Commitment-Flexibility Trade-off and Withdrawal Penalties*

Attila Ambrus†
Duke University, Department of Economics

Georgy Egorov‡
Northwestern University, Kellogg-MEDS

September 2012

Abstract

Using the consumption-savings model proposed by Amador, Werning and Angeletos in their 2006 Econometrica paper (henceforth AWA), in which individuals face the trade-off between flexibility and commitment, we show that withdrawal penalties can be part of the optimal contract from an ex ante perspective, despite involving money-burning. For the case of two states (which we interpret as "normal times" and a "negative liquidity shock"), we provide a full characterization of the optimal contract, and show that within the parameter region where the first best is unattainable, the likelihood that withdrawal penalties are part of the optimal contract is decreasing in the probability of a negative liquidity shock and increasing in the severity of the shock. We also show that contracts with the same qualitative feature (withdrawal penalties for high types) arise in continuous state spaces, too. Our conclusions differ from AWA because the analysis in the latter implicitly assumes that the optimal contract is interior (the amount withdrawn from the savings account is strictly positive in each period in every state). Our results are consistent with empirical evidence on withdrawals from individual retirement accounts and time deposit contracts.

Keywords: Commitment, flexibility, self-control, money-burning
JEL Classification: D23, D82, D86.

1 Introduction

In an important paper, Manuel Amador, Iván Werning and George-Marios Angeletos (2006, from now on AWA) study the optimal savings rule in a model where people are tempted to consume earlier, along the line of Strotz (1956), Phelps and Pollack (1968) and Laibson (1997, 1998), but full commitment is undesirable as it does not allow for incorporation of new information, such as taste shocks and income shocks.[1] They provide an optimal rule in two broad situations: if

*We thank Manuel Amador for useful suggestions, and Alexander Groves for valuable research assistance.

†E-mail: aa231@duke.edu

‡E-mail: g-egorov@kellogg.northwestern.edu

[1] See also Gul and Pesendorfer (2001) and Dekel, Lipman and Rustichini (2001) for axiomatic foundations for preferences that imply temptation by present consumption and relatedly demand for commitment. For other papers studying optimal contracts with agents who suffer from self-control problems, see, e.g., DellaVigna and Malmendier (2004), Eliaz and Spiegler (2006) and Esteban and Miyagawa (2005).

the shock variable can only take two values, and if the shock variable is continuous but a simple regularity condition on the density holds. An important feature of the optimum in the above characterization results is that there is no money burning from the consumer's perspective: in every state total consumption over time is equal to total endowment.[2]

We revisit the model of AWA and first analyze the case of two possible taste shocks. We show that money-burning may be used in equilibrium, imposed on the impatient type, in order to provide incentives for the more patient type not to imitate the impatient type. This is in contrast with Proposition 1 in AWA. The reason is that the arguments in AWA implicitly assume that the optimal contract involves allocating strictly positive amounts of the good to be consumed at both time periods, in every state. However, we show that there is an open set of parameter values for which the optimal contract involves 0 consumption in the second time period in case of a negative liquidity shock in the first time period. There is a natural way to modify the framework in AWA that allows for utility functions that indeed guarantee that the optimal contract always specifies an interior consumption plan, and hence the analysis in AWA is valid: adding utility functions that at consumption level 0 take a value of $-\infty$. However, as we show it in the paper, imposing Inada conditions on the utility functions does not rule out corner solutions and money-burning in the optimal contract.

For utility functions in the original AWA framework, we show that money burning becomes part of the optimal contract when the probability of the impatient type is not too large, and when the negative liquidity shock is severe enough. The intuition behind the first feature is that money burning becomes a realized loss in the high liquidity shock state, hence it can only be optimal ex ante if the high liquidity state is not too likely.

For continuous type spaces, we provide a similar result. Take any distribution of shocks satisfying the regularity condition in AWA, and assume that the high type is present-biased enough to prefer to consume everything in the first period.[3] Consider a distribution which is obtained as the original distribution with probability $1-\varepsilon$, and the same distribution shifted to the right by z with probability ε. The interpretation is that on top of the normal shocks, there is an ε likelihood of a catastrophic liquidity shock. Then whenever ε is small enough and z is large enough, the optimal contract involves money-burning and 0 second period consumption in all states following the catastrophic shock. This in particular implies that in the model with a

[2] Analogously, Athey et al. (2004) and Athey et al. (2005) show in various contract theory settings (that are technically connected to the original models they are primarily interested in) that money burning is not part of the optimal contract. Ambrus and Egorov (2009), in a principal-agent setting different from the one in the current paper, characterize cases when money burning can be part of an optimal delegation scheme. See also Amador and Bagwell (2011).

[3] The latter assumption can be relaxed, but it makes the proof much easier. Characterizing the optimal contract in the continuum types setting is difficult in general.

continuum of types, money burning can be imposed in optimum on the highest types. This is not consistent with Proposition 2 from AWA, for reasons similar to why money-burning can be optimal in the model with two states.

The cases in which we find money-burning to be optimal admit a natural withdrawal penalty interpretation: in high liquidity-shock states the agent is allowed to withdraw everything from the savings account in the first period, at the cost of a withdrawal penalty. This is roughly consistent with the rationale behind early withdrawal penalties associated with individual retirement accounts (IRA) and 401(k) accounts in the US, expressed, for example in Thaler (1994): consumers should be allowed to withdraw their savings in case of severe need, but "...the absence of a withdrawal penalty would imply that whatever funds get contributed would be more at risk to a spending spree."[4] The empirical findings of Holden and Schrass (2008), in that only 5% of IRA withdrawals occur before age $59\frac{1}{2}$ (the period for which withdrawal penalties apply) are consistent with early withdrawals only being chosen in rare high liquidity shock states, while Amronim (2002, 2003) provides evidence that curtailing the liquidity of IRAs and 401(k) plans is particularly attractive for people with weaker self-control.[5] Similarly, our analysis also provides a possible new explanation for withdrawal penalties being common features of time deposit contracts offered by commercial banks, complementing the supply-side explanation that they make the task of liquidity management easier for banks.[6] Empirical evidence in this context also supports that early withdrawals are associated with high liquidity shocks of depositors (Amromin and Smith (2003)), and that their likelihood is nonzero but relatively small (Gilkeson (1999)).

Our point that withdrawal penalties might be part of the optimal contract receives empirical support from Beshears et al. (2011), who in a field experiment offer subjects to split their savings between a savings account with no early withdrawal penalties, and a savings account with positive withdrawal penalties (and the same interest rate). They find that subjects contribute a significantly positive amount to the savings account with withdrawal penalties, and that this contribution is significantly higher when the withdrawal penalty is 20% versus when it is 10%.[7]

[4] The idea that people have a tendency to undersave, and that it is related to self-control problems and hence can lead to preference for commitment is originally brought up in Diamond (1977).

[5] Incidentally, Amromin (2003), using a completely different modeling framework than we do, also conclude that catastrophic income shocks are needed to explain empirical regularities regarding deposits to and withdrawals from tax-deferred withdrawal accounts.

[6] The contracting problem in AWA is not embedded in a market environment, hence it is not a good fit for analyzing savings contracts offered by commercial banks, although it could possibly be reinterpreted as the type of contract that would emerge with homogenous consumers and a perfectly competitive banking sector. We do not pursue this extension here. For a recent paper on a monopolist bank contracting with heterogenous and possibly time-inconsistent agents, see Galperti (2012).

[7] For further empirical evidence for demand for ex ante commitment in contracts, see Ashraf et al. (2006), DellaVigna and Malmendier (2006), and Bryan et al. (2010).

2 The model

The setup reintroduces the model from AWA, and we preserve the notation. There are two periods and a single good. A consumer has a budget y and chooses his consumption in periods 1 and 2, c and k, respectively, so his budget set B is defined by $c \geq 0$, $k \geq 0$, $c + k \leq y$ (the interest rate is normalized to 0). The utility of *self-0* (the individual before the consumption periods) is given by:

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

where $U, W : \mathbb{R}^+ \to \mathbb{R} \cup \{-\infty\}$ are two strictly increasing, strictly[8] concave and continuously differentiable functions, and $\theta \in \Theta$ is a taste shock which is realized in period 1. AWA assume that U, W map $\mathbb{R}^+$ to $\mathbb{R}$, thus ruling out the possibility of, say, $U(c) = \log c$. We extend their framework as allowing for this does not complicate the analysis, and as we show leads to some new insights regarding the possibility of money-burning in optimum.

We assume that Θ is bounded and normalized so that $\mathbb{E}\theta = 1$. Denote the c.d.f. of θ by $F(\cdot)$ and the p.d.f. of θ by $f(\cdot)$. The utility of *self-1* is given by

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

where $0 < \beta \leq 1$ captures the degree of agreement between self-0 and self-1 (and $1 - \beta$ captures the strength of temptation towards earlier consumption). The goal is to characterize the optimal contract with self-0 as the principal and self-1 as the agent, i.e., the consumption scheme that self-0 would choose from behind the veil of ignorance about the realization of the taste shock θ.

Hereinafter, we find it convenient to characterize contracts in terms of utilities rather than allocations (each is a monotone transformation of the other). We let $C(u)$ and $K(w)$ be the inverse functions of $U(c)$ and $W(k)$, respectively, and we let set A be given by

$$A = \left\{(u, w) \in \mathbb{R}^2 : u \geq U(0), w \geq W(0), C(u) + K(w) \leq y\right\}.$$

Since $C(u)$ and $K(w)$ are convex functions, the set A is convex. Define function $z(\cdot)$ by

$$z(x) = W(y - C(x));$$

then $z(\cdot)$ is decreasing and strictly concave. The set $\{u, w : w = z(u)\}$ is the frontier of the set

[8] Assuming strict concavity rules out linear utility functions, but simplifies characterization a lot. Clearly, any linear function may be approximated by strictly concave functions, so the results may be applied to characterize the properties of optimal contracts with linear utility functions as well.

A where there is no money-burning: $C(u) + K(w) = y$. Thus, self-0 solves:

$$\max_{(u(\theta),w(\theta))_{\theta\in\Theta}} \int_{\theta\in\Theta} (\theta u(\theta) + w(\theta))\, dF(\theta) \tag{1}$$

$$\text{subject to} \qquad (u(\theta), w(\theta)) \in A \text{ for every } \theta \in \Theta, \tag{2}$$

$$\theta u(\theta) + \beta w(\theta) \geq \theta u(\theta') + \beta w(\theta') \text{ for every } \theta, \theta' \in \Theta. \tag{3}$$

Finally, let $\left(u^{fb}(\theta), w^{fb}(\theta)\right) = \arg\max_{(u,w)\in A} (\theta u + w)$ denote the first best allocation.

3 Two types

Here we consider the case of two types, so that $\Theta = \{\theta_l, \theta_h\}$ with $0 < \theta_l < \theta_h$ (and given the normalization $E\theta = 1$, $\theta_l < 1 < \theta_h$). This setup can be interpreted such that state θ_l represents "normal times", while state θ_h represents a negative liquidity shock, such as a job loss.

If we denote the probability that $\theta = \theta_l$ by μ, we must have

$$\mu\theta_l + (1-\mu)\theta_h = 1. \tag{4}$$

We are thus solving the problem

$$\max_{(u_l,w_l),(u_h,w_h)\in A} (\mu(\theta_l u_l + w_l) + (1-\mu)(\theta_h u_h + w_h)) \tag{5}$$

$$\text{subject to} \qquad \theta_l u_l + \beta w_l \geq \theta_l u_h + \beta w_h, \tag{6}$$

$$\theta_h u_h + \beta w_h \geq \theta_h u_l + \beta w_l. \tag{7}$$

Throughout this section, we use subscripts l and h to denote the values at θ_l and θ_h, respectively, e.g., $u_l \equiv u(\theta_l)$, etc.

AWA, as part of Proposition 1 in this paper, characterizes the parameter regions in which (i) the optimal contract achieves the first best; (ii) does not achieve the first best but implies separation of the two types; and (iii) implies pooling of the two types. Parts of this proof relied on an argument that there is no money-burning in the optimal contract. We show that this need not hold without additional assumptions, and provide the complete proof of this result in the Appendix, even though Part 1 of Proposition 1 of AWA is correct as stated.

Proposition 1 *Suppose $\Theta = \{\theta_l, \theta_h\}$ with $\theta_l < \theta_h$. Suppose that $\theta_l < \left|\frac{dz}{du}|_{u=U(y)}\right|$ and $\theta_h > \left|\frac{dz}{du}|_{u=U(0)}\right|$.*[9] *Then there exists $\beta^* \in (\theta_l/\theta_h, 1)$ such that for $\beta \in [\beta^*, 1]$ the first-best allocation is implementable.*

[9] This requirement ensures that the first best contract is not pooling, and should have been included in Proposition 1 of AWA as well. If $\theta_l \geq \left|\frac{dz}{du}|_{u=U(y)}\right|$, then the optimal contract is $c_l^{fb} = c_h^{fb} = y$, $k_l^{fb} = k_h^{fb} = 0$, and if $\theta_h \leq \left|\frac{dz}{du}|_{u=U(0)}\right|$, then the optimal contract is $c_l^{fb} = c_h^{fb} = 0$, $k_l^{fb} = k_h^{fb} = y$. In either of these cases, the first best is implementable for all β. (If $z(u)$ does not have a left derivative at $u = U(0)$, or $U(0) = -\infty$, then $\left|\frac{dz}{du}|_{u=U(0)}\right|$ is 0.)

If $\beta \leq \theta_l/\theta_h$, then pooling is optimal, i.e., $u_h = u_l$ and $w_h = w_l$; moreover, there is no money-burning in this case: $w_l = z(u_l)$.

If, however, $\beta \in \left(\frac{\theta_l}{\theta_h}, \beta^*\right)$, then separation is optimal, i.e., $u_h > u_l$ and $w_h < w_l$. In this last case, $w_l = z(u_l)$, but both $w_h = z(u_h)$ and $w_h < z(u_h)$ are possible. In either case, the IC constraint of the low type (6) *is binding and the IC constraint of the high type* (7) *is not.*

Proposition 1 of AWA also claims that money burning is never part of the optimal contract, which, as we find, does not have to hold in general. Our next result below gives a necessary and sufficient condition for money burning to be part of the optimal contract. The proof of Proposition 1 in AWA is invalid without further assumptions at the point where the authors write "Then an increase in $c(\theta_h)$ and a decrease in $k(\theta_h)$ that holds $(\theta_l/\beta) U(c(\theta_h)) + U(k(\theta_h))$ unchanged...", which implicitly assumes that a decrease in $k_h = k(\theta_h)$ is possible. If $k_h = 0$, so type θ_h consumes only in period 1, then such a decrease is clearly impossible. We prove that this is the only possible case consistent with money burning (i.e., money-burning implies $c_h < y$, $k_h = 0$), and it is only possible if $W(0) \neq -\infty$ (Proposition 2). In fact, if $k_h > 0$ in the optimal contract then the argument in AWA goes through, ruling out the possibility of money burning.

As a prelude to the next result, the following figures illustrate the two types of separating contracts that are possible in optimum. Note that if the IC constraint is binding for the low type then the line connecting (u_l, w_l) and (u_h, w_h) has to have a slope of $-\theta_l/\beta_l$. Below we refer to this line as the IC_l line. Figure 1 (left) represents a possibility such that at the optimum the IC_l line intersects set A twice at the Pareto frontier. This corresponds to a separating equilibrium with no money burning, as in AWA. Figure 1 (right) represents a different possibility, when at the optimum the IC_l line crosses the horizontal boundary of set A (on the $w = W(0)$ line), implying that there is money burning in equilibrium. Below we show that both of these cases can indeed occur at the optimum.

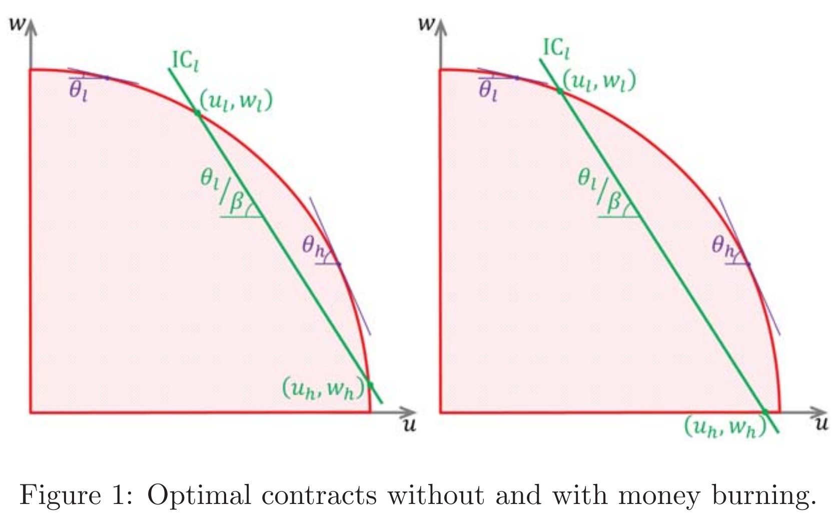

Figure 1: Optimal contracts without and with money burning.

In order to give a precise characterization of when money-burning is part of a separating optimal contract, we need to introduce some further notation. Proposition 1 implies that the IC constraint of type θ_l is binding; let us denote, for any $K \in \mathbb{R}$,

$$\lambda^K = \left\{ (u, w) \in A : u + \frac{\beta}{\theta_l} w = K \right\}. \tag{8}$$

For any K, the above set of points is either a line segment, a point, or the empty set, although for simplicity we just refer to it as the IC_l line. Whenever $\lambda^K \neq \varnothing$, let $\lambda_l^K = \left(u_l^K, w_l^K\right)$ and $\lambda_h^K = \left(u_h^K, w_h^K\right)$ be the points of λ^K that minimize and maximize u, respectively. Fixing $K = u_l + \frac{\beta}{\theta_l} w_l = u_h + \frac{\beta}{\theta_l} w_h$, we observe that $(u_l, w_l) = \lambda_l^K$ and $(u_h, w_h) = \lambda_h^K$ (if it were not the case, then moving (u_l, w_l) north-west along the IC_l line would not violate (6) or (7) and would increase (5), as $\theta_l < \frac{\theta_l}{\beta}$, and moving (u_h, w_h) along the same line would have the same effect as $\theta_h > \frac{\theta_l}{\beta}$). Let us now take a particular value of K,

$$K_0 = U(y) + \frac{\beta}{\theta_l} W(0); \tag{9}$$

then K_0 is finite if $W(0) \neq -\infty$ and $K_0 = -\infty$ otherwise. In the case K_0 is finite, notice that $\lambda_h^{K_0} = (U(y), W(0))$ by definition. The leftmost point of intersection of λ^{K_0} with A, $\lambda_l^{K_0}$, plays a critical role in the following formulation, and we let $u_0 \equiv u_l^{K_0}$.

Proposition 2 *Suppose $\frac{\theta_l}{\theta_h} < \beta < \beta^*$, so the optimal contract is separating. Money-burning will be used as part of the optimal contract if and only if (i) $W(0) \neq -\infty$, (ii) $u_0 > U(0)$, where u_0 is defined as $u_l^{K_0}$ for $K_0 = U(y) + \frac{\beta}{\theta_l} W(0)$, and (iii) for the following inequality holds:*

$$\mu \frac{1 - \beta}{\frac{1}{\left|\frac{dz}{du}\right|_{u=u_0}\right|} - \frac{\beta}{\theta_l}} > 1. \tag{10}$$

While the formal proof is in the Appendix, here we provide a brief intuition for it. The key insight is that when reformulating the optimization program in terms of K, the maximand is strictly concave. Hence, money burning is optimal if and only if starting from an IC allocation that specifies a consumption vector $(y, 0)$ in θ_h, the effect of a marginal amount of money-burning (decreasing first period consumption marginally while keeping second period consumption at 0) in the high state plus the implied consumption vector change through the IC constraint yields a strictly positive expected utility change for the consumer.

The next corollary, which follows directly from the proof described above, further clarifies the set of cases where money-burning may be used.

Corollary 1 *If the optimal contract requires money-burning, then self-1 with type θ_l is impatient enough to prefer allocation $(y, 0)$ to $\left(c^{fb}(\theta_l), k^{fb}(\theta_l)\right)$, i.e.,*

$$\theta_l U(y) + \beta W(0) > \theta_l u_l^{fb} + \beta w_l^{fb}. \tag{11}$$

In particular, $W(0)$ must be finite, so $W(k)$ must be bounded away from $-\infty$. Moreover, whenever the optimal contract requires money-burning, we must have $k_h = 0$.[10]

Note that the case $W(0) = -\infty$ is only realistic if the consumer literally keeps all her resources in the savings account, and has no other source of consumption, and this is unlikely to hold for savings accounts in practice. Putting it differently, it is reasonable in real life savings situations that self-1 is impatient enough (β is low) to want to withdraw all money from the account in period 1, if this option is feasible.

We also want to point out that imposing a condition that $W'(0) = \infty$ (commonly referred to as Inada condition) does not rule out the possibility that the optimal contract involves money-burning and 0 second period consumption in the high state. The intuition is that the IC constraint for the low type is binding. Therefore while marginally increasing second-period consumption in the high state, starting from 0, increases the consumer's expected utility at an infinite rate, this also makes the temptation of the low type to pretend to be a high type, tightening the IC constraint and decreasing utility in the low state at an infinite rate.

The following is an example in which money burning is part of the optimal contract.

Example 1 *Suppose $U(c) = \sqrt{c}$, $W(k) = \sqrt{k}$.*[11] *In this case, $z(u) = \sqrt{y - u^2}$ and $u_0 = \frac{(\theta_l/\beta)^2 - 1}{(\theta_l/\beta)^2 + 1}\sqrt{y}$, $\frac{dz}{du}(u_0) = -\frac{u_0}{\sqrt{y - u_0^2}} = -\frac{1}{2}\left(\frac{\theta_l}{\beta} - \frac{\beta}{\theta_l}\right)$, and the condition (10) becomes*

$$\mu(1-\beta)\frac{\theta_l}{\beta}\frac{(\theta_l/\beta)^2 - 1}{(\theta_l/\beta)^2 + 1} > 1.$$

Now, if we take $\theta_l = \frac{1}{10}$, $\theta_h = 10$, $\mu = \frac{10}{11}$, $\beta = \frac{1}{20}$, the left-hand side equals $\frac{57}{55} > 1$. One can check that the optimal contract is $c_l = \frac{121}{346}$, $k_l = \frac{225}{346}$, $c_h = \frac{1369}{1384}$, $k_h = 0$, and indeed involves money-burning. (The optimal contract with the constraint that money-burning is not allowed would be $c_l = \frac{9}{25}$, $k_l = \frac{16}{25}$, $c_h = 1$, $k_h = 0$, and the ex-ante expected utilities in the two contracts are $\frac{3257}{220\sqrt{346}} = 0.795897$ and $\frac{87}{110} = 0.790909$, respectively, with the difference of 0.005.) In a working paper version, we provided examples with different (power) utility functions, where the use of money-burning increased the gain in ex-ante welfare by more than 36%.

Example 1, which shows that money-burning is possible, is not atypical. In particular, this has nothing to do with the choice of utility functions (except that we need $W(0) \neq -\infty$ to get money-burning). Moreover, as long as the utility functions in both periods are the same, one can find an open set of parameter values (relative to the possible set of parameter values

[10] If $\beta = \theta_l/\theta_h$, then the optimal contract is not uniquely defined, and among them there may be contracts with money-burning and $k_h > 0$. But in this case we can find an optimal contract without money burning.

[11] We thank an anonymous referee for a suggestion that made the example simpler.

defined in the model) for which having money burning is optimal, i.e., (10) is satisfied. In other words, such situations are not knife-edge cases. We formalize this result in the Supplementary Appendix, and also show that for utility functions satisfying $\left|\frac{dz}{du}|_{u=u_0}\right| \geq 1$, the optimal contract involves money-burning for all $\beta \in \left(\frac{\theta_l}{\theta_h}, \beta^*\right)$, i.e., whenever the optimal contract is separating.

Next we examine comparative statics for the region where money burning is part of the optimal contract.

Proposition 3 *Suppose that μ, θ_l, θ_h, β are such that (4) holds and $\beta \in \left(\frac{\theta_l}{\theta_h}, \beta^*\right)$ (so the optimal contract is separating but not the first best) and it involves no money burning. Then, for a fixed θ_l and β, a decrease in μ involves no money-burning either. For a fixed $\frac{\theta_l}{\beta}$ and μ, a higher β implies no money-burning.*

This means that within the parameter region which imply separation, but not the first-best, money-burning is part of the optimal contract when the high state is sufficiently rare. Intuitively, if μ is high enough, then committing to money burning in the state θ_h does not affect the expected utility of self-0 too negatively. Note that the second part of the result, since increasing θ_l for a fixed μ implies decreasing θ_h, can be reworded such that given $\frac{\theta_l}{\beta}$ constant, money burning is part of the optimal contract when the high liquidity shock is severe enough.

We finish the section by pointing out that clearly, with two types, all contracts that involve money-burning have a "withdrawal fee" interpretation. Indeed, they may be implemented as follows. The agent can withdraw up to c_l in period 1 free of charge. Withdrawal of any larger amount is possible, but requires paying a fee of $y - c_h$. In equilibrium then, type θ_l will withdraw c_l, and type θ_h will withdraw the full amount but consume only $c_h < y$.

4 Continuum of types

Let us restrict attention to the case where the support of θ is a compact segment $\Theta = \left[\underline{\theta}, \bar{\theta}\right]$, and that $f(\theta)$ is positive on Θ. Denote

$$G(\theta) = F(\theta) + \theta (1 - \beta) f(\theta), \tag{12}$$

and let θ_p be the lowest $\theta \in \Theta$ such that

$$\int_{\hat{\theta}}^{\bar{\theta}} (1 - G(\tilde{\theta})) d\tilde{\theta} \leq 0 \text{ for all } \hat{\theta} \geq \theta_p.$$

Since $F\left(\bar{\theta}\right) = 1$ and $f\left(\bar{\theta}\right) > 0$, we must have $\theta_p < \bar{\theta}$. The following proposition proves that there is "bunching at the top", i.e., all types $\theta > \theta_p$ get the same allocation.

Proposition 4 *An optimal allocation* $\{(u(\theta), w(\theta))\}_{\theta\in\Theta}$ *satisfies* $u(\theta) = u(\theta_p)$ *and* $w(\theta) = w(\theta_p)$ *for* $\theta \geq \theta_p$*. Both* $w(\theta) = z(u(\theta))$ *and* $w(\theta) < z(u(\theta))$ *are possible for* $\theta \geq \theta_p$*.*

This proposition corrects Proposition 2 in AWA. Like AWA, we claim that the types $\left[\theta_p, \bar{\theta}\right]$ are pooled. Unlike AWA, we do not claim that the budget constraint holds with equality for these types and there is no money-burning at the top. On the contrary, we show that it is possible that types $\left[\theta_p, \bar{\theta}\right]$ will have to burn money. The difference in the conclusions again arises because of the possibility that the optimal contract does not specify an interior consumption plan. In particular, in the proof of Proposition 2 AWA suggest that if θ_p is interior (i.e., $\theta_p \in \left(\underline{\theta}, \bar{\theta}\right)$), then $u(\theta_p)$ can be increased in a way that the IC constraint is preserved and the objective function does not decrease. However, preserving the IC constraint for type θ_p necessarily implies that $w(\theta_p)$ must be decreased, which is impossible if $w(\theta_p) = 0$. As in the case with two types, therefore, we only can have money-burning at the top if $w(\theta) = 0$ for high types.

Next we explicitly characterize cases in which there is money-burning imposed in the optimal contract on the highest types. We do this by starting from any distribution F satisfying the regularity conditions of AWA, and a condition that is satisfied when the degree of present bias is high enough. Then we consider distributions which can be obtained as F with probability $1-\varepsilon$, and F shifted to the right by z with probability ε (we actually need to shift F to the left by $\frac{\varepsilon z}{1-\varepsilon}$ in order to keep the mean one condition). The interpretation is that on top of the normal shocks, there is an ε likelihood of a catastrophic liquidity shock. We then show that whenever ε is small enough and z is large enough, the optimal contract involves money-burning in all states following the catastrophic shock.

Proposition 5 *Take distribution* F *with finite support* $\left[\underline{\theta}, \bar{\theta}\right]$ *and mean* 1*, and let* $u(\theta)$*,* $w(\theta)$ *be the optimal contract. Suppose that in period* 1*, type* $\bar{\theta}$ *would prefer to consume everything immediately rather than stay with the contract he pre-committed to:*

$$\bar{\theta}U(y) + \beta W(0) > \bar{\theta}u\left(\bar{\theta}\right) + \beta w\left(\bar{\theta}\right). \tag{13}$$

Consider now a family of distibutions $F_{z,\varepsilon}$*, given by:*

$$F_{z,\varepsilon}(x) = (1-\varepsilon) F\left(\theta + \frac{\varepsilon z}{1-\varepsilon}\right) + \varepsilon F(\theta - z).$$

Then there exist $\overline{z} > 0$ *such that for every* $z > \overline{z}$ *there is* $\overline{\varepsilon} > 0$ *such that for* $\varepsilon < \overline{\varepsilon}$*, the optimal contract implies* $c(\theta) < y$ *and* $k(\theta) = 0$ *for every* τ *in the support of* $F(\theta - z)$*.*

As AWA demonstrates, with more than two types in general the optimal contract can specify money-burning in a complicated pattern that not have a withdrawal fee interpretation.[12]

[12] A full characterization of the optimal contract in the continuous setting is a difficult problem and it is beyond the reach of the current paper. An alternative direction of future investigation is characterizing the optimal

However, the above result shows that even with a continuum of states, the possibility of a rare catastrophic shock can lead to the optimal contract resembling one withdrawal fees. Therefore the possibility that the optimal contract specifies money-burning and 0 second period consumption (hence has a withdrawal penalty interpretation) is not tied to the two-state case, instead to its qualitative essence of catastrophic shocks.

Lastly, we note that Proposition 6 in AWA, which generalizes Proposition 2 there, is also incorrect in claiming the absence money-burning, for the same reason as Proposition 2. However, under the additional assumption that guarantees that the optimal contract takes the form of minimal savings requirement, the result goes through and there is no money-burning (Proposition 7 in AWA). We omit the details here.

5 Appendix

Proof of Proposition 1. First note that the first-best allocation is implementable if $\beta \geq \beta^*$, where

$$\beta^* = \theta_l \frac{u_h^{fb} - u_l^{fb}}{w_l^{fb} - w_h^{fb}},$$

and moreover $\beta^* > \frac{\theta_l}{\theta_h}$. This is correctly proven in AWA.

From now on, consider the case $\beta < \beta^*$. Adding the incentive constraints (6) and (7) implies $\theta_h (u_h - u_l) \geq \theta_l (u_h - u_l)$, which implies $u_h \geq u_l$. Trivially, if (6) holds with equality, then (7) holds as well. Let us prove that (6) binds (so we could forget about (7)), and also $(u_l, w_l) \in \partial A$, $(u_h, w_h) \in \partial A$.

To see that $(u_l, w_l) \in \partial A$, assume the contrary. Indeed, if $(u_l, w_l) \notin \partial A$, then we can use the reasoning analogous to AWA: we can lower u_l and raise w_l slightly while holding $\theta_l u_l + \beta w_l$ unchanged, so that the modified contract is still in A; this would not change (6), will relax (7), and will increase the objective function (5), which contradict optimality of the initial contract.

To prove that $(u_h, w_h) \in \partial A$, suppose that $(u_h, w_h) \notin \partial A$, and consider the following three cases separately. If $\beta > \frac{\theta_l}{\theta_h}$, then a slight increase in u_h and a corresponding decrease in w_h that holds $\theta_l u_h + \beta w_h$ unchanged will not change (6), will relax (7), and will increase the objective function (5). If $\beta < \frac{\theta_l}{\theta_h}$, then a slight decrease in u_h and a corresponding increase in w_h will do the same. Finally, if $\beta = \frac{\theta_l}{\theta_h}$, then moving (u_h, w_h) to ∂A while preserving $\theta_l u_h + \beta w_h$ will not violate any constraint and will preserve the objective function, so without loss of generality we may assume that $(u_h, w_h) \in \partial A$ in the optimal contract in this case as well.

savings contract in the model framework of Fudenberg and Levine (2006), which is a major alternative to the quasi-hyperbolic approach used in AWA and the current paper. The specification of the above model with linear costs of self-control is particularly tractable, potentially facilitating a more complete analysis.

Let us now prove that (6) holds with equality in the optimal contract. Denote $\partial_f A = \{(u,w) \in \partial A : C(u) + K(w) = y\}$ (i.e., the frontier, where there is no money-burning), $\partial_c A = \{(u,w) \in \partial A : C(u) = 0\}$ (there is no consumption in period 1) and $\partial_k A = \{(u,w) \in \partial A : K(u) = 0\}$ (no consumption in period 2). The latter two may be empty if $U(0) = -\infty$ or $W(0) = -\infty$, respectively, but in any case $\partial A = \partial_f A \cup \partial_c A \cup \partial_k A$.

Suppose, to obtain a contradiction, that (6) is not binding; this already implies that the optimal contract is separating. We must have $(u_h, w_h) \in \partial_f A$, for otherwise we would be able to increase u_h slightly without violating either of the constraints and increasing the objective function). Second, we must have $(u_l, w_l) \in \partial_f A$. Indeed, suppose not, then either $(u_l, w_l) \in \partial_c A$ or $(u_l, w_l) \in \partial_k A$. Notice that (7) must bind, for if (7) did not bind, we could increase c_l to increase the objective function. Now, if $(u_l, w_l) \in \partial_k A$, then we must have $w_l \leq w_h$ (w_l is the lowest possible), we also have $u_l \leq u_h$ and if the contract is separating, one of the inequalities is strict, but then (7) cannot be binding. The remaining case is $(u_l, w_l) \in \partial_c A \setminus \partial_f A$. Since (7) binds, we must have $\left|\frac{dz}{du}|_{u=u_h}\right| > \frac{\theta_h}{\beta}$. But then slightly increasing w_l coupled with moving (u_h, w_h) along $\partial_f A$ so as to preserve (7) would unambiguously increase the objective function. This means that if (6) is not binding, then $(u_l, w_l) \in \partial_f A$, $(u_h, w_h) \in \partial_f A$, and also $u_l < u_h$ (otherwise the contract would be pooling, not separating). Again, suppose first that (7) binds; then $u_l < u_h$ means that (u_h, w_h) is the rightmost point of intersection of the line corresponding to (7) and $\partial_f A$, and so $\left|\frac{dz}{du}|_{u=u_h}\right| > \frac{\theta_h}{\beta}$; in this case, moving (u_h, w_h) slightly in the direction of $\left(u_h^{fb}, w_h^{fb}\right)$ would relax (7) and increase the objective function. The last possibility is that (7) does not bind. Then we could move either (u_h, w_h) slightly in the direction of $\left(u_h^{fb}, w_h^{fb}\right)$ or (u_l, w_l) slightly in the direction of $\left(u_l^{fb}, w_l^{fb}\right)$ so as to increase the objective function without violating any of the non-binding constraints. The only case where such deviation would not be possible is where $(u_h, w_h) = \left(u_h^{fb}, w_h^{fb}\right)$ and $(u_l, w_l) = \left(u_l^{fb}, w_l^{fb}\right)$. But this is not an incentive compatible contract if $\beta < \beta^*$ by the definition of β^*. This contradiction proves that (6) binds.

Consider the case $\frac{\theta_l}{\theta_h} < \beta < \beta^*$. Let us prove that the contract is separating. Indeed, if it were pooling, then, first of all, $(u_l, w_l) = (u_h, w_h) \in \partial_f A$. If this contract is λ_r^K (but not λ_l^K) for $K = u_l + \frac{\beta}{\theta_l} w_l$, then we can lower u_l and raise w_l slightly while holding $\theta_l u_l + \beta w_l$ unchanged; this would not change (6), will relax (7), and will increase the objective function (5). If this contract corresponds to λ_l^K (but not λ_r^K), then we can raise u_h and lower w_h slightly while holding $\theta_l u_h + \beta w_h$ unchanged with similar effects. The remaining case is where $\lambda_l^K = \lambda_r^K$; this means that $\left|\frac{dz}{du}|_{u=u_h}\right| = \frac{\theta_l}{\beta}$, and then moving (u_l, w_l) in the direction of $\left(u_l^{fb}, w_l^{fb}\right)$ and moving (u_h, w_h) in the direction of $\left(u_h^{fb}, w_h^{fb}\right)$ in a way that (6) continues to bind will relax (7) and will increase the objective function. Consequently, the optimal contract is separating. This implies $u_l < u_h$, and thus (7) does not bind. From this one can easily prove that $(u_l, w_l) \in \partial_f A$

(otherwise slightly increasing $\theta_l u_l + \beta w_l$ would create an incentive compatible contract which yields a higher ex-ante payoff) and, moreover, $\left|\frac{dz}{du}|_{u=u_l}\right| \in \left[\theta_l, \frac{\theta_l}{\beta}\right]$ (in particular, $u_l \in \left[u_l^{fb}, u_h^{fb}\right]$). Indeed, if $\left|\frac{dz}{du}|_{u=u_h}\right| < \theta_l$, then moving (u_l, w_l) in the direction of $\left(u_l^{fb}, w_l^{fb}\right)$ would increase the ex-ante payoff, and $\left|\frac{dz}{du}|_{u=u_l}\right| > \frac{\theta_l}{\beta}$ makes $(u_l, w_l) \in \partial_f A$ and (6) binding incompatible with $u_h > u_l$. As for (u_h, w_h), we can rule out $(u_h, w_h) \in \partial_c A$ (as then $u_l < u_h$ is impossible), but as we show, both $(u_h, w_h) \in \partial_f A$ and $(u_h, w_h) \in \partial_k A$ is possible.

Now consider the case $\beta < \frac{\theta_l}{\theta_h}$. Let us prove that the contract is pooling. If it were separating, then we can lower u_h and raise w_h slightly while holding $\theta_l u_h + \beta w_h$ unchanged (the fact that $(u_l, w_l) \in A$ ensures that such deviation results in a contract within A, but it also preserves (6), (7) and increases the ex-ante payoff (5). Hence, the contract is pooling. This means that $(u_l, w_l) = (u_h, w_h) \in \partial_f A$ and also $u_l \in \left[u_l^{fb}, u_h^{fb}\right]$, for otherwise moving the pooled contract along $\partial_f A$ in the direction of the first-best contract would increase the ex-ante payoff.

We thus showed that the contract is separating if $\frac{\theta_l}{\theta_h} < \beta < \beta^*$, pooling if $\beta < \frac{\theta_l}{\theta_h}$, and money-burning is possible only in the separating case and for type θ_h only. The possibility of money-burning for type θ_h is established by Example 1; the construction of an example without money-burning at optimum is trivial. This completes the proof. ■

Proof of Proposition 2. Take $\beta \in \left(\frac{\theta_l}{\theta_h}, \beta^*\right)$. From the proof of Proposition 1, if money-burning is part of the optimal contract, then $(u_h, w_h) \in \partial_k A \setminus \partial_f A$, so $k_h = 0$, $c_h < y$. This already implies $W(0) > -\infty$.

Also, by Proposition 1 we know that (6) is binding. Consequently, if (u_l, w_l, u_h, w_h) is the optimal contract, then $u_h + \frac{\beta}{\theta_l} w_h = u_l + \frac{\beta}{\theta_l} w_l$, which we denote by K. This means that $(u_l, w_l), (u_h, w_h) \in \lambda^K$. Moreover, from the proof of Proposition 1 we know that $(u_l, w_l) = \lambda_l^K$, $(u_h, w_h) = \lambda_h^K$. This proves that the optimal contract solves the following problem (formulated in terms of K, which remains the only degree of freedom):

$$\max_{K : \lambda^K \neq \varnothing} \left(\mu\left(\theta_l u_l^K + w_l^K\right) + (1-\mu)\left(\theta_h u_h^K + w_h^K\right)\right). \tag{14}$$

Indeed, the constraints (6) and (7) would then hold automatically: The IC constraint of type θ_l (6) would hold as equality because $\left(u_l^K, w_l^K\right)$ and $\left(u_h^K, w_h^K\right)$ lie on the same λ^K, and the IC constraint of type θ_h (7) would follow from the fact that (6) holds with equality and $u_h^K \geq u_l^K$. Moreover, again from the proof of Proposition 1, we have $(u_l, w_l) \in \partial_f A$, and also $u_l \geq u_l^{fb}$, so it is suffices to optimize over $K \geq u_l^{fb} + \frac{\beta}{\theta_l} w_l^{fb}$ only.

Let us first establish that (14) is strictly concave in K. Take two values of K, K_1 and K_2, and denote the value of the maximand in (14) by $v(K_1)$ and $v(K_2)$, respectively. Now take any $\delta \in (0,1)$. Given the linearity of the objective function (14) and the constraints (6) and (7), the contract given by $u_l' = \delta u_l^{K_1} + (1-\delta)\, u_l^{K_2}$, $u_h' = \delta u_h^{K_1} + (1-\delta)\, u_h^{K_2}$, $w_l' = \delta w_l^{K_1} + (1-\delta)\, w_l^{K_2}$, $w_h' = \delta w_h^{K_1} + (1-\delta)\, w_h^{K_2}$ satisfies the constraints and yields the value of (14) v' equal to $\delta v(K_1) + (1-\delta)\, v(K_2)$; moreover, it lies in A due to convexity of A. Since we proved that we can only improve by moving (u_l, w_l) to the upper-left and (u_l, w_l) to the lower-right, we get that $v(\delta K_1 + (1-\delta) K_2) > \delta v(K_1) + (1-\delta)\, v(K_2)$ (to see that the inequality is strict, notice that at least (u_l', w_l') necessarily lies in the interior of A. Hence we established that (14) is strictly concave in K.

We now see that money-burning is optimal if and only if (14) increases if we decrease K a little bit from the value $K_0 \geq U(y) + \frac{\beta}{\theta_l} W(0)$. If $u(0) = U(0)$, then doing so decreases the value of the objective function, because both the low type and the high type will get a smaller payoff. Now consider two cases. Suppose first that $K_0 > u_l^{fb} + \frac{\beta}{\theta_l} w_l^{fb}$; then the formula (10) is derived in the main text. If $K_0 \leq u_l^{fb} + \frac{\beta}{\theta_l} w_l^{fb}$, then $\left|\frac{dz}{du}|_{u=u_0}\right| \leq \theta_l$ as $u_0 \leq u_l^{fb}$. But then the right-hand side of (10) does not exceed $\mu\theta_l < 1$, so the formula is correct in this case as well. ■

Proof of Corollary 1. From Proposition 2, we have $W(0) > -\infty$, so $W(\cdot)$ is bounded away from $-\infty$. Now, we have a combination of (in)equalities:

$$\begin{aligned} \theta_l U(y) + \beta W(0) &> \theta_l u_h + \beta w_h; && (15) \\ \theta_l u_h + \beta w_h &= \theta_l u_l + \beta w_l; && (16) \\ \theta_l u_l + \beta w_l &\geq \theta_l u_l^{fb} + \beta w_l^{fb} && (17) \end{aligned}$$

Indeed, $u_h < U(0)$ and $w_h = W(0)$ imply (15); (16) follows because (6) is binding (as shown in the proof of Proposition 1), and (17) holds because $\left|\frac{dz}{du}|_{u=u_l}\right| \in \left[\theta_l, \frac{\theta_l}{\beta}\right]$ (again from the proof of Proposition 1), so $\left|\frac{dz}{du}|_{u=u_l}\right| \leq \frac{\theta_l}{\beta}$, and now $u_l^{fb} < u_l$ and $\left(u_l^{fb}, w_l^{fb}\right) \in \partial_f A$ imply the required condition. This implies $\theta_l U(y) + \beta W(0) > \theta_l u_l^{fb} + \beta w_l^{fb}$, which completes the proof. ■

Proof of Proposition 3. For fixed β and θ_l, u_0 is also fixed. Then the left-hand side of (10) is clearly increasing in μ, so if (10) did not hold for a given μ, then it would not hold for a lower μ. Hence, if the new optimal contract is separating, there will be no money-burning because (10) would not hold, and if the new optimal contract is not separating, then such optimal contract never involves money-burning. Finally, a decrease in θ_h, for a fixed θ_l, implies a lower μ, and we can use the previous reasoning. This completes the proof. ■

Proof of Proposition 4. The proof that $u(\theta) = u(\theta_p)$ for $\theta \geq \theta_p$ in AWA is correct, and is omitted here. Trivially, we must have $w(\theta) = w(\theta_p)$ for $\theta \geq \theta_p$ as well (otherwise, only the contracts with the highest w will be chosen). This proves the first part of the Proposition.

The proof that $w(\theta_p) < z(u(\theta_p))$ is possible follows from Proposition 5. ■

Proof of Proposition 5. The first step in the proof is to show that for z large, the types in the support of $F(\theta - z)$ are bunched. Indeed, we have

$$\begin{aligned}
\int_{\hat{\theta}}^{\bar{\theta}} \left(1 - G\left(\tilde{\theta}\right)\right) d\tilde{\theta} &= \int_{\hat{\theta}}^{\bar{\theta}} \left(1 - F\left(\tilde{\theta}\right) - \tilde{\theta}(1-\beta) f\left(\tilde{\theta}\right)\right) d\tilde{\theta} \\
&= \int_{\hat{\theta}}^{\bar{\theta}} \int_{\tilde{\theta}}^{\bar{\theta}} f(\theta') \, d\theta' d\tilde{\theta} - \int_{\hat{\theta}}^{\bar{\theta}} \tilde{\theta}(1-\beta) f\left(\tilde{\theta}\right) d\tilde{\theta} \\
&= \int_{\hat{\theta}}^{\bar{\theta}} \left(\tilde{\theta} - \hat{\theta}\right) f\left(\tilde{\theta}\right) d\tilde{\theta} - \int_{\hat{\theta}}^{\bar{\theta}} \tilde{\theta}(1-\beta) f\left(\tilde{\theta}\right) d\tilde{\theta} \\
&= \int_{\hat{\theta}}^{\bar{\theta}} \left(\beta\tilde{\theta} - \hat{\theta}\right) f\left(\tilde{\theta}\right) d\tilde{\theta} < \left(\beta\bar{\theta} - \hat{\theta}\right)\left(1 - F\left(\hat{\theta}\right)\right).
\end{aligned}$$

Now, if z is large enough, then $\beta(\theta_h + z) < \theta_l + z$, and therefore $\int_{\hat{\theta}}^{\theta_h + z} \left(1 - G\left(\tilde{\theta}\right)\right) d\tilde{\theta} \leq 0$ for all $\theta \geq \theta_l + z$. Consequently, $\theta_p \leq \theta_l + z$ and the "catastrophic" types are bunched.

The second step is to notice that we must have $W(0) > -\infty$ (otherwise type $\bar{\theta}$ would not prefer to consume everything in period 1, so (13) would be violated).

The third step is to notice that for z sufficiently large (at least $z > \bar{\theta} - \underline{\theta}$), as ε tends to 0, the optimal allocation for the types in the support of $F\left(\theta + \frac{\varepsilon z}{1-\varepsilon}\right)$ must converge (in distribution) to the optimal contract for the original distribution F. This is true because for any given z, the "catastrophic" types from the support of $F(\theta - z)$ have a vanishing impact as $\varepsilon \to 0$, and the optimal contract is unique.

Now assume, to obtain a contradiction, that the statement is false: i.e., there is a monotonically increasing sequence z_n tending to $+\infty$ for each of which there is a monotonically decreasing sequence $\varepsilon_{n,m}$ tending to 0 for which there types in the support of $F(\theta + z)$ are pooled, but either $c_{n,m} = y$ or $k_{n,m} > 0$ for these types. Let us show that none of these points may have $c_{n,m} + k_{n,m} < y$ and $k_{n,m} > 0$ simultaneously (i.e., if there is money-burning, then there must be zero consumption in period 1). Indeed, if these inequalities were true for $(c_{n,m}, k_{n,m})$ where $\varepsilon_{n,m}$ is small enough, then, as in the proof of Proposition 1, we could increase $c_{n,m}$ and decrease $k_{n,m}$ without violating the IC constraint for the type $\bar{\theta} - \frac{\varepsilon_{n,m} z_n}{1-\varepsilon_{n,m}}$ and below, and improving the payoffs for self-0 of every "catastrophic" type (this is true whenever $\left(\bar{\theta} - \frac{\varepsilon_{n,m} z_n}{1-\varepsilon_{n,m}}\right)/\beta < \underline{\theta} + z$, i.e., true for large z and small ε). Consequently, $(c_{n,m}, k_{n,m})$ must satisfy $c_{n,m} + k_{n,m} = y$, in other words, $(U(c_{n,m}), W(k_{n,m})) \in \partial_f A$.

Since the set of allocations is compact, we can, for every z_n, pick a limit point of $(c_{n,m}(\theta), k_{n,m}(\theta))$ and denote it $(c_n(\theta), k_n(\theta))$, and then pick a limit point of this sequence and denote it (c^*, k^*). Consider two possibilities. First, assume that $c^* \neq c\left(\bar{\theta}\right)$, where $c\left(\bar{\theta}\right)$ is the first-period consumption in the optimal contract for type $\bar{\theta}$. If so, monotonicity implies $c^* > c\left(\bar{\theta}\right)$. However, since the set A is convex and $(U(c^*), W(c^*)) \in \partial_f A$, type $\bar{\theta} - \frac{\varepsilon_{n,m} z_n}{1-\varepsilon_{n,m}}$ will prefer taking $(c_{n,m}, k_{n,m})$ for n and m large. This would violate his IC constraint.

The second possibility is that $c^* = c(\bar{\theta})$; this would be true, for example, if the "catastrophic" types were not given any options other than what "normal" types have. Notice that $c(\bar{\theta}) < y$ (otherwise, the condition (13) would not hold). Then for large z_n and small $\varepsilon_{n,m}$ there would be a profitable deviation as before: increase $c_{n,m}$ and decrease $k_{n,m}$ without violating the IC constraint for the type $\bar{\theta} - \frac{\varepsilon_{n,m} z_n}{1-\varepsilon_{n,m}}$ and below, and improving the payoffs for self-0 of every "catastrophic" type. Since such a deviation may not be possible in an optimal contract, we get to a contradiction, which completes the proof. ■

6 References

Amador, Manuel and Kyle Bagwell (2011): "The theory of optimal delegation with an application to tariff caps," mimeo Stanford University.

Amador, Manuel, Iván Werning and George-Marios Angeletos (2006): "Commitment vs Flexibility," Econometrica, 74, vol.2, 365-396.

Ambrus, A. and G. Egorov (2009): "Delegation and nonmonetary incentives," mimeo, Harvard University.

Amromin, G. (2002): "Portfolio allocation choices in taxable and tax-deferred accounts: An empirical analysis of tax efficiency," mimeo Federal Reserve Bank of Chicago.

Amromin, G. (2003): "Household portfolio choices in taxable and tax-deferred accounts: Another puzzle?," European Finance Review, 7, 547-582.

Amromin, G. and P. Smith (2003): "What explains early withdrawals from retirement accounts? Evidence from a panel of taxpayers," National Tax Journal, 56, 595-612.

Ashraf, N., D. Karlan and W. Yin (2006): "Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines:

Athey, S., A. Atkeson and P. Kehoe (2005): "The optimal degree of discretion in monetary policy," Econometrica, 73, 1431-1475.

Athey, S., K. Bagwell and C. Sanchirico (2004): "Collusion and price rigidity," Review of Economic Studies, 71, 317-349.

Beshears, J., J. Choi, D. Laibson, B. Madrian and J. Sakong (2011): "Self Control and Liquidity: How to Design a Commitment Contract," mimeo Harvard University.

Bryan, G., D. Karlan and S. Nelson (2010): "Commitment devices," Annual Review of Economics, 2, 671-698.

Dekel, E., B. Lipman and A. Rustichini (2001): "Representing preferences with a unique subjective state space," Econometrica 69, 891-934.

DellaVigna, S. and U. Malmendier (2004): “Contract design and self-control: Theory and evidence,” Quarterly Journal of Economics, 119, 353-402.

DellaVigna, S. and U. Malmendier (2006): “Paying not to go to the gym,” American Economic Review, 96, 694-719.

Diamond, P. (1977): “A framework for social security analysis,” Journal of Public Economics, 8, 275-298.

Eliaz, K. and R. Spiegler (2006): “Contracting with diversely naive agents,” Review of Economic Studies, 73, 689-714.

Esteban, S. and E. Miyagawa (2005): “Optimal menu of menus with self-control preferences,” mimeo Universitat Autonoma de Barcelona.

Fudenberg, D. and D. Levine (2006): “A dual-self model of impulse control,” American Economic Review, 96, 1449-1476.

Galperti, S. (2012): “Commitment, flexibility, and optimal screening of time inconsistency,” mimeo Northwestern University.

Gilkeson, J., J. List and C. Ruff (1999): “Evidence of early withdrawal in time deposit portfolios,” Journal of Financial Services Research, 15, 103-122.

Gilkeson, J. and C. Ruff (1996): “Valuing the withdrawal option in retail CD portfolios,” Journal of Financial Services Research, 10, 333-358.

Gul, F. and W. Pesendorfer (2001): “Temptation and self-control,” Econometrica, 69, 1403-1435.

Holden, S. and D. Schrass (2008): “The role of IRAs in U.S. households’ saving for retirement,” Investment Company Institute, 18, 1.

Laibson, D. (1997): “Golden eggs and hyperbolic discounting,” Quarterly Journal of Economics, 112, 443-477.

Laibson, D. (1998): “Life-cycle consumption and hyperbolic discount functions,” European Economic Review, 42, 861-871.

Phelps, E. and R. Pollack (1968): “On second best national savings and game-equilibrium growth,” Review of Economic Studies, 35, 185-199.

Strotz, R. (1956): “Myopia and inconsistency in dynamic utility maximization,” Review of Economic Studies, 23, 165-180.

Thaler, R. (1994): “Psychology and savings policies,” American Economic Review, 84, 186-192.