma}(\theta)$.

Proof By Lemma 3, the optimization problem with non-sophisticated agents is the same as the original problem except the ex-ante incentive constraints are non-binding, so $\mu = 0$. As a result, the first order conditions are

$$u'(c_0(H)) = u'(c_0(L)) = \phi,$$

and for all γ,

$$\pi_{\gamma}\delta_1(e_{\gamma}) f(\theta|\kappa_{\gamma}) - \xi'_{\gamma}(\theta) = \lambda_{\gamma}(\theta),$$

$$(1-\beta)\pi_{\gamma}\delta_1(e_{\gamma}) f(\theta|\kappa_{\gamma}) + \beta\lambda_{\gamma}(\theta) = \frac{\phi\pi_{\gamma}\delta_1(e_{\gamma}) f(\theta|\kappa_{\gamma})}{u'(c_2(\gamma,\theta))},$$

$$\lambda_{\gamma}(\theta) u'(c_1(\gamma,\theta)) = \phi\pi_{\gamma}\delta_1(e_{\gamma}) f(\theta|\kappa_{\gamma}),$$

$$\xi_{\gamma}(\underline{\theta}) = \xi_{\gamma}(\overline{\theta}) = 0,$$

$$\lambda_{\gamma}(\theta)\frac{1}{\theta}h'\left(\frac{y(\gamma,\theta)}{\theta}\right) + \xi_{\gamma}(\theta)\left[\frac{1}{\theta^2}h'\left(\frac{y(\gamma,\theta)}{\theta}\right) + \frac{y(\gamma,\theta)}{\theta^3}h''\left(\frac{y(\gamma,\theta)}{\theta}\right)\right] = \phi\pi_{\gamma}\delta_1(e_{\gamma}) f(\theta|\kappa_{\gamma}).$$

By rearranging the first order conditions, the results follow. ■

Proposition 6 demonstrates the government's ability to fully insure agents against differences in innate ability γ. This is not surprising, because Lemma 3 stated that the government can screen innate ability without distortions by utilizing PAM. As a result, the only distortions in the economy stem from the unobserved productivity θ realized in $t = 1$.

Since innate ability is screened for free but productivity is not, Proposition 6 shows that the intertemporal wedge $\tau_0^k(\gamma)$ is characterized by the standard inverse Euler equation for all innate ability types. This is because the government no longer needs the additional intertemporal distortions illustrated in Proposition 1 on $\tau_0^k(\gamma)$ to incentivize investment in human capital. The imaginary allocations are sufficient for that purpose. However, productivity remains unobservable by the government, so savings in $t = 0$ is still restricted and shaped by the inverse Euler equation to relax the ex-post incentive constraints.

More interestingly, Proposition 6 shows that all agents are provided with a commitment device: for any γ and θ, $\frac{u'(c_1(\gamma,\theta))}{u'(c_2(\gamma,\theta))} > \beta$. With non-sophisticated agents, the government can focus on its paternalistic goals since it no longer needs to manipulate retirement savings to screen innate ability. More specifically, notice for any γ, the expected intertemporal distortion is

$$\mathbb{E}\left[\tau_1^k(\gamma,\theta)\right] = (1-\beta)\left[1 - \frac{\mathbb{E}\left[u'(c_1(\gamma,\theta))\right]}{u'(c_0(\gamma))}\right].$$

Since the inverse Euler implies $u'(c_0(\gamma)) < \mathbb{E}[u'(c_1(\gamma,\theta))]$, we have $\mathbb{E}\left[\tau_1^k(\gamma,\theta)\right] < 0$. In essence, agents over-save for retirement in expectation. This is due to the disagreement between the paternalistic government and present-biased agents at $t = 1$. The government uses this disagreement to deter downward deviations in productivity by back-loading the consumption of lower productivity types. At the same time, more productive agents who produce higher output are rewarded with a more front-loaded consumption path. Since it is more cost effective to increase production efficiency by decreasing working-period consumption c_1 of low productivity agents than to increase it for high productivity agents, the consumption path is slightly back-loaded on average. Without this disagreement, it is most efficient to motivate agents to work by respecting their intertemporal preferences and screen through the standard downward labor distortion.

Finally, Proposition 6 shows that the optimal labor distortion is determined solely by the intratemporal component. This means the economic forces that shape the labor wedge are essentially static. Recall from Proposition 2 that both the intertemporal and present-bias components are integral to the optimal provision of dynamic incentives through labor distortion. Since the ex-ante incentive constraint is non-binding, the forces that determine the provision of dynamic incentives are absent from the labor wedge. As a result, the intertemporal and present-bias components no longer influence labor distortion.