American Economic Journal: Macroeconomics 2023, 15(3): 153–190
https://doi.org/10.1257/mac.20200181

Efficient Consolidation of Incentives for Education and Retirement Savings†

By Radoslaw Paluszynski and Pei Cheng Yu*

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.* (*JEL* G51, H21, H24, I26, J24, J26)

The average cost of higher education in the United States 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 United States that suggest making retirement savings contingent on student loan repayment, which establishes a link between retirement and education policies.[1] These proposals are based

*Paluszynski: University of Houston (email: rpaluszynski@uh.edu); Yu: University of New South Wales (email: pei-cheng.yu@unsw.edu.au). Virgiliu Midrigan was coeditor for this article. 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, as well as seminar participants at Academia Sinica, Chinese University of Hong Kong, Hong Kong Baptist University, National Tsing Hua University, National Sun Yat-sen University, Shih Hsin University, and University of New South Wales. We also received helpful feedback at the 2018 Texas Theory Camp, 2019 Asia Econometric Society meeting, 2019 North American Econometric Society summer meeting, Nineteenth Annual SAET Conference, and One Hundred Twelfth Annual Congress of the National Tax Association. Pei Cheng Yu would like to thank the Institute of Economics at Academia Sinica for their hospitality during the pandemic.

†Go to https://doi.org/10.1257/mac.20200181 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[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 One Hundred Sixteenth Congress, which met from January 3, 2019, to January 3, 2021. These bills allow employer 401(k) matching based on student loan payments.

on a pathbreaking Internal Revenue Service (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. 2015). In our framework, agents initially differ in their innate ability, which could be either high or low. Based on their innate ability, agents choose their level of education: college or high school. Afterward, 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' 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 undersaving 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 noncollege 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 noncollege graduates who earn sufficiently high income. This difference in how commitment is provided between college graduates and noncollege 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 noncollege 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 front-loaded for college graduates and more back-loaded for noncollege

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

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 noncollege 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 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 percent 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.

A. *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

[3] We use an extended definition of college that includes master's, doctoral, and professional degrees.

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 (2023) 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 affect 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 Kapička (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 (2021) examine the learning-by-doing aspect of human capital accumulation, so human capital is acquired stochastically as a by-product from labor effort. Kapička 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 (2020) 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 (2021) focus on the design of retirement

[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.

savings policies for time-inconsistent agents in a Mirrlees setting by examining a multidimensional screening environment.

The rest of this paper is organized as follows. Section I presents our model, and Section II characterizes the optimal wedges. In Section III, we calibrate the model and perform the quantitative analysis. Section IV demonstrates a policy that decentralizes the optimum and examines the quantitative effects of the policy proposed in the US Congress. Section V discusses some extensions. The replication package is available as Paluszynski and Yu (2023).

I. 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 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 (invest e_H) or not (invest e_L) in higher education. 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 as well as 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 $\theta \in \Theta = [\underline{\theta}, \bar{\theta}] \subset \mathbb{R}_+$. Productivity is drawn from a differentiable distribution with c.d.f. $F(\theta \mid \kappa)$, which depends on human capital κ and is ranked according to first-order stochastic dominance: if $\kappa > \kappa'$, $F(\theta \mid \kappa) < F(\theta \mid \kappa')$, $\forall \theta \in \Theta$. Also, let $f(\theta \mid \kappa)$ denote the p.d.f. and assume $f(\theta \mid \kappa) > 0$ 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

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

differentiable, strictly increasing, and strictly convex: $h', h'' > 0$, with $h(0) = 0$. γ-agents have the following utility at $t = 0$:

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

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 III.

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 $t = 1$ undersave for retirement. Section VB considers an economy with nonsophisticated agents.

A. *Planning Problem*

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

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

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

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

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

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

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

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

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

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

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

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

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

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

It is worth emphasizing that apart from the inherent investment risk, education is costly for two additional reasons. First, it is costly in terms of resources. Second, it is costly in terms of time, because receiving education delays entry into the labor market.

B. *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 nondecreasing 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 2020). In essence, in contrast to the literature with time-consistent agents, ensuring that 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.

The second difficulty lies in the direction of the relevant deviation at $t = 0$. Usually, the relevant deviation is downward 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 (2020) showed that the optimal allocations are usually nonmonotonic with respect to ex ante information. The nonmonotonicity 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 IIID.

C. *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.

[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 Hao (2000). However, finding suitable conditions that guarantee the sufficiency of local incentive constraints for present-biased agents is difficult (Galperti 2015; Yu 2020).

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

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

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

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

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

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

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

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

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 $\hat{\tau}^k$: $1 - \hat{\tau}_t^k = (1/\beta)(1 - \tau_t^k)$. 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

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

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 undersupply (oversupply) of labor given the market wage.

D. *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.

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

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

By Jensen's inequality, the inverse Euler equation implies $u'\big(c_0(\gamma)\big) < \mathbb{E}_\theta\big[u'\big(c_1(\gamma,\theta)\big)\big]$ 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 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:

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

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.

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,

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

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

[8] See Golosov, Kocherlakota, and Tsyvinski (2003) for more on the inverse Euler equation for time-consistent agents.

advantage of the present bias by back-loading 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. Online Appendix H.3 provides a detailed characterization of the wedges when γ is observable.

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 2021, 2020). To see how, suppose that only H-agents can have productivities greater than θ_H where $\underline{\theta} < \theta_H < \bar{\theta}$, so $f(\theta \mid \kappa_H), f(\theta \mid \kappa_{L,H}) > 0$ and $f(\theta \mid \kappa_L) = 0$ for any $\theta \in (\theta_H, \bar{\theta}]$.[9] In this environment, it is possible to detect misreports on innate ability from some H-agents. Thus, the government can deter misreporting by punishing agents who are caught lying—those who reported $\gamma = L$ and $\theta \in (\theta_H, \bar{\theta}]$. Dishonest agents are punished with a more front-loaded consumption path that exacerbates their present bias: high $c_1(L, \theta)$ and low $c_2(L, \theta)$ for any $\theta \in (\theta_H, \bar{\theta}]$. Present-biased agents are tempted by the front-loaded 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 $\theta \in (\theta_H, \bar{\theta}]$, so that actual L-agents ($\theta \leq \theta_H$) would not be tempted by the front-loaded consumption path to misreport upward as $\theta \in (\theta_H, \bar{\theta}]$. As a result, the set of allocations $\left\{\left(c_1(L, \theta), c_2(L, \theta), y(L, \theta)\right)\right\}_{\theta \in (\theta_H, \bar{\theta}]}$ 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 ID 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 2021). More details are provided in online 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.

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

[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.

II. Theoretical Results

This section presents the optimal intertemporal and labor wedges, which provide the foundations for conditioning retirement savings on education investment. The wedges are derived from the optimization problem and optimality conditions presented in Appendix A.

A. *Intertemporal Wedges*

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

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

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

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

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

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

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

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

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

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:

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

Comparing it with the standard inverse Euler equation (3), the consumption for H-agents is even more front-loaded, while the consumption is relatively back-loaded 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 front-loaded consumption path for L-agents helps discourage downward deviations. As a result, the best the government can do is choose consumption such that the inverse marginal utility is 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

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

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

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

Notice that the retirement savings for L-agents depend on the likelihood ratio $f(\theta\,|\,\kappa_{L,H})/f(\theta\,|\,\kappa_L)$. If $f(\theta\,|\,\kappa_{L,H})/f(\theta\,|\,\kappa_L)$ 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 underinvesting 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 $f(\theta\,|\,\kappa_{L,H})/f(\theta\,|\,\kappa_L)$ 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.

[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 underinvesting in education is complementary to shirking in future periods. Hence, in addition to the usual distortion to deter oversaving, the optimal policy makes the intertemporal distortion worse at the education stage to deter underinvesting in education. If education investment was observable in their environment, like ours, then the intertemporal distortion disappears.

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

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

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

Inequality (8) shows that it is optimal for the government to back-load the consumption of H-agents more than L-agents at $t = 1$. The tightness of (8) increases as $f(\theta | \kappa_{L,H})/f(\theta | \kappa_L)$ decreases. In essence, the degree of back-loading 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.

Now, suppose f satisfies the monotone likelihood ratio property (MLRP): $f(\theta | \kappa)/f(\theta | \kappa')$ 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. Online Appendix I extends our analysis to a model with multiple working periods and finds that optimal retirement savings incentives still depend on education investments, though this effect is now weaker.

B. *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

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

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

[14] See online Appendix H.3 for details.

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

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

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

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

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

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

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

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

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 $\theta f(\theta \mid \kappa)$, while the incentive constraints for higher productivity agents of mass $1 - F(\theta \mid \kappa)$ 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 $1/\phi$. Hence, if the cost of increasing average utility is small relative to the cost of increasing the utility of (γ, x) agents $(1/\phi < 1/u'(c_1(\gamma, x)))$, 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 $u'(c_1(\gamma, \theta))/u'(c_1(\gamma, x))$.

When there is dynamic information and agents are time consistent, then the labor wedge is shaped by both the intratemporal and the 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 $\frac{1 - F(\theta \mid \kappa_{L,H})}{1 - F(\theta \mid \kappa_L)}$ 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 $\frac{1 - F(\theta \mid \kappa_{L,H})}{1 - F(\theta \mid \kappa_L)}$. This shows how the optimal labor distortion for noncollege 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 that of 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.

Online Appendix D quantifies the decomposition of the optimal labor wedge. We show that 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.

III. 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 II, as well as the welfare gains under the optimal tax system. Online Appendix I shows how the dependence of retirement savings subsidies on education investments decreases with finer time periods, so our quantitative estimates are likely loose upper bounds on the significance of education-dependent retirement savings policies.

A. *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, $h(\ell) = \frac{\ell^{1+\frac{1}{\eta}}}{1+1/\eta}$. 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. (2015). 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

[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 nonnegative. 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 IIIC, 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.

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

former with all individuals who hold an associate's degree or less. The share of such respondents 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 to be \$15,700. Online Appendix B.3 discusses the details of our calibration.

In order to calibrate the distributions of skills for agents of different innate ability and education, we create a separate model that we refer to as the "current policies" world. This model is described in detail in online 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 underrepresentation of high-end earnings in the data, we add an upper Pareto tail to each distribution such that the upper 10 percent

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

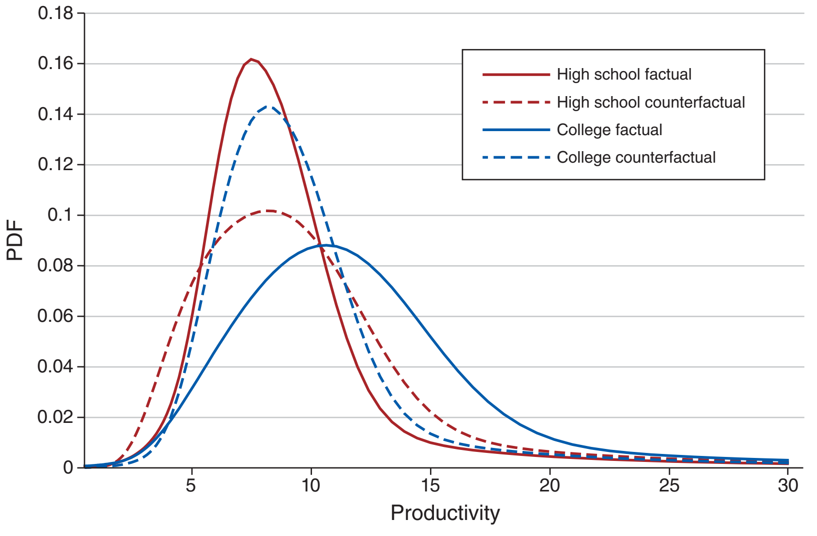

Figure 1. Calibrated Distributions of Skills for the Four Groups of Agents

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.

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 II.

B. *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]

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

TABLE 2—INTERTEMPORAL WEDGES IN PERIOD ZERO: PRESENT-BIAS VERSUS 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

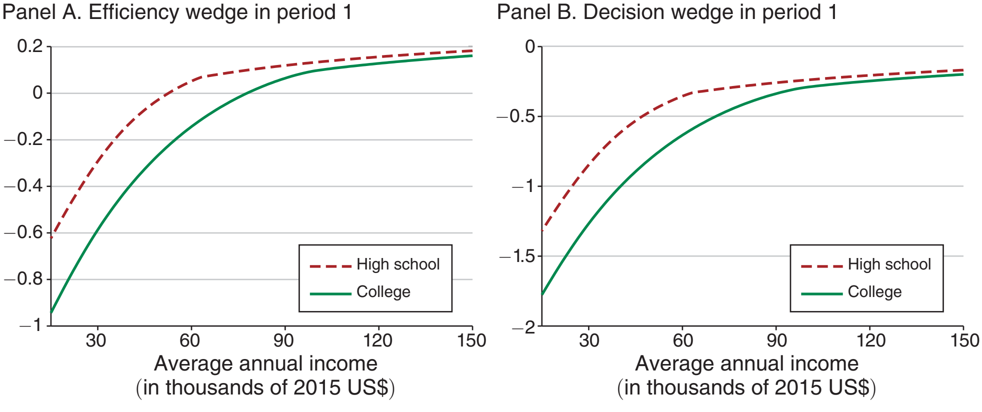

FIGURE 2. INTERTEMPORAL WEDGES IN THE MODEL WITH PRESENT-BIASED AGENTS

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 back-loading 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 is more back-loaded than that of 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 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

[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.

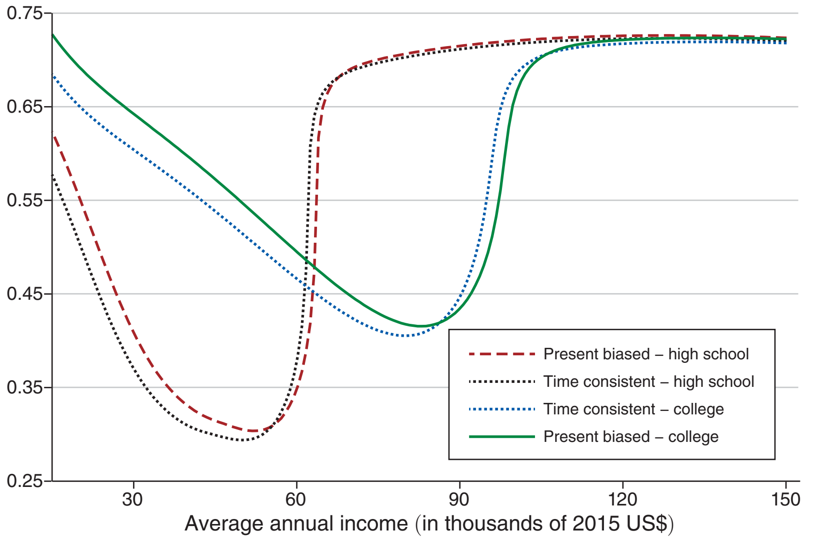

Figure 3. Labor Wedge in the Model with Present-Biased Agents

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. Online Appendix D reinforces this point by showing that the present-bias component of the optimal labor wedge, as introduced in Section IIB, is in general small quantitatively and declines monotonically with income.

Online 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 that the agents are.

C. *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.

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	$\beta = 0.5$	5.76	5.94	5.33
$\beta = 0.7$	1.00	**1.36**	1.32	$\beta = 0.7$	1.75	**1.97**	1.84
$\beta = 0.9$	0.27	0.36	0.37	$\beta = 0.9$	0.32	0.40	0.41

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 that 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 percent of lifetime consumption. The gains come mostly 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 percent of lifetime consumption, but still significant. Since the policy of mandatory savings already forces agents to smooth their consumption, this implies that 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 nonmonotonic 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 online

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

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

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. Online 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.

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 $y\left(H,\tilde{\theta}\right) = y\left(L,\hat{\theta}\right)$, we have

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

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 nontrivial 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 online 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 baseline parametrization (bold numbers), the corresponding gain in lifetime consumption amounts to 0.02 percent. 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.

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* postwar business cycle fluctuations in the United States would be equivalent to a 0.05 percent 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 percent of lifetime consumption. Such parameter values are not empirically implausible as evidenced by the latest estimates of Laibson et al. (2015).

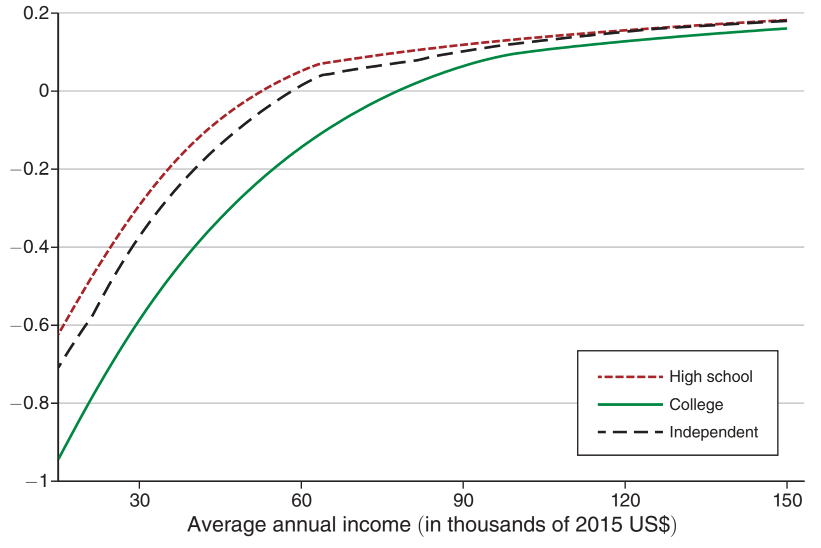

Figure 4. Optimal Education-Dependent and Independent Efficiency Wedge

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
$\beta = 0.7$	0.0104	**0.0167**	0.0262
$\beta = 0.9$	0.0010	0.0016	0.0019

D. *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, one where only L-agents receive it, and one where no agents do.[22]

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 0 up to US$40,000 (the actual calibrated cost, as Table 1 shows, is US$15,700). 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 $\delta_1(e)$.

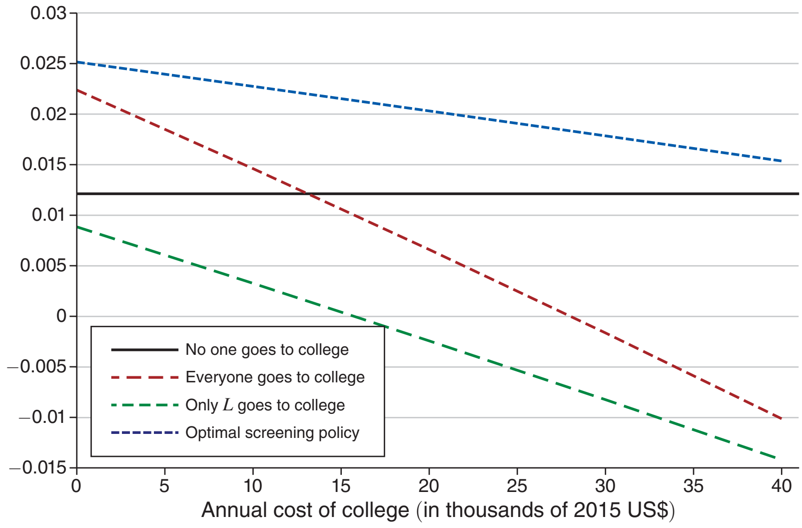

Figure 5. Comparing Optimal Policy to Alternative Screening Policies

IV. 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 $\left\{c_0(\gamma), c_1(\gamma,\theta), c_2(\gamma,\theta)\right\}_{\gamma,\theta\in\Theta}$ can be expressed as a function of income y and education e. It is immediate that by

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

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* $t \geq 1$.

A. *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 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 is 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 pretax 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

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

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

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

[24] The bond savings tax helps the government deter agents from oversaving while simultaneously undersupplying labor (Werning 2010).

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

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

subject to

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

Let $P^{ra} = \left\{ \left[L(e), r(e,y)\right], a(y), [\alpha, \bar{c}], \left[T(y), T^k(b), T^{ra}, g(r)\right] \right\}$ 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 graduates and noncollege 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 that 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

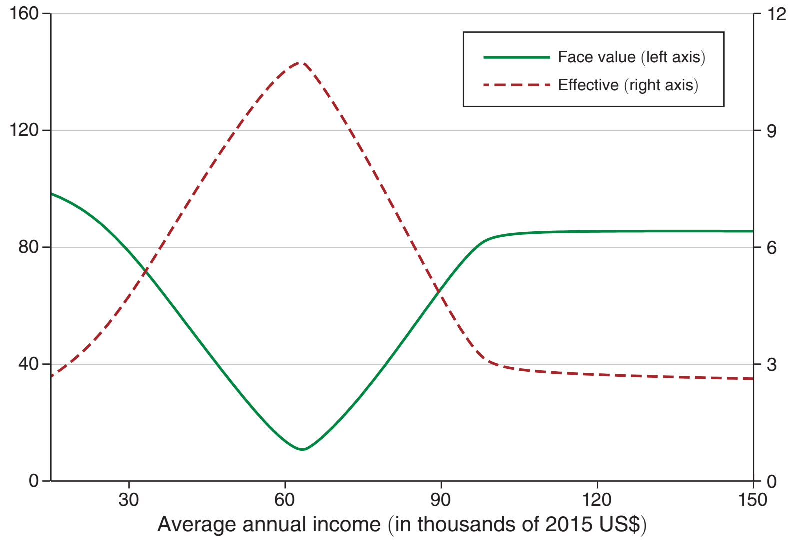

Figure 6. Optimal Student Loan Repayment Schedule

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

B. *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 III 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. Online 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, panel 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

[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.

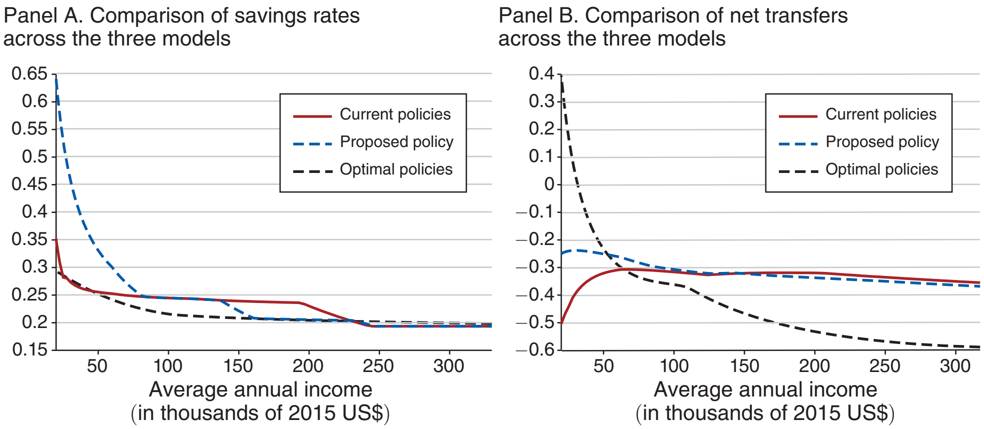

Figure 7. Savings Rates and Transfers of College Graduates across the Three Models

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 ($s_2 > 0$). 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 percent, which is very close to the full information efficient rate of 22.3 percent (and mostly aligns with the constrained efficient rate from our optimal model). In contrast, notice that 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 back-loaded relative to the optimal policies.

Figure 7, panel 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, especially so at the lowest income levels. This contrasts sharply with the optimal Mirrleesian allocations, which redistribute resources toward agents with low income and away from agents with high income. As is evident from Figure 7, panel 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 the 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

[27] Notice that here, in contrast to our previous exercises in Section IIIC, the welfare function is that of the present-biased agents in period $t = 0$, not the one of a paternalistic government.

reform is welfare improving, equivalent to a gain in lifetime consumption of 0.18 percent. 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 percent of lifetime consumption.

V. Extensions

A. *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 multidimensional 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 the following:*

(i) *The inverse Euler equations* (4), (5), *and for any* $\theta \in \Theta$,

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

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

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

(ii) *The labor wedge for H-agents satisfies* (9), *and for L-agents,*

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

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

[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 front-loading consumption for high-income L-agents. Here, this front-loading could be more perverse. It takes advantage of the fact that H-agents value retirement consumption more than L-agents, so a restriction on retirement savings further deters downward deviations by H-agents. This logic is similar to the key finding in Golosov et al. (2013), which shows that discouraging the consumption of a good preferred by high types among low types raises welfare. The labor wedge for L-agents $\tau^w(L,\theta)$ also differs from the case with homogeneous β. Recall that 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 that this case is a theoretical curiosity since empirical studies have demonstrated pervasive present-biased behavior among college students (Steel 2007).

B. *Nonsophistication*

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 undersaving. 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 nonsophisticated agents.

For nonsophisticated agents, the government can use off-path policies to take advantage of their incorrect beliefs. Following Yu (2021), 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 back-loaded consumption path—high retirement consumption but lower working period consumption—compared to the on-path option. At $t = 0$, the nonsophisticated agents underestimate their present bias and thus overestimate the value of retirement consumption to their future selves. As a result, they mispredict that they will select the decoy option in $t = 1$. In reality, their future selves prefer the more front-loaded on-path option instead. Therefore, the government can exploit this incorrect belief by promising college graduates with high retirement benefits—which never needs to be implemented on the equilibrium path—to induce investment in higher education. In other words, the inclusion of a decoy option in the menu can relax the ex ante incentive constraint. In fact, Yu (2021) 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 online 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 nonsophisticated 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.

C. *Nonpaternalism*

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

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

where $\hat{U}_1(c_1, c_2, y; \theta) = u(c_1) - h(y/\theta) + \delta_2 u(c_2)$ and $U_1(c_1, c_2, y; \theta) = u(c_1) - h(y/\theta) + \beta \delta_2 u(c_2)$. 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 nonpaternalistic government that adopts the $t = 0$ agents' preferences, and (iii) a nonpaternalistic government

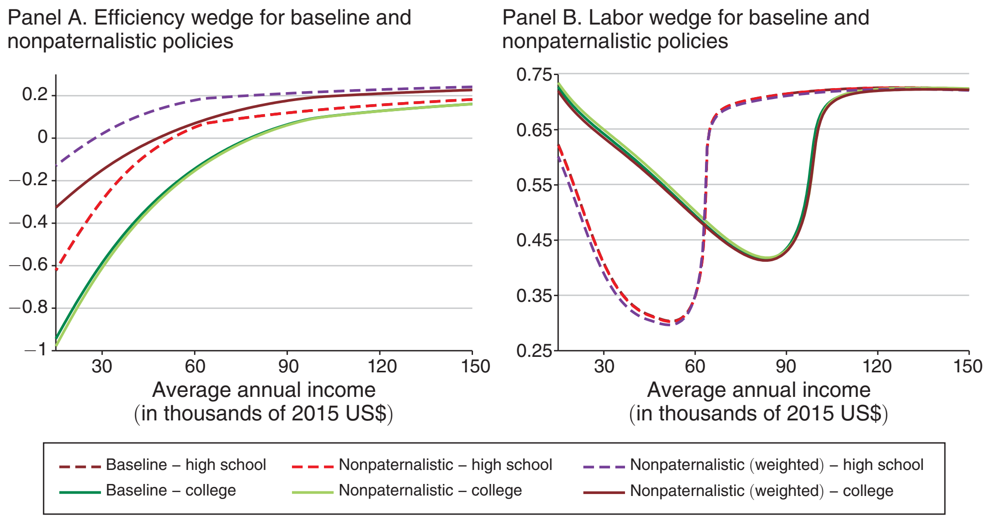

Figure 8. Optimal Wedges with a Non-Paternalistic Government

that puts $\chi = 0.5$ weight on the $t = 1$ agents' preferences. A few observations are noteworthy. First, as evident in Figure 8, panel A, the nonpaternalistic 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 nonpaternalistic government adopts the $t = 1$ agents' preferences, the optimal wedges are smaller in absolute value and involve more front-loading of consumption, a result that aligns with basic intuition: a nonpaternalistic 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, Figure 8, panel B shows that the labor wedges are virtually unaffected by any nonpaternalism considerations.

D. *Alternative Assumptions on Timing*

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

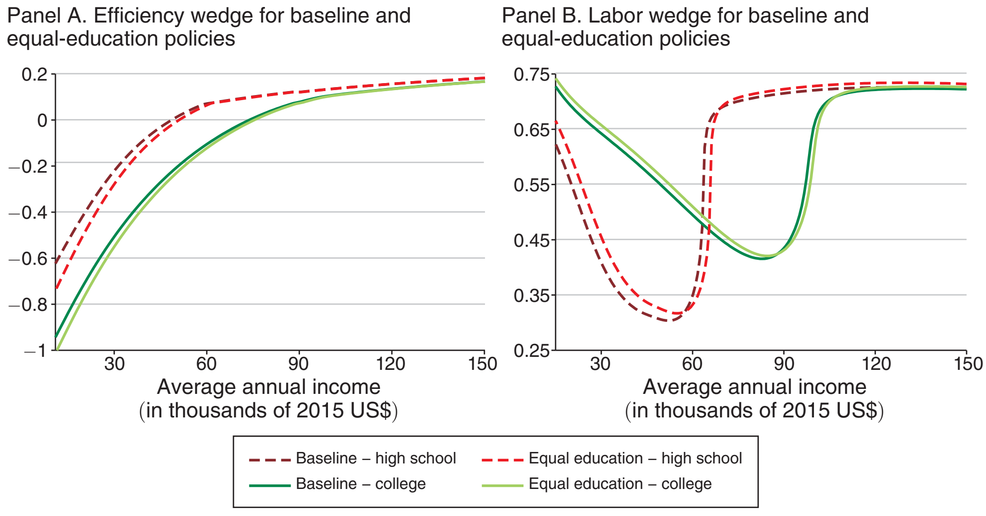

FIGURE 9. OPTIMAL WEDGES WITH EQUAL EDUCATION PERIOD LENGTHS

Yu (2021). 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 IB, 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 online 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.

Difference in Length of Education Periods.—Our baseline parametrization assumes a difference in length of the education period of around five 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 downward, 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 upward for the lowest and highest incomes, and downward for the middle range, but any differences relative to the baseline are small.

VI. 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 United States—for parents.

Appendix A. The Optimization Problem

Given Lemma 1, the relaxed optimal tax problem is

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

subject to

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

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

$$\delta_0(e_H) u\big(c_0(H)\big) + \beta\delta_1(e_H) \int_{\underline{\theta}}^{\overline{\theta}} \big[U_1(H,\theta) + (1-\beta)\delta_2 u\big(c_2(H,\theta)\big)\big] f(\theta \mid \kappa_H) d\theta$$
$$\geq \delta_0(e_L) u\big(c_0(L)\big) + \beta\delta_1(e_L) \int_{\underline{\theta}}^{\overline{\theta}} \big[U_1(L,\theta) + (1-\beta)\delta_2 u\big(c_2(L,\theta)\big)\big] f(\theta \mid \kappa_{L,H}) d\theta,$$

and the resource constraint. As is standard, we ignore the monotonicity constraint—$y(\gamma,\theta)$ is nondecreasing 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:

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

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

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

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

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

$$= \frac{\phi\,\pi_L\delta_1(e_L)f\big(\theta\,|\,\kappa_L\big)}{u'\big(c_2(L,\theta)\big)},$$

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

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

$$\lambda_\gamma(\theta)\frac{1}{\theta}h'\left(\frac{y(\gamma,\theta)}{\theta}\right) + \xi_\gamma(\theta)\left[\frac{1}{\theta^2}h'\left(\frac{y(\gamma,\theta)}{\theta}\right) + \frac{y(\gamma,\theta)}{\theta^3}h''\left(\frac{y(\gamma,\theta)}{\theta}\right)\right]$$

$$= \phi\,\pi_\gamma\delta_1(e_\gamma)f\big(\theta\,|\,\kappa_\gamma\big).$$

In online Appendix A, we show how the theoretical results follow from these conditions.

REFERENCES

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

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

Anderberg, Dan. 2009. "Optimal Policy and the Risk Properties of Human Capital Reconsidered." *Journal of Public Economics* 93 (9–10): 1017–26.

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

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

Bergemann, Dirk, and Juuso Välimäki. 2019. "Dynamic Mechanism Design: An Introduction." *Journal of Economic Literature* 57 (2): 235–74.

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

Bovenberg, A. Lans, and Bas Jacobs. 2005. "Redistribution and Education Subsidies Are Siamese Twins." *Journal of Public Economics* 89 (11–12): 2005–35.

Cadena, Brian C., and Benjamin J. Keys. 2015. "Human Capital and the Lifetime Costs of Impatience." *American Economic Journal: Economic Policy* 7 (3): 126–53.

Courty, Pascal, and Li Hao. 2000. "Sequential Screening." *Review of Economic Studies* 67 (4): 697–717.

Craig, Ashley C. 2023. "Optimal Income Taxation with Spillovers from Employer Learning." *American Economic Journal: Economic Policy* 15 (2): 82–125.

Cronqvist, Henrik, Richard H. Thaler, and Frank Yu. 2018. "When Nudges Are Forever: Inertia in the Swedish Premium Pension Plan." *AEA Papers and Proceedings* 108: 153–58.

Cunha, Flavio, and James J. Heckman. 2007. "Identifying and Estimating the Distributions of Ex Post and Ex Ante Returns to Schooling." *Labour Economics* 14 (6): 870–93.

Diamond, Peter A. 1998. "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates." *American Economic Review* 88 (1): 83–95.

Farhi, Emmanuel, and Xavier Gabaix. 2020. "Optimal Taxation with Behavioral Agents." *American Economic Review* 110 (1): 298–336.

Farhi, Emmanuel, and Iván Werning. 2007. "Inequality and Social Discounting." *Journal of Political Economy* 115 (3): 365–402.

Farhi, Emmanuel, and Iván Werning. 2010. "Progressive Estate Taxation." *Quarterly Journal of Economics* 125 (2): 635–73.

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

Gabaix, Xavier. 2014. "A Sparsity-Based Model of Bounded Rationality." *Quarterly Journal of Economics* 129 (4): 1661–1710.

Galperti, Simone. 2015. "Commitment, Flexibility, and Optimal Screening of Time Inconsistency." *Econometrica* 83 (4): 1425–65.

Gary-Bobo, Robert J., and Alain Trannoy. 2015. "Optimal Student Loans and Graduate Tax under Moral Hazard and Adverse Selection." *RAND Journal of Economics* 46 (3): 546–76.

Golosov, Mikhail, Narayana Kocherlakota, and Aleh Tsyvinski. 2003. "Optimal Indirect and Capital Taxation." *Review of Economic Studies* 70 (3): 569–87.

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski. 2016. "Redistribution and Social Insurance." *American Economic Review* 106 (2): 359–86.

Golosov, Mikhail, Maxim Troshkin, Aleh Tsyvinski, and Matthew Weinzierl. 2013. "Preference Heterogeneity and Optimal Capital Income Taxation." *Journal of Public Economics* 97: 160–75.

Grochulski, Borys, and Tomasz Piskorski. 2010. "Risky Human Capital and Deferred Capital Income Taxation." *Journal of Economic Theory* 145 (3): 908–43.

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

Kapička, Marek. 2015. "Optimal Mirrleesean Taxation in a Ben-Porath Economy." *American Economic Journal: Macroeconomics* 7 (2): 219–48.

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

Koeniger, Winfried, and Julien Prat. 2018. "Human Capital and Optimal Redistribution." *Review of Economic Dynamics* 27: 1–26.

Krähmer, Daniel, and Roland Strausz. 2015. "Dynamic Mechanism Design." In *An Introduction to the Theory of Mechanism Design*, edited by Tilman Börgers, 204–34. Oxford, UK: Oxford University Press.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–77.

Laibson, David, Peter Maxted, Andrea Repetto, and Jeremy Tobacman. 2015. "Estimating Discount Functions with Consumption Choices over the Lifecycle." Unpublished.

Lockwood, Benjamin B. 2020. "Optimal Income Taxation with Present Bias." *American Economic Journal: Economic Policy* 12 (4): 298–327.

Lucas, Robert E. 1987. *Models of Business Cycles*. Oxford, UK: Basil Blackwell.

Madrian, Brigitte C., and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116 (4): 1149–87.

Makris, Miltiadis, and Alessandro Pavan. 2021. "Taxation under Learning by Doing." *Journal of Political Economy* 129 (6): 1878–1944.

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

Myerson, Roger B. 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323–58.

Nakajima, Makoto. 2012. "Rising Indebtedness and Temptation: A Welfare Analysis." *Quantitative Economics* 3: 257–88.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *American Economic Review* 89 (1): 103–24.

Paluszynski, Radoslaw, and Pei Cheng Yu. 2023. "Replication data for: Efficient Consolidation of Incentives for Education and Retirement Savings." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.38886/E154761V1.

Pavoni, Nicola, and Hakki Yazici. 2017. "Intergenerational Disagreement and Optimal Taxation of Parental Transfers." *Review of Economic Studies* 84 (3): 1264–1305.

Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *Review of Economic Studies* 68 (1): 205–29.

Stantcheva, Stefanie. 2017. "Optimal Taxation and Human Capital Policies over the Life Cycle." *Journal of Political Economy* 125 (6): 1931–90.

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

Werning, Iván. 2010. "Nonlinear Capital Taxation." Unpublished.

Yu, Pei Cheng. 2020. "Seemingly Exploitative Contracts." *Journal of Economic Behavior and Organization* 176: 299–320.

Yu, Pei Cheng. 2021. "Optimal Retirement Policies with Present-Biased Agents." *Journal of European Economic Association* 19 (4): 2085–2130.