Efficient Consolidation of Incentives for Education and Retirement Savings*

Radoslaw Paluszynski† Pei Cheng Yu‡

This version:
July 16, 2021

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 Tevonne Brookes, Gonzalo Castex, Taha Choukhmane, German Cubas, Mikhail Golosov, Hyunju Lee, 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 retirement policies can be used to increase education investment when people are present-biased.

Recently, there have been 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.[1] 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 actual 401(k) contributions.[2] In essence, individuals automatically save for retirement while repaying student loans. Other private employers have since offered similar benefits. Despite the enthusiasm of policymakers, the benefit of conditioning retirement savings on student loan repayments is not apparent. This paper provides a theoretical foundation for the dependence of retirement savings on education investment—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.

Our theoretical framework shows how a commitment device that offsets the agents'

[1] The Retirement Parity for Student Loans Act, the Retirement Security and Savings Act, and the Securing a Strong Retirement Act were introduced in the 116th Congress, which met from Jan 3, 2019 to Jan 3, 2021. These bills allow employer 401(k) matching based on student loan payments.

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

present bias in retirement savings can encourage education investment. This forms a natural interdependence between the optimal retirement savings and education policies. 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 mitigates their present bias. On the other hand, for non-college graduates, the commitment device is not provided indiscriminately. When high innate ability agents are more likely to earn higher income, the optimal retirement savings policy may even exacerbate the present bias of non-college graduates who earn sufficiently high income. This difference in how commitment is provided between college and non-college graduates helps the government screen innate abilities.

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. On the other hand, for non-college graduates, the inverse marginal utility of consumption is *strictly lower* than the working period's expected inverse marginal utility. This implies that, compared to the time-consistent case, consumption is more frontloaded for college graduates and more backloaded for non-college graduates. As a result, the optimal education policy gratifies the high innate ability agents' present bias to encourage them to invest in college.

We also derive the optimal labor wedge for our environment. In contrast to the time-consistent benchmark, distortions during the working period have less impact on education incentives when agents are present biased. Therefore, we show that the labor wedge has an additional economic force that serves to weaken the provision of dynamic incentives through labor distortions. 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.

To decentralize the constrained efficient allocations, we consider an implementation 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 implementation is inspired by the recent IRS ruling and policy proposals in the US Congress that treat student loan repayments as equivalent to salary reduction contributions to retirement accounts. In addition, 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.[3] We show that our theoretical predictions are quantitatively

[3]We use an extended definition of college that includes Master's, Doctoral, and Professional degrees.

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, resulting in low gains from consumption smoothing. 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. We also depart from the optimal policy analysis and examine the quantitative impact of the contribution matching based on student loan repayment in a life-cycle model that resembles "current policies". We find that the proposed reform raises savings rates and improves income redistribution among college graduates, leading to higher welfare in general.

1.1 Related Literature

This paper contributes to the literature on optimal human capital policies. Bovenberg and Jacobs (2005) study 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) focus 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. Koeniger and Prat (2018) show how optimal policy on human capital investment is different from optimal policies on bequests or savings. 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) consider risky but unobservable human capital investment, so tax policies are not conditional upon this investment. In 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, which deviates from their setup of time-consistent agents. This allows us to demonstrate how the provision of commitment can be used to encourage investment in education.

Our paper contributes to the literature on Mirrlees taxation when agents have behavioral biases.[4] Farhi and Gabaix (2019) use sparse maximization (Gabaix, 2014) to study optimal taxation of behavioral agents in a static setting. Lockwood (2020) 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. In contrast to Lockwood (2020), we focus on an environment with dynamic private information. In our setting, the optimal income tax for present-biased agents is quantitatively similar to the one for time-consistent agents. In contrast to the sequential screening environment adopted in this paper, Moser and de Souza e Silva (2019) and Yu (2020a) focus on the design of retirement savings policies for time-inconsistent agents in a Mirrlees setting by examining a multi-dimensional screening environment.

The rest of the paper is organized as follows. Section 2 presents the 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 a policy that decentralizes the optimum and examines the quantitative effects of the policy proposed in the US Congress. Section 6 discusses some extensions.

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 refer to agents with innate ability γ as γ-agents.[5] The share

[4] The recent literature on optimal bequest and estate taxation considers models with altruistic parents and a welfare criteria that also weighs the child directly (Farhi and Werning, 2007, 2010; Pavoni and Yazici, 2017). Such models can also be interpreted as a planner that disagrees with the agents' discount factor.

[5] Innate ability should be thought of as "college readiness" (Athreya and Eberly, 2021) and not necessarily general aptitude measured by standardized tests.

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 and increases more with education for H-agents. This captures the fact that education helps raise human capital, and 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 $f(\theta|\kappa) > 0$ 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 with productivity θ have the following utility at $t=1$:

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

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 $h(0) = 0$. γ-agents have the following utility at $t=0$:

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

Notice that innate ability γ only affects the agents' human capital κ, which influences the productivity distribution they face in the future.

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. 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. Hence, agents who invested in higher-education enter the workforce later and retire later than those who did not. Under our specification, the flow utility and allocations are in annual terms. For example, (c_1, y) is the annual consumption-output bundle in $t = 1$. More details are provided in Section 4.

Crucially, since $\beta < 1$, present-biased agents discount the immediate future more than the distant future. We consider agents who are fully aware of their present bias, i.e., sophisticated agents. As a result, agents in $t = 0$ dislike the fact that their future selves in $t = 1$ under-save for retirement. Section 6.2 considers an economy with non-sophisticated agents.

2.1 Planning Problem

To characterize the constrained efficient allocation, we analyze a direct mechanism—agents report their private information to the government. In Section 5, we will use it as a blueprint to decentralize the optimum as a competitive equilibrium. The government designs

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

Since agents privately learn their innate ability γ and productivity θ sequentially, by the dynamic revelation principle, it is without loss in requiring P to be incentive compatible for each period.[6] Since allocations depend on the reports in a direct mechanism, to simplify notation, we will express the utilities U_0 and U_1 as functions of an agent's reports and type. Let the utility of a type (γ, θ) agent who reports $\theta' \in \Theta$ in $t = 1$ be denoted as

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

The ex-post incentive compatibility constraints ensure the agents report θ truthfully: for any $\theta, \theta' \in \Theta$,

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

By the dynamic revelation principle, the ex-post incentive compatibility constraints (1) only require truth-telling in $t = 1$ after truth-telling in $t = 0$ (Myerson, 1986). Let the utility in $t = 0$ of γ-agents who reported innate ability γ' be denoted as

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

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',

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

[6] See Chapter 11 by Krähmer and Strausz in Borgers (2015) and Bergemann and Välimäki (2019) for an overview of the dynamic revelation principle.

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.

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.

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, so it is differentiable almost everywhere 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. 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, 2020b). In essence, in contrast to the literature with time-consistent agents, ensuring present-biased agents do not misreport as adjacent types is insufficient to guarantee that they do not have incentives to make larger misreports.[7] This paper simplifies the problem by examining the case with two levels of innate ability.

[7] The literature on dynamic mechanism design with time-consistent agents has typically exploited regularity conditions that guarantee the sufficiency of local incentive constraints, such as Courty and Li (2000). However, finding suitable conditions that guarantee the sufficiency of local incentive constraints for present-biased agents is difficult (Galperti, 2015; Yu, 2020b).

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 (2020b) 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 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, independent of the agents' present bias, whether it is optimal for everyone, no one, or only the H-agents to invest in higher education depends on the cost and differential returns to college. We will quantitatively verify that it is indeed optimal for only the H-agents to invest given the calibrated benefits and cost of college in Section 4.4.

2.3 Wedges

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

Following Moser and de Souza e Silva (2019), we define two types of intertemporal wedges. First, we define the *efficiency wedge*, which captures the intertemporal distortions from the government's perspective. The efficiency wedge in $t = 0$ for innate ability γ is

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

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

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

The efficiency wedge helps us identify deviations from the full information efficient outcome, which is characterized by $\tau_t^k = 0$. If $\tau_t^k > 0$ ($\tau_t^k < 0$), then the agent is undersaving (oversaving) in t relative to the efficient outcome.

Second, we define the *decision wedge*, which captures deviations from the agent's Euler equation. The decision wedge in $t = 0$ for innate ability γ is

$$\hat{\tau}_0^k(\gamma) = 1 - \frac{u'(c_0(\gamma))}{\beta\mathbb{E}_\theta\left[u'(c_1(\gamma, \theta))\,|\gamma\right]},$$

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

$$\hat{\tau}_1^k(\gamma,\theta)=1-\frac{u'\left(c_1\left(\gamma,\theta\right)\right)}{\beta u'\left(c_2\left(\gamma,\theta\right)\right)}.$$

The decision wedge provides the implied tax on savings. If $\hat{\tau}_t^k<0$ ($\hat{\tau}_t^k>0$), then it is optimal to introduce a savings subsidy (tax).

From the definitions, the following relationship holds for efficiency wedge τ^k and decision wedge $\hat{\tau}^k: 1-\hat{\tau}_t^k=\frac{1}{\beta}\left(1-\tau_t^k\right)$. A negative efficiency wedge implies a negative decision wedge, so the government needs to subsidize savings. However, the decision wedge is ambiguous when the efficiency wedge is positive. Specifically, if $\tau^k<1-\beta$, then the decision wedge is negative. Otherwise, when $\tau^k\geq 1-\beta$, the decision wedge is weakly positive.

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 $\tau^w>0$ ($\tau^w<0$), then there is an under-supply (over-supply) of labor given the market wage.

2.4 Benchmarks

In this section, we discuss three benchmark cases: (i) time-consistent agents, (ii) observable innate ability, and (iii) when off-path mechanisms are used. In these settings, distortions to retirement savings are not used to incentivize human capital investment.

2.4.1 Time-Consistent Agents

With time-consistent agents ($\beta=1$), the optimal intertemporal distortion at $t=0$ satisfies the standard inverse Euler equation:

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

By Jensen's inequality, the inverse Euler equation implies $u'\left(c_0\left(\gamma\right)\right)<\mathbb{E}_\theta\left[u'\left(c_1\left(\gamma,\theta\right)\right)\right]$ 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.[8] By restricting savings, the government can induce effort in $t=1$ at a lower cost, which relaxes

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

the ex-post incentive constraint. Therefore, the efficiency wedge τ_0^k is strictly positive for any innate ability γ.

Finally, the retirement savings of time-consistent agents are not distorted. The optimal intertemporal decision at $t = 1$ 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, regardless of their past investment in education: $\tau_1^k(\gamma, \theta) = 0$ for all γ and θ. This is because there is no additional uncertainty beyond $t = 1$, so there is no need to distort the intertemporal margin at $t = 1$, in contrast to (3). Hence, the retirement savings policies do not need to depend on education investment.

2.4.2 Observable Innate Ability

If the government observes γ, the optimal efficiency wedge at $t = 0$ is characterized by $c_0(L) = c_0(H)$ and the inverse Euler equation (3). Furthermore, the optimal efficiency wedge at $t = 1$ is, for any γ and θ,

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

where $c_1(L, \theta) \neq c_1(H, \theta)$ due to the difference in productivity distributions for H-agents and L-agents. Since γ is observable, the distortion in retirement savings is not used to encourage education investment. Instead, the government takes advantage of the present bias by backloading the consumption of lower productivity types to deter downward misreports in θ. Notice that by (3), the efficiency wedge at $t = 1$ is negative in expectation, and the decision wedge is negative for all agents, implying a savings subsidy for all agents. Appendix H.3 provides a detailed characterization of the wedges when γ is observable.

2.4.3 Off-Path Mechanisms

With present-biased agents, it may be optimal to introduce mechanisms with off-path threats when the productivity distributions do not span the whole range of Θ (Yu, 2020a,b). To see how, suppose only H-agents can have productivities greater than θ_H where $\underline{\theta} < \theta_H < \overline{\theta}$, so $f(\theta|\kappa_H), f(\theta|\kappa_{L,H}) > 0$ and $f(\theta|\kappa_L) = 0$ for any $\theta \in \left(\theta_H, \overline{\theta}\right]$.[9] In this environment, it is possible to detect misreports on innate ability from some H-agents. Thus, the government

[9]In general, as long as there is a subset of productivities $\hat{\Theta} \subset \Theta$ with positive measure that satisfies $f(\theta|\kappa_L) = 0$ and $f(\theta|\kappa_{L,H}) > 0$ for any $\theta \in \hat{\Theta}$, then it is optimal to introduce off-path threats.

can deter misreporting by punishing agents who are caught lying—those who reported $\gamma = L$ and $\theta \in \left(\theta_H, \overline{\theta}\right]$. Dishonest agents are punished with a more frontloaded consumption path that exacerbates their present bias: high $c_1(L, \theta)$ and low $c_2(L, \theta)$ for any $\theta \in \left(\theta_H, \overline{\theta}\right]$. Present-biased agents are tempted by the frontloaded consumption path in $t = 1$, but want to avoid it in $t = 0$. H-agents know that they run the risk of being punished in $t = 1$ with less retirement consumption if they misreported in $t = 0$, which relaxes the ex-ante incentive constraint. The government also increases $y(L, \theta)$ for any $\theta \in \left(\theta_H, \overline{\theta}\right]$, so that actual L-agents $(\theta \leq \theta_H)$ would not be tempted by the frontloaded consumption path to misreport upwards as $\theta \in \left(\theta_H, \overline{\theta}\right]$. As a result, the set of allocations $\{(c_1(L, \theta), c_2(L, \theta), y(L, \theta))\}_{\theta \in \left(\theta_H, \bar{\theta}\right]}$ is off-path: It only punishes misreporting H-agents.

In this setting, it may even be possible to fully relax the ex-ante incentive constraint using off-path threats, so γ is de facto public information and the allocations in Section 2.4.2 are implemented. As a special case, if utility u is unbounded below and above and productivity θ is a deterministic function of human capital κ, then the full information efficient allocation is implementable (Yu, 2020a). More details are provided in Appendix H.1.

The above mechanism essentially asks the agents to report their innate ability in $t = 0$ and again in $t = 1$, penalizing those whose reports are inconsistent with off-path punishments.[10] In contrast, such punishments may no longer be off path when productivity distributions span the whole range of Θ : Punishments meant to deter H-agents from misreporting might penalize certain L-agents.[11] Therefore, we do not consider the use of off-path mechanisms in our paper's setting.

3 Theoretical Results

In this section, we derive the optimal intertemporal and labor wedges, which provide the foundations for conditioning retirement savings on education investment.

3.1 Intertemporal Wedges

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

[10] In contrast to the standard dynamic revelation principle presented in Myerson (1986), with time-inconsistent agents, it may be optimal for agents to report both new and past information if off-path punishments can be used to penalize only the misreporting agents (Galperti, 2015).

[11] When productivity distributions span the whole range of Θ, the punishment would have to be designed such that both H-agents with any fixed productivity θ who misreported their innate ability and actual L-agents of productivity θ are indifferent between the punishment and the on-path allocation (Amador et al., 2003; Halac and Yared, 2014). However, there is no obvious equilibrium refinement that has the dishonest agents selecting the punishment and the honest agents selecting the on-path allocations.

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

$$\sum_{\gamma} \frac{\pi_{\gamma}}{u'(c_0(\gamma))} = \sum_{\gamma} \pi_{\gamma} \mathbb{E}_{\theta} \left(\frac{1}{u'(c_1(\gamma,\theta))} \middle| \gamma \right), \tag{4}$$

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

$$\mathbb{E}_{\theta} \left(\frac{1}{u'(c_1(\gamma,\theta))} \middle| \gamma \right) = \mathbb{E}_{\theta} \left(\frac{1}{u'(c_2(\gamma,\theta))} \middle| \gamma \right), \tag{5}$$

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

$$\frac{1}{\beta u'(c_2(H,\theta))} = \frac{1}{u'(c_1(H,\theta))} + \left(\frac{1-\beta}{\beta}\right)\left(\frac{\pi_H + \beta\mu}{\pi_H + \mu}\right)\frac{1}{u'(c_0(H))} \tag{6}$$

$$\frac{1}{\beta u'(c_2(L,\theta))} = \frac{1}{u'(c_1(L,\theta))} + \left(\frac{1-\beta}{\beta}\right)\left(\frac{\pi_L - \beta\mu\left(\frac{f(\theta|\kappa_{L,H})}{f(\theta|\kappa_L)}\right)}{\pi_L - \mu}\right)\frac{1}{u'(c_0(L))}, \tag{7}$$

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 (4) and (5).

Let us first discuss the distortions to savings in $t = 0$. If we take expectation of (6) and (7) with respect to θ, then by (5) we can derive the following inverse Euler inequalities:

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

Comparing it with the standard inverse Euler equation (3), the consumption for H-agents is even more frontloaded while the consumption is relatively backloaded for L-agents.[12] In other words, the government caters to the H-agents' preference for immediate gratification to encourage them to accumulate human capital. Furthermore, the less frontloaded 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

[12] 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.

equalized in aggregate, which is implied by (4).

The main feature of our model is that distortions in retirement savings are used to incentivize investment in education. To see this, first notice that by (6), we have

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

Here, the government is rewarding H-agents for going to college with a commitment device that helps them save more for retirement. This commitment device helps substitute part of the information rent to H-agents, because commitment is not guaranteed for agents who did not invest in college. By (7), 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.[13] 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.

Alternatively, from Proposition 1, we can understand the commitment argument from the following: for every $\theta \in \Theta$,

$$\frac{1}{\beta u'\left(c_2\left(H,\theta\right)\right)} - \frac{1}{u'\left(c_1\left(H,\theta\right)\right)} > \frac{1}{\beta u'\left(c_2\left(L,\theta\right)\right)} - \frac{1}{u'\left(c_1\left(L,\theta\right)\right)}. \tag{8}$$

Inequality (8) shows that it is optimal for the government to backload the consumption of H-agents more than L-agents at $t = 1$. The tightness of (8) increases as $\frac{f\left(\theta|\kappa_{L,H}\right)}{f\left(\theta|\kappa_L\right)}$ decreases. In essence, the degree of backloading for L agents increases when the reported productivity is likely from L-agents. In contrast, when innate ability is observable, then (8) holds with equality for all θ.[14] Hence, (8) demonstrates how commitment helps screen innate ability.

[13] When $f\left(\theta|\kappa_L\right) = 0$ and $f\left(\theta|\kappa_{L,H}\right) > 0$ for a strictly positive measure of productivities, then off-path threats can relax the ex-ante incentive constraint. See Section 2.4.3 and Appendix H.1 for details.

[14] See Proposition 6 in Appendix H.3 for details.

Now, suppose f satisfies the monotone likelihood ratio property (MLRP): $\frac{f(\theta|\kappa)}{f(\theta|\kappa')}$ is increasing in θ for any $\kappa > \kappa'$, which implies that higher productivity θ is more likely to come from higher accumulated human capital κ. Then, 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 MLRP 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.

To summarize, the efficiency wedge for H-agents in $t = 0$ is positive, and the sign of L-agents' efficiency wedge in $t = 0$ is unclear. Thus, the decision wedges in $t = 0$ for both innate ability types are ambiguous. For $t = 1$, though the efficiency wedge for H-agents depends on productivity θ, the decision wedge for all H-agents is negative: $\hat{\tau}_1^k (H, \theta) < 0$. Furthermore, when MLRP holds, both the optimal efficiency and decision wedges at $t = 1$ for L-agents increase with productivity. As a result, the government subsidizes the retirement savings of all college-educated agents, but it only subsidizes the retirement savings of high-school graduates who earn low income.

3.2 Labor Wedge

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

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

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

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

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

$$E_\gamma(\theta) = (1-\beta) D_\gamma(\theta),$$

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

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

$$\frac{\tau^w(H,\theta)}{1-\tau^w(H,\theta)} = A_H(\theta) B_H(\theta) \left[C_H(\theta) - D_H(\theta) + E_H(\theta)\right], \tag{9}$$

$$\frac{\tau^w(L,\theta)}{1-\tau^w(L,\theta)} = A_L(\theta) B_L(\theta) \left[C_L(\theta) - \frac{1-F(\theta|\kappa_{L,H})}{1-F(\theta|\kappa_L)} \left[D_L(\theta) - E_L(\theta)\right]\right], \tag{10}$$

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

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 (9) and (10) to pinpoint each component:

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

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

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 $\theta f(\theta|\kappa)$, while 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.

The present-bias component highlights the additional force that influences the labor wedge when agents are present-biased. Since present-biased agents are less sensitive to future incentives, the intertemporal component is less effective in screening the innate ability of present-biased agents than time-consistent agents. The present-bias component captures how this effect weakens the intertemporal component. The total effect of the labor wedge on education incentives is the sum of the intertemporal component and present-bias component, which is βD_γ. In essence, only the portion of the intertemporal component that present-biased agents internalize relaxes the ex-ante incentive constraints, so only a fraction of dynamic incentives enters the labor distortion. To illustrate the logic, consider the extreme example where β is close to zero—agents almost completely ignore future incentives. In this example, the screening of innate ability and productivity are essentially independent. Therefore, the optimal labor distortion is approximately equal to the static case because changes in the labor wedge do little to encourage past incentives for education.

Appendix D quantifies the decomposition of the optimal labor wedge. As we show there, the labor wedge is mostly determined by the productivity distribution and 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.

4.1 Calibration

Table 1 presents the calibrated parameter values. We assume the CRRA utility of consumption, $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$,[15] and the disutility of labor, $h(\ell) = \frac{\ell^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}$. 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 adopted from Nakajima (2012), who makes similar timing assumptions as we do, and calibrates these parameters to achieve a capital-output ratio of 3, an average value for the US economy, in a general equilibrium life-cycle model with present-biased agents. 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 derived from the annual factor 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 factor 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.[16]

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 43 years[17] and then retire and live for 20 years in retirement. The annual cost of higher education is calculated

[15] In our quantitative implementation, we add a small utility shifter to make sure that the expected flow utility for each period and each agent type is non-negative. This matters for educational incentives because the low- and high-type agents have different life expectancy (see below). Hence, without taking a stand on "the value of life", we assume that living is at least weakly better than not-living (which has the flow utility of zero). In practice, the shifter of 0.3 suffices to achieve this goal.

[16] To conduct a sensitivity analysis with respect to the short-term discount factor in Section 4.3, we also consider β of 0.9 and 0.5. In these cases, we adjust the long-term discount factor with linear interpolation using the points provided by Nakajima (2012). The resulting values of δ are 0.9749 and 0.9963, respectively.

[17] This is to match the average retirement age of college graduates based on CPS data for 2010-2016 of around 66 years.

to be \$15,700. Section B.3 in the Appendix discusses the details of our calibration.

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	0.16
$\delta_1(e_H)$	College period 1 long-term discount factor	0.93
δ_2	Retirement discount factor	0.29
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	0.20
$\delta_1(e_H)$	College period 1 long-term discount factor	0.85
δ_2	Retirement discount factor	0.17

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 (in particular, we assume the same cost of college as in our main model), solve for optimal behavior and simulate a 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). These are the four discrete levels of the human capital function κ. We assume specific functional forms for the distributions of skills and select their parameters 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 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.

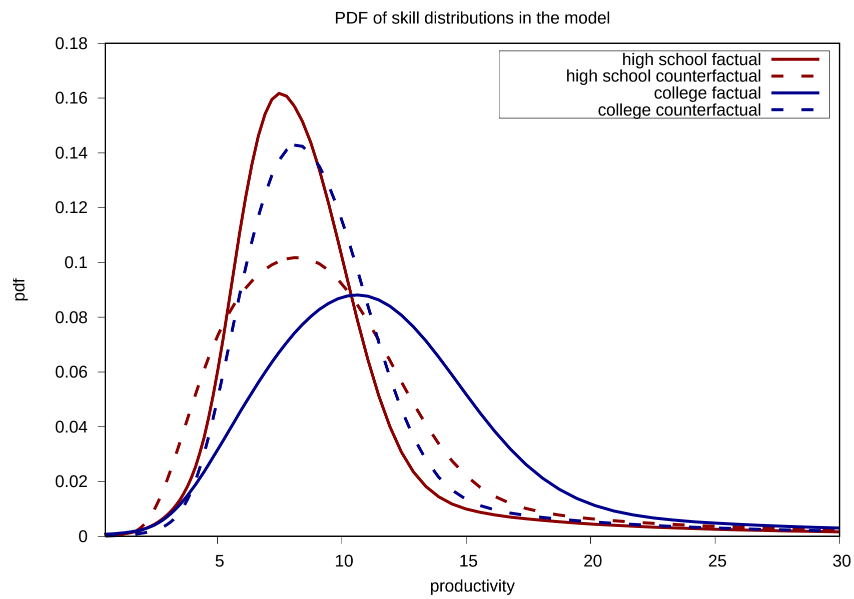

Figure 1: Calibrated distributions of skills for the four groups of agents

It should be emphasized that the quantification of our model is parsimonious and relies on several simplifications. In particular, we assume that the cost of college is equal for all agents (no inheritances or intra-vivo transfers exist), the upper tails of the income distribution across types take the same shape, and agents make their educational choices based exclusively on monetary incentives. At the same time, this parsimony allows us to quantify the main mechanism without losing the clarity of our theoretical analysis in Section 3.

4.2 Optimal Wedges

In what follows, we discuss our quantitative results. We begin with Table 2 which shows the optimal efficiency and decision wedges in $t = 0$. In line with the hallmark dynamic Mirrlees result, the government finds it optimal to restrict savings in $t = 0$ in order to induce higher labor effort from agents in the next period. Notice also that the optimal efficiency 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 efficiency wedge for present-biased agents is slightly higher, raising the consumption of college students and providing additional incentives to make the college investment. Finally, notice that the decision wedge $\hat{\tau}_0^k(H)$ is positive, so it is optimal to introduce a modest savings tax on college students.[18]

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

	Present-biased	Time-consistent
Efficiency wedge $\tau_0^k(H)$	0.33	0.29
Decision wedge $\hat{\tau}_0^k(H)$	0.05	0.29

Figure 2 shows the optimal efficiency and decision wedges in $t = 1$ and conveys a key quantitative result. The efficiency wedges are negative for a wide interval of low incomes, and are always smaller than $1 - \beta$.[19] From the decision wedges, the government introduces a retirement savings subsidy for all agents, and the degree of consumption backloading decreases with income. This is an expected outcome in a model with paternalistic policies and present-biased agents. More importantly, the intertemporal wedges are significantly different for the two education groups. The consumption path of college graduates are more backloaded than high school graduates at all income levels, with the difference eventually disappearing for higher incomes. The government does so in part 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 working-age income across the life cycle. By contrast, notice that in the variant of our model with time-consistent agents, the optimal efficiency and decision wedges in the working-age period are equal to zero for both

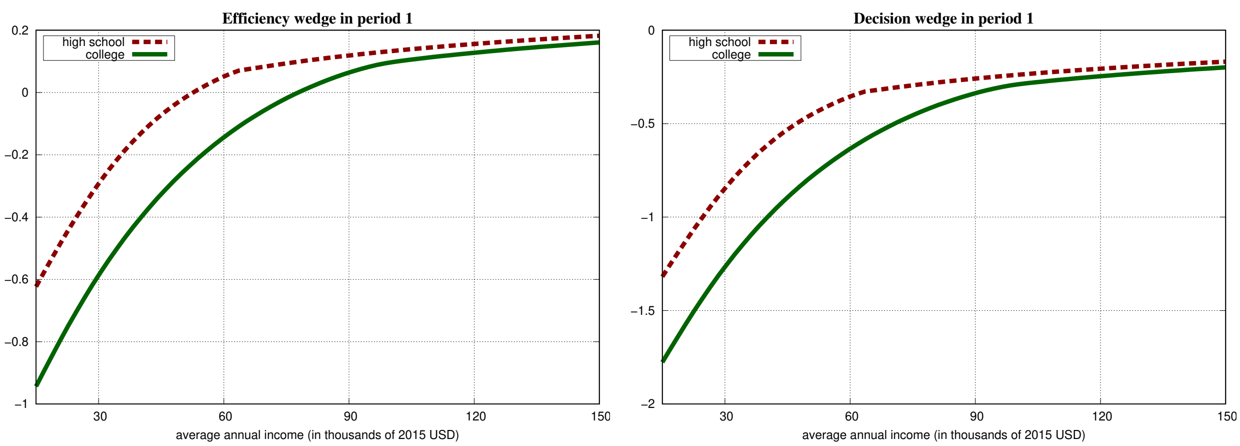

(a) Efficiency wedge (b) Decision wedge

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

[18]The intertemporal wedges for L-agents are not shown since for our quantitative exercise, we assumed L-agents do not have a student period ($\delta_0(e_L) = 0$).

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

education groups. This is because time-consistent agents are able to raise retirement savings on their own.

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 labor wedges mostly decline with income and are significantly different for the two education groups. This resembles the main result of Findeisen and Sachs (2016) which implies 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 labor wedges between the present-biased and time-consistent settings are generally small, and arise predominantly at the lowest incomes. This implies that the presence of present-biased agents may not alter the normative prescriptions in terms of the design of income tax schedules that the literature has established so far. Appendix D 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.

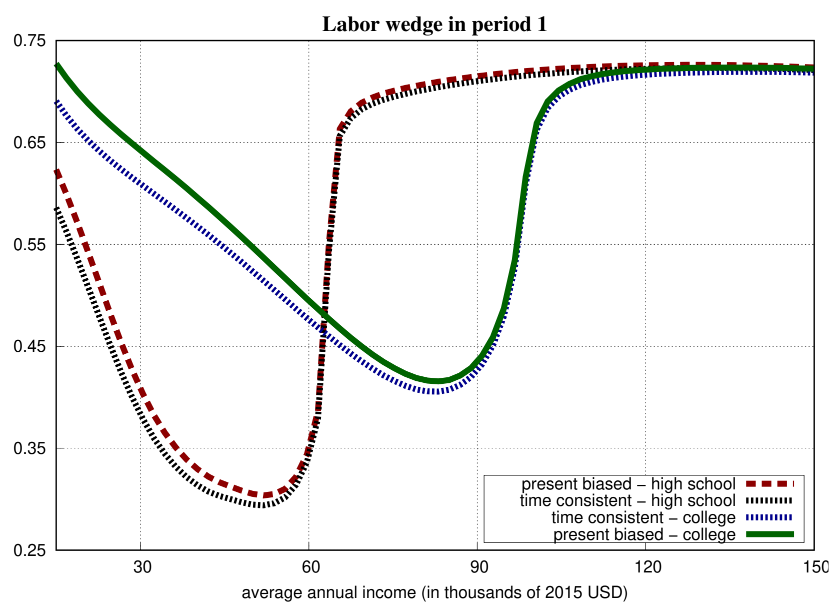

Figure 3: Labor wedge in the model with present-biased agents

Appendix C presents a sensitivity analysis of the efficiency wedge with respect to the main preference parameters, β and σ. In particular, it shows that the wedges for both education groups become steeper with respect to income, the more present-biased and risk-averse the agents are.

4.3 Welfare Gains from Optimal Policies

We now turn our attention to the calculation of potential welfare gains arising from our optimal allocations. We will compare our optimum to three separate benchmarks: optimal policies for time-consistent agents implemented in two ways, as well as the optimum with present-biased agents where the efficiency wedge is restricted to be education-independent.

4.3.1 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.[20] 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 designed 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. Here, agents are forced to smooth their 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, without taking advantage of the education-dependent intertemporal wedge.

Table 3 presents the welfare gains under our baseline parametrization (bold numbers) relative to the two time-consistent benchmarks. In line with our prior expectations, the gains over time-consistent laissez-faire policies are the highest and amount to 1.97% of lifetime consumption. The gains mostly come from increased retirement savings, but also from improved production efficiency. On the other hand, the gains relative to time-consistent policies under mandatory savings are lower, at 1.36% of lifetime consumption, but still

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

	Mandatory savings			Laissez-faire		
	$\sigma = 1.5$	$\sigma = 2$	$\sigma = 2.5$	$\sigma = 1.5$	$\sigma = 2$	$\sigma = 2.5$
$\beta = 0.5$	2.29	3.09	2.90	5.76	5.94	5.33
$\beta = 0.7$	1.00	**1.36**	1.32	1.75	**1.97**	1.84
$\beta = 0.9$	0.27	0.36	0.37	0.32	0.40	0.41

[20] The details of these implementations are presented in Appendix E.

significant. Since the policy of mandatory savings already forces agents to smooth their consumption, this implies the welfare gains of the optimal education-dependent policies largely come from more efficient production. In Table 3, we also conduct a sensitivity analysis with respect to key preference parameters.[21] We find that welfare gains are decreasing in the degree of present bias (since we are getting closer to the time-consistent benchmark) and they are non-monotonic in the degree of risk aversion. This is because there are two forces at play that act in opposite directions. On the one hand, higher risk aversion increases the value of the insurance channel that our mechanism provides, hence increasing potential gains from optimal policies. On the other hand, as we demonstrate in Appendix E.4, higher risk aversion leads to lower efficiency losses from using suboptimal allocations. Hence, the interaction between these two forces leads the overall welfare gains to increase initially, and then decline. Appendix E.4 also shows that most of the difference in welfare gains between the two implementations boils down to laissez-faire agents being unable to smooth consumption over the life cycle.

4.3.2 Welfare Gains from Education-Dependent Savings

We now turn our attention to the benchmark with present-biased agents where the efficiency wedge is restricted to be education-independent. An education-independent wedge is conditioned only on observed income y. Hence, 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

$$\frac{u'\left(c_1(L,\hat{\theta})\right)}{u'\left(c_2(L,\hat{\theta})\right)} = \frac{u'\left(c_1(H,\tilde{\theta})\right)}{u'\left(c_2(H,\tilde{\theta})\right)}. \tag{11}$$

In essence, regardless of education, agents with the same income face an equal decision wedge. Solving for optimal distortions under the set of constraints (11) is non-trivial because these restrictions are contingent on allocations (declared income) rather than the underlying state variable (productivity). We overcome this challenge by designing a computational algorithm, described in Appendix F, which allows us to make the constraints conditional on allocations. Figure 4 presents the optimal efficiency wedge obtained under the set of restrictions (11), along with the education-dependent benchmark.

Table 4 shows welfare gains measured as a corresponding percentage increase in lifetime consumption that would result from moving from the system with an education-independent efficiency wedge to the optimum (where it depends on educational attainment). Under the

[21] When varying the degree of present bias, β, we simultaneously adjust the long term discount factor, as described in Section 4.1.

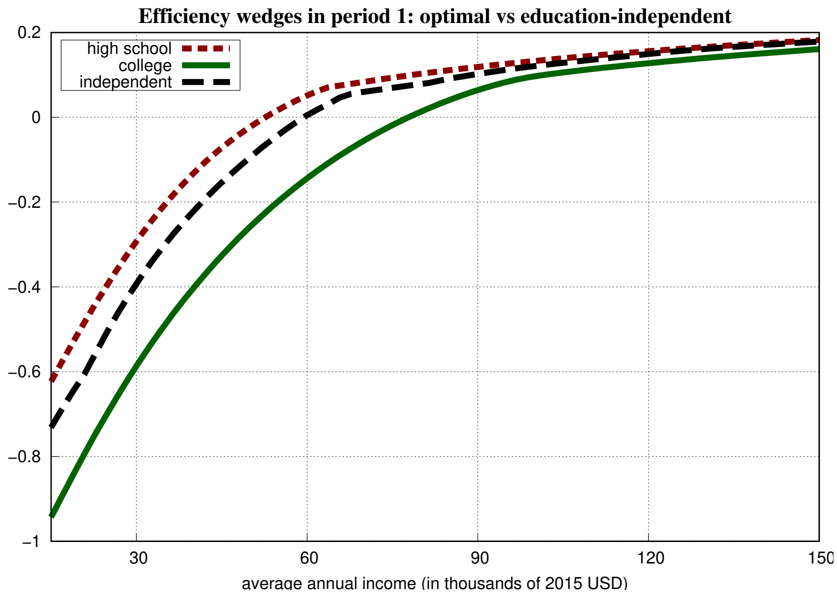

Figure 4: Optimal education-dependent and independent efficiency wedge

baseline parametrization (bold numbers), the corresponding gain in lifetime consumption amounts to 0.02%. Table 4 also conducts a sensitivity analysis of this result with respect to key preference parameters of the model—the degree of risk aversion σ and the short-term discount factor β. As is clear from the table, welfare gains increase in the degree of present bias and the degree of risk aversion.

Table 4: Welfare gains over optimal education-independent savings policies

	$\sigma = 1.5$	$\sigma = 2.0$	$\sigma = 2.5$
$\beta = 0.5$	0.0332	0.0684	0.1031
$\beta = 0.7$	0.0103	**0.0167**	0.0261
$\beta = 0.9$	0.0010	0.0016	0.0019

The fairly small welfare gain that we obtain in the baseline parametrization deserves a comment. First, 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) shows 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. Similarly, Aguiar and Gopinath (2006) show that the threat of financial autarky (and the resulting lack of consumption smoothing) is trivial for borrowing countries and, hence, no realistic amounts of sovereign debt can be sustained in equilibrium without additional sources of default punishment. Second, our sensitivity analysis indicates that this number can be elevated significantly under alternative calibrations. In particular, for a short-term discount factor of 0.5, and risk aversion of 2.5, the welfare gain is equivalent to more than 0.1% of lifetime consumption. Such parameter values are not empirically implausible as evidenced by the latest estimates of Laibson et al. (2017).

4.4 Testing Policies Without Screening

As a final step in our quantitative analysis of the model, we test whether incentivizing only H-agents to attend college is indeed preferable quantitatively to other alternatives. In particular, we calculate the government's value derived under the policy that all agents get higher education, only L-agents receive it, and one where no agents do.[22]

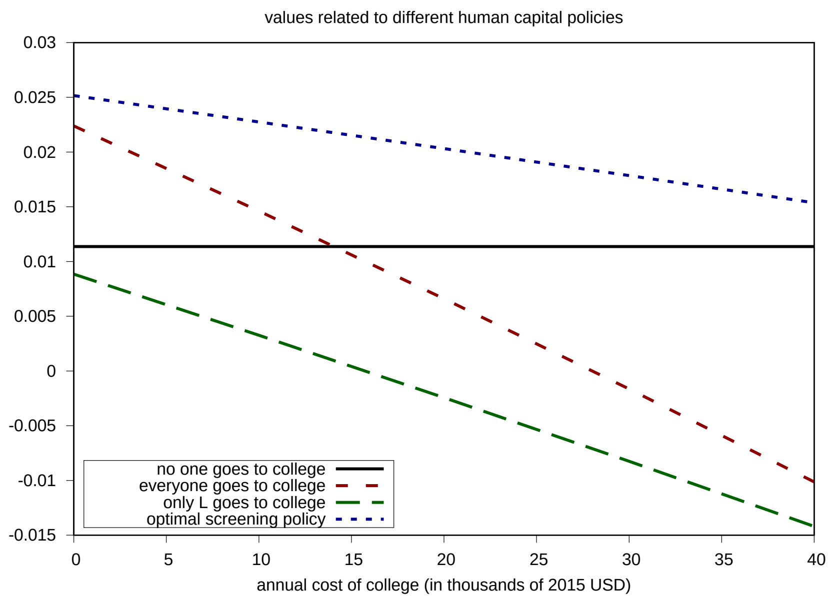

Figure 5: Comparing optimal policy to alternative screening policies

Figure 5 presents the values associated with these alternative policies, along with the optimal screening one. The values are depicted as function of the annual monetary cost of higher education, ranging from zero up to 40,000 USD (the actual calibrated cost, as Table 1 shows, is 15,700 USD). It can immediately be noticed that the optimal Mirrleesian policy dominates the alternatives at all cost values, including when college is free. This is due to the fact that going to college and beyond entails a significant time cost while the expected return to L-agents remains small. It is also worth noticing that for realistic levels of the calibrated cost, sending no one to college weakly dominates the alternative of sending everyone to college, or sending L-agents only.

[22]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 $\delta_1(e)$.

5 Implementation

In this section, we discuss the implications of our findings for the design of student loans, income taxes, and retirement policies. In particular, this section highlights how to decentralize policies where retirement savings can help incentivize education investment, which is the main innovation of the paper. We also provide a quantitative analysis of the policy proposal discussed in the US Congress.

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 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, the Retirement Security and Savings Act, and the Securing a Strong Retirement Act—which intend to qualify student loan repayments for employer matching.[23]

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.

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 Student Loan Payment as Contribution to Retirement Savings

In this section, we consider an implementation with social security benefits and retirement savings accounts that depend on student loan repayments. The advantage of this

[23] We also consider an alternative implementation in Appendix G where the subsidy for retirement savings is both income and education contingent.

decentralization is that it adopts the main features of existing retirement policies. Furthermore, it demonstrates how the retirement bills proposed in the US Congress could be used to implement the optimum.

Agents are offered a student loan $L(e)$ in $t=0$. Agents face an income tax $T(y)$ in $t=1$ that is independent of education. The student loan repayment $r(e,y)$ is tax deductible and reduces income tax by $g(r)$. In each period, agents can save via the risk-free bond b, which are taxed with a history-independent bond savings tax $T^k(b)$.[24]

For the retirement policies, similar to the current system, all agents receive an income-contingent social security benefit $a(y)$ upon retirement. The retirement savings account is defined by the contribution matching rate $\alpha \in [0,1]$ and a contribution limit $\bar{c}$. Retirement account contributions come from pre-tax income (similar to a traditional 401(k)) and are only lump-sum taxed T^{ra} upon withdrawal. Furthermore, similar to current retirement savings accounts, matched contributions are not subject to the contribution limit $\bar{c}$. The novelty of this implementation is that the amount of student loan repaid $r(e,y)$ is considered a contribution, so employers can further contribute $\alpha r(e,y)$ into the account. Let $\omega(s_2,r)$ 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 $\omega(s_2,r)=(1+\alpha)s_2+\alpha r$.

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

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

subject to

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

$$c_2=a(y)+R_2\omega(s_2,r(e,y))+R_2b_2-\mathbf{1}_{\omega>0}T^{ra},$$

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

where $\mathbf{1}_{\omega>0}$ is an indicator function with $\mathbf{1}_{\omega>0}=1$ if and only if there are assets in the account, otherwise $\mathbf{1}_{\omega>0}=0$. Also, $\tilde{R}_1(e)=\frac{R_1(e)}{R_0(e)}$ 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

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

with innate ability γ at $t=0$ is

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

subject to

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

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

Proposition 3 *The optimum can be implemented through P^{ra} where student loan repayments are considered contributions to the retirement savings account.*

Under P^{ra}, both college and high-school graduates face the same income tax and social security policy, while the repayment schedule and corresponding tax deduction generate the different incentives for college and non-college graduates. What is significant is that P^{ra} uses student loan repayments as a retirement savings vehicle for college graduates. At the heart of this implementation is the idea that college graduates can save for retirement while paying off their student loans. Specifically, we construct the social security benefits to match the optimal retirement consumption of high-school graduates. Since college graduates are essentially saving for retirement when they repay their student loans, $\alpha r(e,y)$—the amount of repayment that is being matched—is designed to supplement the social security benefits so college graduates can consume the optimum during retirement. On the other hand, college graduates are poorer at $t=1$ when they repay their student loans, so the tax deduction $g(r(e,y))$ is constructed to ensure they consume the optimum during the working period.

Figure 6 presents the student loan repayment schedule in our implementation. The solid green line shows the face value of the repayment schedule, $r(e,y)$, which starts high and then decreases initially. This allows low-income college-graduates to accumulate additional (and decreasing in annual income) contributions in their 401(k) plans through the match from student loan repayment. The effective loan repayment schedule $r(e,y) - g(r(e,y))$ is represented by the dashed red line as a function of annual income. Notice that the effective repayment schedule increases in income until the Pareto tail for high-school graduates kicks in. Also, except for mid-income agents—who constitute the majority of all agents—higher education is relatively cheap. This implies that the high repayment in face value for low-income agents is mainly for the purpose of increasing retirement savings. It is worth mentioning that the contribution matching rate α that arises in our proposed implementa-

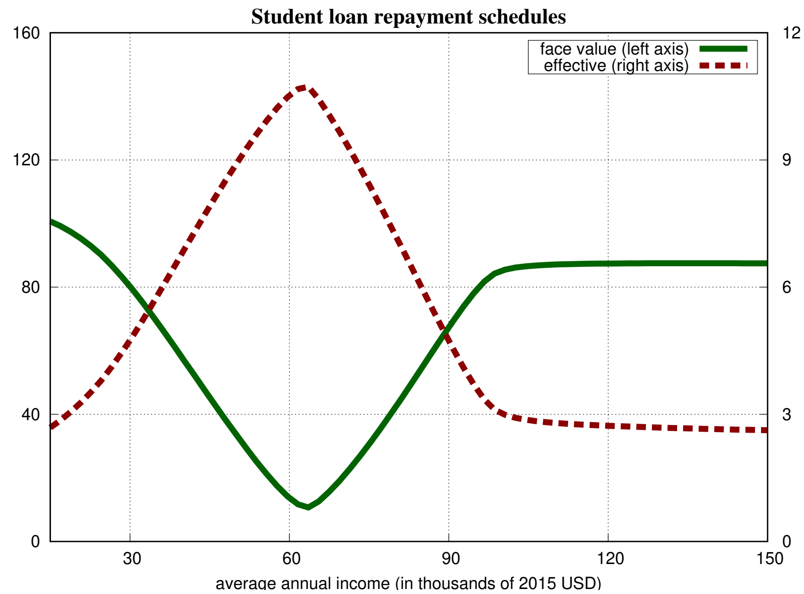

Figure 6: Optimal student loan repayment schedule

tion amounts to 2.00%, which is in the ballpark of the actual rate used by the IRS ruling from May 2018.

5.2 Quantitative Analysis of Reform Proposed by US Congress

This section considers the quantitative impact of the policy reform recently discussed by the US Congress.[25] We deviate from our optimal policy framework and work with the "current policies" life-cycle model, which was developed in Section 4 for the purpose of inferring the productivity distributions. While the "current policies" model plays an auxiliary role in our paper, it is nevertheless instructive to use it to examine its implications for the proposed reform of employer matching based on student loan repayment. Appendix B.2 explains how we incorporate this possibility in the "current policies" model. We introduce the reform in a revenue-neutral way—net revenue remains the same as in the "current policies" model—by simultaneously increasing income taxes on all agents.

Figure 7(a) summarizes the differences in savings between the two variants of the "current policies" model, along with the optimal Mirrlees framework,[26] by plotting the retirement savings rates of college graduates, defined as the ratio $\frac{c_2/R_2}{c_1+c_2/R_2}$, as a function of annual income. Under current policies (without the proposed reform), the savings rate is high for the lowest incomes and then drops fast as agents start actively saving in 401(k) plans ($s_2 > 0$).

[25] We are referring to the proposal that would allow for employer 401(k) matching based on student loan payments, which was included in three recent pieces of legislation, namely the Retirement Parity for Student Loans Act, the Retirement Security and Savings Act, and the Securing a Strong Retirement Act.

[26] To make the models comparable, we re-solve for optimal policies by imposing the same resource constraint imbalance as the one implied by the "current policies" world.

At annual income of around \$200,000, agents' individual savings hit the contributions limit, and agents with annual income just below \$250,000 start holding regular savings ($b_2 > 0$). The savings rate then stabilizes at around 20% which is very close to the full information efficient rate of 22.3% (and mostly aligns with the constrained efficient rate from our optimal model). In contrast, notice the savings rate under the proposed policy achieves this level for a wider interval of incomes, starting at around \$150,000 — the income level when agents choose $s_2 > \bar{c} - i$ (i denotes the annual repayment from traditional non-income contingent student loans). In essence, they forgo a part of the 401(k) matching stemming from student loan repayment to receive more matching on their own deposits s_2. Finally, notice that student loan repayments are independent of income under the proposed reform. As a result, due to the boost in retirement savings from matching on repayments, consumption for low-income college graduates is more backloaded relative to the optimal policies.

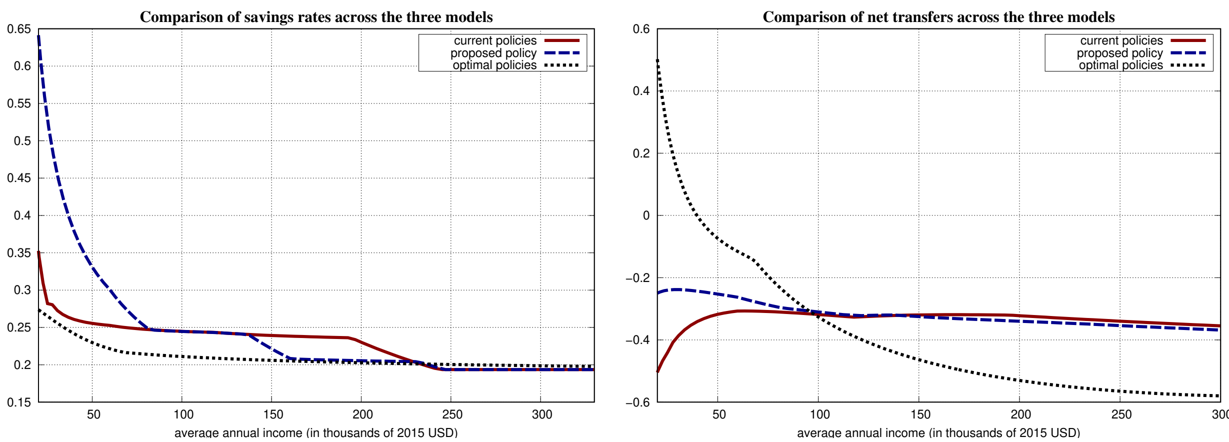

(a) Savings rates

(b) Net transfers

Figure 7: Savings rates and transfers of college graduates across the three models

Figure 7(b) plots the net transfers, defined as $\frac{c_1+c_2/R_2-y}{y}$, for college graduates at different income levels. In the "current policies" model, all college graduates are net contributors, and especially so at the lowest income levels. This contrasts sharply with the optimal Mirrleesian allocations which redistribute resources towards agents with low income and away from agents with high income. As is evident from Figure 7(b), the proposed reform partially achieves this pattern of redistribution. As a result, the policy reform can potentially improve the redistribution of income among college graduates.

To evaluate the welfare implications of the proposed policy, we calculate the percentage gain in lifetime consumption of all agents in the pre-reform economy that produces an aggregate welfare equal to the post-reform one.[27] We find that the reform is welfare-improving,

[27] Notice that here, in contrast to our previous exercises in Section 4.3, the welfare function is that of the

equivalent to a gain in lifetime consumption of 0.18%. Despite the fact that the "current policies" model is very different from our main model, the exercise in this section shows how education-dependent retirement policies can raise welfare in general. On the other hand, moving from "current policies" to the optimal Mirrleesian world yields a substantial welfare gain of 2.04% of lifetime consumption.

6 Extensions

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.[28] Proposition 4 characterizes the distortions and shows that retirement policies are still used to increase education investment when the degree of present bias is heterogeneous.

Proposition 4 *The constrained efficient allocation with heterogeneous present bias satisfies*

i. the inverse Euler equations (4), (5) 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 $\mu = \left[u'\left(c_0\left(L\right)\right) - u'\left(c_0\left(H\right)\right)\right]\left[\frac{u'\left(c_0(L)\right)}{\pi_L} + \frac{u'\left(c_0(H)\right)}{\pi_H}\right]^{-1}$.

ii. the labor wedge for H-agents satisfies (9) and for L-agents:

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

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

present-biased agents in period $t = 0$, not the one of a paternalistic government.

[28] 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.

Though the economic forces determining the wedges for H-agents remain unchanged, Proposition 4 shows us how the optimal policy leverages the difference in β for the L-agents' wedges. For the efficiency wedge $\tau_1^k(L,\theta)$, recall that the optimal policy recommends frontloading consumption for high-income L-agents. Here, this frontloading 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 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 4, 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 encouraging education investment through retirement savings policies 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 (2020a), 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 frontloaded 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 (2020a) 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 H.2.

Off-path policies are powerful, but the optimal policy should still feature the interdependence between retirement savings and education investment discussed in this paper. 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 sufficiently 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 policies 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 be able encourage them to increase investment in education.

6.3 Non-Paternalism

So far, this paper has assumed that the government is paternalistic, i.e., its own preferences over the agents' welfare are time-consistent. In this section, we depart from this assumption by allowing the government to adopt the agents' own present bias when choosing optimal allocations. In doing so, a natural question is whether the government is present-biased only at $t = 0$ or also at $t = 1$. To consider both possibilities, we assume that the government has the objective function given by

$$\sum_{\gamma} \pi_{\gamma} \left\{ \delta_0 \left(e_{\gamma}\right) u \left(c_0 \left(\gamma\right)\right) + \beta \delta_1 \left(e_{\gamma}\right) \int_{\Theta} \left[\chi \hat{U}_1 \left(c_1, c_2, y; \theta\right) + (1-\chi) U_1 \left(c_1, c_2, y; \theta\right) \right] f \left(\theta | \kappa_{\gamma}\right) d\theta \right\}$$

where $\hat{U}_1 \left(c_1, c_2, y; \theta\right) = u \left(c_1\right) - h \left(\frac{y}{\theta}\right) + \delta_2 u \left(c_2\right)$ and $U_1 \left(c_1, c_2, y; \theta\right) = u \left(c_1\right) - h \left(\frac{y}{\theta}\right) + \beta \delta_2 u \left(c_2\right)$. In essence, by setting the parameter χ to a value smaller than one, we allow the government to put some weight on the present-biased agent's preferences at $t = 1$.

Figure 8 presents the optimal efficiency and labor wedges for three alternative cases: i) the baseline paternalistic government; ii) a non-paternalistic government who adopts the $t = 0$ agents' preferences; and iii) a non-paternalistic government who puts $\chi = 0.5$ weight on the $t = 1$ agents' preferences. A few observations are noteworthy. First, as evident in panel 8(a), the non-paternalistic allocations where the entire weight is put on the $t = 0$ agent are essentially the same as the allocations of a paternalistic government. By contrast, when the non-paternalistic government adopts the $t = 1$ agents' preferences, the optimal wedges are smaller in absolute value and involve more frontloading of consumption, a result that aligns with basic intuition: A non-paternalistic government that puts weight on the agents' $t = 1$ preferences feels less need to help the agents save at period $t = 1$ than our baseline paternalistic model. Finally, panel 8(b) shows that the labor wedges are virtually unaffected by any non-paternalistic considerations.

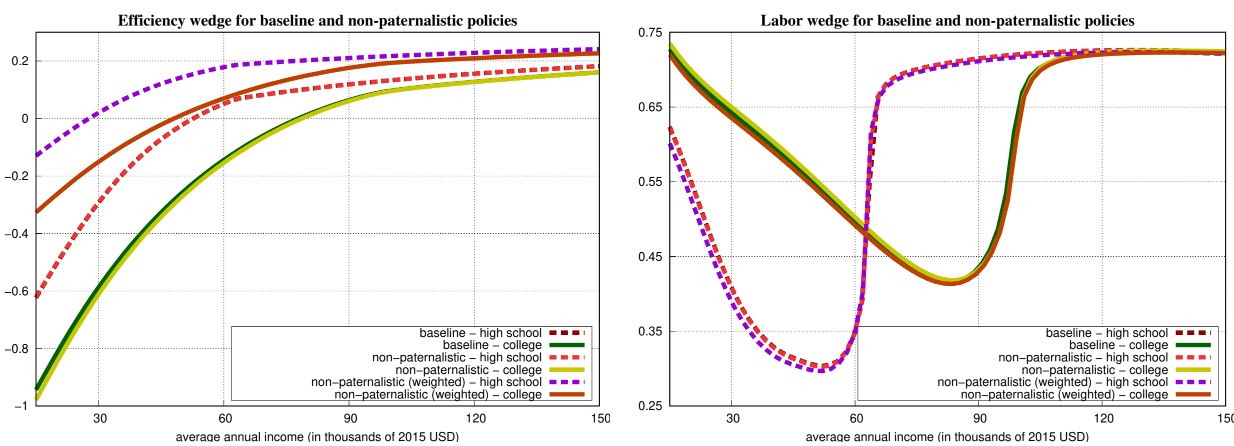

(a) Efficiency wedge

(b) Labor wedge

Figure 8: Optimal wedges with a non-paternalistic government

6.4 Alternative Assumptions on Timing

6.4.1 Length of Period

A possible concern with our three-period model is that agents make a one-time retirement savings decision upon entering the labor force, right after the education period and decades before retiring. In contrast, individuals can continuously save for their retirement in the real world. This could possibly weaken the positive incentive effects of retirement policies on education investment.

We examine a model with a coarse timing for three reasons. First, there is evidence that individuals exhibit inertia in retirement savings decisions (Madrian and Shea, 2001). In particular, evidence suggests that once individuals make a decision on their pension portfolio few ever revisit the decision (Cronqvist et al., 2018). The current timing could be interpreted as modeling the inertia exogenously and highlighting the importance of the agents' initial savings decisions, which could have large ramifications for their retirement welfare. Second, the length of these periods is consistent with other papers on retirement policies that examine a Mirrlees taxation model with present-biased agents, such as Moser and de Souza e Silva (2019) and Yu (2020a). Furthermore, in Moser and de Souza e Silva (2019), the quasi-hyperbolic discounting model is used in a reduced-form way to capture frictions that are not necessarily behavioral. Finally, as discussed in Section 2.2, there are considerable technical difficulties to adopting shorter period lengths with privately informed present-biased agents.

To provide an intuition on how our results may change with finer time periods, in Appendix I, we ignore the technical issues and analyze a four-period model where the working period is split in two. In the second working period, agents draw a new productivity from a distribution that depends on human capital and past productivity. The results of the four-period model align with the main message of our paper: The retirement policies for present-biased agents can help incentivize investments in education.

6.4.2 Difference in Length of Education Periods

Our baseline parametrization assumes a difference in length of the education period of around 5 years. In essence, the college graduates receive education before working, while high school graduates enter the workforce immediately. To show that this assumption is not crucial for the results of our paper, we re-solve the model under the assumption that high school graduates spend an equal amount of time in the initial period (without getting any training). Hence, the life cycles of the two types of agents are perfectly synchronized. Figure 9 presents the impact of lifting this assumption on our model by comparing the resulting two wedges to the benchmark ones. The change results in minor shifts of both wedges for both education groups. The efficiency wedge moves downwards, which implies a higher savings subsidy for both groups. Crucially, college graduates are still subsidized more than high school graduates, by a similar margin as in the baseline. The labor wedge moves upwards for the lowest and highest incomes, and downwards for the middle range, but any differences relative to the baseline are small.

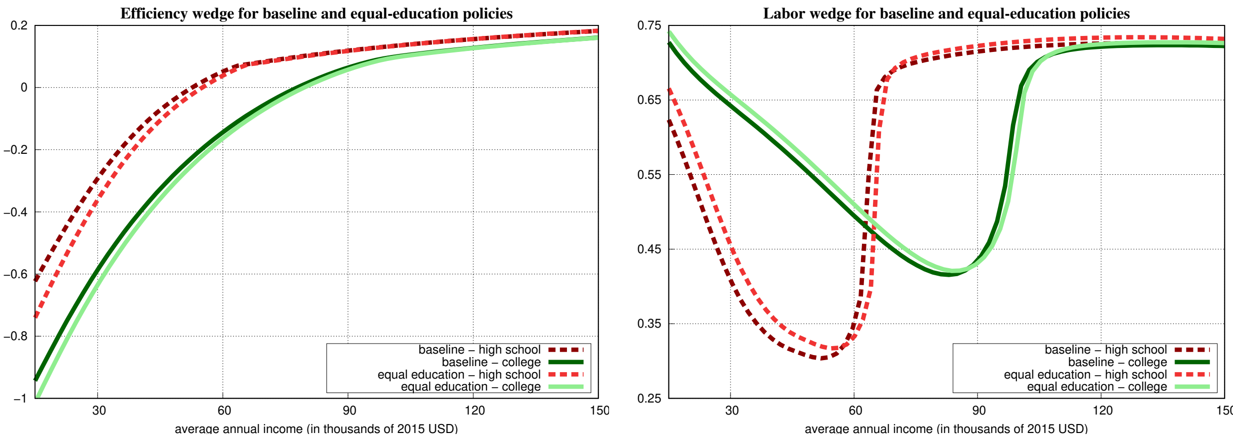

(a) Efficiency wedge

(b) Labor wedge

Figure 9: Optimal wedges with equal education period lengths

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 incentivizes education. Specifically, we show how conditioning retirement savings on student loan repayments, along with some qualitative changes to existing policies, can implement the optimum. We quantify the welfare gains from these policies, and also show that the inverse Euler equation does not hold with present-biased agents, while the labor wedge is quantitatively similar to the case with time-consistent agents.

This paper focuses on the question of how best to design policies for present-biased individuals financing their own education. One potential avenue for future research is to consider parental contributions to human capital investment. With an overlapping generations model, we can potentially analyze a setting where altruistic parents invest in their offspring's education, with both suffering from present bias. Such a richer model may pave the way to a study of optimal college savings policies—such as the 529 plan in the US—for parents.

References

Aguiar, Mark and Gita Gopinath, "Defaultable Debt, Interest Rates, and the Current Account," *Journal of International Economics*, 2006, *69* (1), 64–83.

Amador, Manuel, Ivan Werning, and George-Marios Angeletos, "Commitment vs. Flexibility," *NBER Working Paper 10151*, 2003.

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.

Athreya, Kartik and Janice Eberly, "Risk, the College Premium, and Aggregate Human Capital Investment," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 168–213.

Bergemann, Dirk and Juuso Välimäki, "Dynamic Mechanism Design: An Introduction," *Journal of Economic Literature*, 2019, *57* (2), 235–274.

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

Borgers, Tilman, "An Introduction to the Theory of Mechanism Design," *Oxford University Press*, 2015.

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.

Courty, Pascal and Hao Li, "Sequential Screening," *Review of Economic Studies*, 2000, *67.*

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

Cronqvist, Henrik, Richard Thaler, and Frank Yu, "When Nudges are Forever: Inertia in the Swedish Premium Pension Plan," *AEA Papers and Proceedings*, 2018, *108*, 153–158.

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 Ivan Werning, "Inequality and Social Discounting," *Journal of*

Political Economy, 2007, *115* (3).

— **and** — , "Progressive Estate Taxation," *Quarterly Journal of Economics*, 2010, *125* (2).

— **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.

Koeniger, Winfried and Julien Prat, "Human Capital and Optimal Redistribution," *Review of Economic Dynamics*, 2018, *27.*

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," *American Economic Journal: Economic Policy*, 2020.

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

Madrian, Brigitte and Dennis Shea, "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics*, 2001, *116* (4), 1149–1187.

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.

Milgrom, Paul and Ilya Segal, "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 2002, *70* (2).

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.

Pavoni, Nicola and Hakki Yazici, "Intergenerational Disagreement and Optimal Taxation of Parental Transfers," *Review of Economic Studies*, 2017, *84* (3).

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," *Journal of European Economic Association*, 2020.

__ , "Seemingly Exploitative Contracts," *Journal of Economic Behavior and Organization*, 2020.

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

$$U_1 \left(\gamma, \theta\right) = u \left(c_1 \left(\gamma, \theta\right)\right) - h \left(\frac{y \left(\gamma, \theta\right)}{\theta} \right) + \beta \delta_2 u \left(c_2 \left(\gamma, \theta\right)\right), \tag{12}$$

$$\frac{\partial U_1 \left(\gamma, \theta\right)}{\partial \theta} = \frac{y \left(\gamma, \theta\right)}{\theta^2} h' \left(\frac{y \left(\gamma, \theta\right)}{\theta} \right), \tag{13}$$

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

and the resource constraint. As is standard, we ignore the monotonicity constraint—$y(\gamma, \theta)$ 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 $(\lambda_{\gamma}(\theta), \xi_{\gamma}(\theta), \mu, \phi)$ be the multipliers on (12), (13), 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: $\xi_\gamma(\underline{\theta}) = \xi_\gamma(\overline{\theta}) = 0$, and

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

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

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

A.2 Proofs

Proof of Lemma 1: First, to prove necessity, we have to show that Parts (i.) and (ii.) follow from incentive compatibility. To prove Part (i.), let $\Phi(\gamma,\theta) = u(c_1(\gamma,\theta)) + \beta\delta_2 u(c_2(\gamma,\theta))$. For a fixed γ and productivities θ and θ' with $\theta > \theta'$, incentive compatibility requires

$$\Phi(\gamma,\theta) - h\left(\frac{y(\gamma,\theta)}{\theta}\right) \geq \Phi(\gamma,\theta') - h\left(\frac{y(\gamma,\theta')}{\theta}\right)$$

and

$$\Phi(\gamma,\theta') - h\left(\frac{y(\gamma,\theta')}{\theta'}\right) \geq \Phi(\gamma,\theta) - h\left(\frac{y(\gamma,\theta)}{\theta'}\right).$$

Adding these two inequalities yields

$$h\left(\frac{y(\gamma,\theta)}{\theta'}\right) - h\left(\frac{y(\gamma,\theta)}{\theta}\right) \geq h\left(\frac{y(\gamma,\theta')}{\theta'}\right) - h\left(\frac{y(\gamma,\theta')}{\theta}\right). \tag{14}$$

Since h is strictly increasing and strictly convex, for (14) to hold, it must be the case that $y(\gamma,\theta) \geq y(\gamma,\theta')$.

Next, we will prove Part (ii.). For a fixed γ and θ, the incentive constraint is

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

Theorem 2 of Milgrom and Segal (2002) applies, so $U_1(\gamma,\theta)$ is absolutely continuous in θ. Finally, since $U_1(\gamma,\theta)$ is absolutely continuous in θ, it is differentiable in θ almost everywhere. Hence, we have $\frac{\partial U_1(\gamma,\theta)}{\partial\theta} = \frac{y(\gamma,\theta)}{\theta^2}h'\left(\frac{y(\gamma,\theta)}{\theta}\right)$.

Finally, to prove sufficiency, we have to show that, for a given γ, none of the agents would want to misreport θ when Parts (i.) and (ii.) hold. Consider any productivities θ and θ'

with $\theta > \theta'$, then

$$\int_{\theta'}^{\theta} \frac{y(\gamma, x)}{x^2} h'\left(\frac{y(\gamma, x)}{x}\right) dx \geq \int_{\theta'}^{\theta} \frac{y(\gamma, \theta')}{x^2} h'\left(\frac{y(\gamma, \theta')}{x}\right) dx$$
$$\iff U_1(\gamma, \theta) - U_1(\gamma, \theta') \geq -h\left(\frac{y(\gamma, \theta')}{\theta}\right) + h\left(\frac{y(\gamma, \theta')}{\theta'}\right)$$
$$\iff U_1(\gamma, \theta) \geq \Phi(\gamma, \theta') - h\left(\frac{y(\gamma, \theta')}{\theta}\right).$$

The first inequality comes from Part (i.), $y(\gamma, \theta)$ is non-decreasing in θ, and h is strictly convex. The left-hand side of the second inequality comes from Part (ii.). Rearranging the terms in the second inequality yields the third inequality, which implies incentive compatibility. ■

Proof of Proposition 1: Conditions (6) and (7) and μ follow from the first order conditions. The inverse Euler equations (4) and (5) are derived using the perturbation argument.

Let $P = \left\{c_0(\gamma), [c_t(\gamma, \theta), y(\gamma, \theta)]_{t>0, \theta\in\Theta}\right\}_{\gamma}$ be the allocation that solves the constrained efficient planning problem. We first derive (5) by considering a small increase in $c_2(\gamma, \theta)$ across θ for a fixed γ. That is, for all θ, define $u(\tilde{c}_2(\gamma, \theta)) = u(c_2(\gamma, \theta)) + \Delta$ for some small Δ. We simultaneously decrease $c_1(\gamma, \theta)$ for all θ such that $u(\tilde{c}_1(\gamma, \theta)) = u(c_1(\gamma, \theta)) - \delta_2 \Delta$. 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.,

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

Evaluating the first order condition of this problem at $\Delta = 0$ yields (5).

Similarly, to derive (4), we consider a small decrease in $c_1(\gamma, \theta)$ for all θ and γ such that $u(\tilde{c}_1(\gamma, \theta)) = u(c_1(\gamma, \theta)) - \frac{1}{\delta_1(e_\gamma)}\Delta$ for some small Δ. We simultaneously increase $c_0(\gamma)$ for all γ such that $u(\tilde{c}_0(\gamma)) = u(c_0(\gamma)) + \frac{1}{\delta_0(e_\gamma)}\Delta$. 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 (4). ■

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

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

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

First, we derive (9). Since $\lambda_\gamma\left(\theta\right) = \frac{\phi\pi_\gamma\delta_1(e_\gamma)f(\theta|\kappa_\gamma)}{u'(c_1(\gamma,\theta))}$, we rewrite the first order condition on $y\left(H,\theta\right)$ as

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

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 $\frac{1}{\theta}h'\left(\frac{y(H,\theta)}{\theta}\right)\phi\pi_H\delta_1\left(e_H\right) f\left(\theta|\kappa_H\right)$ yields

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

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

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

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 $\frac{\mu}{\phi\pi_H} = \frac{1}{u'(c_0(H))} - \frac{1}{\phi}$. 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 $E_\gamma(\theta) = (1-\beta)\, D_\gamma(\theta)$, then multiply both sides by $u'(c_1(H,\theta))$ to yield (9).

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

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

Since $\frac{\beta\mu}{\phi\pi_L} = \frac{\mu}{\phi\pi_L} - \frac{(1-\beta)\mu}{\phi\pi_L}$ and from the first order conditions, we get (10). 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 (4) 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 (because $\Phi\left(\gamma,\tilde{\theta}\right) = \Phi\left(\gamma,\hat{\theta}\right)$). 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: By Lemma 2, we can define the optimal consumption derived from the direct mechanism as $(c_0(e), c_1(e,y), c_2(e,y))$. Next, following Werning (2011), we construct bond savings tax $T^k(b)$ such that agents do not double deviate—misreport and buy too much 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

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

where

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

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

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

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

For the other 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

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

Next, we construct the social security benefit $a\left(y\right)=c_2\left(e_L,y\right).$ We set the income tax to be $T\left(y-s_2\right)=y-s_2-c_1\left(e_L,y\right),$ and the tax deduction from student loan repayment is

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

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

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

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

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

We choose T^{ra} such that $r\left(e_H,y\right)$ and $g\left(R_1r\right)$ are weakly positive. Let $T^{ra}\left(y\right)$ be a fictitious tax schedule defined as

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

Observe that given $T^{ra}\left(y\right),$ 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:

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

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 $c_1=c_1\left(e_L,y\right)$ and $c_2=c_2\left(e_L,y\right).$ Next, for agents who invested e_H, given the taxes, $c_1=y-T(y)+g\left(r\left(e_H,y\right)\right)-r\left(e_H,y\right)$ and $c_2=a\left(y\right)+\alpha R_2r\left(e_H,y\right)-T^{ra}$, so they optimally choose $c_1=c_1\left(e_H,y\right)$ and $c_2=c_2\left(e_H,y\right).$ Also, by the taxation principle, agents with productivity θ choose $y=y\left(e,\theta\right).$ Finally, notice that given $L(e)$, agents with innate ability γ optimally choose education level e_γ. ■

Proof of Proposition 4: With heterogeneous β, the government's problem remains the same except (12) 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

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

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. We introduce such loans by assuming that agents may borrow L in period $t = 0$ to finance consumption and tuition. Then, agents start repaying these loans at the beginning of working period $t = 1$, and we assume they repay them in ten years. Loans up to a limit $\bar{L}$ carry the government-subsidized interest rate r_g; any amount above $\bar{L}$ carry a market interest rate $r_m > r_g$. Similarly as in the main calibration in Section 4, we assume that the education, work, and retirement periods last for 5.12, 43, and 20 years, respectively (in the case of high school graduates, the education period lasts for zero years).

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 total loan of L 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

$$c_1 + b_2 + s_2 = \theta l - T\left(\theta l - s_2\right) - T_s\left(\theta l\right) - i,$$

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

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

where $\overline{c}$ is the upper-bound on contributions to the 401(k) account and i is an installment of the student loan defined as follows:

$$i = \frac{1-\delta_a^{10}}{1-\delta_a^{43}}\left[\frac{r_g(1+r_g)^{10}}{(1+r_g)^{10}-1}\min\left\{L,\bar{L}\right\} + \frac{r_m(1+r_m)^{10}}{(1+r_m)^{10}-1}\left(\max\left\{L,\bar{L}\right\}-\bar{L}\right)\right]$$

Agents start repaying their students loans at the beginning of work period and take ten years to pay them down. Loans up to the upper bound of $\bar{L}$ carry a government-subsidized interest rate r_g, while loan amounts above it carry a market interest rate of $r_m > r_g$. The effective installment i is spread out over the entire working-age period using the baseline annual discount factor of δ_a.

To analyze the solution of this model, let $\overline{\chi}_t\left(\theta\right)$ denote the multiplier on the period t budget constraint for agents who invested e_H, and $\underline{\chi}_t\left(\theta\right)$ 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

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

This provides us with the following Euler equation:

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

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

$$h'(l)=\begin{cases}\overline{\chi}_2(\theta)\left\{1.5R_2\left(\frac{\theta l-s_2}{1.5R_2s_2}\right)^{\tau}B(\theta,y,s_2)+0.9\theta\right\} & \text{if } y\leq 9,912\\ \overline{\chi}_2(\theta)\left\{1.5R_2\left(\frac{\theta l-s_2}{1.5R_2s_2}\right)^{\tau}B(\theta,y,s_2)+0.32\theta\right\} & \text{if } 9,912<y\leq 59,760\\ \overline{\chi}_2(\theta)\left\{1.5R_2\left(\frac{\theta l-s_2}{1.5R_2s_2}\right)^{\tau}B(\theta,y,s_2)+0.15\theta\right\} & \text{if } 59,760<y\leq 118,500\\ \overline{\chi}_2(\theta)\,1.5R_2\left(\frac{1}{1.5R_2s_2}\right)^{\tau}\theta\lambda(1-\tau) & \text{if } y>118,500\end{cases},$$

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:

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

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 productivity 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 $t=0$. The total amount of student loans L that carries into the working-age period $t=1$ is defined as

$$L = \left(\frac{1}{\delta_a}\right)^{5.12} b_1 \frac{1-\delta_a^{5.12}}{1-\delta_a}$$

B.2 Deriving Allocations for Proposed Reform

Following Appendix B.1, we derive the allocations from the proposed policy reform that treats student loan repayments as contributions to retirement savings in this section.

In the proposed reform, student loan repayments may qualify as a contribution, even without agents making a direct contribution to their own accounts. Let m denote the match received from student loan repayments. Assuming a 50% matching rate, these agents receive a contribution of $m = 0.5 \min\{i, \overline{c} - s_2\}$. In essence, student loan repayment i is treated as a contribution, but only to the extent that it does not put the total contribution over the limit $\overline{c}$ of what is qualified for matching. Agents can also elect to forgo this option. This happens if agents save $s_2 = \overline{c}$. By contrast, for any $s_2 < \overline{c}$, agents will benefit, at least partially, from receiving a match on their student loan repayments.

During the working period, agents with productivity θ who took out a total loan of L in $t=0$, resulting in an annual loan installment of i (as defined in subsection B.1.1), and invested e in education solve the following problem:

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

subject to

$$c_1 + b_2 + s_2 = \theta l - T(\theta l - s_2) - T_s(\theta l) - i,$$

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

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

There are three cases to consider. In the first case, agents do not benefit from the new policy proposal because they choose to save in the retirement account up to the limit: $s_2 = \bar{c}$ and $b_2 \geq 0$. For the last two cases, $s_2 < \bar{c}$ and student loan repayments act as contributions to their retirement savings. Notice that if agents elect to receive a match on their student loans, they would not save in their regular savings account, so $b_2 = 0$.[29] In the second case, the agents receive a match on the full student loan repayment amount with total contributions below the limit: $i + s_2 \leq \bar{c}$ and $b_2 = 0$. In the last case, agents choose to primarily save on their own and only receive a match on a fraction of the student loan repayment: $i + s_2 = \bar{c}$ and $b_2 \geq 0$.

For the schooling period, by backward induction, agents solve the same problem as the one presented in Appendix B.1.2.

For our quantitative implementation of this model, we assume that the new policy is introduced in a revenue-neutral way. This means that the additional matching provided to employees based on their repayment of student loans is financed by a simultaneous increase in income taxes on everyone. We find that the reform is fully financed by reducing the λ parameter of the tax function from the baseline value of 0.839 to 0.826.

B.3 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 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.[30] The upper bound for 401(k) contributions $\bar{c}$ is set to \$18,000 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 obtain through private market and through a government-subsidized scheme are 10%

[29] It is always optimal to move some of the savings from the regular account into the 401(k) and receive a tax deferral, even when this shift causes a decrease in the match from student loans.

[30] 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.

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	
$\bar{L}$	Cap on government-subsidized student loans	8.75	
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	0.16	
$\delta_1(e_H)$	College period 1 long-term discount factor	0.93	
δ_2	Retirement discount factor	0.29	
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	0.20	
$\delta_1(e_H)$	College period 1 long-term discount factor	0.85	
δ_2	Retirement discount factor	0.17	

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

and 5%, respectively. The amount of subsidized loan is capped at $87,500, in line with the regulations for Stafford loans in the US (weighted by the shares of undergraduate and professional degrees from Table 6). We further assume that an agent takes ten years to repay the student loans.

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[31] as well as relative enrollment data for both types of college.[32] Table 6 presents a breakdown of different higher education outcomes, along with average costs and durations, which we use to calculate this parameter.

In calibrating the short- and long-term discount factors we primarily follow Nakajima

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

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

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.

(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 $\delta_a = 0.9852$ following Nakajima (2012) which we in turn use to calculate effective discount factors across the three periods in our model. These effective discount factors 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 43 years in period 1. This yields $\delta_0(e_H) = \frac{1-\delta_a^{5.12}}{1-\delta_a^{43}} = 0.16$ and $\delta_1(e_H) = \frac{\delta_a^{5.12}-\delta_a^{48.12}}{1-\delta_a^{43}} = 0.93$. We assume that both education types spend 43 years working and 20 years in retirement. This yields a common retirement period discount factor of $\delta_2 = \frac{\delta_a^{43}-\delta_a^{63}}{1-\delta_a^{43}} = 0.29$.[33]

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 factor $\delta_a = 0.9698$. Then, with the same reasoning we assume $\delta_0(e_L) = 0$ and $\delta_1(e_L) = 1$ for high school graduates, compared to $\delta_0(e_H) = 0.20$ and $\delta_1(e_H) = 0.85$ for college graduates. The discount factor for retirement amounts to $\delta_2 = 0.17$.

[33] Because the college type first spends around five years on education before they start to work, we assume that they also retire later and live longer for the same number of years. This is consistent with a significant body of research which shows college graduates live longer than non-college graduates (Meara et al., 2008).

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	205.5	269.2	315.4	223.2
Scale parameter	63.4	86.0	104.5	69.5

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 function of our mixture is then given by

$$
\begin{aligned}
f(\theta) =& p \times \left[\frac{1}{2\pi\sigma_1} \exp\left\{\frac{(\theta-\mu_1)^2}{2\sigma_1^2}\right\}\right] \qquad (15) \\
&+(1-p) \times \begin{cases} \frac{\rho}{\Phi\left(\alpha s(\alpha,\rho)\right)} \frac{1}{\sqrt{2\pi} s(\alpha,\rho)\theta} \exp\left\{-\frac{1}{2}\left(\alpha s(\alpha,\rho) - \frac{\log(\theta^T)-\log(\theta)}{s(\alpha,\rho)}\right)\right\}, & \text{if } \theta \in (0,\theta^T) \\ (1-\rho)\frac{\alpha(\theta^T)^\alpha}{\theta^{\alpha+1}}, \quad \text{if } \theta \in [\theta^T,\infty) \end{cases}
\end{aligned}
$$

In equation (15), 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

Table 8: Parameters of productivity distributions

Symbol	Meaning	Value			
		HS fact.	*HS counter.*	*COL fact.*	*COL counter.*
μ	Mean of normal	8.05	9.74	11.33	8.81
σ	St. dev. of normal	2.19	2.74	3.66	2.29
θ^T	Threshold for Pareto	7.79	9.97	14.70	8.60
ρ	Fraction of lognormal	0.35	0.61	0.63	0.43
p	Probability of normal	0.62	0.39	0.53	0.56

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

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 (15) matter quantitatively and altogether result in a good fit for model-derived distributions of earnings in each group. Figures 10-11 depict the PDFs 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 perhaps a slight exception for 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 report truthfully.

C Sensitivity analyses

In this section, we conduct various sensitivity checks.

C.1 Efficiency wedges for alternative parameters

Here, we show how the efficiency wedge, varies with respect to the two main preference parameters, namely the short-term discount factor and the risk aversion. Figures 12(a) and 12(b) present the efficiency wedges (analogous to Figure 2(a)) when $\beta = 0.5$ and $\beta = 0.9$, respectively. To make the comparison easier, we keep the scale of the y-axes unchanged.

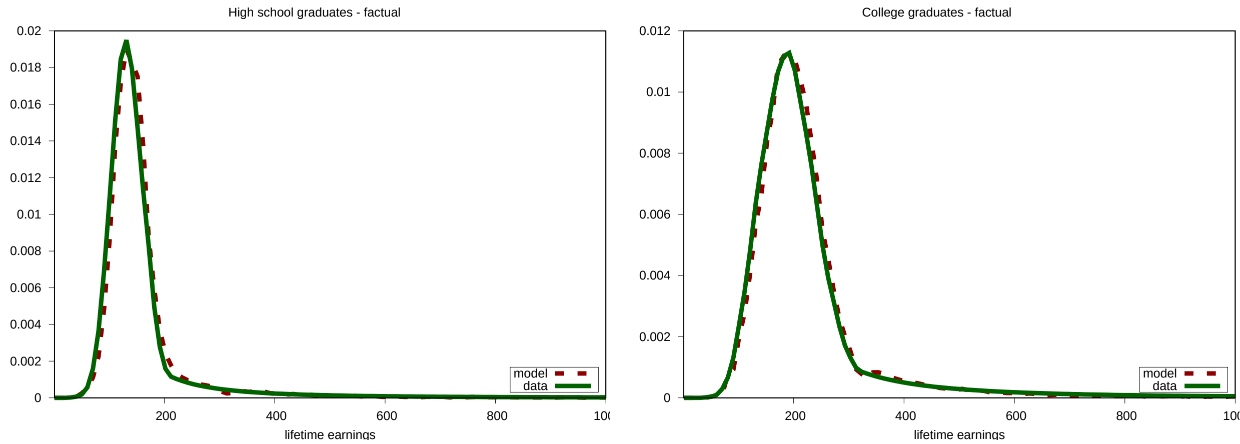

Figure 10: Distributions of lifetime income - factual

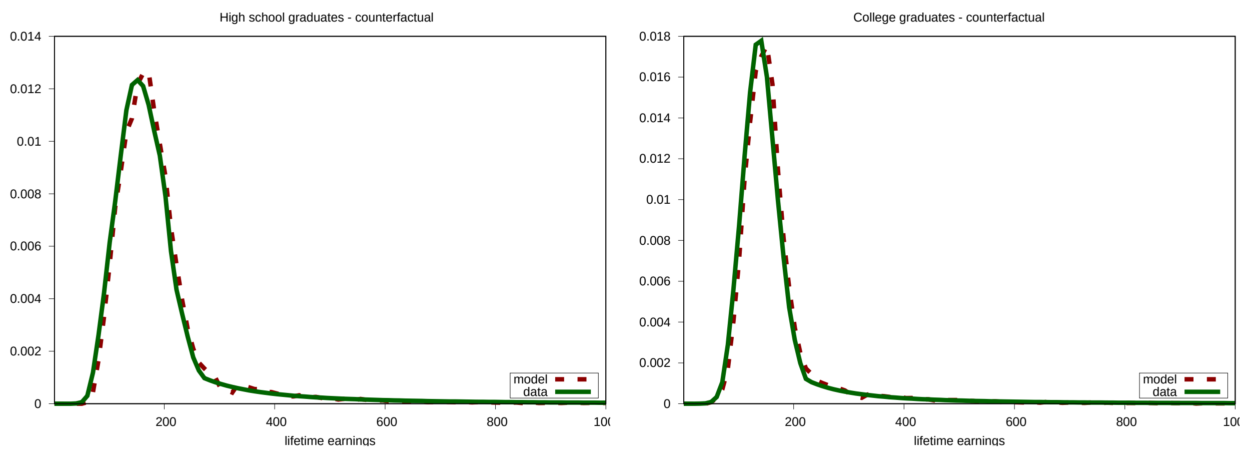

Figure 11: Distributions of lifetime income - counterfactual

The results are quite intuitive in that the two wedges become much steeper (and separated from each other) as β decreases. On the other hand, for a small degree of present bias, the two wedges flatten out and converge to zero.

Figures 13(a) and 13(b) show a similar sensitivity analysis for the efficiency wedges with respect to risk aversion. Intuitively, as σ moves towards risk neutrality the wedges become flatter, while higher risk aversion makes them steeper.

Finally, it should be emphasized that the labor wedge remains virtually unaffected by varying the degree of present bias, which is reminiscent of the result in Figure 3. By contrast, the labor wedge is impacted by the degree of risk aversion, but this issue has been studied extensively by the previous literature and is not the main topic of interest in this paper.

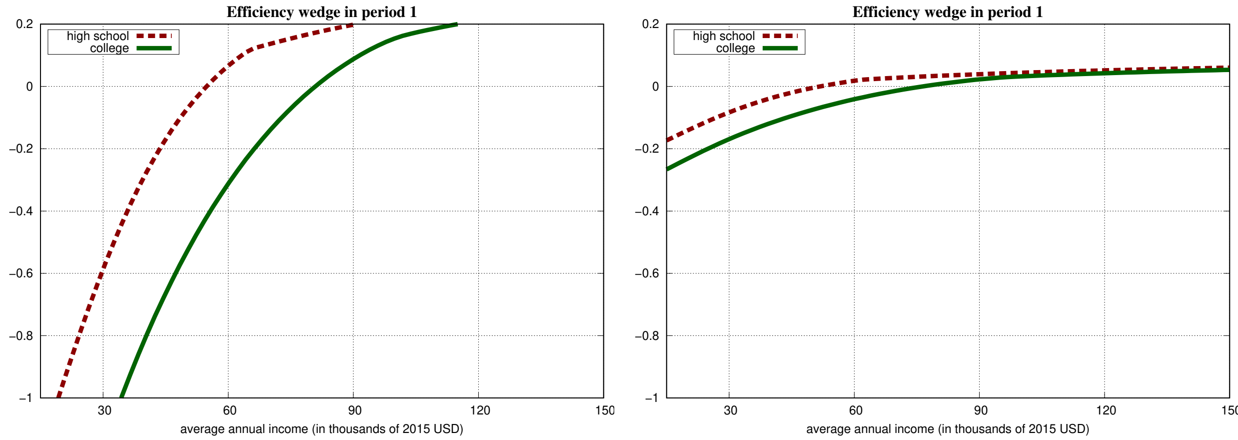

(a) Case $\beta = 0.5$

(b) Case $\beta = 0.9$

Figure 12: Efficiency wedges for alternative values of short-term discount factor

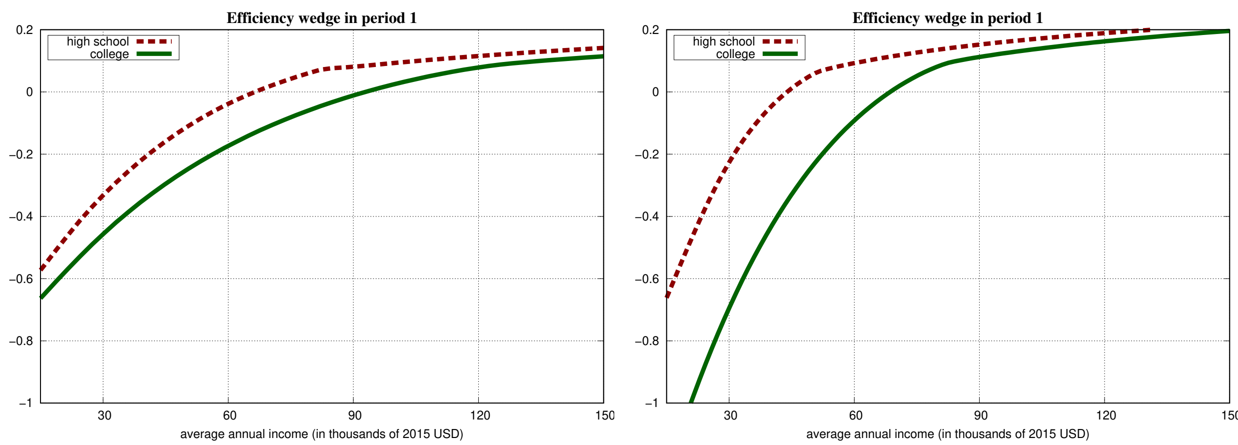

(a) Case $\sigma = 1.5$

(b) Case $\sigma = 2.5$

Figure 13: Efficiency wedges for alternative values of relative risk aversion

D Decomposition of the Labor Wedge

In this section, we quantify the decomposition of the labor wedge introduced in Section 3.2. Figure 14 presents the numerical approximation of the labor wedge components A,C,D and E as function of income for the high innate ability type.[34] 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

[34]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.

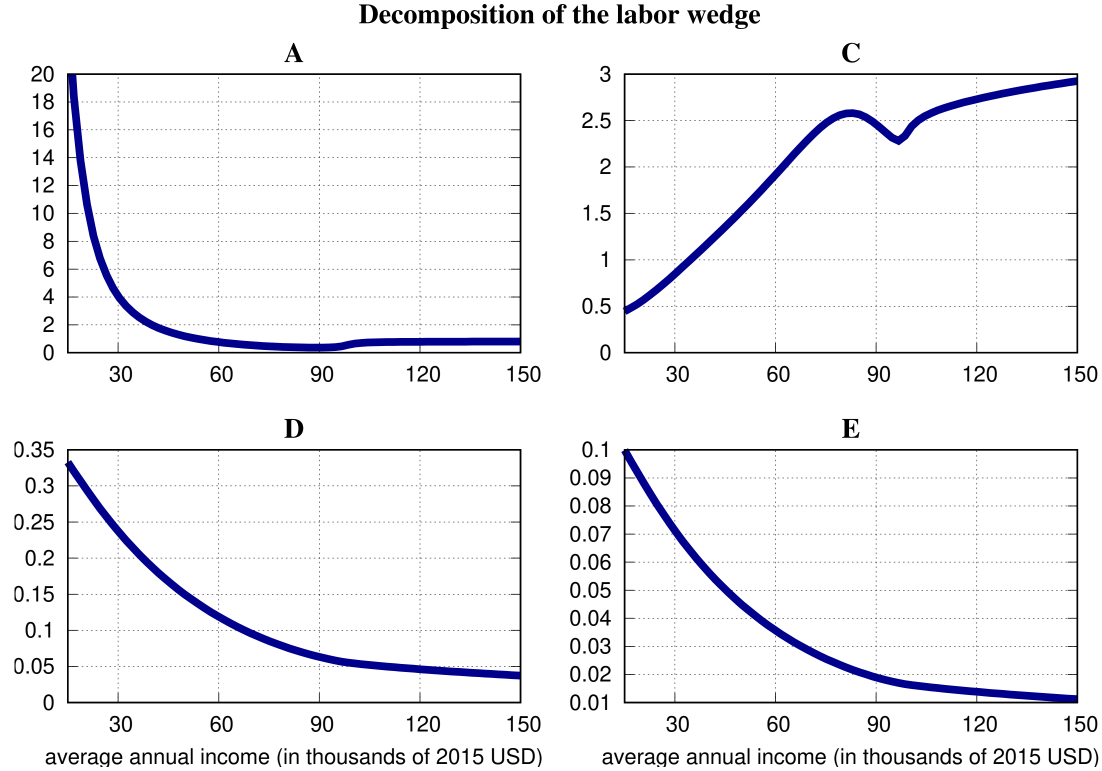

Figure 14: Labor wedge components for the high innate ability type

The novel aspect of our paper is the introduction of E, the present bias component. Since $E = (1 - \beta) D$, this component 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.

E 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 4.3. 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:

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

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

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

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

E.1 Mandatory Savings

Consider a mandatory minimum savings rule that forces agents to smooth consumption: $\tilde{c}_1\left(\gamma,\theta\right)=\tilde{c}_2\left(\gamma,\theta\right)$. 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, $u\left(c\right)=\frac{c^{1-\sigma}}{1-\sigma}$ 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 $t=1$. 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

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

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

E.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 $\tilde{T}\left(y\right)$ that implements the optimum for time-consistent agents. At $t = 1$, time-consistent agents solves the following:

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

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 time-consistent 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:

$$\tilde{T}(y) = \begin{cases} y & \text{if } y \notin \tilde{Y} \\ y - \tilde{c}_1(y) - \frac{1}{R_2}\tilde{c}_2(y) & \text{if } y \in \tilde{Y}. \end{cases}.$$

For simplicity, we assume that if the government observes an off-path income level that it did not 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

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

and

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

It is obvious that there will be intertemporal inefficiencies. Specifically, given $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$, we have

$$\hat{c}_1(y) = \beta^{-\frac{1}{\sigma}} \left(\frac{1+\delta_2}{\beta^{-\frac{1}{\sigma}} + \delta_2} \right) \tilde{c}(y) \text{ and } \hat{c}_2(y) = \left(\frac{1+\delta_2}{\beta^{-\frac{1}{\sigma}} + \delta_2} \right) \tilde{c}(y).$$

In addition to the intertemporal inefficiencies, the agents might also choose suboptimal output. The choice in output y is equivalent to a choice in the report $\hat{\theta}$ in a direct mechanism. Given the savings decision derived above, the agent would solve the following problem

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

which can be rewritten as

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

We can compare this problem to the time-consistent agent's problem:

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

Since the allocations are incentive compatible for time-consistent agents, the time-consistent agents choose $\hat{\theta} = \theta$. This implies that we can rewrite the problem for the present-biased agent as follows:

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

where

$$T = \left[\beta \left(\frac{\beta^{-\frac{1}{\sigma}} + \delta_2}{1+\delta_2} \right)^{\sigma}\right]^{\frac{1}{1+\frac{1}{\eta}}}.$$

This implies that present-biased agents with productivity θ would misreport as

$$T\theta = \left[\beta \left(\frac{\beta^{-\frac{1}{\sigma}} + \delta_2}{1+\delta_2} \right)^{\sigma}\right]^{\frac{1}{1+\frac{1}{\eta}}} \theta.$$

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.

E.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

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

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:

$$\sum_{\gamma} \pi_{\gamma} \left\{ -\frac{\hat{c}_0(\gamma)}{R_0(e_\gamma)} - e_\gamma + \frac{1}{R_1(e_\gamma)} \int_{\Theta} \left[\hat{y}(\gamma,\theta) - \hat{c}_1(\gamma,\theta) - \frac{1}{R_2} \hat{c}_2(\gamma,\theta) \right] f(\theta|\kappa_\gamma)\, d\theta \right\} \geq 0. \quad (16)$$

E.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 (16) 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 which balances the resource constraint under present-biased agents, under both implementations and for all parameter combinations considered in Table 3.

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

	Mandatory savings			Laissez-faire		
	$\sigma = 1.5$	$\sigma = 2$	$\sigma = 2.5$	$\sigma = 1.5$	$\sigma = 2$	$\sigma = 2.5$
$\beta = 0.5$	27.05	21.14	17.52	27.29	21.46	17.86
$\beta = 0.7$	13.16	**10.15**	8.31	13.23	**10.23**	8.40
$\beta = 0.9$	3.60	2.74	2.23	3.61	2.75	2.23

To gain a better understanding of where the difference in welfare gains between the two implementations comes from, Table 10 calculates the welfare gains relative to the laissez-faire allocations, where we assume that agents are able to smooth consumption over the life cycle. In other words, the government spending amount G is the same as in the right panel of Table 9 (because agents report their productivity as in the laissez-faire implementation), but consumption is smoothed over time as in the mandatory savings world. Comparing Table 10 to Table 3, we learn that while the laissez-faire allocations lead to a minor efficiency loss, most of the additional welfare loss in this scenario is due to the agents' inability to smooth consumption over time.

Table 10: Welfare gains relative to time-consistent laissez-faire allocations with perfect consumption smoothing

	Laissez-faire with smoothing		
	$\sigma = 1.5$	$\sigma = 2$	$\sigma = 2.5$
$\beta = 0.5$	2.51	3.41	3.26
$\beta = 0.7$	1.04	**1.43**	1.40
$\beta = 0.9$	0.27	0.37	0.38

F 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 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 the wedge is imposed for each observable income $y(\gamma,\theta)$ which itself is an allocation.

To overcome this challenge we adopt the following approach:

1. Consider a generic set of allocations $\{c_1(\gamma,\theta), c_2(\gamma,\theta), y(\gamma,\theta)\}_{\gamma\in\{H,L\}}$. For each type $\gamma \in \{H, L\}$, and for each productivity θ, find $\hat{\theta}_\gamma$ such that:

$$y(\gamma,\theta) = y(\hat{\gamma},\hat{\theta}_\gamma)$$

   Here, $\hat{\gamma}$ is the innate ability type other than γ. In essence, for each type and productivity level we find an off-grid productivity level that yields the same output for the other innate ability type. We use linear interpolation to evaluate income at off-grid productivity values.

2. Solve for the optimal allocations under an additional set of $2N$ constraints, one for each pair of innate ability and productivity, such that for each γ and θ

$$\frac{u'\Big(c_1(\gamma,\theta)\Big)}{u'\Big(c_2(\gamma,\theta)\Big)} = \frac{u'\Big(c_1(\hat{\gamma},\hat{\theta}_\gamma)\Big)}{u'\Big(c_2(\hat{\gamma},\hat{\theta}_\gamma)\Big)}$$

   Once again, $\hat{\gamma}$ is the innate ability type other than γ and we use linear interpolation to evaluate consumption in both periods at off-grid values of productivity. In essence, for each productivity level we require that the efficiency wedge be equalized with the efficiency wedge of the other innate ability type, at the productivity level which yields the same output.

G Education-Contingent Retirement Savings Subsidy

In this section, we consider an alternative implementation with income and education contingent retirement savings subsidies. 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. 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 $T^k(b)$.

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 $t = 1$. 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 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 5 *The optimum can be implemented with P^{ss}.*

Proof Similar to Section 5.1, we focus on an implementation where agents do not double deviate (misreport and purchase bonds) due to the bond savings tax $T^k(b)$, which is constructed in the proof of Proposition 3.

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

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

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

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

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 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. ■

Figure 15 presents the optimal student loan repayment and retirement savings subsidies for the two education groups as function of income. Panel 15(a) shows that for the H-agents with income below 60,000 in present value, the repayment subsidy starts at over 60% and decreases with income. Once the Pareto tail kicks in, the trend reverts and the optimal subsidy increases and then drops again before settling at around 80%. Panel 15(b) shows the savings subsidy schedules. The optimal savings subsidies are chosen such that they are the negative of the corresponding decision wedges $\hat{\tau}_1^k$, which is why its shape is similar to the efficiency 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.

It is worth pointing out that the optimal student loans subsidy is determined by the labor

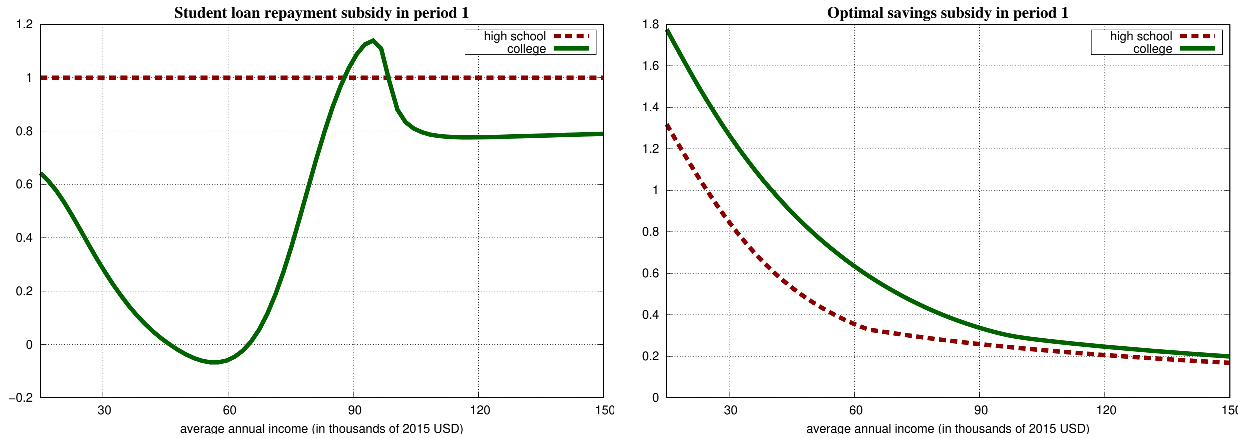

(a) Optimal student loan subsidy $\tau^e(e, y)$

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

Figure 15: Optimal education-contingent subsidies

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 replicate 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 $60,000$.

H Off-Path Mechanisms

In this section, we show how mechanisms with off-path options can relax the ex-ante incentive constraint. First, we discuss how threats can be constructed off the equilibrium path when agents are sophisticated and the government can perfectly identify some of the agents who misreported in the previous period. Then, we illustrate the off-path mechanism for non-sophisticated agents. Finally, we characterize the optimum when the ex-ante incentive constraint is fully relaxed.

H.1 Off-Path Mechanism for Sophisticated Agents

Our paper focuses on a setting where all productivity distributions span the whole range of Θ regardless of agents' human capital κ. Here, we will show how off-path threats can fully relax the ex-ante incentive constraint when there exists a positive measure of productivities that only H-agents can have. Though off-path threats can also be introduced in a setting where all the productivity distributions span the whole range of Θ, it is not apparent why only the dishonest agents choose the off-path threats while the honest agents choose the on-path allocations. This is because, for the off-path threats to work, both H-agents with productivity θ who misreported as low innate ability and actual L-agents with productivity θ would have to be indifferent between the off-path threat and the on-path allocation (Amador et al., 2003; Halac and Yared, 2014). Finally, for off-path threats to be effective, agents have to be aware of their present bias, which is the case with sophisticated agents.

We will assume that only H-agents can have productivities greater than θ_H where $\underline{\theta} < \theta_H < \overline{\theta}$, so $f(\theta|\kappa_H), f(\theta|\kappa_{L,H}) > 0$ and $f(\theta|\kappa_L) = 0$ for any $\theta \in \left(\theta_H, \overline{\theta}\right]$. In other words, only agents with high innate ability can achieve the productivity levels above θ_H. To illustrate how the off-path threats weaken the ex-ante incentive constraint, we also assume that u is unbounded below and above ($u(\mathbb{R}_+) = \mathbb{R}$). When u is unbounded below and above, the ex-ante incentive constraint can be fully relaxed by the off-path threats. In general, off-path threats can weaken the ex-ante incentive constraint—though perhaps not fully—as long as there is a set of productivities with positive measure that only H-agents can achieve.

Following Yu (2020a), the government can introduce a conditional commitment mechanism (CCM). CCM features off-path allocations used to exploit the sophisticated present-biased agents' demand for commitment, which are referred to as threat allocations. For L-agents, the government designs the menu

$$\tilde{P}_L = \left\{c_0(L), [c_1(L,\theta), y(L,\theta), c_2(L,\theta)]_{\theta\in\left(\theta_H,\overline{\theta}\right]}, [c_1(L,\theta), y(L,\theta), c_2(L,\theta)]_{\theta\in[\underline{\theta},\theta_H]}\right\}.$$

Since none of the L-agents would end up with a productivity greater than θ_H, the set of allocations $[c_1(L,\theta), y(L,\theta), c_2(L,\theta)]_{\theta\in\left(\theta_H,\overline{\theta}\right]}$ only punishes H-agents who misreported in $t=0$, so it is the set of threat allocations. The threat allocations are located in the menu for L-agents to deter H-agents from misreporting. Without loss of generality, it is sufficient to assume the government designs a single threat allocation: $c_t(L,\theta) = c_t^T$ and $y(L,\theta) = y^T$ for any $\theta \in \left(\theta_H, \overline{\theta}\right]$. In contrast, we assumed that L-agents do not have an incentive to misreport upwards, so the menu for H-agents stays the same: $\tilde{P}_H = \left\{c_0(H), [c_1(H,\theta), y(H,\theta), c_2(H,\theta)]_{\theta\in\Theta}\right\}$.

The consumption path of the threat allocations are frontloaded to exacerbate the agents' present bias: $c_1^T > c_2^T$. Since present-biased agents at $t = 0$ hope to prevent their future selves from saving too little for retirement, the frontloaded threat consumption path can help deter H-agents from misreporting downwards. Furthermore, to prevent the actual L-agents from selecting the threat allocations, the threat output y^T is increased. Formally, the threat allocations are disciplined by the threat constraints: for any $\theta \in \left(\theta_H, \overline{\theta}\right]$ and $\theta' \in [\underline{\theta}, \theta_H]$,

$$u\left(c_1^T\right) - h\left(\frac{y^T}{\theta}\right) + \beta\delta_2 u\left(c_2^T\right) \geq u\left(c_1\left(L, \theta'\right)\right) - h\left(\frac{y\left(L, \theta'\right)}{\theta}\right) + \beta\delta_2 u\left(c_2\left(L, \theta'\right)\right),$$

and the executability constraints: for any $\theta \in [\underline{\theta}, \theta_H]$,

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

The threat constraints guarantee that H-agents who misreported downwards in $t = 0$ would end up consuming the threat allocation if their productivity is greater than θ_H. The executability constraints ensure that none of the L-agents would consume the threat allocation. Therefore, the ex-ante incentive constraint is

$$\begin{aligned} U_0\left(H\right) \geq \delta_0\left(e_L\right) u\left(c_0\left(L\right)\right) + \beta\delta_1\left(e_L\right) \Bigg[& \int_{\underline{\theta}}^{\theta_H} \left[U_1\left(L, \theta\right) + \left(1 - \beta\right)\delta_2 u\left(c_2\left(L, \theta\right)\right)\right] f\left(\theta|\kappa_{L,H}\right) d\theta \\ & + \int_{\theta_H}^{\overline{\theta}} \left[u\left(c_1^T\right) - h\left(\frac{y^T}{\theta}\right) + \delta_2 u\left(c_2^T\right)\right] f\left(\theta|\kappa_{L,H}\right) d\theta \Bigg]. \end{aligned}$$

Finally, the on-path allocations need to be incentive compatible.

Lemma 3 *When $f\left(\theta|\kappa_H\right), f\left(\theta|\kappa_{L,H}\right) > 0$ and $f\left(\theta|\kappa_L\right) = 0$ for any $\theta \in \left(\theta_H, \overline{\theta}\right]$ and u is unbounded above and below, the ex-ante incentive compatibility constraint is non-binding at the optimum.*

Proof First, for any on-path allocations, let

$$\hat{\theta} = \arg\max_{\theta' \in [\underline{\theta}, \theta_H]} \left\{u\left(c_1\left(L, \theta'\right)\right) - h\left(\frac{y\left(L, \theta'\right)}{\overline{\theta}}\right) + \beta\delta_2 u\left(c_2\left(L, \theta'\right)\right)\right\}.$$

In other words, an agent with the highest productivity $\overline{\theta}$ would misreport as $\hat{\theta} \in [\underline{\theta}, \theta_H]$ if it were restricted to choosing the on-path allocations. This implies that $U_1\left(\hat{\theta}; L, \overline{\theta}\right)$ is the highest attainable utility for the right-hand side of the threat constraint for any agent who reported $\gamma = L$ in $t = 0$. As a result, the threat constraint can be rewritten as

$$u\left(c_1^T\right) + \beta\delta_2 u\left(c_2^T\right) \geq U_1\left(\hat{\theta}; L, \overline{\theta}\right) + \sup_{\theta \in \left(\theta_H, \overline{\theta}\right]} h\left(\frac{y^T}{\theta}\right).$$

This construction would imply that all of the agents with $\theta \in \left(\theta_H, \overline{\theta}\right]$ prefer the threat allocation. We can set c_1^T such that

$$u\left(c_1^T\right) = U_1\left(\hat{\theta}; L, \overline{\theta}\right) + \sup_{\theta \in \left(\theta_H, \overline{\theta}\right]} h\left(\frac{y^T}{\theta}\right) - \beta\delta_2 u\left(c_2^T\right). \tag{17}$$

Next, let

$$\begin{aligned} &U_0\left(L; H, \theta \leq \theta_H\right) \\ &\quad = \delta_0\left(e_L\right) u\left(c_0\left(L\right)\right) + \beta\delta_1\left(e_L\right) \int_{\underline{\theta}}^{\theta_H} \left[U_1\left(L, \theta\right) + \left(1-\beta\right)\delta_2 u\left(c_2\left(L, \theta\right)\right)\right] f\left(\theta | \kappa_{L,H}\right) d\theta, \end{aligned}$$

then the ex-ante incentive constraint can be rewritten as

$$u\left(c_1^T\right) + \delta_2 u\left(c_2^T\right) \leq \frac{\frac{1}{\beta\delta_1(e_L)}\left[U_0\left(H\right) - U_0\left(L; H, \theta \leq \theta_H\right)\right] + \int_{\theta_H}^{\overline{\theta}} h\left(\frac{y^T}{\theta}\right) f\left(\theta | \kappa_{L,H}\right) d\theta}{1 - F\left(\theta | \kappa_{L,H}\right)}.$$

By (17), the ex-ante incentive constraint is

$$\begin{aligned} u\left(c_2^T\right) \leq &\frac{\frac{1}{\beta\delta_1(e_L)}\left[U_0\left(H\right) - U_0\left(L; H, \theta \leq \theta_H\right)\right] + \int_{\theta_H}^{\overline{\theta}} h\left(\frac{y^T}{\theta}\right) f\left(\theta | \kappa_{L,H}\right) d\theta}{\left(1-\beta\right)\delta_2\left[1 - F\left(\theta | \kappa_{L,H}\right)\right]} \\ &- \frac{U_1\left(\hat{\theta}; L, \overline{\theta}\right) + \sup_\theta H\left(\frac{y^T}{\theta}\right)}{\left(1-\beta\right)\delta_2}. \end{aligned}$$

Since u is unbounded below and above, we can decrease c_2^T such that the ex-ante incentive constraint always holds for any on-path allocation and increase c_1^T such that (17) is satisfied.

Finally, to satisfy the executability constraints, notice that for any on-path and threat allocations, we can increase y^T such that the executability constraints hold. ■

Lemma 3 shows how the off-path threat can help relax the ex-ante incentive compatibility constraint. With a positive probability, H-agents who misreported downwards in $t = 0$ are caught lying. Those who are caught are punished with less retirement savings. As a result, H-agents voluntarily report truthfully and surrender their information rent in exchange for commitment. A special case of Lemma 3 is when the innate ability and productivity are the same or when private information is not dynamic. In this case, the government learns the agents' productivity when they report their innate ability truthfully in $t = 0$. As a result, the full information efficient optimum is implementable when the off-path threat fully relaxes the ex-ante incentive constraint on innate ability (Yu, 2020a).

H.2 Off-Path Mechanism 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 $W_1\left(c_1, c_2, y; \theta, \hat{\beta}\right)$ denote the non-sophisticated agents' perceived utility in $t = 1$:

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

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 (2020a) 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 that is implemented on-path is called the *real allocations* denoted as (c, y). To illustrate how PAM works, we assume that u is unbounded below and above ($u(\mathbb{R}_+) = \mathbb{R}$). We will show how the ex-ante incentive constraint can be fully relaxed in this setting. In general, PAM weakens the ex-ante incentive constraint, though perhaps not fully, whenever agents are non-sophisticated.

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 allocation 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 allocation. However, they end-up selecting the real allocation instead. Since we assumed the ex-ante incentive constraints are non-binding for L-agents, the government does not need to design imaginary allocations for them, so $\hat{P}_L = \left\{c_0\left(L\right), \left[c_1\left(L,\theta\right), y\left(L,\theta\right), c_2\left(L,\theta\right)\right]_{\theta\in\Theta}\right\}$. Similar to Yu (2020a), 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 4 *For non-sophisticated present-biased agents, when u is unbounded above and below, the ex-ante incentive compatibility constraint is non-binding at the optimum.*

Proof From Yu (2020a), we first choose the imaginary allocation for a fixed $\hat{\beta}$ 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 allocation has to satisfy the executability constraints to make sure that agents actually choose the real allocation at $t = 1$: for any θ,

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

Thus, the ex-ante incentive compatibility constraint is

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

Next, we show how the imaginary allocation can be designed such that the ex-ante incentive constraint is non-binding for all sophistication levels. Without loss of generality, set $y^I = 0$. Choose the imaginary allocation such that $u\left(c_1^I\right) + \beta\delta_2 u\left(c_2^I\right) = \min_{\tilde{\theta}} U_1\left(H,\tilde{\theta}\right)$, so the executability constraints are non-binding except for H-agents with the lowest utility

U_1. Hence, the preference arbitrage constraints can be expressed as

$$u\left(c_2^I\right) \geq J\left(\hat{\beta}\right) \equiv \frac{\max_{\hat{\theta}}\left\{u\left(c_1\left(H,\hat{\theta}\right)\right) - h\left(\frac{y\left(H,\hat{\theta}\right)}{\theta}\right) + \hat{\beta}\delta_2 u\left(c_2\left(H,\hat{\theta}\right)\right)\right\} - \min_{\tilde{\theta}} U_1\left(H,\tilde{\theta}\right)}{\left(\hat{\beta}-\beta\right)\delta_2}.$$

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

$$u\left(c_2^I\right) \geq \max_{\hat{\beta}} J\left(\hat{\beta}\right).$$

Similarly, the ex-ante incentive constraint can be rewritten as

$$u\left(c_2^I\right) \geq K \equiv \frac{1}{(1-\beta)\delta_2}\left\{\frac{U_0\left(L;H\right) - \delta_0\left(e_H\right)u\left(c_0\left(H\right)\right)}{\beta\delta_1\left(e_H\right)} - \min_{\tilde{\theta}} U_1\left(H,\tilde{\theta}\right)\right\}.$$

Since u is unbounded above, for any real allocation, the imaginary retirement consumption c_2^I can be chosen to satisfy

$$u\left(c_2^I\right) \geq \max\left\{\max_{\hat{\beta}} J\left(\hat{\beta}\right), K\right\}.$$

Also, since u is unbounded below, it is possible to adjust c_1^I so that $u\left(c_1^I\right) = \min_{\tilde{\theta}} U_1\left(H,\tilde{\theta}\right) - \beta\delta_2 u\left(c_2^I\right)$. As a result, it is always possible to find a single set of imaginary allocation 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 4, 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.

H.3 Optimum when Ex-Ante Incentive Constraint is Non-Binding

Lemmas 3 and 4 showed how the government can relax the ex-ante incentive constraint with off-path mechanisms when agents are non-sophisticated or when there are productivities that only H-agents can have. The following proposition describes the optimal wedges when the ex-ante incentive constraint is non-binding.

Proposition 6 *When the ex-ante incentive constraint does not bind, the constrained efficient allocation 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 When the ex-ante incentive constraint is non-binding, the optimization problem is the same as the original problem except that $\mu = 0$. 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, since the ex-ante incentive constraint is non-binding. 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 efficiency 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 when the ex-ante incentive constraint is slack. 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, $\frac{u'(c_1(\gamma,\theta))}{u'(c_2(\gamma,\theta))} > \beta$. When the ex-ante incentive constraint is non-binding, the government can focus on its paternalistic goals since it no longer needs to manipulate retirement consumption to screen innate ability.

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.

I Model with Multiple Working Periods

In this section, we divide the working period in half and allow for stochastic changes in productivity. In this four-period model, the agent is a student at $t = 0$ and then works for two periods at $t = 1$ and $t = 2$. The agent retires at $t = 3$.

Similar to the three-period model, agents learn their innate ability $\gamma \in \{H, L\}$ and choose their education investment $e \in \{e_L, e_H\}$ at $t = 0$. Human capital κ depends on both γ and e, as before. At $t = 1$, agents privately learn their productivity $\theta_1 \in \left[\underline{\theta}, \overline{\theta}\right]$. Productivity θ_1 is drawn from c.d.f. $F_1(\theta_1|\kappa)$ with p.d.f. $f_1(\theta_1|\kappa)$. Also, as before, F_1 is ranked according to first order stochastic dominance and f_1 is strictly positive for any θ_1 and κ. The innovation here is that agents privately draw a new productivity $\theta_2 \in \left[\underline{\theta}, \overline{\theta}\right]$ at $t = 2$, from c.d.f. $F_2(\theta_2|\theta_1, \kappa)$ with p.d.f. $f_2(\theta_2|\theta_1, \kappa)$. In essence, the distribution of θ_2 depends on past productivity θ_1 and human capital κ. We will assume that f_2 is strictly positive for any θ_2 and past history (θ_1, κ). At $t = 3$, agents retire and consume their savings. We will continue to assume that agents only differ in the number of years spent as a student at $t = 0$, and the length of other periods are the same for all agents.

I.1 The Mechanism and Incentive Compatibility

The government designs the following direct mechanism:

$$P = \{c_0(\gamma), [c_1(\gamma, \theta_1), y_1(\gamma, \theta_1)], [c_2(\gamma, \theta_1, \theta_2), c_3(\gamma, \theta_1, \theta_2), y_2(\gamma, \theta_1, \theta_2)]\}.$$

Following the analysis for the three-period model, we require the mechanism P to be incentive compatible for every period. Let the utility of a type $(\gamma, \theta_1, \theta_2)$ agent who reports $\theta_2' \in \Theta$ in

$t = 2$ be denoted as

$$U_2(\theta_2'; \gamma, \theta_1, \theta_2) = u(c_2(\gamma, \theta_1, \theta_2')) - h\left(\frac{y_2(\gamma, \theta_1, \theta_2')}{\theta_2}\right) + \beta\delta_3 u(c_3(\gamma, \theta_1, \theta_2')).$$

The incentive compatibility constraints in $t = 2$ ensure the agents report θ_2 truthfully: for any $\theta_2, \theta_2' \in \Theta$,

$$U_2(\gamma, \theta_1, \theta_2) \equiv U_2(\theta_2; \gamma, \theta_1, \theta_2) \geq U_2(\theta_2'; \gamma, \theta_1, \theta_2).$$

Let the utility of a type (γ, θ_1) agent who reports $\theta_1' \in \Theta$ in $t = 1$ be denote as

$$U_1(\theta_1'; \gamma, \theta_1) = u(c_1(\gamma, \theta_1')) - h\left(\frac{y_1(\gamma, \theta_1')}{\theta_1}\right) + \beta\delta_2 \int_\Theta \left[U_2(\gamma, \theta_1', \theta_2) + (1-\beta)\delta_3 u(c_3(\gamma, \theta_1', \theta_2))\right] dF_2(\theta_2|\theta_1, \kappa_\gamma).$$

The incentive compatibility constraints in $t = 1$ ensure the agents report θ_1 truthfully: for any $\theta_1, \theta_1' \in \Theta$,

$$U_1(\gamma, \theta_1) \equiv U_1(\theta_1; \gamma, \theta_1) \geq U_1(\theta_1'; \gamma, \theta_1).$$

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

$$U_0(\gamma'; \gamma) = \delta_0(e_{\gamma'}) u(c_0(\gamma')) + \beta\delta_1(e_{\gamma'}) \int_\Theta \Bigg[u(c_1(\gamma', \theta_1)) - h\left(\frac{y_1(\gamma', \theta_1)}{\theta_1}\right) + \delta_2 \int_\Theta \left[U_2(\gamma', \theta_1, \theta_2) + (1-\beta)\delta_3 u(c_3(\gamma', \theta, \theta_2))\right] dF_2(\theta_2|\theta_1, \kappa_{\gamma',\gamma})\Bigg] dF_1(\theta_1|\kappa_{\gamma',\gamma}).$$

The incentive compatibility constraints at $t = 0$ ensure that the agents report γ truthfully: for any innate ability γ, γ',

$$U_0(\gamma) \equiv U_0(\gamma; \gamma) \geq U_0(\gamma'; \gamma).$$

The following lemma characterizes the set of allocations that are incentive compatible at $t = 2$. Its proof is similar to the proof of Lemma 1 so it is omitted.

Lemma 5 *For any γ and θ_1, P is incentive compatible at $t = 2$ if and only if (i.) $y_2(\gamma, \theta_1, \theta_2)$ is non-decreasing in θ_2, and (ii.) $U_2(\gamma, \theta_1, \theta_2)$ is absolutely continuous in θ_2, so it is differentiable almost everywhere with*

$$\frac{\partial U_2(\gamma, \theta_1, \theta_2)}{\partial \theta_2} = \frac{y_2(\gamma, \theta_1, \theta_2)}{\theta_2^2} h'\left(\frac{y_2(\gamma, \theta_1, \theta_2)}{\theta_2}\right). \tag{18}$$

Unfortunately, it is difficult to simplify the incentive compatibility constraints at $t = 1$.

With a continuum of productivities in $t = 1$, local incentive compatibility may not imply global incentive compatibility when agents are time inconsistent (Halac and Yared, 2014; Galperti, 2015; Yu, 2020b). This issue and others were discussed in Section 2.2. We will not focus on the theoretical properties that guarantee the sufficiency of local incentive compatibility.[35] Instead, we will follow the standard procedure and replace the incentive constraints in $t = 1$ with the following envelope condition:

$$\begin{aligned}\frac{\partial U_1(\gamma,\theta_1)}{\partial\theta_1} = \frac{y_1(\gamma,\theta_1)}{\theta_1^2}h'\left(\frac{y_1(\gamma,\theta_1)}{\theta_1}\right) \\ - \beta\delta_2\int_\Theta \frac{\partial F_2(\theta_2|\theta_1,\kappa_\gamma)}{\partial\theta_1}\frac{y_2(\gamma,\theta_1,\theta_2)}{\theta_2^2}h'\left(\frac{y_2(\gamma,\theta_1,\theta_2)}{\theta_2}\right)d\theta_2 \\ + \beta(1-\beta)\delta_2\delta_3\int_\Theta u(c_3(\gamma,\theta_1,\theta_2))\frac{\partial f_2(\theta_2|\theta_1,\kappa_\gamma)}{\partial\theta_1}d\theta_2,\end{aligned} \tag{19}$$

which is derived with the help of Lemma 5 by assuming P is incentive compatible at $t = 2$. It is important to note that (19) is a necessary condition for incentive compatibility, not sufficient.

Finally, similar to the three-period model, it is difficult to assess which of the incentive constraints on innate ability are relevant at $t = 0$. Following the analysis of the three-period model, we will focus on the case where only the H-agents go to college, so the relevant incentive constraint at $t = 0$ is

$$U_0(H) \geq U_0(L;H). \tag{20}$$

I.2 The Planning Problem

The government maximizes the sum of long-run utilities:

$$\begin{aligned}\sum_\gamma \pi_\gamma\Bigg\{\delta_0(e_\gamma)u(c_0(\gamma)) + \delta_1(e_\gamma)\int_\Theta\Bigg[U_1(\gamma,\theta) \\ + (1-\beta)\delta_2\int_\Theta[U_2(\gamma,\theta_1,\theta_2) + (1-\beta)\delta_3 u(c_3(\gamma,\theta,\theta_2))]\,dF_2(\theta_2|\theta_1,\kappa_\gamma)\Bigg]dF_1(\theta_1|\kappa_\gamma)\Bigg\}\end{aligned}$$

subject to

$$U_2(\gamma,\theta_1,\theta_2) = u(c_2(\gamma,\theta_1,\theta_2)) - h\left(\frac{y_2(\gamma,\theta_1,\theta_2)}{\theta_2}\right) + \beta\delta_3 u(c_3(\gamma,\theta_1,\theta_2)), \tag{21}$$

[35] See Appendix B in Yu (2020b) for more information on the conditions that guarantee the sufficiency of local incentive compatibility constraints in an environment with quasi-linear utility.

$$U_1(\gamma,\theta_1) = u(c_1(\gamma,\theta_1)) - h\left(\frac{y_1(\gamma,\theta_1)}{\theta_1}\right) + \beta\delta_2 \int_\Theta \left[U_2(\gamma,\theta_1,\theta_2) + (1-\beta)\delta_3 u(c_3(\gamma,\theta_1,\theta_2))\right] dF_2(\theta_2|\theta_1,\kappa_\gamma), \quad (22)$$

the incentive constraints (18), (19), and (20), and the resource constraint

$$\sum_\gamma \pi_\gamma \left\{ \frac{-c_0(\gamma) - e_\gamma}{R_0(e_\gamma)} + \frac{1}{R_1(e_\gamma)} \int_\Theta \left[y_1(\gamma,\theta_1) - c_1(\gamma,\theta_1) + \frac{1}{R_2}\int_\Theta \left[y_2(\gamma,\theta_1,\theta_2) - c_2(\gamma,\theta_1,\theta_2) - \frac{c_3(\gamma,\theta_1,\theta_2)}{R_3}\right] dF_2(\theta_2|\theta_1,\kappa_\gamma) \right] dF_1(\theta_1|\kappa_\gamma) \right\} \geq 0.$$

We will assume that $\delta_t R_t = 1$ for all t.

I.2.1 The Optimality Conditions

Let $(\lambda_\gamma(\theta_1,\theta_2), \lambda_\gamma(\theta_1), \xi_\gamma(\theta_1,\theta_2), \xi_\gamma(\theta_1), \mu, \phi)$ be the multipliers on (21), (22), (18), (19), (20), and the resource constraint respectively. Using standard Hamiltonian techniques, we derive the following necessary conditions for optimality

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

$$\lambda_H(\theta_1) + \beta\mu\delta_1(e_H) f_1(\theta_1|\kappa_H) = \frac{\phi\pi_H\delta_1(e_H) f_1(\theta_1|\kappa_H)}{u'(c_1(H,\theta_1))},$$

$$\lambda_L(\theta_1) - \beta\mu\delta_1(e_L) f_1(\theta_1|\kappa_{L,H}) = \frac{\phi\pi_L\delta_1(e_L) f_1(\theta_1|\kappa_L)}{u'(c_1(L,\theta_1))},$$

$$(1-\beta)\left(\pi_H + \frac{\beta\mu}{1-\beta}\right)\delta_1(e_H)\delta_2 f_1(\theta_1|\kappa_H) f_2(\theta_2|\theta_1,\kappa_H) + \beta\lambda_H(\theta_1)\delta_2 f_2(\theta_2|\theta_1,\kappa_H) - \frac{\partial\xi_H(\theta_1,\theta_2)}{\partial\theta_2} = \lambda_H(\theta_1,\theta_2),$$

$$(1-\beta)\left[\pi_L - \frac{\beta\mu}{1-\beta}\left(\frac{f_1(\theta_1|\kappa_{L,H}) f_2(\theta_2|\theta_1,\kappa_{L,H})}{f_1(\theta_1|\kappa_L) f_2(\theta_2|\theta_1,\kappa_L)}\right)\right]\delta_1(e_L)\delta_2 f_1(\theta_1|\kappa_L) f_2(\theta_2|\theta_1,\kappa_L) + \beta\lambda_L(\theta_1)\delta_2 f_2(\theta_2|\theta_1,\kappa_L) - \frac{\partial\xi_L(\theta_1,\theta_2)}{\partial\theta_2} = \lambda_L(\theta_1,\theta_2),$$

$$
\begin{aligned}
&(1-\beta)^2\left(\pi_H+\frac{\beta\mu}{1-\beta}\right)\delta_1(e_H)\,\delta_2 f_1(\theta_1|\kappa_H)\,f_2(\theta_2|\theta_1,\kappa_H)\\
&\quad+\beta(1-\beta)\,\lambda_H(\theta_1)\,\delta_2 f_2(\theta_2|\theta_1,\kappa_H)+\beta\lambda_H(\theta_1,\theta_2)-\beta(1-\beta)\,\xi_H(\theta_1)\,\delta_2\frac{\partial f_2(\theta_2|\theta_1,\kappa_H)}{\partial\theta_1}\\
&\qquad=\frac{\phi\pi_H\delta_1(e_H)\,\delta_2 f_1(\theta_1|\kappa_H)\,f_2(\theta_2|\theta_1,\kappa_H)}{u'(c_3(H,\theta_1,\theta_2))},
\end{aligned}
$$

$$
\begin{aligned}
&(1-\beta)^2\left[\pi_L-\frac{\beta\mu}{1-\beta}\left(\frac{f_1(\theta_1|\kappa_{L,H})\,f_2(\theta_2|\theta_1,\kappa_{L,H})}{f_1(\theta_1|\kappa_L)\,f_2(\theta_2|\theta_1,\kappa_L)}\right)\right]\delta_1(e_L)\,\delta_2 f_1(\theta_1|\kappa_L)\,f_2(\theta_2|\theta_1,\kappa_L)\\
&\quad+\beta(1-\beta)\,\lambda_L(\theta_1)\,\delta_2 f_2(\theta_2|\theta_1,\kappa_L)+\beta\lambda_L(\theta_1,\theta_2)-\beta(1-\beta)\,\xi_L(\theta_1)\,\delta_2\frac{\partial f_2(\theta_2|\theta_1,\kappa_L)}{\partial\theta_1}\\
&\qquad=\frac{\phi\pi_L\delta_1(e_L)\,\delta_2 f_1(\theta_1|\kappa_L)\,f_2(\theta_2|\theta_1,\kappa_L)}{u'(c_3(L,\theta_1,\theta_2))},
\end{aligned}
$$

$$
\begin{aligned}
&[\lambda_H(\theta_1)+\beta\mu\delta_1(e_H)\,f_1(\theta_1|\kappa_H)]\,\frac{1}{\theta_1}h'\left(\frac{y_1(H,\theta_1)}{\theta_1}\right)\\
&\quad+\xi_H(\theta_1)\left[\frac{1}{\theta_1^2}h'\left(\frac{y_1(H,\theta_1)}{\theta_1}\right)+\frac{y_1(H,\theta_1)}{\theta_1^3}h''\left(\frac{y_1(H,\theta_1)}{\theta_1}\right)\right]=\phi\pi_H\delta_1(e_H)\,f_1(\theta_1|\kappa_H),
\end{aligned}
$$

$$
\begin{aligned}
&[\lambda_L(\theta_1)-\beta\mu\delta_1(e_L)\,f_1(\theta_1|\kappa_{L,H})]\,\frac{1}{\theta_1}h'\left(\frac{y_1(L,\theta_1)}{\theta_1}\right)\\
&\quad+\xi_L(\theta_1)\left[\frac{1}{\theta_1^2}h'\left(\frac{y_1(L,\theta_1)}{\theta_1}\right)+\frac{y_1(L,\theta_1)}{\theta_1^3}h''\left(\frac{y_1(L,\theta_1)}{\theta_1}\right)\right]=\phi\pi_L\delta_1(e_L)\,f_1(\theta_1|\kappa_L),
\end{aligned}
$$

and for all γ,

$$
\pi_\gamma\delta_1(e_\gamma)\,f_1(\theta_1|\kappa_\gamma)-\xi_\gamma'(\theta_1)=\lambda_\gamma(\theta_1),
$$

$$
\lambda_\gamma(\theta_1,\theta_2)\,u'(c_2(\gamma,\theta_1,\theta_2))=\phi\pi_\gamma\delta_1(e_\gamma)\,\delta_2 f_1(\theta_1|\kappa_\gamma)\,f_2(\theta_2|\theta_1,\kappa_\gamma),
$$

$$
\begin{aligned}
&\lambda_\gamma(\theta_1,\theta_2)\,\frac{1}{\theta_2}h'\left(\frac{y_2(\gamma,\theta_1,\theta_2)}{\theta_2}\right)+\left[\xi_\gamma(\theta_1,\theta_2)-\beta\delta_2\xi_\gamma(\theta_1)\,\frac{\partial F_2(\theta_2|\theta_1,\kappa_\gamma)}{\partial\theta_1}\right]\\
&\qquad\times\left[\frac{1}{\theta_2^2}h'\left(\frac{y_2(\gamma,\theta_1,\theta_2)}{\theta_2}\right)+\frac{y_2(\gamma,\theta_1,\theta_2)}{\theta_2^3}h''\left(\frac{y_2(\gamma,\theta_1,\theta_2)}{\theta_2}\right)\right]\\
&\qquad\qquad=\phi\pi_\gamma\delta_1(e_\gamma)\,\delta_2 f_1(\theta_1|\kappa_\gamma)\,f_2(\theta_2|\theta_1,\kappa_\gamma),
\end{aligned}
$$

and the following boundary conditions hold: for all γ,

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

and for any γ and $\theta_1\in\Theta$,

$$\xi_\gamma\left(\theta_1,\underline{\theta}\right)=\xi_\gamma\left(\theta_1,\overline{\theta}\right)=0.$$

Next, using the conditions above, we present the optimal labor and intertemporal distortions in the four-period life-cycle model. We will discuss the labor wedges before the intertemporal wedges, because, surprisingly, the intertemporal wedges depend on the distortions in labor when agents are present biased.

I.3 Labor Wedges

To separate the economic forces that determine the optimal labor distortions, we first define the elements that influence the wedges in $t=1$:

$$A_\gamma(\theta_1)=\frac{1-F_1\left(\theta_1|\kappa_\gamma\right)}{\theta f_1\left(\theta_1|\kappa_\gamma\right)},$$

$$B_\gamma(\theta_1)=1+\frac{\frac{y_1(\gamma,\theta_1)}{\theta_1}h''\left(\frac{y(\gamma,\theta_1)}{\theta_1}\right)}{h'\left(\frac{y_1(\gamma,\theta_1)}{\theta_1}\right)},$$

$$C_\gamma(\theta_1)=\int_{\theta}^{\overline{\theta}}\frac{u'\left(c_1\left(\gamma,\theta_1\right)\right)}{u'\left(c_1\left(\gamma,x\right)\right)}\left[1-\frac{u'\left(c_1\left(\gamma,x\right)\right)}{\phi}\right]\frac{f_1\left(x|\kappa_\gamma\right)}{1-F_1\left(\theta_1|\kappa_\gamma\right)}dx,$$

$$D_\gamma(\theta_1)=u'\left(c_1\left(\gamma,\theta_1\right)\right)\left[\frac{1}{u'\left(c_0\left(\gamma\right)\right)}-\frac{1}{\phi}\right],$$

$$E_\gamma(\theta_1)=\left(1-\beta\right)D_\gamma\left(\theta_1\right).$$

Also, for the wedges in $t=2$, we define $A_\gamma\left(\theta_1,\theta_2\right)$, $B_\gamma\left(\theta_1,\theta_2\right)$, and $C_\gamma\left(\theta_1,\theta_2\right)$ analogously, and let

$$D_\gamma(\theta_1,\theta_2)=u'\left(c_2\left(\gamma,\theta_1,\theta_2\right)\right)\left[\frac{1}{u'\left(c_1\left(\gamma,\theta_1\right)\right)}-\frac{1}{\phi}\right],$$

$$E_\gamma(\theta_1,\theta_2)=\left(1-\beta\right)D_\gamma\left(\theta_1,\theta_2\right),$$

$$\tilde{E}_\gamma(\theta_1,\theta_2)=\beta\left(1-\beta\right)\frac{u'\left(c_2\left(\gamma,\theta_1,\theta_2\right)\right)}{u'\left(c_1\left(\gamma,\theta_1\right)\right)}D_\gamma\left(\theta_1\right).$$

Most of the components are the same as the three-period model and represent the same forces, except for $\tilde{E}_{\gamma}(\theta_1, \theta_2)$. Similar to the present-bias component $E_{\gamma}(\theta_1, \theta_2)$, notice that $\tilde{E}_{\gamma}(\theta_1, \theta_2)$ is also zero when agents are time consistent, so it is unique to our environment with present-biased agents.

The following proposition characterizes the optimal labor distortion in the four-period model.

Proposition 7 *The labor wedge at $t = 1$ for any $\theta_1 \in \Theta$ satisfies*

$$\frac{\tau^w(H, \theta_1)}{1 - \tau^w(H, \theta_1)} = A_H(\theta_1) B_H(\theta_1) \left[C_H(\theta_1) - D_H(\theta_1) + E_H(\theta_1)\right],$$

$$\frac{\tau^w(L, \theta_1)}{1 - \tau^w(L, \theta_1)} = A_L(\theta_1) B_L(\theta_1) \left[C_L(\theta_1) - \frac{1 - F_1(\theta_1|\kappa_{L,H})}{1 - F_1(\theta_1|\kappa_L)} \left[D_L(\theta_1) - E_L(\theta_1)\right]\right],$$

and the labor wedge at $t = 2$ for any $\theta_1, \theta_2 \in \Theta$ satisfies

$$\begin{aligned}\frac{\tau^w(H, \theta_1, \theta_2)}{1 - \tau^w(H, \theta_1, \theta_2)} &= A_H(\theta_1, \theta_2) B_H(\theta_1, \theta_2) \\ &\times \left\{ C_H(\theta_1, \theta_2) - D_H(\theta_1, \theta_2) + E_H(\theta_1, \theta_2) - \tilde{E}_H(\theta_1, \theta_2) \right. \\ &\left. - \frac{\beta u'(c_2(H, \theta_1, \theta_2))}{u'(c_1(H, \theta_1))} \frac{\frac{\partial F_2(\theta_2|\theta_1, \kappa_H)}{\partial \theta_1}}{f_2(\theta_2|\theta_1, \kappa_H)} \frac{f_2(\theta_2|\theta_1, \kappa_H)}{1 - F_2(\theta_2|\theta_1, \kappa_H)} \frac{1 - F_1(\theta_1|\kappa_H)}{f_1(\theta_1|\kappa_H)} \frac{\frac{\tau^w(H, \theta_1)}{1 - \tau^w(H, \theta_1)}}{A_H(\theta_1) B_H(\theta_1)} \right\},\end{aligned}$$

$$\begin{aligned}\frac{\tau^w(L, \theta_1, \theta_2)}{1 - \tau^w(L, \theta_1, \theta_2)} &= A_L(\theta_1, \theta_2) B_L(\theta_1, \theta_2) \\ \times \Bigg\{ C_L(\theta_1, \theta_2) - D_L(\theta_1, \theta_2) &+ E_L(\theta_1, \theta_2) - \frac{f_1(\theta_1|\kappa_{L,H})}{f_1(\theta_1|\kappa_L)} \left[\frac{\frac{1 - F_1(\theta_1|\kappa_{L,H})}{1 - F_1(\theta_1|\kappa_L)} - \beta}{1 - \beta}\right] \tilde{E}_L(\theta_1, \theta_2) \\ &- \frac{\beta u'(c_2(L, \theta_1, \theta_2))}{u'(c_1(L, \theta_1))} \frac{\frac{\partial F_2(\theta_2|\theta_1, \kappa_L)}{\partial \theta_1}}{f_2(\theta_2|\theta_1, \kappa_L)} \frac{f_2(\theta_2|\theta_1, \kappa_L)}{1 - F_2(\theta_2|\theta_1, \kappa_L)} \frac{1 - F_1(\theta_1|\kappa_L)}{f_1(\theta_1|\kappa_L)} \frac{\frac{\tau^w(L, \theta_1)}{1 - \tau^w(L, \theta_1)}}{A_L(\theta_1) B_L(\theta_1)} \Bigg\},\end{aligned}$$

where $\frac{1}{\phi} = \mathbb{E}_{\gamma}\left[\mathbb{E}_{\theta_1}\left(\left.\frac{1}{u'(c_1(\gamma, \theta_1))}\right|\gamma\right)\right]$.

Proof The results follow from the first order conditions by using the techniques shown in the proof of Proposition 2. ■

Proposition 7 shows that the labor wedges at $t = 1$ are the same as the ones in the three-period model. However, the labor wedges at $t = 2$—the period just before retirement—contains new economic forces. First, due to dynamic productivity, the labor wedge at $t = 2$ depends on the previous period's labor distortion. This dependence on past labor distortions is closely related to the intertemporal component charaterized in Golosov et al. (2016), which serves to relax the incentive constraints in $t = 1$. Furthermore, this dependence is weaker than when agents are time consistent, since present-biased agents are less sensitive to future incentives. Finally, the labor wedge at $t = 2$ also directly depends on education investment through the new present-bias component $\tilde{E}_\gamma(\theta_1, \theta_2)$, which depends on $D_\gamma(\theta_1)$.

I.4 Intertemporal Wedges

Before we present the optimal intertemporal distortions, let us define the elasticity of the density f_2 with respect to θ_1 as

$$\epsilon_\gamma(\theta_1, \theta_2) = \frac{\partial f_2(\theta_2|\theta_1, \kappa_\gamma)}{\partial \theta_1} \frac{\theta_1}{f_2(\theta_2|\theta_1, \kappa_\gamma)}.$$

The following proposition provides the inverse Euler equations for present-biased agents in the four-period model.

Proposition 8 *The constrained efficient allocation satisfies (i.) the inverse Euler equation in aggregate:*

$$\sum_\gamma \frac{\pi_\gamma}{u'(c_0(\gamma))} = \sum_\gamma \pi_\gamma \mathbb{E}_{\theta_1}\left(\frac{1}{u'(c_1(\gamma, \theta_1))}\middle|\gamma\right),$$

$$\mathbb{E}_{\theta_1}\left(\frac{1}{u'(c_1(\gamma, \theta_1))}\middle|\gamma\right) = \mathbb{E}_{\theta_1, \theta_2}\left(\frac{1}{u'(c_2(\gamma, \theta_1, \theta_2))}\middle|\gamma\right) = \mathbb{E}_{\theta_1, \theta_2}\left(\frac{1}{u'(c_3(\gamma, \theta_1, \theta_2))}\middle|\gamma\right),$$

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

$$\frac{1}{\beta u'(c_3(H, \theta_1, \theta_2))} = \frac{1}{u'(c_2(H, \theta_1, \theta_2))} + \left[1 - \frac{\epsilon_H(\theta_1, \theta_2)\frac{\tau^w(H, \theta_1)}{1 - \tau^w(H, \theta_1)}}{B_H(\theta_1)}\right] \frac{1 - \beta}{u'(c_1(H, \theta_1))} + \frac{(1 - \beta)^2}{\beta}\left(\frac{\pi_H + \beta\mu}{\pi_H + \mu}\right)\frac{1}{u'(c_0(H))},$$

$$\frac{1}{\beta u'(c_3(L,\theta_1,\theta_2))} = \frac{1}{u'(c_2(L,\theta_1,\theta_2))} + \left[1 - \frac{\epsilon_L(\theta_1,\theta_2)\frac{\tau^w(L,\theta_1)}{1-\tau^w(L,\theta_1)}}{B_L(\theta_1)}\right]\frac{1-\beta}{u'(c_1(L,\theta_1))}$$

$$+\frac{(1-\beta)^2}{\beta}\left\{\frac{\left[\pi_L - \beta\mu\frac{f_1(\theta_1|\kappa_{L,H})f_2(\theta_2|\theta_1,\kappa_{L,H})}{f_1(\theta_1|\kappa_L)f_2(\theta_2|\theta_1,\kappa_L)}\right] - \frac{\beta^2\mu}{1-\beta}\frac{f_1(\theta_1|\kappa_{L,H})}{f_1(\theta_1|\kappa_L)}\left[\frac{f_2(\theta_2|\theta_1,\kappa_{L,H})}{f_2(\theta_2|\theta_1,\kappa_L)} - 1\right]}{\pi_L - \mu}\right\}\frac{1}{u'(c_0(L))},$$

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

Proof The results follow from the first order conditions by using the techniques shown in the proof of Proposition 1. ■

Proposition 8 shows that, similar to the three-period model, the best the government can do is to choose consumption such that the inverse marginal utility is equalized in aggregate when agents are present biased. Similarly, the intertemporal distortions at $t = 0$ are also characterized by the inverse Euler inequalities. However, Proposition 8 also shows that there are new economic forces that determine the intertemporal wedges at $t = 1$ and $t = 2$.

From Proposition 8, the optimal intertemporal wedges at $t = 1$ are characterized by

$$\frac{1}{u'(c_1(H,\theta_1))} + \frac{1-\beta}{\beta}\left(\frac{\pi_H + \beta\mu}{\pi_H + \mu}\right)\frac{1}{u'(c_0(H))} = \mathbb{E}_{\theta_2}\left(\frac{1}{\beta u'(c_2(H,\theta_1,\theta_2))}\middle|\theta_1\right),$$

$$\frac{1}{u'(c_1(L,\theta_1))} + \frac{1-\beta}{\beta}\left[\frac{\pi_L - \beta\mu\frac{f_1(\theta_1|\kappa_{L,H})}{f_1(\theta_1|\kappa_L)}}{\pi_L - \mu}\right]\frac{1}{u'(c_0(L))} = \mathbb{E}_{\theta_2}\left(\frac{1}{\beta u'(c_2(L,\theta_1,\theta_2))}\middle|\theta_1\right).$$

Notice that these wedges look similar to the three-period model's inverse Euler equations (6) and (7), with a key difference—future productivity θ_2 is unknown. Due to this uncertainty, it is optimal to restrict savings at $t = 1$ to relax the incentive constraints at $t = 2$. On the other hand, the main mechanism of this paper also exists: A commitment device that helps agents save at $t = 1$. With these two opposing forces, it is unclear whether the optimal efficiency wedge at $t = 1$ is smaller than $1 - \beta$.

Finally, Proposition 8 demonstrates that the retirement savings of present-biased agents at $t = 2$ depends on the labor distortion at $t = 1$ and education investment. The government uses reports on θ_2 to detect possible prior misreports and distorts the allocations according to that likelihood to relax past incentive constraints. Crucially, this dependence is not presesnt for time-consistent agents and when θ_2 is independent of past productivity θ_1 or human capital κ.

To see how this mechanism works, first notice that the degree of dependence on previous period's labor distortion is affected by the elasticity $\epsilon_\gamma(\theta_1, \theta_2)$, which measures the percentage change of density f_2 in response to a change in θ_1. The consumption path is more frontloaded when the elasticity is positive compared to when it is negative. This is because a positive elasticity implies that the current productivity θ_2 was more likely to have come from a slightly higher past productivity θ_1. Since the relevant deviation is for agents to misreport downwards, a frontloaded consumption path for agents with $\epsilon_\gamma(\theta_1, \theta_2) > 0$ exacerbates their present bias, which helps deter them from misreporting θ_1. Furthermore, the degree of frontloading or backloading is increasing in the absolute value of $\epsilon_\gamma(\theta_1, \theta_2)$, the labor wedge at $t = 1$, and consumption $c_1(\gamma, \theta_1)$.

Most importantly, Proposition 8 shows that the retirement savings of present-biased agents also depend directly on past investments in education through the consumption at $t = 0$ and indirectly through the labor wedge and consumption at $t = 1$. For the four-period model, this direct link between education investment and retirement savings works in a similar fashion as in the three-period model, albeit with two differences. First, the mechanism is the same for H-agents, but for L-agents the savings commitment depends on the whole productivity profile (θ_1, θ_2). Similar to the three-period model, the commitment is stronger for agents' whose productivity profile is more likely to have come from genuine L-agents. Second, this direct effect of education investment on retirement savings is weaker in the four-period model than in the three-period model by a factor of $1 - \beta$. However, together with the indirect effect—through the labor wedge and consumption at $t = 1$, the role of education-dependent policies on retirement savings remains important.