

Optimal Life-cycle Capital Taxation under Self-Control Problems*

Nicola Pavoni[†] and Hakki Yazici[‡]

March 22, 2015

Abstract

We study optimal taxation of savings in an economy where agents face self-control problems, and we allow the severity of self-control to change over the life cycle. We focus on quasi-hyperbolic discounting with constant elasticity of intertemporal substitution utility functions and linear Markov equilibria. We derive explicit formulas for optimal taxes that implement the efficient (commitment) allocation. We show, analytically, that if agents' ability to self-control increases concavely with age, then savings should be subsidized and the subsidy should decrease with age. We also study the quantitative effects of age-dependent self-control problems and find that the optimal subsidies in our environment are much larger than those implied by models with constant self-control. Finally, we compare our optimal subsidies with those implied by the 401(k) plan. Although the subsidy levels in the two cases are of comparable magnitudes, the 401(k) plan implies an increasing pattern of subsidies while the optimal subsidies decrease over the life cycle.

JEL classification: E21, E62, D03.

Keywords: Self-control problems, Linear Markov equilibrium, Life-cycle taxation of savings.

*We thank the editor, Morten O. Ravn, two anonymous referees, as well as Per Krusell, John Leahy, and seminar participants at Bogazici University, Goethe University in Frankfurt, IIES, the SED meetings in Ghent, University of Alicante, University of Bologna, University of Oxford for their comments and suggestions. Yazici gratefully acknowledges financial support from the European Community Framework Programme through the Marie Curie International Reintegration Grant #268457.

[†]Bocconi University, IGIER, IFS, and CEPR.

[‡]Sabanci University.

1 Introduction

Economists traditionally assume that people discount streams of utility over time exponentially. An important implication of exponential discounting is that under this assumption people have time-consistent intertemporal preferences: How an individual feels about a given intertemporal tradeoff is independent of when he is asked. However, laboratory and field studies on intertemporal choice have cast doubt on this assumption.¹ This evidence suggests that discounting between two future dates gets steeper as we get closer to these dates. Such time-inconsistent intertemporal preferences capture self-control problems. Naturally, all this evidence on self-control problems have led many economists to model this phenomenon and study its positive and normative implications.²

In this paper, we study optimal capital taxation over the life cycle in the presence of self-control problems. A common modeling assumption in the literature on self-control problems is that the degree of self-control problem is constant over time. This contrasts with the significant body of empirical research indicating that, like many other personality traits, people's ability to self-control changes as they age. Using an experimental approach to measuring self-control problems, one of the conclusions Ameriks, Caplin, Leahy, and Tyler (2007) reach is summarized in the following quotation: *"A particularly interesting finding in Table 4 is the profound reduction in the scale of self-control problems as individuals age, which shows up only when one uses the absolute value of the self-control measure. Older individuals experience fewer self-control problems, either of overconsumption or underconsumption, than do their younger counterparts. This finding is certainly consistent with the common view that temptation falls with age, and has important connections with actual consumption behavior over the life cycle. Models that allow for such a time-changing self-control parameter retirement may be necessary to explain the absence of a spike in consumption spending at the point when retirement assets become fully liquid."*

Another set of evidence for changing levels of self-control problems comes from literature

¹See DellaVigna (2009) for a survey of field studies and Frederick, Loewenstein, and O'Donoghue (2002) for a survey of experimental studies. Also, see Laibson, Repetto, and Tobacman (2007) for evidence of self-control problems in consumption asset holdings panel data.

²Three main models that have been proposed to capture self-control problems are the hyperbolic discounting model of Laibson (1997), the temptation model of Gul and Pesendorfer (2001), and the planner doer model of Thaler and Shefrin (1981).

that investigates intertemporal discounting over the life span. This research has shown that the discrepancy between short term and long term discount rates falls with age predicting a life-cycle developmental trend toward increased self-control. Read and Read (2004) is an experimental study which estimates hyperbolic discount functions for three different groups of individuals: young adults, middle-aged, and old adults, with average ages of 25, 44, and 75. The authors find that the discrepancy between the short-term and the long-term discount rates is significantly higher for the young adults relative to the middle-aged. Similarly, this discrepancy is significantly higher for middle-aged adults compared to the old adults group, for which, they conclude the discrepancy virtually vanishes. Green, Myerson, and O’Staszewski (1999) estimates hyperbolic discount functions for two groups of adults, young and old, with mean ages of 20 and 70 respectively, and finds that the discrepancy between short-term and long-term discount rates that exists for young adults disappears as people get old.³

A third set of evidence for changing self-control problems comes from personality psychology. As Ameriks, Caplin, Leahy, and Tyler (2007) states “personality psychologists associate self-control with conscientiousness, one of the ‘big five’ personality factors.”⁴ There is a long list of empirical studies in personality psychology that show that conscientiousness, and in particular its lower-level facet, self-control, changes with age.⁵ For example, in their survey article on personality development in adulthood, Caspi, Roberts, Robins, and Trzesniewski (2003) conclude that: “it appears that the increase in conscientiousness is one of the most robust patterns in personality development, especially in young adulthood.”

In this paper, we extend the traditional models of self-control to allow for varying degrees of self-control problem over the life cycle, and study optimal capital taxation. In our model, agents make consumption and savings decisions facing self-control problems at all ages. In the last period of their lives, people make consumption and bequest decisions knowing that

³For more empirical work on life-span analysis of intertemporal discounting, see also Green, Fry, and Myerson (1994) and Green, Myerson, Lichtman, Rosen, and Fry (1996).

⁴Actually, Ameriks, Caplin, Leahy, and Tyler (2007) validates this relationship between conscientiousness and the measure of self-control used in the experiment (the EI gap) and finds that “the data reveal a strong relationship between the conscientiousness questions and the absolute value of the EI gap.”

⁵For example, see John, Gosling, Potter, and Srivastava (2003) and Helson, Jones, and Kwan (2002). Ameriks, Caplin, Leahy, and Tyler (2007) also, through their experimental finding, show that there is a profound reduction in the scale of self-control and conscientiousness problems as individuals age.

they are going to be replaced by their offspring next period. We model preferences that exhibit self-control problems through the quasi-hyperbolic discounting framework of Laibson (1997), which builds on the seminal works of Strotz (1955) and Phelps and Pollak (1968). We extend the Laibson (1997) model by allowing for the degree of self-control problem to change over time. We assume people are sophisticated, meaning they anticipate future self-control problems.

In this environment, we define *efficient (or commitment)* allocation as the allocation that would arise in the absence of self-control problems. It is given by the solution to a fictitious social planner's consumption-saving problem where the planner future utilities *exponentially*. In our environment, this preference corresponds to the preference of an initial generation parent. The main exercise in this paper is to examine the optimal tax policy that implements the efficient allocation. In this sense, this paper is a normative exploration of optimal paternalistic tax policy regarding life-cycle saving behavior. It is well-known that in models of quasi-hyperbolic discounting there is multiplicity of equilibria.⁶ We restrict attention to the (unique) linear Markov equilibrium of our economy.

We derive closed form formulas for optimal age-dependent capital taxes. Our closed form solution represents the equilibrium obtainable as the limit of the equilibria of finite-period economies. We show that optimal capital taxes can be positive as well as negative in different periods of life and they can be increasing, decreasing, or changing non-monotonically with age, depending on what we assume about the evolution of self-control problem over the life cycle. This ambiguity result about the qualitative properties of optimal taxes shows that researchers who take self-control problems seriously should also take the evolution of self-control problems over the life cycle seriously before making policy suggestions. This result also questions the basic presumption in the literature that self-control problems always imply optimality of saving subsidies, which, as we show, arises from the assumption of constant self-control over age.

Our closed form tax formulas are obtained assuming agents have constant elasticity of intertemporal substitution (CEIS) preferences. The formulas are valid (i) if CEIS coefficient is one (the utility function is logarithmic) or (ii) for any CEIS coefficient if the economy is in a steady-state. Using these formulas, we prove that if, as suggested by the available empirical evidence, the degree of self-control increases with age, then capital should indeed be *subsidized* in all periods. We also prove that, if self-control increases concavely with age, then optimal

⁶For discussions of multiplicity of equilibria, see, among others, Laibson (1994) and Krusell and Smith (2003).

capital subsidies should *decrease* with age.

We study the quantitative effects of age-dependent self-control problems in a calibrated version of our model and find that the optimal subsidies in our model with decreasing self-control problems are *much larger than those implied by a model with constant self-control problems*. We also compare our optimal subsidies with those implied by the 401(k) plan. If we exclude the very last periods before retirement - where the subsidy rate in the 401(k) essentially mimics the employer matching rate - the subsidy levels in the two cases are of comparable magnitudes. A marked difference emerges however: the 401(k) plan implies an increasing pattern of subsidies while the optimal subsidies decrease over the life cycle.

Our benchmark model assumes perfectly altruistic parents, making it equivalent to a standard infinite horizon framework. Section 4.1 allows for imperfect altruism and generalizes our optimal tax formulas to take into account the finite life time effects of Krusell, Kuruscu, and Smith (2010). We find that, in our model, the effect of age-dependent self-control dominates the finite time effect induced by imperfect altruism: optimal capital subsidies decrease over the life cycle even when parents do not care at all about their offsprings (i.e., when finite life time effect is the strongest). In Section 4.2, we extend the finite horizon version of the model developed in Section 4.1 with borrowing constraints to study the life-cycle consumption implications of our model. We find that, even though our model abstracts from important life cycle issues such as child-rearing and health, life-cycle consumption profiles implied by our calibrated self-control patterns capture key features of empirical life-cycle consumption profiles fairly well.

Related Literature. Our paper is related to a number of recent papers that have explored the implications of self-control problems for optimal paternalistic taxation. O'Donoghue and Rabin (2003) and O'Donoghue and Rabin (2006) analyze models of paternalistic taxation of unhealthy goods. More closely related is Krusell, Kuruscu, and Smith (2010), which analyzes optimal taxation of savings in an economy where agents live finitely many periods and have self-control problems à la Gul and Pesendorfer (2001).⁷ First, they prove that the optimal policy prescriptions of the quasi-hyperbolic model and the temptation model are identical when the utility function is logarithmic or when it is CEIS and the temptation parameter goes to infinity. Second, they show that savings should be subsidized and that this *subsidy should be increasing with*

⁷Krusell, Kuruscu, and Smith (2002) also analyze optimal taxation of savings under self-control problems but their main focus is on an environment where the government as well as the people face issues of time-inconsistency.

time due to finite life time effect.⁸ Our work differs from this paper most importantly by allowing for changing levels of self-control problems over the life cycle: like all papers prior to ours, Krusell, Kuruscu, and Smith (2010) assume that the degree of self-control problem is constant over time. The implications of modeling age-dependent self-control problems turn out to be significant. First, by assuming empirically plausible patterns of self-control problems over the life cycle, we show, analytically, that capital subsidies should actually be decreasing with age. Section 4.1 shows that this result continues to hold even when we extend the model to take into account the finite life time effects of Krusell, Kuruscu, and Smith (2010). Second, we find that, for an agent with a given level of self-control problem, the age-dependence of self-control model imply much higher levels of optimal subsidies relative to the ones implied by the constant self-control model.

Another important paper that is related to ours is Imrohroglu, Imrohroglu, and Joines (2003) which studies the role of social security in a model where agents have self-control problems. They consider a rich overlapping generations model with uninsurable unemployment shocks and liquidity constraints, and find that social security is not very useful in helping agents solve their self-control problems. Ours is a theory of capital subsidies under complete markets. One advantage of our analysis is that, whenever utility is logarithmic, our results are robust to many dimensions of heterogeneity - such as the life-cycle wage profile and the wealth distribution - whereas the normative predictions in models with incomplete markets may obviously depend on all these features.

As discussed above, an immediate implication of age-dependent self-control problems is that capital taxes should be age-dependent. The age-dependence result is also a feature of two sets of earlier contributions that analyze benefits of age-dependent capital income taxes with time-consistent agents. First, in the Ramsey taxation tradition, Erosa and Gervais (2002) shows that, in life-cycle economies, if the government has access to age-dependent linear capital and labor income taxes, the resulting optimal tax system features age-dependence both for capital and labor income. Second, the New Dynamic Public Finance literature calls for age-dependence in optimal capital and labor income tax codes (e.g., Farhi and Werning (2013) and Golosov, Troshkin, and Tsyvinski (2011)). The forces generating age-dependence in the current

⁸It is indeed straightforward to show that, in the infinite horizon version of their model, the subsidies would be constant.

paper, however, are completely different from the forces in these papers.⁹ Therefore, our paper complements this literature by providing a new mechanism through which capital taxes should depend on age. As such, in this paper, the life-cycle pattern of optimal capital taxes depends on features of the environment that are neglected by these papers.

The rest of the paper is organized as follows. Section 2 lays down the baseline model whereas Section 3 analyzes optimal subsidies both qualitatively and quantitatively. In Section 4, we consider imperfect altruism extension and the life-cycle consumption implications of the calibrated self-control model. Section 5 concludes. The proofs of propositions, the procedure we employ to approximate quasi-hyperbolic discount functions, and the analysis of optimal subsidies under partial sophistication and in the presence of an illiquid asset are presented in a separate online Appendix.

2 Model

The economy is populated by a continuum of a unit measure of dynasties who live for a countable infinity of periods, $t = 1, 2, \dots$, where each agent within a dynasty is active for $I + 1$ periods. In the first I periods, agents make consumption saving decisions facing different degrees of self-control problems at different ages. In the last period of their lives, agents decide how much to consume and bequeath to their offspring, who replace them in next period. People are altruistic and they anticipate their offspring's self-control problems. We use quasi-hyperbolic discounting framework formalized by Laibson (1997) to model self-control problems as follows. An agent in his ultimate period of life (we refer to this agent as the parent hereafter) has the following preferences over dynastic consumption stream:

$$u(c_0) + \delta u(c_1) + \delta^2 u(c_2) + \dots + \delta^I u(c_I) + \delta^{I+1} u(c'_0) + \dots$$

where c_0 , c_i , and c'_0 refer to the the consumption levels of the current parent, the offspring at age i , and the offspring when he becomes a parent, respectively. u is the instantaneous utility

⁹The optimality of age-dependence in Erosa and Gervais (2002) is a direct implication of time-dependent consumption and labor plans that is valid even in the steady state due to life cycle changes in people's productivity. In the New Dynamic Public Finance models, capital is taxed in order to deter people from joint deviation of saving and shirking. Since people at different ages (and contingencies) have different levels of accumulated wealth and future prospects, they have different tendencies to save, and hence, the corrective taxes depend on age.

function and δ refers to both the discount factor and the altruism factor. The offspring has different preferences at different ages. His preferences are given by

$$u(c_i) + \beta_i \delta \left[u(c_{i+1}) + \delta u(c_{i+2}) + \dots + \delta^{I-i} u(c'_0) + \dots \right], \text{ at age } i \in \{1, \dots, I-1\},$$

$$u(c_I) + \beta_I \delta \left[u(c'_0) + \dots \right], \text{ at age } I.$$

When $\beta_i = 1$ for all i , all agents at all ages are time-consistent as there is no self-control problem. Throughout the paper we will assume that $\beta_i < 1$, meaning individuals postpone their planned savings when the date of saving comes. Prior to the current paper, the literature has assumed $\beta_i = \beta$ for all i , meaning the degree of self-control problem is constant as people age. Following the large body of empirical findings provided by personality psychologists and experimental studies, we allow for the severity of self-control problems, β_i to depend on i . In the baseline model, we assume people are fully sophisticated, meaning they fully anticipate the self-control problems faced by future selves and descendants. In Appendix C, we analyze a version of the model in which people are only partially aware of their self-control problems.¹⁰

The instantaneous utility function, u , is of the CEIS form with elasticity parameter $\sigma > 0$:

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}, \text{ for } \sigma \neq 1;$$

$$= \log(c), \text{ else.}$$

Production takes place at the aggregate level according to the function $F(k, l)$, where k is aggregate capital and l is aggregate labor. The production function satisfies the usual neoclassical properties together with the Inada conditions:

$$F_1, F_2 > 0 ; F_{11}, F_{22} \leq 0; \quad \text{and} \quad \lim_{k \rightarrow 0} F_1 = \infty; \quad \lim_{k \rightarrow \infty} F_1 = 0.$$

Labor is inelastically supplied, so at all dates $l = 1$. Define $f(k) = F(k, 1) + (1-d)k$, where d refers to the fraction of capital that is forgone due to depreciation. There is a credit market in which agents can trade one period risk-free bonds and capital as perfectly substitutable assets. Since at any given date there is not cross-sectional heterogeneity, all agents have the same level of asset holdings. Hence, letting b_t be the amount of asset holdings of the agent alive in period t , the credit market clearing condition is $k_t = b_t$.

¹⁰See Ariely and Wertenbroch (2002) for behavioral evidence on partial sophistication.

2.1 The Efficient Allocation

The *efficient* or – as we use interchangeably throughout the paper – the *commitment* allocation is the allocation that would arise in the absence of self-control problems. It is given by the solution to a fictitious social planner’s consumption-saving problem where the planner discounts exponentially with discount factor δ . In our environment, this preference corresponds to the preference of an initial generation parent. By taking a long-term perspective and evaluating welfare according to the initial parent’s preference, we are following much of the literature.¹¹ The following Euler Equations characterize the efficient allocation, which we denote with an asterisk throughout the paper:¹²

$$\begin{aligned} u'(c_i^*) &= \delta f'(k_i^*) u'(c_{i+1}^*), \text{ for } i \in \{0, 1, \dots, I-1\}, \text{ and} \\ u'(c_I^*) &= \delta f'(k_I^*) u'(c_0^*), \dots \end{aligned} \tag{1}$$

2.2 Implementing the Efficient Allocation

Since people in this economy face self-control problems, laissez-faire market equilibrium cannot attain the commitment allocation. Our main interest in this paper is to find and characterize a tax system that implements the commitment allocation in the market environment. We call such a tax system *optimal*. We proceed by defining a market equilibrium with taxes. It is important to note that from the outset we restrict the set of taxes that are available to the government to linear taxes on savings coupled with lump-sum rebates (throughout the paper we call this the set of linear taxes). In general, it is not obvious that there is a linear tax system that implements the efficient allocation. However, since we focus our attention to linear equilibria, a linear tax system that implements the efficient allocation exists. We will verify this claim in Section 3.

2.3 Markov Equilibrium with Taxes

For notational simplicity, here in the main text, we only present the stationary version of the model where the level of aggregate capital stock starts from its steady-state level, k . The prices

¹¹See DellaVigna and Malmendier (2004), Gruber and Koszegi (2004) and O’Donoghue and Rabin (2006), for example.

¹²We do not state the transversality condition but the commitment allocation will converge to a steady state with positive capital as long as $k_0 > 0$.

at the steady state are given by

$$R = f'(k) \quad \text{and} \quad w = f(k) - f'(k)k. \quad (2)$$

In such a world, the only index we need to carry around is the age index i . In Appendix A.1, we provide the general setup where the economy starts from an arbitrary level of capital stock and prices change over time. We prove our main result, Proposition 1, for the general case, and show that if the utility function is logarithmic, then optimal taxes do not depend on whether the economy is at a steady state or in a transition.

Let τ_i be the savings (capital) tax agent $i = 0, 1, \dots, I$ pays. Tax proceeds are rebated in a lump-sum manner in every period. Denote the lump-sum rebate in period i by T_i and let $\tau = \{\tau_i, T_i\}_{i=0}^I$. For each set of taxes, we define the policy functions $b_i(\cdot; \tau)$ for $i = 0, 1, \dots, I$, describing the optimal level of asset holdings of agent i given prices and taxes. When agent of age n is deciding b_n , his evaluation of the effect of his choice on $b_i, i > n$ will be described by the (nested) function $b_i(b_{i-1}(\dots b_{n+1}(b_n; \tau)\dots; \tau); \tau)$, which will be referred to as $b_i(\dots(b_n)\dots)$ so as to simplify notation. In addition, in order to only deal with functions, we assume each agent's solution is unique, a property satisfied by our closed form solution involving linear policies. Of course, in case of multiple solutions, our policy functions correspond to appropriate selections from the policy correspondences.

A Stationary Markov equilibrium with taxes τ consists of a level of capital k , prices R, w , value function $V(\cdot; \tau)$ and policy functions $\{b_i(\cdot; \tau)\}_{i=0}^I$ such that: (i) the prices satisfy (2); (ii) the value functions and the policies are consistent with the parent's problem described below; (iii) the government budget is satisfied period-by-period and markets clear: $T_i = R\tau_i b_i(k; \tau)$ and $b_i(k; \tau) = k$ for all i .

We now formally define the parent's problem. Let $V(b; \tau)$ be the value of a parent's problem who had b units of assets in his last period before parenthood and faces the tax system τ . The parent chooses his bequest b_0 and does not have any direct control over b_1, \dots, b_I . Note that his preferences are not aligned with his offspring's (in a given period i , parent's discount factor is δ whereas offspring's is $\beta_i \delta$). The parent foresees this misalignment of preferences, and correctly forecasts future policies.

The parent's value and the policies are given by the solution to the following problem:

$$V(b; \tau) = \max_{b_0} u(R(1 - \tau_I) b + w + T_I - b_0)$$

$$+ \delta \left\{ \sum_{i=0}^{I-1} \delta^i u(R(1 - \tau_i) b_i(\dots(b_0)\dots) + w + T_i - b_{i+1}(\dots(b_0)\dots)) + \delta^I V(b_I(\dots(b_0)\dots); \tau) \right\}$$

s.t. for all b_0

$$b_1(b_0; \tau) = \arg \max_{\hat{b}_1} u(R(1 - \tau_0) b_0 + w + T_0 - \hat{b}_1) + \delta \beta_1 u(R(1 - \tau_1) \hat{b}_1 + w + T_1 - b_2(\hat{b}_1))$$

$$+ \delta \beta_1 \left\{ \sum_{i=2}^{I-1} \delta^{i-1} u(R(1 - \tau_i) b_i(\dots(\hat{b}_1)\dots) + w + T_i - b_{i+1}(\dots(\hat{b}_1)\dots)) + \delta^{I-1} V(b_I(\dots(\hat{b}_1)\dots); \tau) \right\}$$

s.t. for all b_1

$$b_2(b_1; \tau) = \arg \max_{\hat{b}_2} u(R(1 - \tau_1) b_1 + w + T_1 - \hat{b}_2) + \delta \beta_2 u(R(1 - \tau_2) \hat{b}_2 + w + T_2 - b_3(\hat{b}_2))$$

$$+ \delta \beta_2 \left\{ \sum_{i=3}^{I-1} \delta^{i-2} u(R(1 - \tau_i) b_i(\dots(\hat{b}_2)\dots) + w + T_i - b_{i+1}(\dots(\hat{b}_2)\dots)) + \delta^{I-2} V(b_I(\dots(\hat{b}_2)\dots); \tau) \right\}$$

...

$$b_{I-1}(b_{I-2}; \tau) = \arg \max_{\hat{b}_{I-1}} u(R(1 - \tau_{I-2}) b_{I-2} + w + T_{I-2} - \hat{b}_{I-1})$$

$$+ \delta \beta_{I-1} \left\{ u(R(1 - \tau_{I-1}) \hat{b}_{I-1} + w + T_{I-1} - b_I(\hat{b}_{I-1})) + \delta V(b_I(\hat{b}_{I-1}); \tau) \right\} \quad (3)$$

s.t. for all b_{I-1}

$$b_I(b_{I-1}; \tau) = \arg \max_{\hat{b}_I} u(R(1 - \tau_{I-1}) b_{I-1} + w + T_{I-1} - \hat{b}_I) + \delta \beta_I V(\hat{b}_I; \tau). \quad (4)$$

To understand the nested nature of policies better, let us analyze the definition of policies in (3) and (4). First, constraint (4) describes how self I chooses b_I . This agent chooses b_I anticipating correctly that next period when he becomes a parent he will face an offspring with self-control problems, and the offspring will face an offspring with self-control problems, and so on. Second, consider constraint (3) which defines how self $I - 1$ chooses b_{I-1} . Being sophisticated, self $I - 1$ knows that his followers will have self-control problems. In particular, self $I - 1$ knows that self I chooses b_I according to (4). We have just seen that the last constraint, (4), enters the parent's problem in at least two ways: first, in the definition of self I 's policy function and then as a constraint in the definition of self $I - 1$'s policy function. These two different constraints are represented by a single constraint, (4), because the parent and self $I - 1$'s sophisticated belief agree about how self I will behave. Similarly, the constraint describing self $I - 1$'s policy is also a constraint in the constraint that describes self $I - 2$'s policy, and self

$I - 2$'s policy is also a constraint of self $I - 3$'s, and so on. Thus, actually the constraint that describes the policy of self i enters parent's problem in i different places but since these are all identical constraints, we represent them with just one constraint that describes self i 's policy.

We restrict attention to linear equilibria, meaning equilibria in which policy functions are linear in net present value of current wealth. Mathematically, defining $\Gamma_{i-1}(b; \tau)$ to be the net present value of wealth available to an agent at the beginning of age i with asset level b and under tax system τ , we derive closed form solutions of the form $c_i(b; \tau) = M_i(\tau)\Gamma_{i-1}(b; \tau)$.

3 Optimal Taxes

In this section we analyze optimal capital taxes in the model introduced in Section 2. Proposition 1 characterizes optimal taxes when utility is logarithmic.

Proposition 1 *Suppose $u(c) = \log(c)$. The optimal taxes are given by:*

$$1 - \tau_i^* = \frac{1}{\beta_i} (1 - \delta + \beta_{i+1}\delta), \text{ for } i \in \{0, 1, \dots, I\},$$

with $\beta_0 = \beta_{I+1} = 1$.

Proof. Relegated to Appendix A.1. ■

The main task in proving Proposition 1 is to compute linear policies, $M_i(\tau)$, which we do as follows. The assumption of linearity of policy functions implies that offspring's problems at all ages are strictly concave maximization problems. This implies that the first-order optimality conditions are not only necessary but also sufficient. As a result, in the parent's problem, we can replace the constraints that take the form of maximization problems with the corresponding first-order optimality conditions. Using the the first-order approach and a version of a guess and verify method, we find analytic expressions for the value function V and the vector of constants $M_i(\tau)$ describing equilibrium linear policies. Then, we use the policy functions to compute Euler equations that describe people's optimality conditions regarding savings at different ages in equilibrium. The comparison of these equilibrium Euler equations with the planner's Euler equations, given by (1), gives the optimal distortions (taxes) that need to be created to implement the commitment allocation.

It might be important to stress that - as shown in Appendix A.1. - Proposition 1 holds regardless of whether the economy is in a steady state or in a transition. In particular, since agents do not face binding liquidity constraints, *the expressions for taxes hold for any life-cycle path of wages*. In Appendix C, we also prove that, under the logarithmic utility assumption, the optimal taxes given by Proposition 1 are valid independent of whether we assume people are aware of their future self-control problems or not.

Proposition 2 shows that if the economy is in a steady state, then the optimal taxes characterized in Proposition 1 for the $\sigma = 1$ case are valid for any $\sigma > 0$.

Proposition 2 *Assume k is such that $\delta f'(k) = 1$. Then, for any $\sigma > 0$, optimal taxes are given by the exact same expressions as in Proposition 1.*

Proof. Relegated to Appendix A.2. ■

The tax formula for in Proposition 1 consists of two components. The first component, $\frac{1}{\beta_i}$, is easier to understand. Because of his *current* self-control problem, self i discounts tomorrow by an extra β_i and hence wants to undersave relative to the efficient allocation. By factoring the after tax return with $\frac{1}{\beta_i}$, we can exactly offset the extra discounting, thereby getting rid of this undersaving motive. We call this first part of the tax formula the *current component*. Clearly, the current component is always greater than one, i.e., it always calls for a subsidy.

The second component of the tax formula, $(1 - \delta + \beta_{i+1}\delta)$, is there to correct deviations in current savings caused by suboptimal actions of future selves. This part the tax formula, which we refer to as the *future component*, is always less than one, meaning it calls for a tax on savings. Intuitively, from self i 's perspective, self $i+1$ is undersaving due to his self-control problems ($\beta_{i+1} < 1$.) Thus, self i 's welfare would go up if he can make self $i+1$ increase his savings, which self i can achieve by increasing his own savings since, under our parametric assumptions, self $i+1$'s savings is strictly increasing in self i 's savings. This discussion implies that self i has an additional marginal benefit of saving in equilibrium when $\beta_{i+1} < 1$ relative to the case where $\beta_{i+1} = 1$. It is this extra benefit that makes self i save more relative to the commitment level. The future component of the tax formula calls for a tax in order to offset this extra return, and ensure self i does not oversave. The reader might still feel puzzled by our argument: after all, when facing the optimal taxes, self $i + 1$ saves the efficient amount. Notice, however, that from self i 's perspective, self $i + 1$ is still undersaving (at the new price that is inflated by the subsidy),

and hence future component of the tax is still needed.

Obviously, the sign of the optimal capital tax depends on whether the current or the future component dominates. For constant self-control (i.e., $\beta_i \equiv \beta < 1$), the current component always dominates, implying an optimal negative tax (i.e., optimality calls for a saving subsidy). We will see below that when β_i changes with age, depending on the pattern of change, either component may dominate, and the optimal tax can in general be positive or negative.

Notice that τ_0^* is only shaped by the future component. It is hence always positive. Since it is applied to the wealth transferred to future generations, τ_0^* can be interpreted as a bequest tax. In this paper, we do not analyze taxation of wealth transferred across generations. We study this topic in detail in Pavoni and Yazici (2013b).

3.1 Lessons for Capital Taxation

Propositions 1 and 2 imply several general lessons for capital taxes which are summarized below in a series of corollaries.

Corollary 3 (Age-dependence) *Optimal capital taxes are age-dependent. In particular, depending on how the degree of self-control changes with age: (i) Optimal capital taxes might be positive or negative at different ages. (ii) Optimal capital taxes might be increasing or decreasing with age at different ages.*

Proof. (i) For an example of $\tau_i > 0$, set $\beta_{i+1} \approx 0$ and $\beta_i > 1 - \delta$. For an example of $\tau_i < 0$, set $\beta_i = \beta_{i+1} = \beta < 1$. (ii) See the brown line with crosses in Figure 1 for an example. ■

The optimal capital taxes depend on age as long as the degree of self-control problems depend on age. Figure 1 plots various life cycle profiles of self-control on the left panel and the corresponding optimal taxes on the right panel. The figure shows that, under constant self-control of $\beta = 0.5$, depicted by blue dots, the subsidy is also constant over the life cycle at 4%. Corollary 3 also states that, unlike the common presumption in the literature, it might be optimal to tax people with self-control problems. However, although a theoretical possibility, optimal capital taxes can be positive in our model only under parameter specifications that seem to be inconsistent with data. Figure 1 shows that the model can generate a few periods of positive optimal capital taxes only when β_i declines sharply with age, as depicted by the brown crosses. As we discuss in Section 3.2, increasing level of self-control problems with age (β_i declining with i) is at odds with empirical findings. The orange dashed line and the light

blue solid line in Figure 1 display self-control patterns that are increasing with age. In both cases, capital should be subsidized. Then, whether subsidies should increase or decrease with age depends on the curvature of β_i .




Figure 1: *Left panel: Examples of constant, decreasing, concave increasing and convex increasing patterns self-control problems over the life cycle. Right panel: Corresponding optimal capital subsidies.*

Corollaries 4 and 5 below characterize quite sharply the sign and the monotonicity properties of optimal capital taxes over the life cycle under the assumption that self-control problems decrease concavely with age (β_i is increasing and concave in i). This pattern is in line with the available empirical evidence and agrees with the self-control profiles we calibrate in Section 3.2.¹³ Corollary 4 shows that if the severity of self-control problems decline with age, then

¹³Several different strands of literature provide evidence on this pattern of self-control problems. First, research on intertemporal discounting over the life span has shown that short term discount rates fall with age predicting a life-cycle developmental trend toward increased self-control. See, in particular, Green, Myerson, and Ostadzowski (1999), and Read and Read (2004). Second, Ameriks, Caplin, Leahy, and Tyler (2007) finds that ‘EI gap,’ the measure of self-control problem used in the paper’s experiment, decreases concavely with age. Finally, personality psychologists associate self-control with conscientiousness, one of the ‘big five’ personality factors, and in the words of Caspi, Roberts, Robins, and Trzesniewski (2003) ‘it appears that the increase in conscientiousness is one of the most robust patterns in personality development, especially in young adulthood.’ John, Gosling, Potter, and Srivastava (2003) estimates conscientiousness as a quadratic function of age and finds that the quadratic age term has a negative coefficient ‘indicating that the rate of increase [in conscientiousness] was greater at younger ages than at older

capital should be subsidized at all ages.

Corollary 4 (Optimality of Capital Subsidies) *If $\beta_{i+1} \geq \beta_i$, for all $i \geq 1$, optimal capital tax is negative for all ages.*

Proof. $1 - \tau_i^* = \frac{1}{\beta_i}(1 - \delta + \beta_{i+1}\delta) > \frac{\beta_{i+1}}{\beta_i} \geq 1$. ■

Corollary 5 shows that, if people's ability to self-control increases concavely with age, then capital subsidies should decrease with age.

Corollary 5 (Decreasing Capital Subsidies) *If $0 \leq \beta_{i+1} - \beta_i \leq \beta_i - \beta_{i-1}$ for all $i \geq 1$ (concavity), then optimal capital subsidies decrease with age.*

Proof. $1 - \tau_{i-1}^* = \frac{1-\delta}{\beta_{i-1}} + \frac{\beta_i\delta}{\beta_{i-1}} > \frac{1-\delta}{\beta_i} + \frac{\beta_i\delta}{\beta_{i-1}} > \frac{1-\delta}{\beta_i} + \frac{\beta_{i+1}\delta}{\beta_i} = 1 - \tau_i^*$, where the first and second inequalities follow from $\beta_{i-1} < \beta_i$ and $\beta_{i+1} - \beta_i \leq \beta_i - \beta_{i-1}$, respectively. ■

The result of Corollary 5 differs from Krusell, Kuruscu, and Smith (2010) which concludes that in any finite economy with constant self-control, capital subsidies should be increasing with age. The optimality of increasing subsidies in their case is due to the finite life time people face, and this element is missing from our analysis due to our assumption of perfect altruism. In Section 4.1, we show that the finite life time effect is quantitatively small within the relevant parameter space, implying that the optimality of decreasing subsidies with age is generally optimal.

3.2 Quantitative Analysis

In this section, we quantitatively analyze optimal capital taxation over the life cycle assuming either one of the justifications of the tax formulas in Proposition 1 hold: either utility is logarithmic or the economy is at a steady state. In order to conduct a numerical analysis, we have to choose values for the parameters of the model. Individuals are assumed to be born at the real time age of 25 and they live 51 years, so they die at the end of age 75. Observe that the tax formulas do not depend on the constant relative risk aversion coefficient σ , the shape of the production function F , or the depreciation rate, d . Therefore, we do not specify values for

ages.' Roberts, Walton, and Viechtbauer (2006) also estimates a concave conscientiousness pattern.

these parameters.¹⁴ The only parameter values needed for the analysis are the true yearly discount factor δ and the evolution of self-control with age, $\{\beta_i\}_{i=1}^{51}$. Here, β_i represents the level of self-control problem at real age of $i + 24$.

We set the long-run yearly discount factor $\delta = 0.96$. This value corresponds to the benchmark estimate in a constant self-control model with $\sigma = 1$ by Laibson, Repetto, and Tobacman (2007). The self-control vector, $\{\beta_i\}_{i=1}^{51}$, is calibrated as follows. We assume the relationship $i \rightarrow \beta_i$ takes the following functional form:

$$\beta_i = a - d \exp \left\{ \frac{51 - i}{b} \right\}. \quad (5)$$

An advantage of this functional form over some other - perhaps simpler - forms is that it is relatively easy to ensure that it satisfies a key condition of our model, namely $\beta_i \leq 1$ for all $i \in \{1, \dots, 51\}$.¹⁵ This form is also flexible in the sense that it allows for both concave and convex and decreasing and increasing patterns of self-control over the life cycle. This flexibility is important as we do not want to - a priori - put any restrictions on the self-control pattern. We calibrate the parameters of the self-control function in (5) in two alternative ways.

RR Calibration. Read and Read (2004) conducts a survey with 129 respondents between the ages of 19 and 89 in which respondents are asked to make a large number of time discounting decisions on both computerized and paper-and-pencil questions. The study estimates hyperbolic discount functions for three age groups: young, middle-aged, and older adults, with mean ages of 25, 44, and 75, respectively. Unlike quasi-hyperbolic discount functions, which have one short-term and one long-term discount factors (β and δ), hyperbolic discount functions allow many different discount factors depending on the length of delay for future reward. This implies that the hyperbolic discount functions estimated in Read and Read (2004) are not readily available for our study. Therefore, our calibration strategy works in two stages. First, for each age group, we find the β that best approximates the hyperbolic discount function for that group. We assume that the β at the mean age of a group is equal to the β that is approximated for that

¹⁴Observe that if we want the taxes computed using the formulas in Proposition 1 to be valid under any σ , then we need to assure that the interest rate R or the deeper parameters of the production function F and d satisfy the steady-state condition $R = f'(k) = F'(k, 1) - d = \delta^{-1}$, where k refers to steady-state level of capital stock.

¹⁵To be precise, β_i is bounded above by 1 with this functional form under the assumption that $a = 1 + d$. As we see when we discuss our calibration exercises in the next section, this assumption is empirically supported.

group (for example, β at age 25 equals β of young adults). Second, we use the β for ages 25, 44, and 75 (or correspondingly for periods $i = 1, 20,$ and 51) to pin down the parameters $a, b,$ and d of the functional form in (5).

For the older adults group, the approximation stage is quite simple because Read and Read (2004) finds that the older adults group display exponential discounting. The authors write: “Green et al’s major result- that younger people show hyperbolic discounting while older people show exponential discounting - is supported by our data.” This implies that a single (long-term) discount factor, $\delta,$ is sufficient to describe the behavior of older adults, implying this group has $\beta = 1.$ Since age 75 corresponds to period 51 in our period, this gives us our first calibration target, $\beta_{51} = 1,$ which implies $a = 1 + d$ when plugged in (5).

For the young adults and middle-aged groups, the approximation procedure chooses β using a least squares procedure: that is, β is chosen to minimize the sum of the squares of errors between the yearly discount factors that are implied by the hyperbolic discount function and our (δ, β) model. The details of our approximation procedure can be found in Appendix B. The β ’s that come out of our approximation procedure are 0.525 and 0.732 for the age 25 and age 44 groups, respectively. Thus, we set $\beta_1 = 0.525$ and $\beta_{20} = 0.732$ as the two other calibration targets. These two targets allow us to identify the remaining parameters of the self-control function (5), b and $d.$ The calibration targets and the calibrated parameter values are reported in the upper panel of Table 1.

GMO Calibration. This calibration uses a strategy that is quite similar to RR calibration except that the calibration targets come from two other studies. Two of the targets of GMO calibration come from Green, Myerson, and O’Donoghue (1999) which is an earlier paper that also uses an experimental approach to estimate hyperbolic discount functions. They do so for two groups of adults: the young and the old, with average ages of 20 and 70, respectively. Similar to Read and Read (2004), Green, Myerson, and O’Donoghue (1999) also finds that the intertemporal discounting behavior of their old adults group can be best described by exponential discounting function. Therefore, we set β for this group to be 1. Using the exact same approximation procedure we use in RR calibration, we find β for the age 20 group to be 0.362. Then, we identify (5) for the 51 periods starting at age 20 and ending at age 70. Thus, we set $\beta_1 = 0.362$ and $\beta_{51} = 1.$ The latter again implies that $a = 1 + d$ as in the RR calibration. We have two

RR Calibration		
Parameters	Targets	Source
d=0.795	β at age 25 = 0.525	RR
b=106.74	β at age 44 = 0.732	RR
a=1+d	β at age 75 = 1	RR
GMO Calibration		
Parameters	Targets	Source
d=0.055	β at age 20 = 0.362	GMO
b=19.73	Average β = 0.818	LRT
a=1+d	β at age 70 = 1	GMO

Table 1: *This table reports the two alternative calibration exercises that we conduct for self-control pattern over the life cycle. The acronyms GMO, LRT, and RR stand for Green, Myerson, and Ostadzewski (1999), Laibson, Repetto, and Tobacman (2007), and Read and Read (2004), respectively. a, b and d are the parameters of the self-control function in expression (5).*

parameters, b and d to be identified and only one target $\beta_1 = 0.362$. We still need one more calibration target to identify b and d . We use the average level of self-control problems in the economy, call it β_{avg} , as an additional target. We take β_{avg} to be 0.818 in our benchmark analysis following the estimate of Laibson, Repetto, and Tobacman (2007) for $\sigma = 1$ for a constant self-control model. The calibration targets and calibrated parameter values are reported in the lower panel of Table 1.

It is important to notice that period 1 in GMO calibration does not refer to period 1 in our model since in our model people are 25 in period 1 whereas people are 20 in period 1 of GMO calibration. The period 1 in the calibration rather refers to the first period of calibrated self-control function (5). Thus, the β of the agent in the first period of his life (at age 25) is equal to β_6 according to (5), which equals 0.517. Similarly, we need to determine β for agents between the ages of 70 and 75 since the function (5) defines self-control problems only for the ages between 20 and 70. Given that according to Green, Myerson, and Ostadzewski (1999) self-control problems already vanish by age 70, we set $\beta = 1$ for agents at ages between 70 and 75.

Now, we summarize our results. First, in both of our calibrations, self-control problems decrease with age, and thus, it is always optimal to subsidize life-cycle savings. The red dots

and the solid green line on the left panel of Figure 2 represent the RR and GMO calibrations of the evolution of self-control problems over the life cycle, respectively. The corresponding lines on the right panel display the corresponding optimal subsidies. In the RR calibration, the optimal subsidies start about 6% and decrease slowly to about 1% at the end of the life cycle. In the GMO calibration, the subsidies start at a significantly higher 8.5% but decrease more sharply to 0% at the end of the life cycle. The optimality of declining subsidies in both cases is expected given Corollary 4 and Corollary 5 and the concavely increasing pattern of β_i with i in both calibrations. The fact that GMO optimal subsidies start higher and decline more steeply is due to the fact that the GMO self-control pattern is substantially more concave than that of RR.




Figure 2: *Left panel: Life cycle patterns of self-control problems under Read and Read (2004) (RR) and Green, Myerson, and Ostaszewski (1999) (GMO) calibrations. Right panel: Corresponding optimal subsidies.*

We conduct sensitivity analysis with respect to our choice of functional form for the self-control function given in (5). We do so by assuming a quadratic form: $\beta_i = Ai^2 + Bi + C$. For the RR calibration targets, the quadratic calibration gives a β_i pattern that is virtually identical to the original calibration. The implied optimal subsidies are virtually identical as well. For the GMO calibration targets, the implied β_i pattern rises strictly above one toward the end of the life cycle and stays above one until the end. This is inconsistent with the evidence found

in both papers which calls for short term discount rates that are (weakly) larger than long term discount rates at all ages. The fact that quadratic function cannot be bounded above by one is one of the reasons why we have adopted the functional form in (5).¹⁶

We conclude this section by summarizing the general pattern of optimal taxes that emerges in virtually all of our simulations: optimal taxes are negative, - i.e., they are in fact subsidies, and these subsidies typically decline with age.

3.3 How Large Are the Optimal Subsidies Relative to Existing Ones?

Observe that in our computations the tax base is the gross return on asset holdings. Most actual tax systems, however, tax *asset income*. If we translate our numbers into subsidies on capital income, we find that optimal subsidy at age 25 is 217%, and subsidies decline to 10% by age 65.¹⁷ These are obviously large numbers. In this section, we compare these numbers with existing saving subsidies in the United States.

Tax-deferred retirement accounts are the main channel through which savings are subsidized in the United States. Probably, the most well-known tax-deferred saving account in the United States is the classical 401(k). Each dollar invested into a classical 401(k) is deductible from taxable income. Moreover, the taxes on the returns to 401(k) are deferred till retirement. This means, instead of paying taxes on interest or dividend income earned every year, the person pays tax based only on the income generated at the date of retirement, according to the tax rate faced by the agent at that date. Moreover, it is quite likely that, at retirement age, contributors face lower marginal taxes than when they invested into the plan. As we see below, this feature may generate considerable saving subsidies, and, importantly, these subsidies depend on where the agent is over her life cycle. Consider an agent who is at age i and is facing a marginal income tax rate τ_C^y based on the income tax bracket she falls in currently. Suppose there are N periods before she retires. If she invests \$1 today in 401(k), with the current tax deduction, this is as if she invests $\$ \frac{1}{1-\tau_C^y}$. If τ_R^y is the income marginal tax rate at retirement age, the agent will receive

¹⁶See Pavoni and Yazici (2013a) for an extensive sensitivity analysis of optimal subsidies regarding the levels of β_1 and β_{avg} for the case of Green, Myerson, and Ostry (1999) calibration.

¹⁷Denoting the capital income tax by τ_i^k , the relation between our taxes and tax on capital income is given by: $1 - \tau_i^k = \frac{R(1-\tau_i^y)-1}{R-1}$. As a consequence, for $\tau_1^y = -8.67\%$ and $\tau_{41}^y = -0.41\%$, and under $R = 1/1.04$, we have $\tau_1^k = -217\%$ and $\tau_{41}^k = -10\%$.

$\$R^N \frac{1-\tau_R^y}{1-\tau_C^y}$ at retirement. This implies a yearly saving subsidy of¹⁸

$$\left(\frac{1-\tau_R^y}{1-\tau_C^y} \right)^{\frac{1}{N}} - 1. \quad (6)$$

This saving subsidy depends on a person's age in two ways. First, people's income tend to depend on where they are on their life cycle, which implies the tax deduction they receive, τ_C^y , effectively depends on their age. Second, a person's age determines how far away she is from retirement, N , which clearly affects the subsidy rate in (6). Observe that in computing the implied 401(k) subsidy rate in (6), we take the tax base as the gross return on asset holdings to make it comparable to the optimal subsidies we have computed earlier in this section.

The 401(k) also allows employers to contribute to the worker retirement saving plans. The most common methods of employer matching are the \$1 per \$1 up to 6% of pay and the \$0.5 per \$1 up to 6% of pay.¹⁹ According to these options, if the saving rate is below 6%, then for each dollar that a worker contributes to the 401(k) account, the employer contributes one (or, respectively, one half) dollar. This means a worker investing one dollar is effectively investing $\frac{\$2}{(1-\tau_C^y)}$ into the plan when the employer matching is 1-to-1 and $\frac{\$1.5}{(1-\tau_C^y)}$ when the employer matching is 0.5-to-1. The formula in (6) can then be straightforwardly adapted to compute implied 401(k) subsidies in the presence of employer matching.²⁰

¹⁸The formula in (6) indicates that the implied subsidy rate of the 401(k) scheme is independent of the agent's saving rate. The 401(k) scheme, however, puts a cap on the amount agents can invest. In 2012, the maximal amount for agents aged 50 or below was \$17,000; older contributors faced an higher cap. The cap, however, is unlikely to be binding for the median household with an average saving rate.

¹⁹According to a 2009 survey conducted by Hewitt Associates, \$1 for \$1 up to 6% pay is the most common matching plan and is offered by 27% of all employers in their sample while \$0.5 per \$1 is the second most common matching plan. "Trends and Experience in 401(k) Plans." Retrieved from <http://www.retirementmadesimpler.org/Library/Hewitt-Research-Trends-in-401k-Highlights.pdf>

²⁰Let a person's annual income be w and his amount contributed to 401(k) be x . If $x > 0.06w$, meaning the person is contributing more than 6% of his income, then in a one-to-one matching plan, the employer contributes $0.06w$ dollars, which implies that for each dollar he invests he is effectively investing

$$\frac{x + 0.06w}{x(1-\tau_C^y)}.$$

Therefore, in this case, the implied subsidy depends on the amount contributed. However, Thaler and Benartzi (2004) report that the average saving rates into the SMarT plan (for the 'control group') are between 4.4% and 6.6% (see page S174). Thus, in our computations of the implied 401(k) subsidies, we assume that contribution rate is

Now, we compare the saving subsidies implied by a typical 401(k) plan to the optimal subsidies implied by our model. In Figure 3, right panel, we report the life-cycle profile of the median income per household head - between 25 and 66 years of age - in the period 2000-2006 and the corresponding marginal tax rates implied by the 2006 income tax code.²¹ In Figure 3, left panel, we report the implied saving subsidies for several 401(k) plans together with the optimal subsidies given by our model under the two calibrations. Three observations are immediate. First, interestingly, the range of values for the subsidies implied by the 401(k) plan are not very different from the optimal ones. Second, the subsidies implied by the 401(k) plan are very much age-dependent. Third, the life-cycle pattern of the 401(k) subsidies is qualitatively very different from the optimal ones as they are *increasing over the life cycle*. Existing subsidies appear too low for young individuals and too high for individuals close to retirement.




Figure 3: *Left panel: Optimal saving subsidies and subsidies implied by the 401(k) plan according to marginal income tax rates in 2006, at different levels of employer matching. Right panel: Median income of U.S. household head over life cycle in 2000-2006 and implied marginal income taxes in the year 2006.*

less than 6% and use the formula explained in the main text.

²¹The data for the life-cycle profile of the median income per household head in the period 2000-2006 is taken from Heathcote, Perri, and Violante (2010).

3.4 Comparison to the Constant Self-Control Model

In this section, we compare the optimal capital subsidies obtained in our model to those one would obtain in a model where self-control problems are constant over the life cycle. For the sake of brevity, we only discuss the comparison for the GMO calibration. The thick green line on the left panel of Figure 4 displays the GMO calibration whereas the thin (and flat) green line displays the corresponding constant self-control pattern in which people have $\beta = 0.818$ at all ages, which is the value that corresponds to average β in the GMO calibration.

The thin green line on the right-hand panel of Figure 4 displays that the optimal subsidies in the constant self-control model are constant at slightly below 1%, at 0.89% to be precise. In our GMO simulation, optimal capital subsidies start as high as 8.5% at the beginning of the life cycle. The fact that our subsidies are higher than those in the constant self-control model in the early years of the life cycle might not be very surprising since people have bigger self-control problems at earlier ages in our model. What is perhaps more surprising is that the level of optimal subsidies remain higher than the ones implied by the constant self-control model until as late as age 55 even though the left panel of Figure 4 shows that at 55 the agent in our model has significantly more self-control than the agent in the constant self-control model (with a β difference of about 0.1). This indicates that our model implies larger subsidies for similar levels of current self-control problems. The comparison at age 42 makes this point clearer: at 42, agents in both models have virtually the same level of $\beta \approx 0.82$. As the right panel of Figure 4 displays, the optimal subsidy in our model at this age, 2.14%, is much higher than the optimal subsidy in the constant self-control model, 0.89% (in terms of subsidies on capital income the comparison is 54% vs. 22%).

To see why the optimal subsidies in our model are significantly larger than the ones in constant self-control model, rewrite the optimal tax formula given by Proposition 1 as

$$-\tau_i^* = (1 - \delta) \left(\frac{1}{\beta_i} - 1 \right) + \delta \frac{\beta_{i+1} - \beta_i}{\beta_i}. \quad (7)$$

The expression (7) decomposes optimal subsidy formula into two components. The first component is the optimal subsidy that arises in a model if the self-control problem remains constant at the current level, β_i . The second component is the additional amount of subsidy needed due purely to the change in the level of self-control problems. Obviously, as long as self-control problem is decreasing with age, $\beta_{i+1} > \beta_i$, this term calls for additional subsidization of sav-

ings. Since δ is typically close to one, the second component plays a quantitatively important role in shaping capital subsidies and taxes.




Figure 4: *Left panel: Life cycle patterns of self-control problems under Read and Read (2004) (RR) and Green, Myerson, and Ostadzewski (1999) (GMO) calibrations, and under constant self-control model at $\beta = 0.732$ and $\beta = 0.818$. Right panel: Corresponding optimal subsidies.*

To grasp the intuition why our model implies higher subsidies, remember the decomposition of optimal taxes into the *current* and the *future* components, which we discuss following Proposition 2. The current component is related to an agent’s current degree of self-control problem and calls for undersaving. For agent at age 42, this component is the same between our model and the constant self-control model. The future component summarizes how much a person oversaves to compensate for future self’s undersaving. Since people’s degree of self-control improves with age in our model, the future component makes the agent save more today in the constant self-control model relative to ours. As a result, an agent with the same level of current self-control problem saves more in the constant self-control model, which implies the required subsidy to make him save the right amount is going to be lower. Due to this future component, optimal subsidies in our model remain higher than the one in the constant self-control model even long after age 42 (until age 55).

4 Imperfect Altruism and Life-Cycle Consumption Profile

In the main body of the paper, we have assumed that people are perfectly altruistic towards their descendants. In Section 4.1, we remove this assumption and investigate the quantitative importance of imperfect altruism on optimal subsidies. The main conclusion is that imperfect altruism has little effect on both the level and pattern of optimal saving subsidies. In Section 4.2, we extend the finite life time version of the model developed in Section 4.1 by allowing for borrowing constraints and use this model to study the life-cycle consumption implications of our model. We find that, even though our model abstracts from important life cycle issues such as child-rearing and health, life-cycle consumption profiles implied by our model under calibrated self-control patterns capture the key properties of empirical life-cycle consumption patterns fairly well.

4.1 Imperfect Altruism

Using a *constant self-control* model with sophisticated agents Krusell, Kuruscu, and Smith (2010) find that optimal saving subsidies should be increasing with age if agents face finite life times. In our baseline model, the finite life time channel, which calls for increasing subsidies with age, is shut down by the perfect altruism assumption. We now consider an extended version of our model allowing for imperfect altruism and assess its quantitative importance in shaping optimal subsidies over the life cycle.

A parent has the following preferences over dynastic consumption streams

$$u(c_0) + \gamma \left[\delta u(c_1) + \delta^2 u(c_2) + \dots + \delta^I u(c_I) + \delta^{I+1} u(c'_0) + \delta^{I+1} \gamma \left[\delta u(c'_1) + \delta^2 u(c'_2) + \dots \right] \dots \right],$$

where this preference specification is equivalent to the one in the baseline model whenever the altruism factor, γ , is equal to 1. When $\gamma \in [0, 1)$, there is imperfect altruism. The finite life time case of Krusell, Kuruscu, and Smith (2010) corresponds to the case of $\gamma = 0$. The rest of the parent's problem is identical to the one in Section 2.3. Proposition 6 generalizes the optimal tax formulas of Proposition 1 to the case with a general altruism factor, γ . In the case of perfect altruism, $\gamma = 1$, these formulas reduce back to the ones in Proposition 1.

Proposition 6 *Suppose $u(c) = \log(c)$. For $\gamma \in (0, 1]$, the optimal taxes are given by:*

$$1 - \tau_i^* = \frac{1}{\beta_i} \frac{1 + \beta_{i+1} \delta [1 + \delta + \dots + \delta^{I-i-2} + \delta^{I-i-1} D]}{1 + \delta + \dots + \delta^{I-i-1} + \delta^{I-i} D}, \text{ for } i \in \{0, 1, \dots, I\},$$

where $\beta_0 = \beta_{I+1} = 1$ and

$$D = \frac{1 + \delta\gamma(1 + \delta + \dots + \delta^{I-1})}{1 - \delta^{I+1}\gamma}.$$

Proof. Relegated to Appendix A.3. ■

Figure 5 displays the quantitative effects of the finite life time channel on the monotonicity properties of optimal subsidies. The red dots represent optimal subsidies under RR calibration with $\gamma = 1$, whereas the dashed red line represents those under RR calibration but with $\gamma = 0$. The comparison of the two lines shows that the finite life time effect is not strong enough to overturn the optimality of subsidies declining with age. Similarly, the solid green line and the dashed green line in Figure 5 display optimal subsidies under GMO calibration with $\gamma = 1$ and $\gamma = 0$, respectively. The comparison in this case supports the earlier conclusion that optimal subsidies decline with age even when we take into account the finite life time effect. Observe that in both robustness exercises we set $\gamma = 0$. For $\gamma \in (0, 1)$, the effect of finite life time on the monotonicity properties of optimal subsidies would be even smaller.

We simulated our model adopting several different parameterizations of γ, δ and $i \rightarrow \beta_i$. The optimality of decreasing saving subsidies is quite robust (details are available upon request). What seems to play an important role to maintain robustness of the decreasing subsidy result under various values of γ is that, in our model, self-control problems vanish towards the end of the life cycle. Notice that the only way in which γ enters the subsidy formula in Proposition 6 is via D . Early in the life cycle, optimal subsidies are not sensitive to γ because the subsidy formula for early periods discounts D repeatedly. The optimal subsidy formula for later periods does not discount D heavily, but this time β_{i+1} converges to one, and the second fraction on the right-hand side of the subsidy formula converges to one, making optimal subsidies insensitive to γ .

Observe that, in Figure 5, for each calibration, the optimal subsidies are uniformly higher in the model with imperfect altruism than those implied by the model with perfect altruism. Mechanically, this is because the parameter γ enters into the formula for $1 - \tau_i^*$ in Proposition 6 only through the constant D . It is easy to show that D is increasing in γ , which implies that the subsidy is decreasing in γ since the subsidy is decreasing in D . Intuitively, the second term in the formula of $1 - \tau_i^*$ in Proposition 6 (the future component) is there because agent i disagrees with agent $i + 1$ regarding how agent $i + 1$ should discount consumption in period $i + 2$ and onwards relative to consumption in $i + 1$: from agent i 's perspective, the correct discount

factor between consumption at date $i + 1$ and $i + s$ is δ^{s-1} , whereas agent $i + 1$ discounts by $\beta_{i+1}\delta^{s-1}$. To correct for the eventual undersaving of self $i + 1$, self i oversaves relative to the efficient allocation, and to prevent this, the government taxes self i . If $\gamma = 0$, the disagreement between self i and self $i + 1$ regarding the discounting between $i + 1$ consumption and future consumption levels stop at the end of the current life cycle ($D = 1$), while for $\gamma = 1$, the disagreement piles up for infinitely many generations ($D = \frac{1}{1-\delta}$). As a result, when $\gamma = 0$, there is less cumulative disagreement, which means self i is less motivated to oversave relative to the efficient allocation, which implies the tax implied by the future component is lower. Therefore, the overall subsidy is larger.




Figure 5: Left panel: Life cycle patterns of self-control problems under Read and Read (2004) (RR) and Green, Myerson, and Ostadzewski (1999) (GMO) calibrations with $\gamma = 0$ and $\gamma = 1$. Right panel: Corresponding optimal subsidies.

4.2 Life-cycle Consumption Implications of Self-Control Calibration

Throughout the paper, we motivate and calibrate the life-cycle pattern of self-control using experimental data. In this section, we would like to document the life-cycle profile of consumption implied by our model as a partial validation of our calibration exercise.

We take Gourinchas and Parker (2002) (GP hereafter) as our empirical reference. GP computes life-cycle profiles of mean annual real disposable income and consumption using Consumer Expenditure Survey (CEX) data, employing synthetic cohort techniques.²² They construct their variables from a sample of roughly 40,000 households from 1980 to 1993. A graphical representation of the income and consumption profiles they compute can be found in Figure 2 (page 67) in GP.




Figure 6: *The blue dashed line and the black circles represent the life-cycle profiles of mean disposable income and consumption between 1980 and 1993 as fitted by Gourinchas and Parker (2002). The red crosses and the green solid line represent the simulated consumption profiles implied by our model under self-control problems given by RR and GMO calibrations, respectively. The dash-dot line represents the life-cycle consumption profile of a planner who discounts exponentially and faces the same liquidity constraints that households face.*

Clearly, our simple model cannot have the ambition of matching the life-cycle pattern of consumption as closely as GP does. In particular, we abstract from income and health uncertainty, and from all life cycle changes other than income and self-control. We assume households face

²²Consumption is obtained by subtracting expenditure on education, medical care, and mortgage interest payments from total household expenditure. Disposable labor income is obtained by subtracting Social Security tax payments, pension contributions, after tax asset and interest income, and those expenditures subtracted from consumption from after-tax family income.

a deterministic path of labor income which corresponds to the life-cycle profile of income constructed in GP. We also impose a constant (and exogenous) gross return to savings, R , and a simple form of liquidity constraints: households cannot borrow at all. To be consistent with GP, and given that our target is to get a life-cycle profile of consumption, we move away from the infinite horizon model and assume households are not altruistic ($\gamma = 0$). We use this model to compute two alternative life-cycle consumption profiles, one assuming self-control problems evolve according to our RR calibration and another one according to GMO calibration. We set the remaining parameters as follows. Agents are assumed to have logarithmic utility and we set $R = 1.0344$ and $\delta = \frac{1}{R}$.²³ Consistently with GP, agents start their working life at age 26 and retire at 65. People begin their lives with zero assets. After retirement, agents are assumed to leave for 10 more years and we set retirement income to zero.²⁴

Figure 6 depicts life-cycle consumption profiles coming from RR and GMO simulations (red dots and solid green line, respectively), together with life-cycle (fitted) income and consumption profiles (blue dashed line and black circles, respectively) constructed in GP. All series are plotted for the age range between 26 and 65.²⁵ Figure 6 shows that both RR and GMO calibrations of self-control profiles imply reasonable life-cycle consumption profiles. In particular, our simple model of self-control is able to generate the concave pattern of consumption we observe in the data. It matches the declining pattern of consumption in the second part of the life cycle fairly well (for the GMO calibration the fit is, in fact, - surprisingly - quite remarkable), while

²³We assume logarithmic utility mainly in order to use our closed form solutions; nonetheless, the relative risk aversion coefficient of one is well in the range estimated by GP, 0.5 to 1.4, depending on the exact specification of the model. The value of R is chosen to be consistent with GP, which estimates an average real return of 3.44% for Moody's AAA municipal bonds over the sample period.

²⁴ GP does not guide us regarding the length of retirement period or retirement income as it reports neither of them, and estimates a value of retirement and the implied 'retirement wealth' to best match moments on consumption and savings. Instead of choosing these values to optimize the model's fit, we set retirement income to zero and retirement length to 10 years. For both calibrations, the consumption profiles simulated using the model with 20 years of post retirement life - not report for ease of graphical exposition - are virtually identical to the profiles obtained from the 10 years model.

²⁵Given the hump shaped pattern of income and the decreasing pattern of self-control, in all parametrizations, agents are liquidity constrained early in life. After they start saving, they never become liquidity constrained again. This characteristic of the model allows us to use the machinery we developed to solve the model without liquidity constraints (leading to the closed form solution for taxes) to solve this model with liquidity constraints. Further details on the numerical algorithm are available upon request.

it is unable to match the flatter pattern of consumption we observe in the data during the first part of the life cycle. GP argues that precautionary savings and measurement error play a crucial role in explaining the consumption profile in the initial segment of the life cycle while our model postulates deterministic income.²⁶

It might be worthwhile discussing the optimal taxes implied by the model with liquidity constraints. The black dash-dot line in Figure 6 represents the life-cycle consumption profile of a planner who discounts exponentially and faces the same liquidity constraints that households face. This is the commitment consumption profile that optimal subsidies target under borrowing constraints. The figure shows that the consumption patterns of the commitment allocation and the self-control model (for both calibrations) coincide until age 34. The optimal subsidies in the borrowing constraint model starting with age 34 coincides with the optimal subsidies given by Proposition 6 (for $\gamma = 0$) since after this period the commitment allocation requires people to save. The optimal subsidies after age 34 corresponds to that shown by the thin dotted (green or red depending on the calibrations) lines in Figure 5. During the ages between 25 and 34, the optimal subsidies are indeterminate. Intuitively, as long as subsidies are not too high during these years, people will be borrowing constrained, and consume exactly their labor income, which is what we see in the commitment allocation. Notice that, the subsidies given by Proposition 6 are optimal for these years as well as they lie within the range of optimal subsidies.

5 Conclusion

This paper studies optimal capital taxation in an economy where agents face self-control problems. In line with evidence suggested by personality psychology and experimental studies, we allow for the severity of the self-control problem to change over the life cycle. We restrict attention to CIES utility functions and focus on linear Markov equilibria. We derive explicit formulas which allow us to compute optimal taxes given the evolution of self-control problem over the life cycle. We show that if agents ability to self-control increases concavely with age,

²⁶On page 49, GP writes: “The importance of the precautionary motive early in life implies that between 60 and 70 percent of non-pension wealth is due to precautionary savings, according to poor estimates and holding the real interest rate fixed.” On page 67, GP also writes: “Consumption lies above income over the late twenties. Given that the CEX wealth data, and better household wealth surveys, show modest increases in liquid wealth over these ranges, this feature seems likely due to misreporting of income or consumption.”

then capital should be subsidized and the subsidy should decrease with age.

In our calibrated quantitative exercises, we find that optimal capital subsidies start somewhere between 6% and 8% at the beginning of the life cycle and decline monotonically with age to somewhere between 0% and 1%. If we translate the subsidies on capital into subsidies on capital income, these are very large numbers. Perhaps more importantly, we show they are much larger than the savings subsidy we would obtain in models with constant self-control, at most ages. Our model is probably too simple for delivering precise policy predictions. Nevertheless, our analysis suggests that researchers who take self-control problems seriously should also carefully measure the evolution of self-control problems over the life cycle before making policy suggestions.

We also compare our optimal subsidies with those implied by the 401(k) plan. If we exclude the very last periods before retirement - where the subsidy rate in the 401(k) essentially mimics the employer matching rate - the subsidy levels in the two cases are of comparable magnitudes. A marked difference emerges, however, in the life-cycle pattern of optimal subsidies: the 401(k) plan implies an increasing pattern of subsidies while the optimal subsidies decrease over the life cycle.

In Appendix C, we extend the model by allowing for different degrees of self awareness (partial sophistication) about the existence of future self-control problems. We first prove that when utility function is logarithmic, the optimal tax formulas are independent of the pattern of partial sophistication. When CEIS coefficient is different from one, closed form solutions for optimal taxes are unavailable. Our numerical experiments show that, as long as the level of sophistication is not changing abruptly from one period to another, the pattern of optimal capital subsidies over the life cycle is surprisingly robust to the degree of sophistication.

The existence of illiquid assets does not change our optimal tax results as long as there are no borrowing constraints. More precisely, in Appendix D we use a three periods example to show that a tax system that is optimal in an environment without illiquid assets is still optimal in the same environment with an illiquid asset as long as we complement the tax system with an appropriate tax on the illiquid asset.

References

- AMERIKS, J., A. CAPLIN, J. LEAHY, AND T. TYLER (2007): "Measuring Self-Control Problems," *American Economic Review*, 97(3), 966–972.
- ARIELY, D., AND K. WERTENBROCH (2002): "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," *Psychological Science*, 13(3), 219–224.
- CASPI, A., B. W. ROBERTS, R. W. ROBINS, AND K. H. TRZESNIEWSKI (2003): *Handbook of Life Course*.
- DELLAVIGNA, S. (2009): "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47(2), 315–372.
- DELLAVIGNA, S., AND U. MALMENDIER (2004): "Contract Design and Self-control: Theory and Evidence," *The Quarterly Journal of Economics*, 119(2), 353–402.
- EROSA, A., AND M. GERVAIS (2002): "Optimal Taxation in Life-Cycle Economies," *Journal of Economic Theory*, 105(2), 338–369.
- FARHI, E., AND I. WERNING (2013): "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, 810(2), 596–635.
- FREDERICK, S., G. LOEWENSTEIN, AND T. O'DONOGHUE (2002): "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40(2), 351–401.
- GOLOSOV, M., M. TROSHKIN, AND A. TSYVINSKI (2011): "Optimal Dynamic Taxes," NBER Working Papers 17642, National Bureau of Economic Research, Inc.
- GOURINCHAS, P.-O., AND J. A. PARKER (2002): "Consumption Over the Life Cycle," *Econometrica*, 70(1), 47–89.
- GREEN, L., A. FRY, AND J. MYERSON (1994): "Discounting of Delayed Rewards: A Life-Span Comparison," *Psychological Science*, 5, 33–36.
- GREEN, L., J. MYERSON, D. LICHTMAN, S. ROSEN, AND A. FRY (1996): "Temporal Discounting in Choice Between Delayed Rewards: The Role of Age and Income," *Psychology and Aging*, 11(1), 79–84.

- GREEN, L., J. MYERSON, AND P. OSTASZEWSKI (1999): "Discounting of Delayed Rewards across the Life-Span: Age Differences in Individual Discounting Functions," *Behavioural Processes*, 46(1), 89–96.
- GRUBER, J., AND B. KOSZEGI (2004): "Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes," *Journal of Public Economics*, 88(9-10), 1959–1987.
- GUL, F., AND W. PESENDORFER (2001): "Temptation and Self-Control," *Econometrica*, 69(6), 1403–1435.
- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006," *Review of Economic Dynamics*, 13(1), 15–51.
- HELSON, R., C. JONES, AND V. S. Y. KWAN (2002): "Personality Change over 40 Years of Adulthood: Hierarchical Linear Modeling Analyses of Two Longitudinal Samples," *Journal of Personality and Social Psychology*, 83(3), 752–766.
- IMROHOROGLU, A., S. IMROHOROGLU, AND D. H. JOINES (2003): "Time-Inconsistent Preferences And Social Security," *The Quarterly Journal of Economics*, 118(2), 745–784.
- JOHN, O. P., S. D. GOSLING, J. POTTER, AND S. SRIVASTAVA (2003): "Development of Personality in Early and Middle Adulthood: Set Like Plaster or Persistent Change?," *Journal of Personality and Social Psychology*, 84(5), 1041–1053.
- KRUSELL, P., B. KURUSCU, AND A. J. SMITH (2002): "Equilibrium Welfare and Government Policy with Quasi-geometric Discounting," *Journal of Economic Theory*, 105(1), 42–72.
- (2010): "Temptation and Taxation," *Econometrica*, 78(6), 2063–2084.
- KRUSELL, P., AND A. A. SMITH (2003): "Consumption–Savings Decisions with Quasi-Geometric Discounting," *Econometrica*, 71(1), 365–375.
- LAIBSON, D. (1994): "Self-Control and Saving," *Ph.D. Dissertation*.
- (1997): "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 112(2), 443–477.

- LAIBSON, D., A. REPETTO, AND J. TOBACMAN (2007): "Estimating Discount Functions with Consumption Choices over the Lifecycle," NBER Working Paper 13314, National Bureau of Economic Research.
- O'DONOGHUE, T., AND M. RABIN (2003): "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes," *American Economic Review*, 93(2), 186–191.
- (2006): "Optimal Sin Taxes," *Journal of Public Economics*, 90(10-11), 1825–1849.
- PAVONI, N., AND H. YAZICI (2013a): "Optimal Life-cycle Capital Taxation under Self-Control Problems," Working paper.
- (2013b): "Present-Bias Preferences and Optimal Taxation of Parental Transfers," mimeo.
- PHELPS, E. S., AND R. A. POLLAK (1968): "On Second-Best National Saving and Game-Equilibrium Growth," *The Review of Economic Studies*, 35(2), 185–199.
- READ, D., AND N. READ (2004): "Time Discounting over the Lifespan," *Organizational Behavior and Human Decision Processes*, 94(1), 22–32.
- ROBERTS, B. W., K. E. WALTON, AND W. VIECHTBAUER (2006): "Patterns of Mean-Level Change in Personality Traits Across the Life Course: A Meta-Analysis of Longitudinal Studies," *Psychological Bulletin*, 132(1), 1–25.
- STROTZ, R. H. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 23(3), 165–180.
- THALER, R. H., AND S. BENARTZI (2004): "Save More Tomorrow (TM): Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, 112(S1), S164–S187.
- THALER, R. H., AND H. M. SHEFRIN (1981): "An Economic Theory of Self-Control," *Journal of Political Economy*, 89(2), 392–406.