Commitment vs. Flexibility*

Manuel Amador†
Iván Werning‡
George-Marios Angeletos§

Abstract

We study the optimal tradeoff between commitment and flexibility in a consumption-savings model. Individuals expect to receive relevant information regarding tastes, and thus value the flexibility provided by larger choice sets. On the other hand, they also expect to suffer from temptation, with or without self-control, and thus value the commitment afforded by smaller choice sets. The optimal commitment problem we study is to find the best subset of the individual's budget set. We show that the optimal commitment problem leads to a principal-agent formulation. We find that imposing a minimum level of savings is always a feature of the solution. Necessary and sufficient conditions are derived for minimum-savings policies to completely characterize the solution. Individuals may perfectly mimic such a policy with a portfolio of liquid and illiquid assets. We discuss applications of our results to situations with similar tradeoffs, such as paternalism, the design of fiscal constitutions to control government spending, and externalities.

*We thank comments and suggestions from Daron Acemoglu, Andy Atkeson, Paco Buera, VV Chari, Peter Diamond, Doireann Fitzgerald, Narayana Kocherlakota and especially Pablo Werning. We also thank seminar and conference participants at the University of Chicago, Rochester, Harvard, M.I.T., Stanford, Pennsylvania, Berkeley, New York, Yale, Maryland, Austin-Texas, Torcuato di Tella, Stanford Institute for Theoretical Economics, American Economic Association, Society of Economic Dynamics, CESifo Venice Summer Institute, the Central Bank of Portugal, and the Federal Reserve Banks of Cleveland and Boston. Emily Gallagher provided valuable editing assistance.

†Stanford Graduate School of Business

‡Massachusetts Institute of Technology, National Bureau of Economic Research and Universidad Torcuato di Tella

§Massachusetts Institute of Technology, National Bureau of Economic Research

Introduction

A commonly articulated justification for government involvement in retirement income is the belief that an important fraction of the population saves inadequately when left to their own devices (Diamond, 1977). From the workers' perspective most pension systems, pay-as-you-go and capitalized systems alike, effectively impose a minimum-savings requirement. One purpose of this paper is to see if such minimum-savings policies are optimal in a model where agents suffer from the temptation to over-consume.

More generally, if people suffer from temptation and self-control problems, what should be done to help them? Current models emphasizing such problems lead to a simple but extreme answer: it is optimal to completely remove all future choices. In particular, in the intertemporal choice framework it is best to commit individuals to a particular consumption path, removing all future savings choices. In these models, the desire to commit is simply overwhelming.

Eliminating all ex post choices is unlikely to be optimal when new information regarding preferences or other variables is expected to arrive in the future. In these circumstances, individuals value the flexibility to act on their information. Indeed, in the absence of temptation or self-control problems full flexibility would be optimal.

This paper studies the non-trivial design of optimal commitment devices in situations where eliminating all choices is not necessarily optimal. We introduce a value for flexibility and study the resulting tradeoff with commitment, defined as the removal of some future choices. Our model combines a preference for flexibility and a preference for commitment by introducing taste shocks into both a time-inconsistent quasi-hyperbolic discounting framework (Phelps and Pollack, 1968; Laibson, 1997) and the temptation model of Gul and Pesendorfer (2001). The resulting preferences belong to a class introduced by Dekel, Lipman and Rustichini (2001).

The individual we model suffers from temptation for higher present consumption. Each period a taste shock is realized that affects the individual's desire for current versus future consumption.[1] Importantly, taste shocks are privately observed by the individual. If, instead, taste shocks were observable and verifiable by an outside party, one could simply contract upon them in a way that avoids all temptation and achieves the unconstrained ex-ante optimum. But when the shocks are private information

[1] Our analysis focuses on taste shocks, but the crucial feature is the arrival of any new information relevant to the savings decision. Flexibility would also be valuable if one modeled health, employment and income shocks. As we later show, with constant absolute risk aversion preferences, a model with income shocks is isomorphic to a model with taste shocks.

only the agent can act upon them, introducing a tradeoff between commitment and flexibility. Commitment is valued because it reduces temptation, while flexibility is valued because it allows the use of valuable private information.

The optimal commitment problem we study selects a subset of the individual's budget set to maximize ex-ante utility, taking into account the ex-post temptation problem facing that set. We show that this commitment problem is equivalent to a principal-agent problem: the principal has the individual's ex-ante preferences, the agent has the ex-post preferences. This formulation makes the analysis tractable and allows us to apply powerful Lagrangian optimization techniques.

A very simple commitment device in this setting is a minimum-savings rule, which restricts individuals to save above some level, allowing complete flexibility otherwise. Facing such a rule, individuals with low enough taste shocks will be unconstrained, while the highest types will be constrained by the minimum-savings level and obtain the same bundle of consumption and savings.

Our first finding is that a minimum-savings rule is always part of the solution to the commitment problem. That is, the optimal allocation necessarily has types above some threshold consuming the same bundle. This result has strong economic intuition: at the very top of the distribution there is no trade-off between commitment and flexibility. For any set of consumption-savings options made available, if the highest types are separated then they are necessarily undersaving from an ex-ante perspective, given these options. Thus, removing a lower segment of these options necessarily improves welfare.

We then turn to the question of whether a minimum-savings policy is all that is needed. Our next result establishes that under a condition for the distribution of taste shocks a minimum-savings rule completely characterizes the optimum. The condition ensures that individuals do not consume in the interior of the budget set, so that 'money burning' is not optimal. It also ensures that there is no gain from removing intervals on the budget line above the minimum-savings level. We show that this condition is also necessary for a minimum-savings rule alone to achieve the optimum.

Our analysis is also useful for other applications, quite different from the consumption-savings model we focus on, and we discuss three examples. The first concerns paternalism, a principal cares about an agent but disagrees with the agent's preferences. Applied to schooling our results may be relevant for thinking about minimum-schooling laws. Second, we discuss fiscal constitutional design, where citizens value government spending, but ruling administrations value it even more. Our results translate to conditions for simple spending caps to be optimal. Finally, we discuss an environment with a continuum of agents imposing externalities on each other, such as pollution. A utilitarian planner maximizes average welfare and internalizes these externalities, but agents acting privately do not. Many other situations feature a similar tradeoff between commitment and flexibility that might be captured by our model.

Models with time-inconsistent preferences solved as a competitive game, as in Strotz (1956), were the first to formalize a value for commitment. In particular, the hyperbolic discounting model has proven useful for modeling the possibility of undersaving and the desirability of commitment devices (Phelps and Pollack, 1968; Laibson, 1997). In a series of recent papers Gul and Pesendorfer (2001, 2002a,b) and Dekel, Lipman and Rustichini (2001, 2004) have provided axiomatic foundations for preferences that value commitment and have derived useful representation theorems. Kreps (1979) provided an early axiomatic foundation for a preference for flexibility, and showed that these preferences can always be represented by including taste shocks in an expected-utility framework.

A large literature on social security policy incorporates a concern for possible undersaving by individuals. Feldstein (1985) models overlapping generations that discount the future at a higher rate than the social planner. Laibson (1998) discusses policies to avoid undersaving by quasi-hyperbolic agents, while Imrohoroglu, Imrohoroglu and Joines (2000) perform a quantitative analysis of the benefits of pay-as-you-go policies in such a setting.

As discussed above, many other applications feature a similar tradeoff to commitment and flexibility. Since Weitzman's (1974) seminal paper on the benefits of prices versus quantities restrictions, there has been interest in the tradeoff between flexibility and control of managers, see Holmstrom (1984). A similar tradeoff arises between discretion and rules to address the time inconsistency of government policy (Kydland and Prescott, 1977). In a recent paper Athey, Atkeson and Kehoe (2004) study the optimal design of monetary policy rules. Benabou and Tirole (2002) consider an individual who manipulates ex post choices by suppressing some information from his memory, instead of using commitment devices.

The rest of the paper is organized as follows. Section 1 lays out the basic model with quasi-hyperbolic preferences. Section 2 studies optimal commitment and derives the main results. Section 3 extends the results to preferences displaying temptation and self-control. We discuss other applications of our results in Section 4. The final section concludes.

1 Basic Consumption-Savings Problem

In this section, we introduce the basic consumption-saving setup with time-inconsistent preferences. There are two periods and a single consumption good each period. We denote first and second period consumption by c and k, respectively. Given total resources y, the consumer is constrained by the budget set $B(y) \equiv \{(c,k) \in \mathbb{R}_+^2 \mid c+k \leq y\}$, where we have normalized the net interest rate to zero.

In the first period individuals receive a taste shock θ from a bounded set Θ with distribution function $F(\theta)$, normalized so that $\mathbb{E}[\theta] = 1$. The taste shock affects the marginal utility of current consumption: higher θ makes current consumption more valuable. Taste shocks are assumed to be the individual's private information.

We follow Strotz (1956), Phelps and Pollack (1968), Laibson (1994, 1997) and others by modeling the agent in each period as different *selves,* with different preferences. For the ensuing games played between *selves* we consider subgame perfect equilibria as our solution concept.

The utility for *self-1* from periods $t = 1, 2$ with taste shock θ is then,

$$\theta U(c) + \beta W(k),$$

where $U : \mathbb{R}_+ \rightarrow \mathbb{R}$ and $W : \mathbb{R}_+ \rightarrow \mathbb{R}$ are increasing, concave and continuously differentiable, and $0 < \beta \leq 1$. Utility for *self-0* from periods $t = 1, 2$ is given by:

$$\mathbb{E}\ [\theta U(c) + W(k)].$$

This setup represents a two-period version of quasi-geometric discounting. We associate β with the strength of temptation towards present consumption.

There is *disagreement* among the different *selves* on discounting but *agreement* regarding taste shocks. The tension is between tailoring consumption to the taste shock and *self-1*'s constant desire for higher current consumption. This tension generates the tradeoff between commitment and flexibility from the point of view of *self-0*. Indeed, this is the central feature of the model, which can be reinterpreted and applied to other situations with similar tradeoffs (see Section 4).

Taste shocks can be interpreted in two ways. Under an objective interpretation, they represent actual differences in ex-post preferences that are evaluated by the expected-utility agent. Under a subjective interpretation, the introduction of taste shocks is an 'as if' construction to represent an ex-ante preference for flexibility, in which case ex-post behavior need not be modeled.

Taste shocks are a tractable way of introducing a value for flexibility. Within an objective interpretation, they can be thought of as representing the significant variation observed in consumption and savings data, after conditioning on all available variables. Other shocks, such as unobservable income or health, can also generate a value for flexibility. Indeed, a model with privately observed income shocks is equivalent to the model with privately observed taste shocks when the utility function is exponential $U(c) = -e^{-\gamma c}$. To see this, note that total consumption in the first period is $c + z$, where z is privately observed income and $c = y - k$ is the observable component of consumption and savings. Then the first period utility is $U\left(c+z\right) = \theta U(c)$ where $\theta \equiv -U\left(z\right)$.

A useful benchmark allocation is the ex-ante *first-best allocation*, $\left(c^{fb}\left(\theta\right), k^{fb}\left(\theta\right)\right)$, defined by the solution to $\max_{(c,k)\in B(y)}\{\theta U\left(c\right) + W\left(k\right)\}$. This allocation would be feasible if taste shocks were not private information and were contractible. Another benchmark allocation is that obtained with full flexibility or no commitment: self-1 is constrained only by the resource constraint and solves $\max_{(c,k)\in B(y)}\left[\theta U(c) + \beta W(k)\right]$. We denote the unique solution to this problem by $(c^f(\theta), k^f(\theta))$.

2 Optimal Commitment without Self-Control

Commitment entails reducing the set of choices available. The optimal commitment problem is to choose the best subset $C \subset B(y)$ of the budget set that maximizes the expected utility of *self-0* given that choices are in the hands of *self-1*, that is, that the allocation is the outcome of a subgame perfect equilibrium. Formally, we choose $C \in B(y)$ so as to maximize $\int[\theta U\left(c(\theta)\right) + W\left(k(\theta)\right)]dF\left(\theta\right)$ subject to $c(\theta), k(\theta) \in \arg\max_{(c,k)\in C}(\theta U\left(c\right) + \beta W\left(k\right))$.

Finding the best subset C is equivalent to the following principal-agent problem directly over allocations $c\left(\theta\right)$ and $k\left(\theta\right)$:

$$\max_{c,\ k} \int \left[\theta U\left(c\left(\theta\right)\right) + W\left(k\left(\theta\right)\right)\right] dF\left(\theta\right)$$

$$\theta U\left(c\left(\theta\right)\right) + \beta W\left(k\left(\theta\right)\right) \geq \theta U\left(c\left(\theta'\right)\right) + \beta W\left(k\left(\theta'\right)\right) \text{ for all } \theta, \theta' \in \Theta \tag{1}$$

$$c\left(\theta\right) + k\left(\theta\right) \leq y \text{ for all } \theta \in \Theta \tag{2}$$

Given total resources y, the problem is to maximize expected utility from the point of view of *self-0* (henceforth: the principal) subject to the constraint that θ is private information of *self-1* (henceforth: the agent). The incentive compatibility constraint

(1) ensures that the agent reports the shock truthfully.

This principal-agent of the commitment problem formulation highlights the restriction that resources cannot be transferred across types. Although this restriction is not without loss of generality, it is a natural starting point for at least three reasons. First, it is useful to isolate the problem of commitment – defined as a reduction of choices from the individual's budget constraint – from the problem of insurance or redistribution, which is beyond the scope of this paper. Second, individuals may have access to commitment technologies, such as an illiquid asset, but not insurance contracts. Thus, it is important to understand what the ideal commitment device, not featuring insurance, looks like. Finally, the possibility of transferring resources across different types is simply absent in some reinterpretations of our model discussed in Section 4.

Two and Three Types

We begin by studying the optimal commitment problem with only two taste shocks and then turn to the case with a continuum. When taste shocks take only two possible values the optimum can be fully characterized as follows.

Proposition 1 *Suppose* $\Theta = \{\theta_l, \theta_h\}$, *with* $\theta_l < \theta_h$. *There exists a* $\beta^* \in (\theta_l/\theta_h, 1)$ *such that for* $\beta \in [\beta^*, 1]$ *the first-best allocation is implementable. Otherwise,*

(a) if $\beta > \theta_l/\theta_h$ *separation is optimal, i.e.* $c^*(\theta_h) > c^*(\theta_l)$ *and* $k^*(\theta_h) < k^*(\theta_l)$,

(b) if $\beta < \theta_l/\theta_h$ *bunching is optimal, i.e.* $c^*(\theta_l) = c^*(\theta_h)$ *and* $k^*(\theta_l) = k^*(\theta_h)$,

(c) if $\beta = \theta_l/\theta_h$ *separating and bunching are both optimal*

In all cases, the optimum can be attained without burning money: $c^*(\theta) + k^*(\theta) = y$ *for* $\theta = \theta_h, \theta_l$.

Proof. In the Appendix. ■

The result that the first-best allocation is incentive compatible for low enough levels of temptation relies on the discrete difference in taste shocks and does not hold with a continuum of shocks. For higher temptation the first-best allocation is no longer incentive compatible and the proposition shows that the solution takes one of two forms. For intermediate levels of temptation it is optimal to separate the agents. To achieve separation the principal must offer bundles that yield to the agent's ex-post desire for higher consumption, giving them higher consumption in the first period than the first-best. For high enough temptation, however, separating the agents requires too much first-period consumption, and bunching both types becomes preferable.

Bunching resolves the commitment problem at the expense of flexibility. The optimal amount of flexibility depends negatively on the degree of disagreement relative to the dispersion of taste shocks. The proposition also shows that the optimum can be attained on the frontier of the budget set, so that 'money burning' is not required.

Unfortunately, with more than two types extending these conclusions is not straightforward. For example, consider three taste shocks, $\theta_l < \theta_m < \theta_h$, with respective probabilities p_l, p_m and p_h. In this case bunching may occur between any consecutive pair of shocks. Money burning for the middle type is optimal if p_m is small enough and $\beta \in (\beta^*, \theta_l/\theta_m)$, where β^* is as defined by the proposition above with two types, θ_l and θ_h. This captures the intuition that if the middle shock occurs with very low probability, money burning is not very costly and might be preferable for incentive purposes. However, this does not exhaust all cases where money burning is optimal. If either $\beta < \beta^*$ or $\beta > \theta_l/\theta_m$, money burning is never optimal for small enough p_m, but may be optimal for sufficiently high p_m.[2] These results help illustrate that money burning is a possible feature of the solution and that conditions on the distribution are required to rule it out.

Continuous Distribution of Types

For the rest of the paper we assume that the distribution of types is represented by a continuous density $f(\theta)$ over the bounded interval $\Theta \equiv [\underline{\theta}, \overline{\theta}]$. It is convenient to change variables from $(c(\theta), k(\theta))$ to $(u(\theta), w(\theta))$ where $u(\theta) \equiv U(c(\theta))$ and $w(\theta) \equiv W(k(\theta))$, and we term either pair of functions an allocation. Let $C \equiv U^{-1}$ and $K \equiv W^{-1}$, which are then increasing and convex functions.

We now characterize the incentive compatibility constraints (1). Facing a direct mechanism given by $(u(\theta), w(\theta))$, an agent with taste shock θ maximizes over the report and obtains utility $V(\theta) \equiv \max_{\theta' \in \Theta}\{(\theta/\beta)\, u(\theta') + w(\theta')\}$. If truth-telling is optimal then $V(\theta) = (\theta/\beta)\, u(\theta) + w(\theta)$, by integrating the envelope condition $V'(\theta) = u(\theta)/\beta$:

$$\frac{\theta}{\beta} u(\theta) + w(\theta) = \int_{\underline{\theta}}^{\theta} \frac{1}{\beta} u(\tilde{\theta}) d\tilde{\theta} + \frac{\underline{\theta}}{\beta} u(\underline{\theta}) + w(\underline{\theta}) \tag{3}$$

(see Milgrom and Segal, 2002). Incentive compatibility of (u, w) also requires u to be a non-decreasing function of θ: agents that are more eager for current consumption cannot consume less. Thus, condition (3) and the monotonicity of u are necessary for

[2] Proofs of these results are available upon request.

incentive compatibility. These two conditions are also sufficient.

The principal's problem is thus to maximize $\int_{\underline{\theta}}^{\overline{\theta}} (\theta u(\theta) + w(\theta)) f(\theta) d\theta$ subject to the budget constraint $C(u(\theta)) + K(w(\theta)) \leq y$, the incentive compatibility constraint (3), and monotonicity $u(\theta') \geq u(\theta)$ for $\theta' \geq \theta$. Note that this problem is convex since the objective function is linear and the constraint set is convex.

Substituting the incentive compatibility constraint (3) into the objective function and the resource constraint and integrating by parts allows us to simplify the problem by dropping the function $w(\theta)$, except for its value at $\underline{\theta}$. Consequently, the principal's problem reduces to finding a function $u : \Theta \to \mathbb{R}$ and a scalar $\underline{w}$ that solves:[3]

$$\max_{\underline{w}, u \in \Phi} \left\{ \frac{\underline{\theta}}{\beta} u(\underline{\theta}) + \underline{w} + \frac{1}{\beta} \int_{\underline{\theta}}^{\bar{\theta}} (1 - G(\theta)) u(\theta) d\theta \right\} \tag{4}$$

$$W(y - C(u(\theta))) + \frac{\theta}{\beta} u(\theta) - \frac{\underline{\theta}}{\beta} u(\underline{\theta}) - \underline{w} - \frac{1}{\beta} \int_{\underline{\theta}}^{\theta} u(\tilde{\theta}) d\tilde{\theta} \geq 0 \text{ for all } \theta \in \Theta \tag{5}$$

where $\Phi = \{\underline{w}, u \mid \underline{w} \in W(\mathbb{R}_+),\ u : \Theta \to U(\mathbb{R}_+) \text{ and } u \text{ is non-decreasing}\}$ and

$$G(\theta) \equiv F(\theta) + \theta(1 - \beta) f(\theta).$$

Any allocation $(\underline{w}, u) \in \Phi$ uniquely determines an incentive compatible direct mechanism using (3). An allocation $(\underline{w}, u)$ is *feasible* if $(\underline{w}, u) \in \Phi$ and the budget constraint (5) holds.

Minimum-Savings

This section shows that minimum-savings rules are necessarily part of the optimum.

Bunching at the top can be achieved by removing bundles previously offered for types above $\hat{\theta}$, who then move to the bundle of $\hat{\theta}$, which is the one still available. That is, for any feasible allocation $(\underline{w}, u)$ and $\hat{\theta} \in \Theta$, take the allocation $(\underline{w}, \hat{u})$ given by $\hat{u}(\theta) = u(\theta)$ for $\theta < \hat{\theta}$, and $\hat{u}(\theta) = u(\hat{\theta})$ for $\theta \geq \hat{\theta}$. Thus, bunching the upper tail is always feasible; we now show that it is also always optimal.

Proposition 2 *An optimal allocation* $(\underline{w}, u^*)$ *satisfies* $u^*(\theta) = u^*(\theta_p)$ *for* $\theta \geq \theta_p$*,*

[3]The objective function and the left-hand side of the constraint are well defined for all $(\underline{w}, u) \in \Phi$ since monotonic functions are integrable and the product of two integrable functions, $1 - G(\theta)$ and $u(\theta)$, is integrable (Rudin, 1976, Theorem 6.9 and 6.13).

where θ_p is the lowest value $\theta \in \Theta$ such that

$$\int_{\hat{\theta}}^{\bar{\theta}} \left(1 - G(\tilde{\theta})\right) d\tilde{\theta} \leq 0$$

for $\hat{\theta} \geq \theta_p$. It is optimal for the budget constraint (5) to hold with equality at θ_p.

Proof. The contribution to the objective function from types with $\theta \geq \theta_p$ is $(1/\beta) \int_{\theta_p}^{\bar{\theta}} (1 - G(\theta)) u(\theta) \, d\theta$. Substituting $u = \int_{\theta_p}^{\theta} du + u(\theta_p)$ and integrating by parts we obtain,

$$u(\theta_p) \frac{1}{\beta} \int_{\theta_p}^{\bar{\theta}} (1 - G(\theta)) \, d\theta + \frac{1}{\beta} \int_{\theta_p}^{\bar{\theta}} \int_{\theta}^{\bar{\theta}} \left(1 - G(\tilde{\theta})\right) d\tilde{\theta} du. \tag{6}$$

Note that, for the second term, $\int_{\theta}^{\bar{\theta}} \left(1 - G(\tilde{\theta})\right) d\tilde{\theta} \leq 0$ for all $\theta \geq \theta_p$. It follows that it is optimal to set $du = 0$, or equivalently $u(\theta) = u(\theta_p)$ for $\theta \geq \theta_p$.

When $\theta_p = \underline{\theta}$, all types are pooled at the same bundle and it is clearly not optimal to be in the interior of the budget set. If θ_p is interior then the first term in 6 is zero, so $u(\theta_p)$ can always be increased up to the point where the budget constraint binds without affecting the objective function. Thus, it is optimal not to have money burning at θ_p. ■

This result states that for any bounded distribution of taste shocks a positive mass of upper agents gets the same bundle of consumption and savings, which lies on the budget line. A minimum-savings rule that binds for some types has the property that top types are bunched. Thus, this section of the allocation can be implemented by a minimum-savings rule that is binding for precisely these agents. Thus, minimum-savings are necessarily part of the optimum.

To gain some intuition for this result, note that *self-1* with taste shock $\theta \leq \beta\bar{\theta}$ shares the preferences of *self-0* with a higher taste shock, equal to θ/β. That is, the indifference curves of $\theta u + \beta w$ and $(\theta/\beta)u + w$ are equivalent. Informally, these types can make a case for their preferences. In contrast, *self-1* types with $\theta > \beta\bar{\theta}$ display a blatant desire for current consumption from *self-0*'s point of view. That is, there is no possible taste shock for *self-0* that justifies *self-1*'s preferences. Separating such types requires consumption to increase with θ, but this cannot be optimal since they are overconsuming from *self-0*'s point of view. Thus, these agents should be bunched. In other words, at the very top of the distribution, for $\theta \geq \beta\bar{\theta}$, there is no trade-off between commitment and flexibility. The Lemma shows that bunching goes further, that in the neighborhood of $\beta\bar{\theta}$ the value of commitment continues to dominate that of flexibility: $\theta_p < \beta\bar{\theta}$.[4]

Simple Minimum-Savings Policies

We showed above that minimum-savings are necessarily part of the optimum. We now investigate whether minimum-savings policies may fully characterize the optimum. The results with discrete types suggest the need for some condition on the distribution of taste shocks. The following condition turns out to be exactly what is needed.

Assumption A: *$G(\theta) \equiv (1-\beta)\,\theta f(\theta) + F(\theta)$ is increasing for all $\theta \leq \theta_p$.*

When the density f is differentiable Assumption A is equivalent to a lower bound on its elasticity

$$\theta \frac{f'(\theta)}{f(\theta)} \geq -\frac{2-\beta}{1-\beta} \ ,$$

The lower bound is negative and continuously decreasing in β. The highest lower bound is attained when $\beta = 0$, while as $\beta \to 1$ the lower bound goes off to $-\infty$. Hence, for any density f with $\theta f'/f$ bounded below, Assumption A is satisfied for β close enough to 1. Moreover, many densities satisfy this condition for all β. For example, it is trivially satisfied for all density functions that are non-decreasing, and holds for the exponential distribution, the log-normal, and the Pareto and Gamma distributions for a subset of their parameters.

It is important to recall the two possible interpretations for taste shocks when interpreting Assumption A. Given the state-dependent utility function, $\theta U(c) + \beta W(k)$, an objective interpretation of the distribution of shocks, $F(\theta)$, implies that it can be identified from ex-post behavior. For example, if individuals have full flexibility and choose freely along the budget constraint, the observed distribution of consumption and savings choices, $c^f(\theta)$ and $k^f(\theta)$, identifies the distribution of taste shocks, given the utility functions and the temptation parameter. In contrast, under a subjective interpretation information regarding taste shocks must be elicited directly ex-ante from the individual.

Our next result shows that under Assumption A agents with $\theta \leq \theta_p$ are offered their ex post unconstrained optimum from the budget line, and agents with $\theta \geq \theta_p$

[4]The assumption that taste shocks are bounded above, equivalent to assuming that consumption in the second period is bounded away from zero under full flexibility, ensures that θ_p is well defined. For example, with a Pareto distribution $F(\theta) = 1 - (b/\theta)^{\alpha}$ for $x \geq b$ implying $G(\theta) = 1 + ((1-\beta)\,a - 1)\,(b/\theta)^{\alpha}$. For $\alpha \geq (1-\beta)^{-1}$ one obtains that $\theta_p = b$ so it is optimal to pool all agents. However, for $\alpha < (1-\beta)^{-1}$ there is no solution to θ_p and, it turns out, it is optimal to provide full flexibility.

are bunched at the unconstrained optimum for θ_p. That is, the optimal mechanism offers the whole budget line to the left of the point $(c^f(\theta_p),\ k^f(\theta_p))$ and corresponds to a simple minimum-savings rule that imposes $k \geq k^f(\theta_p)$. Denote the proposed allocation in terms of utility assignments by $(\underline{w}^*, u^*)$, with $\underline{w}^* = W(k^f(\underline{\theta}))$, $u^*(\theta) = U\left(c^f(\theta)\right)$ for $\theta \leq \theta_p$ and $u^*(\theta) = U(c^f(\theta_p))$ for $\theta > \theta_p$.

We next show that this simple allocation is optimal if and only if Assumption A holds. Our strategy involves applying Lagrangian theorems, which require verifying that our problem is sufficiently convex and differentiable. Once this is established the argument is simple: we impose the necessary and sufficient first-order conditions at the conjectured allocation and back out the implied Lagrangian multipliers; the required non-negativity of these multipliers turns out to be equivalent to Assumption A.[5]

Define the Lagrangian function as

$$L(\underline{w}, u|\Lambda) \equiv \frac{\underline{\theta}}{\beta}u(\underline{\theta}) + \underline{w} + \frac{1}{\beta}\int_{\underline{\theta}}^{\bar{\theta}}(1 - G(\theta))\,u(\theta)\,d\theta$$
$$+ \int_{\underline{\theta}}^{\bar{\theta}}\left(W(y - C(u(\theta))) + \frac{\theta}{\beta}u(\theta) - \left(\frac{\underline{\theta}}{\beta}u(\underline{\theta}) + \underline{w}\right) - \int_{\underline{\theta}}^{\theta}\frac{1}{\beta}u(\tilde{\theta})d\tilde{\theta}\right)d\Lambda(\theta),$$

where the function Λ is the Lagrange multiplier associated with the incentive compatibility constraint.[6] Without loss of generality we set $\Lambda\left(\bar{\theta}\right) = 1$. Note that we do not need to incorporate the monotonicity into the Lagrangian. Instead, we work directly with Φ, which includes the monotonicity condition. Integrating the Lagrangian by parts yields:

$$L(\underline{w}, u|\Lambda) = \left(\frac{\underline{\theta}}{\beta}u(\underline{\theta}) + \underline{w}\right)\Lambda(\underline{\theta}) + \frac{1}{\beta}\int_{\underline{\theta}}^{\bar{\theta}}(\Lambda(\theta) - G(\theta))\,u(\theta)\,d\theta$$
$$+ \int_{\underline{\theta}}^{\bar{\theta}}\left(W(y - C(u(\theta))) + \frac{\theta}{\beta}u(\theta)\right)d\Lambda(\theta).$$

The next lemma exploits the convexity of the problem to show that appropriate

[5]One virtue of this approach is that we do not need to restrict the maximization with ad hoc 'technical conditions'. In contrast, most analyses of optimal contracts proceed under assumptions such as piecewise differentiability or continuity.

[6]Intuitively, the Lagrange multiplier Λ can be thought of as a cumulative distribution function that determines the importance of the resource constraints. If Λ is representable by a density λ, then the constraints can be incorporated as the familiar integral of the product with the density function $\lambda(\theta)$. Although this is a common approach, in general, Λ may have points of discontinuity. Indeed, the multiplier we construct has two points of discontinuity.

first-order conditions are necessary and sufficient for optimality.

Lemma of Optimality. *(a) If an allocation* $(\underline{w_0}, u_0) \in \Phi$ *is optimal with* u_0 *continuous then there exists a non-decreasing* Λ_0 *such that the following first-order conditions in terms of Gateaux differentials:*[7]

$$\partial L \left(\underline{w_0}, u_0; \underline{w_0}, u_0 | \Lambda_0\right) = 0 \tag{7}$$

$$\partial L \left(\underline{w_0}, u_0; h_{\underline{w}}, h_u | \Lambda_0\right) \leq 0 \tag{8}$$

hold for all $(h_{\underline{w}}, h_u) \in \Phi$ *and* h_u *continuous.*

(b) Conversely, if there exists a non-decreasing Λ_0 *such that the first-order conditions (7) and (8) hold, for all* $(h_{\underline{w}}, h_u) \in \Phi$*, then* (u_0, w_0) *is optimal.*

Proof. See appendix. ■

Using the second expression for the Lagrangian the Gateaux differential at the proposed allocation $(\underline{w}^*, u^*)$ is given by:

$$\partial L \left(\underline{w}, u; h_{\underline{w}}, h_u | \Lambda\right) = \left(\frac{\underline{\theta}}{\beta} h_u \left(\underline{\theta}\right) + h_{\underline{w}}\right) \Lambda \left(\underline{\theta}\right) + \frac{1}{\beta} \int_{\underline{\theta}}^{\bar{\theta}} \left(\Lambda \left(\theta\right) - G \left(\theta\right)\right) h_u \left(\theta\right) d\theta \tag{9}$$
$$+ \frac{\theta_p}{\beta} \int_{\theta_p}^{\bar{\theta}} \left(\frac{\theta}{\theta_p} - 1\right) h_u d\Lambda \left(\theta\right)$$

for all $(h_{\underline{w}}, h_u) \in \Phi$. The next proposition uses this Lemma to prove that a minimum-savings rule is the optimum under Assumption A.

Proposition 3 *The minimum-savings allocation* $(\underline{w}^*, u^*)$ *is optimal if Assumption A holds.*

Proof. We show that there exists a non-decreasing multiplier Λ^* such that the proposed $(\underline{w}^*, u^*)$ satisfies the first-order conditions (7) and (8) for all $(h_{\underline{w}}, h_u) \in \Phi$. Let $\Lambda^* \left(\underline{\theta}\right) = 0$, $\Lambda^* \left(\theta\right) = G \left(\theta\right)$ for $(\underline{\theta}, \theta_p]$, and $\Lambda^* \left(\theta\right) = 1$ for $\theta \in \left(\theta_p, \bar{\theta}\right]$. Note that Λ^* is not continuous; it has an upward jump at $\underline{\theta}$ and a jump at θ_p. We need to show

[7] Given a function $T : \Omega \to Y$, where $\Omega \subset X$ and X and Y are normed spaces, if for $x \in \Omega$ and $h \in X$ the limit

$$\lim_{\alpha \downarrow 0} \frac{1}{\alpha} \left[T \left(x + \alpha h\right) - T \left(x\right)\right]$$

exists, then it is called the Gateaux differential at x with direction h and is denoted by $\partial T \left(x; h\right)$.

that the jump at θ_p is upward. Indeed,

$$\lim_{\theta \downarrow \theta_p} \Lambda^*(\theta) - \Lambda^*(\theta_p) = 1 - G(\theta_p) \geq 0,$$

which follows from the definition of θ_p. To see this, note that if $\theta_p = \underline{\theta}$ the result is immediate since then Λ^* would jump from 0 to 1 at $\underline{\theta}$. Otherwise, by definition θ_p is the lowest $\hat{\theta}$ such that $\int_{\hat{\theta}}^{\bar{\theta}} \left(1 - G\left(\tilde{\theta}\right)\right) d\tilde{\theta} \leq 0$ for all $\theta \geq \hat{\theta}$, which implies that $1 - G(\theta_p) \geq 0$.

Substituting the proposed multiplier Λ^* into the Gateaux differential (9),

$$\partial L(\underline{w}, u; h_{\underline{w}}, h_u | \Lambda^*) = \frac{1}{\beta} \int_{\theta_p}^{\bar{\theta}} (1 - G(\theta)) h_u(\theta) d\theta = \frac{1}{\beta} \int_{\theta_p}^{\bar{\theta}} \left[\int_{\theta}^{\bar{\theta}} (1 - G(\tilde{\theta})) d\tilde{\theta} \right] dh_u(\theta),$$

where the last equality follows by integrating by parts, which can be done given the monotonicity of h_u and by the definition of θ_p. This Gateaux differential is zero at the proposed allocation, and by the definition of θ_p it is non-positive for all h_u non-decreasing. It follows that the first-order conditions (7) and (8) are satisfied for all $(h_{\underline{w}}, h_u) \in \Phi$. ■

Proposition 3 shows that the optimal allocation can be very simple and implemented by imposing a minimum level of savings. The next proposition shows that more complicated schemes are optimal if Assumption A does not hold.

Proposition 4 *If assumption A does not hold, no minimum-savings rule is optimal.*

Proof. Let $[a, b] \subset [\underline{\theta}, \theta_p)$, with $a < b$, be an interval where G is strictly decreasing (Assumption A does not hold). Let $(\hat{\underline{w}}, \hat{u})|_{x_p}$ be a minimum-savings allocation indexed by x_p : $\hat{u}(\theta) = u^f(\theta)$ for $\theta < x_p$; $\hat{u}(\theta) = u^f(x)$ for $\theta \geq x_p$, and $\hat{\underline{w}} = w^f(\underline{\theta})$. So, x_p denotes the proposed bunching point.

The proof proceeds by contradiction. Suppose that $(\hat{\underline{w}}, \hat{u})|_{x_p}$ is optimal for some x_p. Then by part (a) of the Lemma of Optimality, there has to exist a non-decreasing Lagrange multiplier $\hat{\Lambda}$ such that the conditions for optimality (8) and (7) are satisfied at the proposed allocation for all $(h_{\underline{w}}, h_u) \in \Phi$ and h_u continuous. Condition (8) with $h_u = 0$ requires that $\hat{\Lambda}(\underline{\theta}) = 0$ since h_w is unrestricted. Using $\hat{\Lambda}(\underline{\theta}) = 0$ and integrating (9) by parts (Theorem 6.20 in Rudin, 1976, guarantees this step given that h_u continuous) leads to:

$$\partial L\left(\hat{\underline{w}}, \hat{u}; h_{\underline{w}}, h_u | \hat{\Lambda}\right) = \gamma(\underline{\theta}) h_u(\underline{\theta}) + \int_{\underline{\theta}}^{\bar{\theta}} \gamma(\theta) dh_u(\theta), \tag{10}$$

with

$$\gamma(\theta) \equiv \frac{1}{\beta}\int_{\theta}^{\bar{\theta}} (\hat{\Lambda}(\tilde{\theta}) - G(\tilde{\theta}))d\tilde{\theta} + \frac{x_p}{\beta}\int_{\max\{\theta,x_p\}}^{\bar{\theta}} \left(\frac{\tilde{\theta}}{x_p} - 1\right) d\hat{\Lambda}(\tilde{\theta});$$

where by condition (8) it follows that $\gamma(\theta) \leq 0$ for all $\theta \in \Theta$ is necessary for optimality.

Then (7) implies that $\gamma(\theta) = 0$ for $\theta \in [\underline{\theta}, x_p]$, i.e. wherever $\hat{u}$ is strictly increasing. It follows then that $\hat{\Lambda}(\theta) = G(\theta)$ for all $\theta \in (\underline{\theta}, x_p]$. The proposed allocation $(\underline{\hat{w}}, \hat{u})$ thus determines a unique candidate multiplier $\hat{\Lambda}$ in the separating region $(\underline{\theta}, x_p]$. This implies that $x_p \leq a$, otherwise, and the associated multiplier $\hat{\Lambda}(\theta)$, which is equal to G in the separating region, would be decreasing for $\theta \in [a, \min\{x_p, b\}]$. Integrating by parts the second term of the γ equation we obtain:

$$\gamma(x_p) = \frac{1}{\beta}\int_{x_p}^{\bar{\theta}} (1 - G(\tilde{\theta}))d\tilde{\theta},$$

which is independent of the choice of the multiplier $\hat{\Lambda}$. But for any $x_p \leq a < \theta_p$, we have that $\gamma(x_p) > 0$ by the definition of θ_p, contradicting a necessary condition for optimality. Hence no minimum-savings rule is optimal. ■

The minimum-savings allocations in both propositions do not entail money burning. Recall that with three types money burning may be optimal. A situation with three types can be approximated by continuous types taking a sequence of continuous densities becoming increasingly peaked around θ_l, θ_m and θ_h. However, the distributions in the sequence would eventually violate Assumption A, which requires a density with bounded slope. Thus, with continuous distributions that violate Assumption A money burning may be optimal.

Even if one restricts attention to allocations that do not involve money burning, an improvement over the minimum-savings policy can be constructed by removing intervals in the separating regions wherever the monotonicity condition in Assumption A fails.[8] Since the resulting allocation does not involve money burning it illustrates that the proposition does more than rule such allocations out.

This construction also yields intuition into Assumption A, for suppose it's condition is not satisfied for $\theta_a < \theta < \theta_b \leq \theta_p$. When one removes the open interval between $c^f(\theta_a)$ and $c^f(\theta_b)$ all types with $\theta \in (\theta_a, \theta_b)$ move from their unconstrained optimum to one of the two extremes, $c^f(\theta_a)$ or $c^f(\theta_b)$. The change in welfare depends critically on how many of such types moved to the left versus the right, since

[8] A formal statement and proof of this "drilling" result is contained in a previous version of this paper (Amador, Werning and Angeletos, 2003).

welfare rises from those moving left and falls from those moving right. The slope of the density function affects precisely this, explaining its role in Assumption A.

Taken together, the previous two propositions imply that minimum-savings policies completely characterize the optimum if and only if the distribution of taste shocks satisfies Assumption A, which is the case for a wide class of distributions. Recall that as temptation vanishes, so that $\beta \to 1$, the condition in Assumption A is essentially satisfied for all distributions. Thus, optimal simple minimum-savings policies seem especially likely for modest levels of temptation. As mentioned above, minimum-savings are a common feature of government retirement policy in developed countries.

Our result is also relevant for thinking about the market's provision of commitment devices. Indeed, simple market arrangements may be able to mimic the optimal one. Under Assumption A, the optimal allocation can be implemented with the use of a particular form of illiquid asset. Suppose the consumer initially divides his wealth between two assets: liquid and illiquid. Both assets have the same rate of return, but funds invested in the illiquid asset cannot be used for consumption at $t = 1$, they can only be consumed at $t = 2$. Thus, investing in the illiquid asset represents a self-imposed minimum level of savings, and in this way the individual can implement the optimal allocation.

In an earlier version of this paper (Amador, Werning and Angeletos, 2003) we show that all our results extend to finitely many periods with i.i.d. taste shocks. By using a dynamic programming argument each stage is similar to the two-period problem in (4)-(5). Minimum-savings are always part of the solution, and a simple minimum-savings policy completely characterizes the solution if and only if Assumption A holds. This result establishes that the commitment afforded by the illiquid asset structure studied by Laibson (1997) may in fact be fully optimal.

We turn next to comparative statics with respect to temptation. As temptation increases, so that β decreases, θ_p decreases so more types are bunched and the minimum-savings level decreases.

Proposition 5 *The bunching point θ_p increases with β. The minimum-savings level, $k_{\min} = y - C\left(u\left(\theta_p\right)\right)$, decreases with β.*

Proof. That θ_p is weakly increasing follows directly from its definition. To see that $k_{\min}$ is decreasing note that it solves $(\theta_p/\beta)U'\left(y - k_{\min}\right)/W'\left(k_{\min}\right) = 1$, and that an interior θ_p solves $\theta_p/\beta = \mathbb{E}\left[\theta \,|\, \theta \geq \theta_p\right]$. Combining these we obtain

$$\mathbb{E}\left[\theta|\theta \geq \theta_p\right] \frac{U'\left(y - k_{\min}\right)}{W'\left(k_{\min}\right)} = 1$$

Since $\mathbb{E}[\theta \mid \theta \geq \theta_p]$ is increasing in θ_p, the result follows from the concavity of U and W. ■

3 Optimal Commitment with Self-Control

In this section we study an individual facing temptation, but with some power of self-control. Under a condition similar to that used for the quasi-hyperbolic case, we show that minimum-savings rules are also optimal in such cases.

Dekel, Lipman and Rustichini (2001, 2004) and Gul and Pesendorfer (2001) consider ex ante preferences defined over choice sets made available to the agent ex post. Our specification for the ex ante utility of set C is

$$P(C) = \mathbb{E}\,[\max_{(c,k)\in C} (\theta U(c) + W(k) + \varphi(\theta U(c) + \beta W(k))) - \varphi \max_{(c,k)\in C} (\theta U(c) + \beta W(k))], \tag{11}$$

which is adopted from Krusell, Kuruscu and Smith (2001). The parameter $\varphi > 0$ captures the cost of self-control, while β captures the strength of the temptation to consume in the current period. Our specification has the convenient property that as $\varphi \to \infty$ preferences converge to the quasi-hyperbolic model.

Up to this point we have allowed only the taste shock θ to be uncertain. We now pursue a generalization, allowing the levels of temptation and self-control to be uncertain as well. There are two motivations for such a generalization. First, in a recent paper Dekel, Lipman and Rustichini (2004) provide plausible economic examples illustrating the need for uncertain temptation and provide axiomatic foundations for it. Second, the generalization also allows us to capture the commonly held view that differences in savings may be partly due to differences in temptation or self-control (Diamond, 1977). We assume θ, β and φ are drawn from a continuous joint distribution over some bounded rectangular support $[\underline{\theta}, \bar{\theta}] \times [\underline{\beta}, \bar{\beta}] \times [\underline{\varphi}, \bar{\varphi}]$ where $\underline{\beta} > 0$.

The optimal commitment problem can be stated as maximizing $P(C)$ by choosing a subset $C \subset B(y)$ where $B(y) = \{(c,k) \mid c + k \leq y\}$ is the budget constraint. As before, we seek to rewrite this as a principal-agent problem. The objective function in (11) can be written as:

$$\mathbb{E}\ \{(1+\beta\varphi) \max_{\upsilon,\omega\in C}[(\theta/\hat{\beta})\upsilon + \omega] - \varphi\beta \max_{\upsilon,\omega\in C}[(\theta/\beta)\upsilon + \omega]\},$$

where we let $\hat{\beta}$ be $(1+\beta\varphi)\,/\,(1+\varphi)$. Define the random variables $\hat{z}$ and z by $\hat{z} \equiv \theta/\hat{\beta}$

and $z \equiv \theta/\beta$, and let the extended support $\hat{\Theta}$ be the union of the supports for z and $\hat{z}$, so that $\hat{\Theta} \equiv [\underline{x}, \bar{x}] \equiv [\underline{\theta}(1+\underline{\phi})/(1+\overline{\beta}\underline{\phi}),\ \bar{\theta}/\underline{\beta}]$. Let an allocation over the extended support $\hat{\Theta}$ be given by a pair of functions $u : \hat{\Theta} \to U\left(R_+\right)$ and $w : \hat{\Theta} \to W\left(R_+\right)$.

The principal-agent formulation of the commitment problem is to find an allocation that maximizes

$$\mathbb{E}\left[(1+\beta\varphi)\left(\hat{z}u(\hat{z}) + w(\hat{z}) - \beta\varphi(zu(z) + w(z)\right)\right] \tag{12}$$

subject to $C\left(u\left(x\right)\right) + K(w(x)) \leq y$ and

$$(u(x), w(x)) \in \arg\max_{\upsilon,\omega \in C}[x\upsilon + \omega] \tag{13}$$

for all $x \in \hat{\Theta}$.

Let $\alpha\left(x\right) \equiv \mathbb{E}\left[(1+\beta\varphi) \mid \hat{z} = x\right]$, and $\kappa\left(x\right) \equiv \mathbb{E}\left[\beta\varphi \mid z = x\right]$. Denote by $h_1\left(\hat{z}\right)$ and $h_2\left(z\right)$ the densities of $\hat{z}$ and z, respectively. By the law of iterated expectations, we have that the new objective function (12) can be written as

$$\int_{\hat{\Theta}} (xu(x) + w(x))\hat{g}\left(x\right) dx, \tag{14}$$

where the density $\hat{g}\left(x\right) \equiv \alpha\left(x\right) h_1\left(x\right) - \kappa\left(x\right) h_2\left(x\right)$ can be negative or positive, and defines a signed measure over the state space $\hat{\Theta}$. This alternative expression for the utility function (11) corresponds to that obtained in the representation theorem by Dekel, Lipman and Rustichini (2001).

The incentive compatibility constraints (13) are equivalent, as before, to

$$xu\left(x\right) + w\left(x\right) = \underline{x}u\left(\underline{x}\right) + \underline{w} + \int_{\underline{x}}^{x} u\left(x'\right) dx' \tag{15}$$

with the associated monotonicity constraint.

We can now substitute (15) into the objective function (14) and the resource constraints. Let $\hat{G}\left(x\right) = \int^{x} \hat{g}\left(z\right) dz$ (where $\hat{G}\left(\bar{x}\right) = 1$). Integrating the objective function by parts then yields the following program,

$$\max_{(\underline{w}, u(\cdot)) \in \hat{\Phi}} \left\{ \underline{x}u\left(\underline{x}\right) + \underline{w} + \int_{\hat{\Theta}} [1 - \hat{G}\left(x\right)] u\left(x\right) dx \right\} \tag{16}$$

subject to the resource constraints:

$$W\left(y - C\left(u\left(x\right)\right)\right) + xu\left(x\right) - \underline{x}u\left(\underline{x}\right) - \underline{w} - \int_{\underline{x}}^{x} u(x')dx' \geq 0, \tag{17}$$

where $\hat{\Phi} \equiv \{\underline{w}, u \mid \underline{w} \in W\left(\mathbb{R}_{+}\right), u : \hat{\Theta} \to U\left(\mathbb{R}_{+}\right)$ and u is non-decreasing$\}$.

With the problem mapped into a version that is formally equivalent to the problem (4)-(5) the following propositions are direct extensions of our previous results.

Proposition 6 *Let x_p be the lowest value in $\hat{\Theta}$ such that for all $\hat{x} \geq x_p$*

$$\int_{\hat{x}} (1 - \hat{G}\left(x\right))dx \ \leq 0$$

An optimal allocation, $\left(\underline{w}^*, u^*\right)$ has $u^*\left(x\right) = u^*\left(x_p\right)$ for $x \geq x_p$.

Define the full flexibility allocation as $u^f(x), w^f(x) \in \arg\max_{u,w}\left\{xu + w\right\}$ subject to $C\left(u\right) + K\left(w\right) \leq y$. Let the proposed allocation be given by $\underline{w} = w^f\left(\underline{x}\right)$, and $u^*\left(x\right) = u^f\left(x\right)$ if $x < x_p$ and $u^*\left(x\right) = u^f\left(x_p\right)$ if $x \geq x_p$. We introduce the following assumption analogous to that of Assumption A,

Assumption B $\hat{G}\left(x\right)$ *is non-decreasing for all $x \leq x_p$.*

The next proposition states that minimum-savings rules are optimal under assumption B.

Proposition 7 *The allocation $\left(\underline{w}^*, u^*\right)$ is optimal if assumption B holds. If assumption B does not hold, no minimum-savings rule is optimal.*

Proof. The proof of the first statement is identical to that in Proposition 3, except that the multiplier Λ does not jump at the bottom $\underline{x}$ (because here $\hat{G}\left(x\right)$ is zero at $\underline{x}$). The second statement follows the proof of Proposition 4. ■

We now discuss two results that are possible when temptation and self-control, β and φ, are not uncertain. Proofs for both results are contained in the appendix. We first connect the condition behind Assumption A, from Section 2, with assumption B. The first result shows that self-control strengthens the case for minimum-savings to be enough for the optimum. The second result shows that higher temptation and lower self-control raises the minimum-savings level, and thus increases commitment at the expense of flexibility.

Proposition 8 ($A \Rightarrow B$) *When β and φ are certain, assumption B holds if $G\left(\theta\right)$ is non-decreasing on $[\underline{\theta}, \hat{\beta}x_p]$.*

Proposition 9 *The bunching point x_p increases and the minimum savings decreases with β and decreases with φ.*

4 Other Applications

Our model can be reinterpreted and applied to other situations that also feature a tradeoff between commitment and flexibility, and that are unrelated to the intertemporal temptation model we have focused on. In this section we discuss three such examples.

4.1 Optimal Paternalism

Few would argue that parents are not, at times, literally paternalistic towards their children. Much government regulation – such as minimum schooling laws, drinking and drug restrictions or prohibitions, etc – is also largely justified on paternalistic grounds. Paternalism involves disagreement regarding preferences *between* individuals, instead of *within* an individual as is the case with temptation. The crucial feature of our model, however, is disagreement, regardless of its source, so it can be applied to these situations.

Consider, for example, the case of a child who must divide time between schooling, s, and leisure, l, constrained by a time endowment, $s + l \leq 1$. The child has utility function $\theta U(l) + \beta W(s)$ with $\beta < 1$. The parameter θ affects the relative valuation of schooling versus leisure and is private information. The parent cares about the child but has a different preference over his allocation of time. In particular, the parent has utility $\theta U(l) + W(s)$. She values schooling more than the child does.

The problem faced by the parent maps directly into our time-inconsistent setup. Our result then provides conditions under which minimum-schooling rules are optimal.

4.2 Optimal Fiscal Constitutions

Consider an economy where a ruling government decides the allocation of resources between private and public consumption. Ex post, the government obtains valuable information regarding the social value of public services, but is biased towards higher public spending. Ex ante society faces the constitutional problem of deciding the restrictions to place on the fiscal choices of the government.

We map this problem into our framework as follows. Citizens' welfare is given by

$$\theta U(g) + W(c)$$

and the resource constraint is the set $B(y) = \left\{(c,g) \in \mathbb{R}_+^2 \mid c + g \leq y\right\}$, where c denotes private consumption and g public services. The realization of the value of public services θ is private information of the government. The government can finance g with lump-sum taxes, in which case any allocation $(c,g) \in B(y)$ can be implemented as a competitive equilibrium. A fiscal constitution is a subset $C \subseteq B(y)$. Given $C \subseteq B(y)$ the government chooses $(c,g) \in C$ to maximize $\beta^{-1}\theta U(g) + W(c)$, where $\beta^{-1} > 1$ parameterizes the government's bias towards public spending.

The optimal constitution is the set $C \subseteq B(y)$ that maximizes citizens' welfare given the behavior of the government. Proposition 2 shows that it is always optimal to limit government spending. Proposition 3 implies that under Assumption A only an upper cap on government spending is needed.

4.3 Externalities

There are two consumption goods, c and k. The population is composed of a continuum of agents indexed by θ, distributed according to $F(\theta)$. The utility of agent-θ is given by

$$V(\theta) \equiv \theta U(c(\theta)) + \beta W(k(\theta)) + (1-\beta) \int W(k(\tilde{\theta})dF(\tilde{\theta}),$$

where $\beta < 1$ and $(c(\theta),\ k(\theta))$ represent the allocation in the population. The last term captures a positive externality generated by the consumption of good k.

Agents do not internalize the externality and maximize $\theta U(c) + \beta W(k)$. A utilitarian planner, however, maximizes:

$$\int V(\theta)\, dF(\theta) = \int \left[\theta U(c(\theta)) + W(k(\theta))\right] dF(\theta).$$

This welfare function is equivalent to one without externalities but where a utilitarian planner assigns utility $\theta U(c) + W(k)$ to agent θ. Suppose that the only instrument available to the government is the removal of consumption opportunities. This maps directly into our framework, and our main result provides conditions for the optimality of a rule that imposes a minimum level of consumption for the good generating positive externalities.

5 Conclusions

Our consumer values commitment to avoid the temptation of current consumption, and flexibility to respond to taste shocks. The resulting tradeoff makes the design of an optimal commitment device non-trivial.

We find that a minimum-savings rule is always part of the optimal commitment policy. Moreover, a minimum-savings rule completely characterizes the optimum when a condition on the distribution of taste shocks is satisfied. The minimum-savings level then increases with the strength of temptation. These results are robust to the way temptation is modeled and can be extended to situations with uncertain levels of temptation and self-control, as well as to longer time horizons.

Our model and results can be applied other situations featuring similar trade-offs between commitment and flexibility, such as paternalism, the design of fiscal constitutions to control government spending, and externalities. Another potential application is to problems of time inconsistency of government policy, to examine the tradeoff of rules vs. discretion.

To isolate the problem of commitment as one reducing available choices from the budget set this paper ignored the possibility of transfers across types. An interesting direction for future research is to consider insurance and taxes that allow these transfers in order to provide a more complete characterization of the optimal tax and social-security policies for the class of environments we have considered in this paper.[9]

Appendix

Proof of Proposition 1

With $\beta = 1$ the incentive constraints are slack at the first-best allocation. Define $\beta^* < 1$ to be the value of β for which the incentive constraint of agent-θ_l holds with equality at the first-best allocation. Then for $\beta > \beta^*$ both incentive constraints are

[9]See Amador, Angeletos, and Werning (2004).

slack at the first best allocation and $\beta^* > \theta_l/\theta_h$ follows since

$$\beta^* \equiv \theta_l \frac{U\left(c^{fb}\left(\theta_h\right)\right) - U\left(c^{fb}\left(\theta_l\right)\right)}{W\left(y - c^{fb}\left(\theta_l\right)\right) - W\left(y - c^{fb}\left(\theta_h\right)\right)}$$

$$> \theta_l \frac{U'\left(c^{fb}\left(\theta_h\right)\right)\left(c^{fb}\left(\theta_h\right) - c^{fb}\left(\theta_l\right)\right)}{W'\left(y - c^{fb}\left(\theta_h\right)\right)\left(c^{fb}\left(\theta_h\right) - c^{fb}\left(\theta_l\right)\right)} = \theta_l \frac{U'\left(c^{fb}\left(\theta_h\right)\right)}{W'\left(y - c^{fb}\left(\theta_h\right)\right)} = \frac{\theta_l}{\theta_h}$$

Now, consider the case where $\beta > \theta_l/\theta_h$ and suppose that $c\left(\theta_h\right) + k\left(\theta_h\right) < y$. Then an increase in $c\left(\theta_h\right)$ and a decrease in $k\left(\theta_h\right)$ that holds $\left(\theta_l/\beta\right) U\left(c\left(\theta_h\right)\right) + U\left(k\left(\theta_h\right)\right)$ unchanged increases $c\left(\theta_h\right) + k\left(\theta_h\right)$ and the objective function. Such a change is incentive compatible because it strictly relaxes the incentive constraint of the high type pretending to be a low type, leaving the incentive constraint of the low type unchanged. It follows that we must have $c\left(\theta_h\right) + k\left(\theta_h\right) = y$ at an optimum. This also shows that separating is optimal in this case, proving part (a). Analogous arguments establish parts (b) and (c).

Finally, $c\left(\theta_l\right) + k\left(\theta_l\right) < y$ cannot be optimal since lowering $c\left(\theta_l\right)$ and raising $k\left(\theta_l\right)$ holding $\theta_l U(c(\theta_l)) + \beta W(k(\theta_l))$ constant would then be feasible. Such a variation does not affect the incentive constraint of the low type and relaxes the incentive constraint of the high type, yet it increases the objective function since $\theta_l U(c(\theta_l)) + W(k(\theta_l))$ increases.

Lemma of Optimality and First-Order Conditions

We first show that the maximization of the Lagrangian is a necessary and sufficient condition for optimality of an allocation. This is stated in the following two results:

Result (a'). *Necessity.* If an allocation $\left(\underline{w_0}, u_0\right) \in \Phi$ with u_0 continuous is optimal then there exists a non-decreasing Λ_0 such that the Lagrangian is maximized:

$$L\left(\underline{w_0}, u_0; \underline{h_w}, h_u | \Lambda_0\right) \leq L\left(\underline{w_0}, u_0; \underline{w_0}, u_0 | \Lambda_0\right) \text{ for all } \left(\underline{h_w}, h_u\right) \in \Phi, h_u \text{ continuous} \quad (18)$$

Result (b'). *Sufficiency.* An allocation $\left(\underline{w_0}, u_0\right) \in \Phi$ is optimal if there exists a non-decreasing Λ_0 such that

$$L\left(\underline{w_0}, u_0; \underline{h_w}, h_u | \Lambda_0\right) \leq L\left(\underline{w_0}, u_0; \underline{w_0}, u_0 | \Lambda_0\right) \text{ and all } \left(\underline{h_w}, h_u\right) \in \Phi. \quad (19)$$

Proof. Our optimization problem maps into the general problem studied in Section

8.3-8.4 by Luenberger (1969): $\max_{x \in X} Q(x)$ subject to $x \in \Omega$ and $G(x) \in P$, where Ω is a subset of the vector space X, $Q : \Omega \to \mathbb{R}$ and $G : \Omega \to Z$, where Z is a normed vector space, and P is a positive non-empty convex cone in Z.

For Result (b'), set:

$$
\begin{aligned}
&X = \{\underline{w}, u \mid \underline{w} \in W(\mathbb{R}_+) \text{ and } u : \Theta \to \mathbb{R}\}, \\
&\Omega = \{\underline{w}, u |\ \underline{w} \in W(\mathbb{R}_+),\ u : \Theta \to U(\mathbb{R}_+) \text{ and } u \text{ is non-decreasing}\} \equiv \Phi, \\
&Z = \left\{ z \mid z : \Theta \to \mathbb{R} \text{ with } \sup_{\theta \in \Theta} |z(\theta)| < \infty \right\} \text{ with the norm } \|z\| = \sup_{\theta \in \Theta} |z(\theta)|, \\
&P = \{z \mid z \in Z \text{ and } z(\theta) \geq 0 \text{ for all } \theta \in \Theta\}.
\end{aligned}
$$

We let the objective function be Q and the left-hand side of the resource constraint be defined as G. Result (b') then follows immediately since the hypothesis of Theorem 1, pg. 220 in Luenberger (1969) are met.

For Result (a'), modify Ω and Z to require continuity of u:

$$
\begin{aligned}
&\Omega = \{\underline{w}, u |\ \underline{w} \in W(\mathbb{R}_+),\ u : \Theta \to U(\mathbb{R}_+), \text{and } u \text{ is continuous and non-decreasing}\} \\
&Z = \{z \mid z : \Theta \to \mathbb{R} \text{ and } z \text{ is continuous}\}, \text{ with the norm } \|z\| = \sup_{\theta \in \Theta} |z(\theta)|
\end{aligned}
$$

with X, P, Q and G as before. Note that Q and G are concave, Ω is convex, P contains an interior point (e.g. $z(\theta) = 1$ for all $\theta \in \Theta$) and that the positive dual of Z is isomorphic to the space of non-decreasing functions on Θ by the Riesz Representation Theorem (see Chapter 5, pg. 113 in Luenberger (1969)). Finally, if $\underline{w_0}, u_0$ is optimal within Φ and $\underline{w_0}, u_0 \in \Phi \cap \{u \textit{ is continuous}\}$ then $\underline{w_0}, u_0$ is optimal within the subset $\Phi \cap \{u \textit{ is continuous}\} \equiv \Omega$. Result (a') then follows since the hypotheses of Theorem 1, pg. 217 in Luenberger (1969) are met. ■

Once we have obtained results (a') and (b'), to prove the Lemma of Optimality, we need to show that the maximization conditions in (18) and (19) are equivalent to the appropriate first-order conditions. We first show that these first-order conditions can indeed be computed. The following Lemma helps do this.

Lemma A.1. (Differentiability of integral functionals with convex integrands). *Given a measure space* (Θ, Θ, μ) *and a function* $\psi : X \times \Theta \to R$, *where* $X \subset R^n$, *suppose the functional* $T : \Omega \to R$, *where* Ω *is some subset of the space of all functions mapping* Θ *into* X, *is given by* $T(x) = \int_\Theta \psi(x(\theta), \theta)\, \mu(d\theta)$.

Suppose that (i) for each $\theta \in \Theta$, $\psi(\cdot, \theta) : X \to R$ *is concave; (ii) that the derivative*

ψ_x exists and is a continuous function of (x,θ); and that (iii) $x+\alpha h \in \Omega$ for $\alpha \in [0,\varepsilon]$ for some $\varepsilon > 0$.

Then the h-directional Gateaux differential, $\partial T(x;h)$ exists and is given by

$$\partial T(x;h) = \int_{\Theta} \psi_x(x(\theta),\theta)\, h(\theta)\, \mu(d\theta),$$

if the right hand side expression is well defined.

Proof. Adding and subtracting $\int_{\Theta} \psi_x(x(\theta),\theta)\, h(\theta)\, \mu(d\theta)$ from the definition of the Gateaux differential,

$$\begin{aligned}\partial T(x;h) = &\int_{\Theta} \psi_x(x(\theta),\theta)\, h(\theta)\, \mu(d\theta) \\ &+ \lim_{\alpha\downarrow 0} \int_{\Theta} \left[\frac{1}{\alpha}\left[\psi(x(\theta)+\alpha h(\theta),\theta) - \psi(x(\theta),\theta)\right] - \psi_x(x(\theta),\theta)\, h(\theta)\right] \mu(d\theta).\end{aligned}$$

We seek to show that the last term is well defined and vanishes.

For $\alpha < \varepsilon$ one can show that,

$$\begin{aligned}&\left|\frac{1}{\alpha}\left[\psi(x(\theta)+\alpha h(\theta),\theta) - \psi(x(\theta),\theta)\right] - \psi_x(x(\theta),\theta)\, h(\theta)\right| \\ &\leq \left|\frac{1}{\varepsilon}\left[\psi(x(\theta)+\varepsilon h(\theta),\theta) - \psi(x(\theta),\theta)\right] - \psi_x(x(\theta),\theta)\, h(\theta)\right|, \end{aligned} \tag{20}$$

by concavity of $\psi(\cdot,\theta)$. Given that $\psi(x(\theta)+\varepsilon h(\theta),\theta)$, $\psi(x(\theta),\theta)$ and $\psi_x(x(\theta),\theta)\, h(\theta)$ are all integrable by hypothesis, it follows that $\frac{1}{\varepsilon}\left[\psi(x(\theta)+\varepsilon h(\theta),\theta) - \psi(x(\theta),\theta)\right] - \psi_x(x(\theta),\theta)\, h(\theta)$ is also integrable. Since a function is integrable if and only if its absolute value is integrable, then (20) provides the required integrable bound to apply Lebesgue's Dominated Convergence Theorem (see Theorem 7.10, pg. 192, Stokey and Lucas with Prescott, 1989) implying:

$$\begin{aligned}&\lim_{\alpha\downarrow 0} \int_{\Theta} \left[\frac{1}{\alpha}\left[\psi(x(\theta)+\alpha h(\theta),\theta) - \psi(x(\theta),\theta)\right] - \psi_x(x(\theta),\theta)\, h(\theta)\right] \mu(d\theta) \\ &= \int_{\Theta} \left[\lim_{\alpha\downarrow 0} \frac{1}{\alpha}\left[\psi(x(\theta)+\alpha h(\theta),\theta) - \psi(x(\theta),\theta)\right] - \psi_x(x(\theta),\theta)\, h(\theta)\right] \mu(d\theta) = 0\end{aligned}$$

by definition of ψ_x. It follows that $\partial T(x;h) = \int_{\Theta} \psi_x(x(\theta),\theta)\, h(\theta)\, \mu(d\theta)$. ■

We can apply the lemma A.1 because the Lagrangian functional is the sum of three terms that can be expressed as integrals with concave differentiable integrands. Since the Lagrangian functional is defined over a convex cone Φ, the hypothesis (iii) of the lemma is met with any $\varepsilon \leq 1$ for any $x \in \Phi$ and $h = y - x$ for $y \in \Phi$.

Furthermore, in our case $\int \psi_u (u(\theta), \theta) h_u(\theta) d\Lambda(\theta)$ is well defined for any u and h_u such that $(\underline{w}, u) \in \Phi$ and $(h_{\underline{w}}, h_u) \in \Phi$, for some $\underline{w}, h_{\underline{w}} \in \mathbb{R}$. This follows since u and h_u are non-decreasing on Θ, they are measurable and bounded; and by standard arguments $\psi_u (u(\theta), \theta) h_u(\theta)$ is also measurable and bounded, and thus integrable.

These arguments establish that we can write the Gateaux differential of the Lagrangian for $(\underline{w}, u)$, $(h_{\underline{w}}, h_u) \in \Phi$ as

$$\begin{aligned} \partial L(\underline{w}, u; h_{\underline{w}}, h_u | \Lambda) = & \left(\frac{\underline{\theta}}{\beta} h_u(\underline{\theta}) + h_{\underline{w}} \right) \Lambda(\underline{\theta}) + \frac{1}{\beta} \int_{\underline{\theta}}^{\bar{\theta}} (\Lambda(\theta) - G(\theta)) h_u(\theta) d\theta \\ & + \int_{\underline{\theta}}^{\bar{\theta}} \left[\frac{\theta}{\beta} - W'(y - C(u(\theta))) C'(u(\theta)) \right] h_u d\Lambda(\theta) \end{aligned}$$

which collapses to (9) at the proposed allocation.

Finally, the following Lemma, which is a simple extension of a result in Luenberger (Lemma 1, pg. 227, 1969), allows us to characterize the maximization conditions of the Lagrangian (obtained in results (a') and (b')) by the appropriate first-order conditions.

Lemma A.2. (Optimality and first-order conditions) *Let f be a concave functional on P, a convex cone in X. Take $x_0 \in P$ and define $H(x_0) \equiv \{h : h = x - x_0 \text{ and } x \in P\}$. Then $\delta f(x_0, h)$ exists for $h \in H(x_0)$. Assume that $\delta f(x_0, \alpha_1 h_1 + \alpha_2 h_2)$ exists for $h_1, h_2 \in H(x_0)$ and $\delta f(x_0, \alpha_1 h_1 + \alpha_2 h_2) = \alpha_1 \delta f(x_0, h_1) + \alpha_2 \delta f(x_0, h_2)$ for all $\alpha_1, \alpha_2 \in R$.*

A necessary and sufficient condition that $x_0 \in P$ maximizes f is that

$$\begin{aligned} \delta f(x_0, x) & \leq 0 \text{ for all } x \in P \\ \delta f(x_0, x_0) & = 0 \end{aligned}$$

In our case, all the hypotheses of Lemma A.2 are met for the Lagrangian, because it is a convex functional over a convex cone, and because Lemma A.1 verifies the differentiability requirement, as discussed above. Thus, we obtain that a necessary and sufficient condition for the Lagrangian to be maximized at $\left(u_0, \underline{w_0}\right)$ over Φ is

$$\begin{aligned} \partial L\left(\underline{w_0}, u_0; \underline{w_0}, u_0 | \Lambda_0\right) = 0, \\ \partial L\left(\underline{w_0}, u_0; h_{\underline{w}}, h_u | \Lambda_0\right) \leq 0, \end{aligned}$$

for all $(h_{\underline{w}}, h_u) \in \Phi$.

Given results (a') and (b'), the proof of the Lemma of Optimality follows.

Proof of Proposition 8

Writing $1-\hat{G}(x)=(1+\beta\varphi)(1-F(\hat{\beta}x))-\beta\varphi(1-F(\beta x))$, integrating and rearranging:

$$\begin{aligned}
\int_{x_0}^{\bar{x}}\left(1-\hat{G}(z)\right)dz &= (1+\varphi)\,\hat{\beta}\int_{x_0}^{\bar{x}}\left(1-F(\hat{\beta}x)\right)dx-\beta\varphi\int_{x_0}^{\bar{x}}(1-F(\beta x))\,dx \\
&= (1+\varphi)\int_{\hat{\beta}x_0}^{\bar{\theta}}(1-F(\theta))\,d\theta-\varphi\int_{\beta x_0}^{\bar{\theta}}(1-F(\theta))\,d\theta \\
&= \int_{\beta x_0}^{\bar{\theta}}(1-F(\theta))\,d\theta-(1+\varphi)\int_{\beta x_0}^{\hat{\beta}x_0}(1-F(\theta))\,d\theta \\
&= \int_{\beta x_0}^{\bar{\theta}}(1-F(\theta))\,d\theta-\int_{0}^{x_0(1-\beta)}\left(1-F\left(\frac{y}{1+\varphi}+\beta x_0\right)\right)dy
\end{aligned}$$

The second equality uses the change in variables $\theta=\hat{\beta}x$ for the first integral, $\theta=\beta x$ for the second, and the fact that $\hat{\beta}\bar{x}>\beta\bar{x}=\bar{\theta}$. The third equality simply rearranges the integrals. The fourth equality performs the change of variables $y=(1+\varphi)(\theta-\theta_0)$ using the fact that $1+\varphi=(1-\beta)/(\hat{\beta}-\beta)$.

The comparative static with respect to φ is now straightforward: an increase in φ raises the integrand $1-F(y/(1+\varphi)+\beta x_0)$ so that x_p must fall with φ. To obtain the comparative static with respect to β we differentiate the last expression:

$$\frac{\partial}{\partial\beta}\int_{x_0}^{\bar{x}}\left(1-\hat{G}(z)\right)dz=\left[F(\hat{\beta}x_0)-F(\beta x_0)+\int_{0}^{x_0(1-\beta)}f\left(\frac{1}{1+\varphi}y+\beta x_0\right)dy\right]x_0>0,$$

implying that x_p rises with β.

Finally, note that the minimum-savings $k_{\min}$ is defined as the solution to:

$$x_p\frac{U'(y-k_{\min})}{W'(k_{\min})}=1,$$

so that comparative statics for x_p translate directly into $k_{\min}$. In particular, $k_{\min}$ is increasing in φ and decreasing in β.

Proof of Proposition 9

Let $F(\cdot)$ be the c.d.f. of the taste shocks. We want to show that if $G(x) \equiv F(x) + x(1-\beta) f(x)$ is non-decreasing, then $\hat{G}(x) = (1+\beta\varphi) F(\hat{\beta}x) - \beta\varphi F(\beta x)$ is non-decreasing. After letting $\lambda = 1/\varphi$ and differentiating, we obtain

$$\Delta(x,\lambda) \equiv \Big(\frac{\lambda+\beta}{\beta}\Big)\hat{\beta}(\lambda) f\Big(\hat{\beta}(\lambda)x\Big) - \beta f(\beta x) \geq 0,$$

and note that $\Delta(x,0) = 0$. Substituting the definition of $G(\cdot)$ yields the alternative expression,

$$\Delta(x,\lambda) = \frac{\lambda+\beta}{\beta(1-\beta)x}\Big[G(\hat{\beta}(\lambda)x) - F(\hat{\beta}(\lambda)x)\Big] - \beta f(\beta x).$$

Define,

$$\tilde{\Delta}(x,\lambda,z) \equiv \frac{\lambda+\beta}{\beta(1-\beta)x}\Big[G(z) - F(\hat{\beta}(\lambda)x)\Big] - \beta f(\beta x). \tag{21}$$

Note that $\tilde{\Delta}(x,\lambda,z)$ increases in z and that $\tilde{\Delta}(x,\lambda,\hat{\beta}(\lambda)x) = \Delta(x,\lambda)$.

To prove $\Delta(x,\lambda) \geq 0$ we write,

$$\Delta(x,\lambda) = \tilde{\Delta}(x,\lambda,\hat{\beta}(\lambda)x) = \tilde{\Delta}(x,0,\hat{\beta}(\lambda)x) + \int_0^\lambda \tilde{\Delta}_\lambda(x,\tilde{\lambda},\hat{\beta}(\lambda)x)d\tilde{\lambda} \tag{22}$$

and proceed to show that both the terms on the right-hand side are non-negative.

To see the sign of the first term in (22) note that since $\tilde{\Delta}$ is increasing in z,

$$\tilde{\Delta}(x,0,\hat{\beta}(\lambda)x) \geq \tilde{\Delta}(x,0,x) = \tilde{\Delta}(x,0,\hat{\beta}(0)x) = \Delta(x,0) = 0.$$

For the integral term in (22) we compute the integrand by differentiating (21) and rearranging using the definition of $G(\cdot)$:

$$\tilde{\Delta}_\lambda(x,\lambda,z) = \frac{1}{\beta(1-\beta)x}\left[G(z) - G(\hat{\beta}(\lambda)x) + \frac{\lambda}{1+\lambda}\hat{\beta}(\lambda)x(1-\beta)f(\hat{\beta}(\lambda)x)\right].$$

Thus, for $z \geq \hat{\beta}(\tilde{\lambda})x$ we have $\tilde{\Delta}_\lambda(x,\tilde{\lambda},z) \geq 0$. It follows that for $\tilde{\lambda} \in [0,\lambda]$ we have $\hat{\beta}(\lambda)x \geq \hat{\beta}(\tilde{\lambda})x$, and therefore $\tilde{\Delta}_\lambda(x,\tilde{\lambda},\hat{\beta}(\lambda)x) \geq 0$. Thus, the integral term in (22) is non-negative. Given that $\hat{\beta}(\lambda)x_p(\lambda)$ is non-decreasing in λ we need $G(x)$ to be non-decreasing up to $\hat{\beta}x_p$.

References

[1] **Amador, Manuel, Iván Werning, and George-Marios Angeletos** (2003) "Optimal Commitment", NBER Working Paper #10151.

[2] **Amador, Manuel, George-Marios Angeletos, and Iván Werning**(2004) "Corrective vs. Redistributive Taxation", MIT working paper.

[3] **Athey, Susan, Andy Atkeson, and Patrick Kehoe** (2003) "On the Optimality of Transparent Monetary Policy", Federal Reserve Bank of Minneapolis Staff Report 326.

[4] **Benabou, Roland, and Jean Tirole** (2002), "Self-Confidence and Personal Motivation," Quarterly Journal of Economics 117 (3), 871-915.

[5] **Dekel, Eddie, Barton L. Lipman, and Aldo Rustichini** (2001) "Representing Preferences with a Unique Subjective State Space", *Econometrica,* 69, n4, pp. 891-934.

[6] ———————— (2004) "Temptation Driven Preferences", Boston University working paper.

[7] **Diamond, Peter** (1977) "A Framework for Social Security Analysis", *Journal of Public Economics*, Vol. 8, n3, pp. 275-98.

[8] **Feldstein, Martin S.** (1985)"The Optimal Level of Social Security Benefits", *Quarterly Journal of Economics,* Vol. 10, n2, pp. 303-20.

[9] **Gul, Faruk, and Wolfgang Pesendorfer** (2001) "Temptation and Self-Control", *Econometrica*, Vol. 69, n6, pp. 1403-35.

[10] ———————— (2002a) "Self-control and the Theory of Consumption", forthcoming in *Econometrica.*

[11] ———————— (2002b) "Self-control, Revealed Preference and Consumption Choice", prepared for Society for Economic Dynamics.

[12] **Holmstrom, Bengt** (1984) "On the Theory of Delegation", in *Bayesian Models in Economic Theory, Studies in Bayesian Econometrics*, Vol. 5, edited by Boyer, Marcel and Kihlstrom, Richard E.: 115-41.

[13] **Imrohoroglu, Ayse, Selahattin Imrohoroglu, and Douglas H. Joines** (2000) "Time Inconsistent Preferences and Social Security", Federal Reserve of Minneapolis Discussion Paper Number 136.

[14] **Kreps, David**, (1979) "A Preference for Flexibility", *Econometrica*, Vol. 47, n5, pp. 65-576.

[15] **Krusell, Per, Burhanettin Kuruscu, and Anthony Smith** (2001) "Temptation and Taxation", mimeo.

[16] **Laibson, David (**1994) "Self-Control and Saving", unpublished M.I.T. PhD thesis.

[17] ——————— (1997) "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics,* v112, n2, pp. 443-77.

[18] ——————— (1998) "Life-Cycle Consumption and Hyperbolic Discount Functions", *European Economic Review*, Vol. 42, nos. 3-5, pp. 861-871.

[19] **Luenberger, David G.** (1969) *Optimization by Vector Space Methods*, John Wiley & Sons, Inc.

[20] **Milgrom, Paul, and Ilya Segal** (2002) "Envelope Theorems for Arbitrary Choice Set", *Econometrica*, Vol. 70, n2, pp. 583-601.

[21] **Phelps, Edmund S., and R.A. Pollack** (1968) "On Second Best National Savings and Game-Equilibrium Growth", *Review of Economic Studies*, Vol. 35, n2, pp. 185-199.

[22] **Rudin, Walter** (1976) "Principles of Mathematical Analysis", 3rd Edition, McGraw-Hill.

[23] **Stokey, Nancy, and Robert E. Lucas Jr., with Edward C. Prescott** (1989) *Recursive Methods in Economic Dynamics*, Harvard University Press.

[24] **Strotz, Robert H** (1956) "Myopia and Inconsistency in Dynamic Utility Maximization", *Review of Economic Studies*, Vol. 23, pp. 165-180.

[25] **Weitzman, Martin L.** (1974) "Prices vs. Quantities", *Review of Economic Studies*, v41, n4, pp. 477-91.