Commitment contracts*

Philip Bond† Gustav Sigurdsson‡

October 2015

Abstract

We analyze a consumption-saving problem in which time-inconsistent preferences generate demand for commitment, but uncertainty about future consumption needs generates demand for flexibility. We characterize in a standard contracting framework the circumstances under which this combination is possible, in the sense that a *commitment contract* exists that implements the desired state-contingent consumption plan, thus offering both commitment and flexibility. The key condition that we identify is a *preference reversal* condition: high desired consumption today should be associated with low marginal utility at future dates. We argue that there are conditions under which this preference reversal condition is naturally satisfied. The key insight of our paper is that time-inconsistent preferences not only generate commitment problems, but also allow their possible solution, since the preferences of later selves can be exploited to punish overconsumption by earlier selves.

*We thank Andres Almazan, Greg Fischer, David Laibson, Dilip Mookherjee, Andrew Postlewaite, Debraj Ray, and Andrzej Skrzypacz for their helpful comments. We also thank seminar and conference audiences at UCLA, El Colegio de Mexico, the London School of Economics, the University of Minnesota, the University of North Carolina, the University of Pennsylvania, the Stanford Institute for Theoretical Economics, the "Townsend" Conference, the Finance Theory Group, the Western Finance Association meeting, and the Econometric Society's Latin American Workshop in Economic Theory.

†University of Washington.

‡Citigroup Global Markets.

1 Introduction

Preferences with hyperbolic time discounting, introduced by Strotz (1956), are widely used to model individual behavior in a variety of settings.[1] In his original article, Strotz observed that hyperbolic discounting generates demand for *commitment*.[2] But in addition to commitment, individuals value the *flexibility* to respond to economic shocks. For example, an individual is likely to be uncertain about his future consumption needs. In such cases, an individual will be reluctant to commit to future consumption levels that are state-independent. In other words, there is a tension between commitment and flexibility (Amador et al 2006). In this paper, we analyze the extent to which commitment and flexibility can be successfully combined. When this is possible, hyperbolic discounting has no impact on equilibrium consumption.

In our setting, an individual would like to commit at date 0 to a consumption plan that may depend on unverifiable shocks that are realized in the future. To this end, the individual can enter into a *commitment contract* with the aim of implementing self 0's[3] desired consumption plan. The key contracting difficulty is that the shocks are realized only after the contract is signed, and since they are unverifiable, the contract cannot directly condition the individual's consumption on their realization. Rather, a commitment contract must provide the individual both with flexibility to respond to these shocks, and with incentives to adhere to self 0's desired consumption plan.

Our results characterize conditions under which the tension between commitment and flexibility can be resolved. Our key condition is a *preference reversal* condition, which states

[1]See Frederick, Loewenstein and O'Donoghue (2002) for a review of models of time discounting. Applications of hyperbolic discounting include consumer finance (e.g., Laibson 1996 on savings behavior in general; Laibson, Repetto, and Tobacman 1998 on retirement planning; DellaVigna and Malmendier 2004 and Shui and Ausubel 2004 on credit card usage; Skiba and Tobacman 2008 on payday lending; and Jackson 1986 on bankruptcy law), asset pricing (e.g., Luttmer and Mariotti 2003), and procrastination (e.g., O'Donoghue and Rabin 1999a, 1999b, 2001).

[2]See Ariely and Wertenbroch (2002) for direct evidence of demand for commitment.

[3]We follow the literature and refer to the individual at date t as *self* t.

that desired consumption at date 1 is negatively correlated with marginal utility (MU) at date 2. When this condition is satisfied, it is often possible to design a commitment contract in which an individual is deterred from overconsumption at date 1 by the prospect that future selves will engage in more costly forms of overconsumption at subsequent dates.

The key insight of our paper is that time-inconsistent preferences are not only the source of the individual's commitment problem, but also allow its possible solution. With time-inconsistent preferences, the individual's different selves have different preferences but still share knowledge of the shock realizations. This opens up the possibility of later selves punishing prior selves for deviating from self 0's desired consumption plan, which would be impossible if their preferences were the same. In essence, time-inconsistent preferences turn a single-agent contracting problem into a multi-agent mechanism design problem. As is well known from the implementation theory literature,[4] this can dramatically expand the set of outcomes that are attainable in equilibrium.

1.1 Illustrative examples

We illustrate our main results with three examples. In each example, there are three dates, and quasi-hyperbolic time preferences over these dates, with a hyperbolic discount factor of $\beta = \frac{1}{2}$ and no regular time discounting. The first two examples concern procrastination in performing a task, a case heavily studied in the literature (see footnote 1).

Example 1: An individual must complete a task by the end of date 3. The task requires 3 units of time. At each of dates 1, 2 and 3, the individual has a time endowment of 2 units. At date 1, the individual learns whether he must perform an essential chore (say, going to the doctor) at date 2. If it arises, the chore takes $\frac{1}{2}$ unit of time. The individual has log preferences over leisure at each date, but derives no utility from time spent on either the task or the chore.

[4]See Maskin and Sjöström (2002), Palfrey (2002), and Serrano (2004) for surveys of implementation theory.

In order to equalize MU of leisure across dates, the individual would like to work 1 unit of time on the task at each of the three dates if the chore is not required, but to work $\frac{7}{6}$, $\frac{2}{3}$ and $\frac{7}{6}$ units at each of the three dates if the chore is required.

However, at date 1, hyperbolic discounting implies that the individual wants to procrastinate and choose the work schedule $(1,1,1)$ over $\left(\frac{7}{6},\frac{2}{3},\frac{7}{6}\right)$, even if he knows he will have to do the date 2 chore.[5]

Suppose, however, that the individual agrees to interim deadlines to submit work corresponding to $\frac{7}{6}$ time units at date 1, and then $\frac{2}{3}$ at date 2, and then $\frac{7}{6}$ at date 3. Moreover, at date 0 he also agrees that the deadlines can be extended. First, at date 1 he can request a partial extension, which reduces the work he needs to submit to just 1 time unit, but at the same time changes the subsequent deadlines to 1 unit of work being required at each of dates 2 and 3. Second, if the date 2 deadline is for 1 unit of work (as it is after a date 1 extension), then at date 2 he can request a further partial extension, which reduces the required submission of work at date 2 to $\frac{2}{3}$ units, but in return he must submit a further $\frac{3}{2}$ units of work at date 3. Note that these deadlines allow for one costless extension, while the second extension is associated with an increase in total work demanded.

These deadlines allow the individual to commit to self 0's desired work schedule and overcome the procrastination problem. If self 1 requested an extension, then self 2 requests a further (work-increasing) extension if and only if the date 2 chore arose.[6] Given this, if self 1 learns the chore is necessary, he avoids procrastination and does not request an extension—because if he were to do so, he anticipates that self 2 will then also request an extension, increasing the individual's total workload.[7] If instead self 1 learns that he does

[5]Formally, $\log(2-1)+\beta\log\left(2-1-\frac{1}{2}\right)+\beta\log(2-1)>\log\left(2-\frac{7}{6}\right)+\beta\log\left(2-\frac{2}{3}-\frac{1}{2}\right)+\beta\log\left(2-\frac{7}{6}\right)$ since $\left(\frac{1}{2}\right)^{\frac{1}{2}}>\frac{5}{6}\frac{5}{6}$.

[6]Formally, $\log\left(2-\frac{2}{3}-\frac{1}{2}\right)+\beta\log\left(2-\frac{3}{2}\right)>\log\left(2-1-\frac{1}{2}\right)+\beta\log(2-1)$ since $\frac{5}{6}\left(\frac{1}{2}\right)^{\frac{1}{2}}>\frac{1}{2}$ (date 2 chore), and $\log(2-1)+\beta\log(2-1)>\log\left(2-\frac{2}{3}\right)+\beta\log\left(2-\frac{3}{2}\right)$ since $1>\frac{4}{3}\left(\frac{1}{2}\right)^{\frac{1}{2}}$ (no date 2 chore).

[7]Formally, $\log\left(2-\frac{7}{6}\right)+\beta\log\left(2-\frac{2}{3}-\frac{1}{2}\right)+\beta\log\left(2-\frac{7}{6}\right)>\log(2-1)+\beta\log\left(2-\frac{2}{3}-\frac{1}{2}\right)+\beta\log\left(2-\frac{3}{2}\right)$ since $\left(\frac{5}{6}\right)^2>\left(\frac{5}{6}\right)^{\frac{1}{2}}\left(\frac{1}{2}\right)^{\frac{1}{2}}$.

not need to do the date 2 chore, he requests an extension, allowing him to better smooth his effort over time, secure in the knowledge that self 2 will not request an extension.[8]

Example 2: In Example 1, the individual is able to commit to self 0's desired consumption plan. But suppose now that, if the chore arises, it takes $\frac{1}{4}$ units of the individual's time at both dates 1 and 2. Consequently, self 0 would like to commit to spend time $\left(\frac{11}{12}, \frac{11}{12}, \frac{7}{6}\right)$ on the task if the chore arises, but to $(1, 1, 1)$ if the chore does not arise. Similar to Example 1, hyperbolic discounting causes self 1 to prefer the work schedule $\left(\frac{11}{12}, \frac{11}{12}, \frac{7}{6}\right)$ regardless of whether the chore arises, because it offers high date-1 leisure.[9]

The key to attaining commitment Example 1 is that it was possible to offer self 2 an alternative work schedule after self 1 picks low date 1 work that self 2 chooses if and only if self 1 overconsumed leisure (i.e., there is no chore), and that also hurts self 1. But this is impossible in Example 2, as follows. First, the only work schedules that raise self 2's utility while lowering self 1's utility are those that reduce date 2 work and increase date 3 work. But second, if self 2 prefers such a path when there is no chore—which is when self 1 has overconsumed leisure in Example 2—then he also prefers such a path when there is a chore, since MU of date 2 leisure is higher in this case. Consequently, it is impossible to impose a state-contingent punishment on self 1 for working too little at date 1. This point is formalized in Lemma 2 below; this simple result is in many ways the key to our analysis.

Discussion: The key distinction between the two examples is that in Example 1 self 0's desired date 1 leisure is negatively correlated with MU at date 2, while in Example 2 these same quantities are positively correlated. We refer to the case of negative correlation as *preference reversal*. As Example 1 illustrates, and as we establish in our formal results, under preference reversal it is frequently possible for an individual to fully reconcile the

[8]Formally $\log(2-1) + \beta \log(2-1) + \beta \log(2-1) > \log\left(2-\frac{7}{6}\right) + \beta \log\left(2-\frac{2}{3}\right) + \beta \log\left(2-\frac{7}{6}\right)$ since $1 > \left(\frac{5}{6}\right)^{\frac{3}{2}} \left(\frac{4}{3}\right)^{\frac{1}{2}}$.

[9]Formally, $\log\left(2-\frac{11}{12}\right) + \beta \log\left(2-\frac{11}{12}\right) + \beta \log\left(2-\frac{7}{6}\right) > \log(2-1) + \beta \log(2-1) + \beta \log(2-1)$ since $\left(\frac{13}{12}\right)^{\frac{3}{2}} \left(\frac{5}{6}\right)^{\frac{1}{2}} > 1$.

conflict between commitment and flexibility, and to attain exactly the outcome desired by self 0. In contrast, this is impossible without the preference reversal condition.

When the preference reversal condition is satisfied, the key qualitative feature of contracts that allow self 0 to reconcile commitment with flexibility is that they increase self 2's discretion relative to the discretion that would be granted absent hyperbolic discounting. This is clear in Example 1, where, absent hyperbolic discounting, all decisions could be delegated to self 1, since no new information arrives at date 2. Moreover, and as our formal results establish, even when new information does arrive at date 2, it remains the case that the combination of preference reversal and hyperbolic discounting leads self 0 to increase the discretion granted to self 2.

In Example 1, preference reversal arises because self 1 learns about MU at date 2. Example 1 illustrates this in the context of a procrastination problem, but exactly the same forces operate in a consumption-savings problem in which MU is state-dependent, and an individual learns about MU realizations in advance. Related, Laibson (1996) analyzes a consumption-savings problem in which income and interest rates vary. If an individual learns about income shocks in advance, this setting is again isomorphic to Example 1. Moreover, variation in interest rates also induces correlation between self 0's desired date 1 consumption and MU at date 2, arguably in an even more direct way—the interest rate earned on investments made at date 1 both affects desired date 1 consumption and date 2 income, and hence MU at date 2. We illustrate this in our final example:

Example 3: An individual has wealth 1, and at date 1 has access to an investment maturing at date 2. At date 1 he learns whether the investment has a (gross) return of 1 or 1.1. At date 2, he has access to an investment with a certain return 1. Writing c_t for date t consumption, contemporaneous utility at date t is $-c_t^{-1}$, i.e., relative risk aversion of 2.

It is straightforward to verify that if the date 1 return is 1 then self 0 would like selves

1 and 2 to invest $\frac{2}{3}$ and $\frac{1}{3}$ respectively, delivering $c_t = \frac{1}{3}$ at all dates.[10] Similarly, if the date 1 return is instead 1.1 then self 0 would like self 1 to reduce investment to 0.656 and self 2 to increase investment to .361, delivering $c_1 = .344$ and $c_2 = c_3 = .361$, i.e., higher consumption at all dates.

As in Example 1, hyperbolic discounting means that if investment decisions are fully delegated to self 1, then self 1 prefers the sequence of investment decisions $(.656, .361)$ to $\left(\frac{2}{3}, \frac{1}{3}\right)$, even when the date 1 return is 1, because by reducing date 1 investment he raises date 1 consumption.[11]

However, and again as in Example 1, self 0 is able to achieve his desired investment levels if he also grants self 2 discretion over investment. Specifically, suppose that self 0 stipulates that, at date 1, self 1 must invest either $\frac{2}{3}$ or .656. If self 1 invests $\frac{2}{3}$, self 2 must invest $\frac{1}{3}$. But if self 1 reduces investment to .656, self 2 can choose between investing .361, or reducing his investment to .3. However, if self 2 chooses the lower investment of .3, he must pay an additional fee of .03.

These investment options deliver self 0's preferred choices, as follows. Self 2 chooses the lower investment level (and associated fee) if and only if self 1 "cheated" and chose the lower investment level when the date 1 return was 1.[12]Anticipating self 2's investment choices, self 1 chooses the higher investment level $\frac{2}{3}$ when the return is 1, because he understands that if he reduces investment to .656, self 2 will in turn reduce investment (and pay the associated fee), which is good for self 2 (because c_2 is higher) but bad for self 1 (because c_3 is lower).[13] On the other hand, self 1 reduces investment to .656 when the return is 1.1, because in this case self 2 picks the higher investment level and avoids the fee.

Note that a key feature of Example 3 is that desired date 1 investment is decreasing in

[10]If the date 1 gross return is R, self 0 desires investment $\frac{2}{2+\sqrt{R}}$ and $\frac{R}{2+\sqrt{R}}$ at dates 1 and 2 respectively.

[11]Formally, $-(1-.656)^{-1} - \beta(.656-.361)^{-1} - \beta(.361)^{-1} > -\left(1-\frac{2}{3}\right)^{-1} - \beta\left(\frac{2}{3}-\frac{1}{3}\right)^{-1} - \beta\left(\frac{1}{3}\right)^{-1}$.

[12]Formally, $-(.656-.3-.03)^{-1} - \beta(.3)^{-1} > -(.656-.361)^{-1} - \beta(.361)^{-1}$ and $-(.656\cdot 1.1-.361)^{-1} - \beta(.361)^{-1} > -(.656\cdot 1.1-.3-.03)^{-1} - \beta(.3)^{-1}$.

[13]Formally, $-\left(1-\frac{2}{3}\right)^{-1} - \beta\left(\frac{2}{3}-\frac{1}{3}\right)^{-1} - \beta\left(\frac{1}{3}\right)^{-1} > -(1-.656)^{-1} - \beta(.656-.3-.03)^{-1} - \beta(.3)^{-1}$.

the return. This is equivalent to the elasticity of intertemporal substitution (EIN) being less than 1, which is consistent with many empirical studies (see also Section 5).

2 Related literature

Central to our analysis is the idea that the commitment contract sets up a game between selves. O'Donoghue and Rabin (1999a) demonstrate that this inter-self game has some surprising properties; for example, "sophistication" may worsen self-control problems relative to "naïveté."[14] This previous paper focuses on a setting in which an individual must take an action exactly once, and takes the costs and rewards of this action as exogenously given. The basic commitment problem confronted by an individual in our paper is covered by their analysis: for instance, in Example 1, the individual can take an immediate reward of $\ln 1 - \ln \frac{5}{6}$ at date 1, with the cost of this reward deferred until the future. Our main results explore whether it is possible to design a contract (which determines costs and rewards) that deters the individual from taking the immediate reward at date 1. When such a contract exists, it gives the individual the possibility of taking two rewards. Although such a contract falls outside O'Donoghue and Rabin's framework, because the number of actions is not fixed, the basic flavor of our contract is related to their Example 4, in which self 1 is deterred from taking the immediate reward by the knowledge that, if he does so, self 2 will then also take an early reward.[15]

Our paper is closely related to Amador et al (2006). Like us, they study a hyperbolic individual who is hit by unverifiable taste shocks, but consider only a two-date version of the problem. This restriction immediately rules out the possibility of self 2 imposing a

[14] Following the literature, *sophistication* refers to the case in which self t correctly understands that selves $s > t$ have present-biased preferences. In contrast, *naïveté* refers to the case in which self t incorrectly believes that selves $s > t$ are not present-biased. See Section 7 for a discussion of partial naïveté.

[15] In addition, O'Donoghue and Rabin observe that present-biased preferences often violate independence of irrelevant alternatives (their Proposition 5), a point they refer to as a "smoking gun." This point—that actions never taken on the equilibrium path may nonetheless affect equilibrium decisions—is central to the design of contracts in our paper.

state-contingent punishment on self 1 for deviating—a key feature of our setting—because with two dates self 1 is effectively the only strategic agent.[16] Consequently, the only way to deter self 1 from deviating is to distort consumption in at least some states; the authors characterize the least costly way to do so.

Like Amador et al (2006), DellaVigna and Malmendier (2004) restrict attention to two dates, again ruling out the possibility of self 2 punishing self 1. Moreover, in their setting self 1 faces a binary choice (e.g., whether or not to go to the gym) and consequently a contract exists under which self 1 acts exactly as self 0 desires. The authors characterize the contract that maximizes the profits of a monopolist counterparty facing a partially naïve agent (Section 7 discusses partial naïveté). In particular, they characterize the combination of flat upfront fees and per-usage fees in the profit-maximizing contract.[17]

O'Donoghue and Rabin (1999b) analyze optimal contracts for procrastinators in a multi-period environment, where the socially efficient date at which a task should be performed is random. They explicitly rule out the use of contracts that induce an agent to reveal his type, which are the focus of our paper. As they observe, this restriction is without loss of generality in the main case they study, that of agents who are completely naïve about their future preferences. By contrast, we study sophisticated agents (again, see Section 7 for a discussion of partial naïveté).

While we examine the use of *external* commitment devices, such as contracts, other research considers what might be termed *internal* commitment devices. Bernheim, Ray, and Yeltekin (2013) and Krusell and Smith (2003) consider deterministic models in which an individual is infinitely lived, and show that Markov-perfect equilibria exist in which he gains some commitment ability from the fact that deviations will cause future selves to punish him. Carrillo and Mariotti (2000) and Benabou and Tirole (e.g., 2002, 2004)

[16] Amador et al (2003) extend the analysis to three or more dates. They assume that shocks are independent across dates; see subsection 4.3 below.

[17] Similarly, Eliaz and Spiegler (2006) analyze profit maximization by a monopolist who deals with a population of time-inconsistent individuals who differ in their degree of sophistication (see Section 7 below).

consider models in which an individual can commit his future selves to some action by manipulating their beliefs, respectively, through the extent of his own information acquisition, through direct distortion of beliefs, or through self-signalling.

3 Model

At each of dates $t = 1, 2, 3$, a single agent consumes c_t. At dates 1 and 2 his contemporaneous utility depends on state variables $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$, realized at dates 1 and 2 respectively, and is given by $u_1(c_1; \theta_1)$ and $u_2(c_2; \theta_2)$. Without loss, $\Pr(\theta_t) > 0$ for all $\theta_t \in \Theta_t$. We write $\Theta \equiv \Theta_1 \times \Theta_2$, and assume Θ is compact. At date 3, contemporaneous utility is $u_3(c_3)$. Note that $u_3(c_3)$ can be interpreted more generally as a value function covering multiple future dates, i.e., the expected discounted future utility of an agent inheriting wealth c_3.

The agent discounts the future quasi-hyperbolically: for $t = 0, 1, 2$, self t's intertemporal utility function is $U^t \equiv u_t + \beta \sum_{s=t+1}^{3} u_s$, so that $\beta \in (0, 1)$ is the hyperbolic discount factor. Note that the regular, i.e., non-hyperbolic, discount rate is normalized to zero; likewise, the risk-free interest rate is zero. Finally, the agent is self-aware (i.e., sophisticated), in the sense that at each date, he correctly anticipates his preferences at future dates (we relax this in Section 7). The contemporaneous utility functions u_t are strictly increasing and strictly concave in c_t. Finally, write $V^t = \sum_{s=t}^{3} u_s$ for utility under exponential discounting.

Write $C(\theta_1, \theta_2) = (C_1(\theta_1), C_2(\theta_1, \theta_2), C_3(\theta_1, \theta_2))$ for a contract, which consists of a sequence of date- and state-contingent consumption levels.

The total resources available for the agent to consume across the three dates is W, and is state-independent.[18] This could either represent an initial endowment of the agent, or the present value of future income. Since our focus is on the effect of hyperbolic discounting

[18] However, the additive shock parameterization of our environment that we introduce below is equivalent to allowing W to vary in an unverifiable way across states.

on intertemporal efficiency, not its effect on insurance across states, we rule out transfers across states and impose the following resource constraint: for all $(\theta_1, \theta_2) \in \Theta$,

$$C_1(\theta_1) + C_2(\theta_1, \theta_2) + C_3(\theta_1, \theta_2) \leq W. \tag{RC}$$

This assumption also facilitates comparison with the existing literature, which like us focuses on intertemporal efficiency.[19] Moreover, it would be hard—and sometimes impossible—to insure the agent if self 0 had private information about the distribution of states.[20] Note that RC covers even zero-probability state realizations (θ_1, θ_2), a point we discuss below.

3.1 Incentive constraints

The central friction in our framework is that the states θ_1 and θ_2 are unverifiable. Unverifiability means that a contract must satisfy the following incentive compatibility (IC) constraints, which ensure that the agent does not gain by misrepresenting the state.[21] At date 2, the IC constraints are: for all $\theta_1 \in \Theta_1$ and $\theta_2, \tilde{\theta}_2 \in \Theta_2$,

$$U^2(C(\theta_1, \theta_2); \theta_2) \geq U^2\left(C\left(\theta_1, \tilde{\theta}_2\right); \theta_2\right). \tag{IC$_2$}$$

At date 1, the IC constraints are: for all $\theta_1, \tilde{\theta}_1 \in \Theta_1$,

$$E_{\theta_2}\left[U^1(C(\theta_1, \theta_2); \theta_1, \theta_2) | \theta_1\right] \geq E_{\theta_2}\left[U^1\left(C\left(\tilde{\theta}_1, \theta_2\right); \theta_1, \theta_2\right) | \theta_1\right]. \tag{IC$_1$}$$

[19] Amador et al (2006) rule out transfers across states. O'Donoghue and Rabin (1999b) and DellaVigna and Malmendier (2004) study risk-neutral agents, and so insurance across states is not a concern.

[20] Note that private information about the distribution of θ_1 would not affect our analysis, which characterizes when intertemporal efficiency is possible.

[21] In principle, a contract could also condition on self 2's report of the date 1 state, say θ_{21}, so that the contract would take the form $C(\theta_1, \theta_2, \theta_{21})$. However, given that self 2's preferences are independent of state θ_1, the only way for a contract with $C(\theta_1, \theta_2, \theta_{21}) \neq C\left(\theta_1, \theta_2, \tilde{\theta}_{21}\right)$ to be incentive compatible is if self 2 is indifferent between $C(\theta_1, \theta_2, \theta_{21})$ and $C\left(\theta_1, \theta_2, \tilde{\theta}_{21}\right)$, and resolves the indifference differently depending on the true realization θ_1. We assume throughout that self 2 resolves indifference in the same way in all states, and accordingly, write the contract and ICs as in the main text. In a discussion of the same issue, Amador et al (2003) show that indifference is only possible in a finite number of states, so that if there are a continuum of states, as in Section 4.5, this assumption is without loss.

3.2 Benchmark: State θ_1 verifiable

Unverifiability of the state induces a potential trade-off between commitment and flexibility for the agent, as discussed in the introduction. Our main results characterize when self 0 can successfully combine commitment with flexibility with respect to date 1 shocks. That is, we characterize when the constraint IC_1 is non-binding in the maximization problem

$$\max_{C \text{ s.t. RC, IC}_1\text{, IC}_2} E_{\theta_1,\theta_2}\left[U^0\left(C\left(\theta_1,\theta_2\right);\theta_1,\theta_2\right)\right]. \quad (1)$$

Formally, IC_1 is non-binding in (1) if self 0's maximal attainable utility coincides with

$$\max_{C \text{ s.t. RC,IC}_2} E_{\theta_1,\theta_2}\left[U^0\left(C\left(\theta_1,\theta_2\right);\theta_1,\theta_2\right)\right], \quad (2)$$

i.e., θ_1 is verifiable. We focus on the reconciliation of commitment and flexibility with respect to date 1 shocks because the analogous trade-off at date 2 is a two-date problem of exactly the type analyzed by Amador et al (2006), and as their results show, in general there is no way to fully resolve the date 2 conflict.

For use throughout the paper, let $C^*(\cdot;\beta)$ be a solution to the relaxed problem (2). We omit the argument β whenever possible. Note that $C^*(\theta_1,\theta_2)$ is indeterminate if $\Pr(\theta_1,\theta_2)=0$. We resolve this indeterminacy by choosing C^* to minimize self 2's discretion, i.e., for any (θ_1,θ_2) such that $\Pr(\theta_1,\theta_2)=0$, let $\tilde{\theta}_2$ be such that $\Pr\left(\theta_1,\tilde{\theta}_2\right)>0$, and let $C^*(\theta_1,\theta_2)=C^*\left(\theta_1,\tilde{\theta}_2\right)$.

3.3 Preference assumptions and a preliminary result

Before formally stating assumptions on how the state (θ_1,θ_2) affects preferences, it is useful to give two leading examples:

Example, multiplicative shocks: $u_t(c_t;\theta_t)=\theta_t u(c_t)$ for $t=1,2$ and $\theta_t\in\Theta_t$.[22]

[22] Amador et al (2006) focus completely on this case.

Example, additive shocks: $u_t(c_t;\theta_t) = u(c_t - f_t(\theta_t))$ for $t = 1, 2$ for some functions f_t and u, where u has NIARA. This shock specification has a natural interpretation as either essential expenditure shocks (e.g., the agent must take his child to the doctor), or income shocks (e.g., the individual receives a bonus).

Motivated by these examples, we make Assumptions 1-3:

Assumption 1 *For any* $\theta_2 \neq \tilde{\theta}_2 \in \Theta_2$*, sign* $\left(u_2'(c_2;\theta_2) - u_2'\left(c_2;\tilde{\theta}_2\right)\right)$ *is independent of* c_2*, and* $u_2'(c_2;\theta_2) \neq u_2'\left(c_2;\tilde{\theta}_2\right)$.

Assumption 1 simply says that Θ_2 is unambiguously ordered by date 2 MU. Accordingly, we write $\tilde{\theta}_2 > \theta_2$ if and only if $u_2'\left(\cdot;\tilde{\theta}_2\right) > u_2'(\cdot;\theta_2)$, and write $\underline{\theta}_2$ and $\bar{\theta}_2$ for the minimal and maximal elements of Θ_2 under this ordering.

In addition, we impose the follow regularity condition, which is easily verified to be satisfied by both multiplicative and additive shocks. It is used only in subsection 4.4.

Assumption 2 *If* $\tilde{\theta}_2 > \theta_2$ *then* $\frac{u_2'(c_2;\theta_2)}{u_2'\left(c_2;\tilde{\theta}_2\right)}$ *is either constant, or strictly increasing in* c_2.

To guarantee interior solutions, we impose the standard Inada condition, modified to allow a state-contingent minimum consumption level (see, e.g., the case of additive shocks):

Assumption 3 *For* $t = 1, 2$ *and* $\theta_t \in \Theta_t$*, there exists* $\underline{c}_t$ *such that* $u_t'(c_t;\theta_t) \to \infty$ *as* $c_t \to \underline{c}_t$*, and moreover,* $u_3'(c_3) \to \infty$ *as* $c_3 \to 0$.

Finally, we note the following straightforward monotonicity result, which is standard to mechanism design problems.[23]

Lemma 1 *If* $\tilde{\theta}_2 > \theta_2$ *and* C *satisfies* IC_2 *then* $C_2\left(\theta_1,\tilde{\theta}_2\right) \geq C_2(\theta_1,\theta_2)$ *for all* $\theta_1 \in \Theta_1$.

[23] See Lemma 2 of Myerson (1981); or Chapter 2.3 of Bolton and Dewatripont (2005).

3.4 Applications

We have described the model in terms of consumption of a good. But the model can also be straightforwardly interpreted as consumption of leisure, enabling us to analyze incentives for procrastinators, as in Examples 1 and 2. In this interpretation, the agent must complete a task that requires a total of h hours of work.[24] His total time endowment across the three dates is $W + h$. The agent decides how much leisure c_t to enjoy at each of dates 1,2,3, subject to the constraint that he completes the task, $\sum_{t=1}^{3} c_t \leq W$.

A second alternative interpretation relates to Amador et al's (2006) analysis of a society that wishes to constrain government spending, while recognizing that in some circumstances higher government spending is socially desirable. Our model extends this setting to cover both federal and local government spending. In this interpretation, the federal government chooses spending c_1; taking federal spending as given, the local government chooses spending c_2; and the private sector consumes c_3. Both federal and local governments want to spend more than is socially optimal, corresponding to hyperbolic discounting.

4 Analysis

We initially assume that Θ_1 and Θ_2 are binary, and then relax this in subsection 4.5. Note that when Θ_2 is binary, $\Theta_2 = \left\{\underline{\theta}_2, \bar{\theta}_2\right\}$.

4.1 Perfect correlation

We start by considering the case of perfect correlation, so that the date 1 state θ_1 perfectly forecasts the date 2 state, which we denote by $\theta_2 = \phi(\theta_1)$. Under perfect correlation, IC_2 does not bind in the relaxed problem (2). Consequently, C^* is independent of β.

Under perfect correlation, the benchmark C^* fully delegates all consumption decisions

[24] In the additive shock parameterization, different states can interpreted as changes in the amount of time required to complete the task.

to self 1. Since self 0 and self 1's preferences coincide at β=1, C^* satisfies IC_1 at $\beta = 1$. Conversely, C^* violates IC_1 at $\beta = 0$, since here self 1 only values date 1 consumption, and so would report $\arg\max_{\tilde{\theta}_1} C^*\left(\tilde{\theta}_1\right)$ regardless of the true state.[25] Given these extremes, define β^* as the hyperbolic discount rate such that C^* just satisfies IC_1:[26]

$$\beta^* = \inf\left\{\beta : C^*(\cdot;\beta) \text{ satisfies IC}_1 \text{ for all } \tilde{\beta} \geq \beta\right\}. \tag{3}$$

Hence hyperbolic discount rates $\beta \geq \beta^*$ do not distort consumption, since self 0 can straightforwardly delegate all decisions to self 1, and attain consumption C^*.

We now turn to the more interesting case of $\beta < \beta^*$.[27] We characterize when self 0 can reconcile commitment and flexibility by granting self 2 some discretion over consumption.

Write $\bar{\theta}_1$ and $\underline{\theta}_1$ for the elements of Θ_1 such that

$$C_1^*\left(\bar{\theta}_1;\beta^*\right) \geq C_1^*\left(\underline{\theta}_1;\beta^*\right). \tag{4}$$

Our main result is Proposition 1, which establishes that the key condition for reconciling commitment and flexibility is that high desired date 1 consumption $C^*\left(\bar{\theta}_1\right)$ is followed by low date 2 MU. We refer to this condition as *preference reversal.*[28]

Proposition 1 *If $C_1^*\left(\bar{\theta}_1\right) > C_1^*\left(\underline{\theta}_1\right)$ then:*

(A, No Preference Reversal): If $\phi\left(\bar{\theta}_1\right) > \phi\left(\underline{\theta}_1\right)$, then for all $\beta < \beta^*$ constraint IC_1 binds in (1), i.e., commitment and flexibility cannot be combined.

[25] In the non-generic case in which $C_1^*(\cdot)$ is constant, IC_1 holds with equality at $\beta = 0$.

[26] We include the argument β in (3) and (4) only so that they remain applicable when we introduce imperfect correlation in subsection 4.4.

[27] It is straightforward to show that, under perfect correlation, C^* violates IC_1 for all $\beta < \beta^*$.

[28] The preference reversal condition may remind readers of Maskin's (1999) monotonicity condition. However, while preference reversal may fail in our setting, monotonicity is trivially satisfied as long as some self's preferences differ across the two states. In our setting, the social choice rule of interest is $F(\theta_1,\theta_2) = C(\theta_1,\theta_2)$. This social choice rule is monotonic if and only if for all (θ_1,θ_2) and $(\theta_1',\theta_2') \neq (\theta_1,\theta_2)$, $U^t\left(C(\theta_1,\theta_2);\theta_1,\theta_2\right) \geq U^t(x;\theta_1,\theta_2)$ and $U^t\left(C(\theta_1,\theta_2);\theta_1',\theta_2'\right) < U^t(x;\theta_1',\theta_2')$ for some self $t \in \{1,2,3\}$ (self 0 is non-strategic) and some $x \in \mathbb{R}^3$. As long as some self's preferences differ across the two states, this condition is satisfied

(B, Preference reversal): If $\phi\left(\bar{\theta}_1\right) < \phi\left(\underline{\theta}_1\right)$, then there exists $\hat{\beta} < \beta^*$ such that for all $\beta \geq \hat{\beta}$ constraint IC_1 is non-binding in (1), i.e., commitment and flexibility can be combined.

Moreover, we explicitly note the following immediate corollary:

Corollary 1 *If θ_1 and θ_2 are perfectly correlated, and IC_1 is non-binding in (1), then IC_2 is likewise non-binding in (1). Consequently, under the conditions of Part (B) of Proposition 1, self 0 attains the same utility as if both θ_1 and θ_2 were verifiable.*

For Part (A), IC_1 in state $\underline{\theta}_1$ for a contract that implements consumption C^* on the equilibrium path is

$$u_1\left(C_1^*\left(\underline{\theta}_1\right);\underline{\theta}_1\right) + \beta V^2\left(C^*\left(\underline{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right) \geq u_1\left(C_1^*\left(\bar{\theta}_1\right);\underline{\theta}_1\right) + \beta V^2\left(C\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right). \quad (5)$$

The key contract terms are $C_2\left(\bar{\theta}_1,\underline{\theta}_2\right)$ and $C_3\left(\bar{\theta}_1,\underline{\theta}_2\right)$, which determine consumption if self 1 falsely claims high consumption $C_1^*\left(\bar{\theta}_1\right)$ in the low date-1 consumption state $\underline{\theta}_1$, which in Part (A) is followed by $\underline{\theta}_2$. Because $\beta < \beta^*$, a necessary condition for satisfying (5) is that $V^2\left(C\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right) < V^2\left(C^*\left(\bar{\theta}_1,\bar{\theta}_2\right);\underline{\theta}_2\right)$. In words, C must punish self 1 for overconsuming at date 1 by delivering date-2 continuation utility strictly below $V^2\left(C^*\left(\bar{\theta}_1,\bar{\theta}_2\right);\underline{\theta}_2\right)$. However, this is impossible. From Lemma 1, IC_2 imply that $C_2\left(\bar{\theta}_1,\underline{\theta}_2\right) \leq C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)$. But because self 2 is hyperbolic, and hence especially values date 2 consumption, a necessary condition for him to pick the punishment $C\left(\bar{\theta}_1,\underline{\theta}_2\right)$ over $C\left(\bar{\theta}_1,\bar{\theta}_2\right) = C^*\left(\bar{\theta}_1,\bar{\theta}_2\right)$ is that $V^2\left(C\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right) \geq V^2\left(C^*\left(\bar{\theta}_1,\bar{\theta}_2\right);\underline{\theta}_2\right)$, which means that $C\left(\bar{\theta}_1,\underline{\theta}_2\right)$ is not an effective punishment. This establishes Part (A).

Summarizing, self 2 cannot be induced to impose an effective punishment on self 1 in state $\underline{\theta}_2$. This is the key observation behind our preference reversal condition, because when the low date-1 consumption state is followed by low date 2 MU-state $\underline{\theta}_2$, it is precisely in state $\underline{\theta}_2$ that a punishment is needed. Because of the centrality of this argument to our analysis, we state the following Lemma for use in subsequent results:

Lemma 2 *If C satisfies IC_2 then* $V^2\left(C\left(\theta_1,\underline{\theta}_2\right);\underline{\theta}_2\right)\geq V^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\underline{\theta}_2\right)$ *for all* $\tilde{\theta}_2\in\Theta_2$.

For Part (B), IC_1 in state $\underline{\theta}_1$ for a contract that implements consumption C^* on the equilibrium path is

$$u_1\left(C_1^*\left(\underline{\theta}_1\right);\underline{\theta}_1\right)+\beta V^2\left(C^*\left(\underline{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)\geq u_1\left(C_1^*\left(\bar{\theta}_1\right);\underline{\theta}_1\right)+\beta V^2\left(C\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right). \quad (6)$$

Here, the key contract terms are $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)$ and $C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)$, which determine consumption if self 1 falsely claims high $C_1^*\left(\bar{\theta}_1\right)$ in the low date-1 consumption state $\underline{\theta}_1$, which in Part (B) is followed by $\bar{\theta}_2$. In this case, it is possible to exploit the agent's time-inconsistency to punish self 1 for overconsumption. To do so, one can raise $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)$ relative to $C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$, while at the same time lowering $C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)$ relative $C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$; note that this is consistent with Lemma 1. Precisely because of hyperbolic discounting, self 2 is motivated to choose $C\left(\bar{\theta}_1,\bar{\theta}_2\right)$ over $C\left(\bar{\theta}_1,\underline{\theta}_2\right)$ in the high MU state $\bar{\theta}_2$. By making $C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)$ low, it is possible to make $V^2\left(C\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)$ low, and thus potentially satisfy (6).

This argument explains why self 1 can be punished, which is necessary in order to reconcile commitment and flexibility. However, one still needs to show that self 1 can be punished *enough*, i.e., $V^2\left(C\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)$ can be made low enough to satisfy (6). In particular, the requirement that RC hold even at the off-equilibrium state $\left(\bar{\theta}_1,\bar{\theta}_2\right)$ precludes punishing self 1 by setting $C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)$ very low while simultaneously raising $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)>W-C_1\left(\bar{\theta}_1,\bar{\theta}_2\right)$ in order to satisfy IC_2. Proposition 1 deals with this complication by establishing a result for β close to β^*. In brief, for all $\beta\leq\beta^*$ one can reduce $V^2\left(C\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)$ below $V^2\left(C^*\left(\bar{\theta}_1,\underline{\theta}_1\right);\bar{\theta}_2\right)$ by an amount that is bounded away from 0. But as β approaches β^*, the size of the punishment needed approaches 0, and so (6) can be satisfied.

To the extent to which empirical studies suggest that hyperbolic discounting is not too extreme,[29] the condition that β is not too low is the empirically relevant case. Moreover,

[29] For example, representative estimates include: Shui and Ausubel (2005), who estimate $\beta\approx 0.8$; Laibson et al (2007), who estimate $\beta\approx$0.7; and Augenblick et al (2015), who estimate $\beta\approx 0.9$.

for some circumstances, part (B) of Proposition 1 can be extended to cover all $\beta < \beta^*$:

Proposition 2 *Suppose the preference reversal condition* $\phi\left(\bar{\theta}_1\right) < \phi\left(\underline{\theta}_1\right)$ *holds. If (I)* $\lim_{c\to 0} u_3(c) = -\infty$*, (II)* $u_2\left(c;\bar{\theta}_2\right) \equiv u_3(c-\kappa)$ *for some* κ*, and (III)* $C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right) \geq C_3^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$*, then for all* β *the constraint* IC_1 *is non-binding in (1).*

In other words, under the conditions stated in Proposition 2, if the preference reversal condition holds then self 0 can attain both flexibility and commitment, independent of the degree of hyperbolic discounting.

Condition (I) of Proposition 2 simply states that it is possible to impose arbitrarily large utility punishments on self 1. It is satisfied by many utility functions, including, for example, CRRA utility functions with a coefficient of risk aversion above 1. Condition (II) holds if shocks are additive. It also holds, trivially, if $u_2\left(c;\bar{\theta}_2\right) \equiv u_3(c)$, which arises if the date 2 shock lowers MU relative to a baseline. Although condition (III) is more demanding, it is nonetheless satisfied in many cases, as discussed in subsection 4.2.

Contracts that reconcile commitment with flexibility when $\beta < \beta^*$ grant self 2 discretion that would be unnecessary absent hyperbolic discounting. In the context of consumption-savings problems, this excess flexibility may take the form of giving self 2 access to additional credit. In the context of procrastination problems, this excess flexibility means that if self 1 misses a deadline, self 2 also has the option to miss a deadline, but this increases the work required in the future. Related, under many circumstances, contracts that deliver equilibrium consumption C^* can be implemented using standard consumer financial products.[30] Similarly, in the procrastination setting, contracts that deliver C^* qualitatively resemble standard workplace rules or norms relating to deadlines (see Example 1).

[30] Formal results are contained in an earlier version of this paper.

4.2 Sources of preference reversal

The preference reversal condition identified in Proposition 1 arises naturally in multiple settings. Example 1 illustrates a leading case, namely *one-period ahead shocks*, i.e., at date 1, self 1 learns about a shock that that affects MU at date 2. This generates preference reversal because high date 2 MU raises C_2^* while reducing both C_1^* and C_3^*. In particular, the reduction in C_1^* means that the preference reversal condition is satisfied. By continuity, preference reversal also holds if date 1 MU is somewhat elevated in states in which self 1 learns date 2 MU will be high, provided the effect on date 1 MU is not too pronounced.

A second leading case in which preference reversal is satisfied is the case of *timing shocks*, in which a shock to MU occurs either at date 1 or date 2. For example, self 0 may know that he will encounter an attractive consumption opportunity at either date 1 or date 2, but not both (e.g., an out-of-town friend will visit at one of dates 1 and 2): that is, the incidence of attractive consumption opportunities is negatively correlated over time. Timing shocks generate preference reversal because if a shock raises MU at date 2, this raises C_2^* while leaving $C_1^* + C_2^*$ unchanged (because it is a timing shock), and hence reduces C_1^*. Note that condition (III) of Proposition 2 is satisfied by both one-period ahead shocks and timing shocks.

In the public finance interpretation discussed in subsection 4.1, preference reversal arises when date 0 uncertainty centers on the efficiency of local government spending. For this case, our analysis suggests that it is possible to design a constitution that controls government spending (at both federal and local levels), while still allowing flexibility to respond to shocks.

Finally, Section 5 shows that preference reversal arises naturally in investment problems, as illustrated by Example 3.

4.3 Zero correlation

Next, consider the case of zero correlation, in which the conditional probability $\Pr(\theta_2|\theta_1)$ of the date 2 state θ_2 is independent of the date 1 state θ_1. For example, this is the case either if date 2 utility u_2 is constant across states, or if θ_1 and θ_2 are uncorrelated: see Amador et al (2006, 2003). Given zero correlation, IC_1 simplifies to

$$u_1(C_1(\theta_1);\theta_1)+\beta E_{\theta_2}\left[V^2(C(\theta_1,\theta_2);\theta_2)\right] \geq u_1\left(C_1\left(\tilde{\theta}_1\right);\theta_1\right)+\beta E_{\theta_2}\left[V^2\left(C\left(\tilde{\theta}_1,\theta_2\right);\theta_2\right)\right]. \tag{IC$_1$}$$

Because of zero correlation, there is a strictly positive probability of all states $(\theta_1,\theta_2)\in\Theta$. So by the definition of β^*, if $\beta<\beta^*$ then constraint IC_1 binds in (1), and commitment and flexibility cannot be fully reconciled, as in Amador et al (2006).

4.4 Imperfect correlation

Under perfect correlation, θ_2 is perfectly forecastable at date 1, and so the only reason to grant self 2 discretion over consumption is so he can punish self 1 for overconsumption. Propositions 1 and 2 give conditions under which self 2 can be induced to punish self 1 so effectively that self 1 is deterred from overconsumption, and so self 0 can reconcile commitment and flexibility.

In contrast, under imperfect correlation, state θ_2 is unknown at date 1. In this case, even the benchmark C^* grants discretion to self 2. Here, we give conditions under which Proposition 1 is robust to imperfect correlation, and we characterize how IC_1 affects the amount of discretion self 0 grants to self 2.

As a preliminary step, recall that Part (B) of Proposition 1 is a local result, established for hyperbolic discount rates in the neighborhood of β^*. The proof of this result relies on the magnitude of self 1's temptation to overconsume under C^* being continuous as a function of β. Although this property is immediate for perfect correlation, it is nontrivial under imperfect correlation. The choice of date 1 consumption $C_1(\theta_1)$ sets the conditions for the subproblem of dividing the remaining resources $W - C_1(\theta_1)$ across dates 2 and 3 while satisfying IC_2. Consumption $C_1^*(\theta_1;\beta)$ depends on β because IC_2 and hence the date 2 subproblem depends on β. Establishing continuity of the magnitude of temptation requires establishing that date 1 consumption $C_1^*(\theta_1;\beta)$ is continuous in β. This in turn requires establishing that the indirect utility function associated with the solution to the subproblem just described is concave in $C_1(\theta_1)$ (see Lemma A-4 in appendix).

Lemma 3 *$C_1^*(\theta_1;\beta)$ and $E\left[V^2\left(C_1^*(\theta_1,\theta_2;\beta);\theta_1,\theta_2\right)|\theta_1\right]$ are continuous in β.*

Parallel to conditions (5) and (6), self 1 is dissuaded from claiming high consumption in the low consumption state $\underline{\theta}_1$ if and only if

$$\begin{aligned} & u_1\left(C_1\left(\underline{\theta}_1\right);\underline{\theta}_1\right)+\beta\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)V^2\left(C\left(\underline{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)+\beta\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)V^2\left(C\left(\underline{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right) \\ \geq\ & u_1\left(C_1\left(\bar{\theta}_1\right);\underline{\theta}_1\right)+\beta\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)V^2\left(C\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)+\beta\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)V^2\left(C\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right). \end{aligned}$$

To delineate the consequences of imperfect correlation, we examine two distinct cases.

First, we consider the case in which the date 2 state following the high-consumption date 1 state $\bar{\theta}_1$ is deterministic. In this case, absent IC_1 there is no need to give self 2 any discretion after self 1 chooses high consumption $C_1\left(\bar{\theta}_1\right)$. However, and as with the case of perfect correlation, allowing for such discretion potentially relaxes IC_1. Our main result for this case is that discretion is useful precisely when the preference reversal condition is satisfied, i.e., high consumption at date 1 is followed by low MU at date 2.

Proposition 3 *If $\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)\in\{0,1\}$ and $C_1^*\left(\bar{\theta}_1;\beta^*\right) > C_1^*\left(\underline{\theta}_1;\beta^*\right)$ then:*
(A, No Preference Reversal) If $\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)=1$, then for all $\beta<\beta^*$ such that $C^*(\cdot;\beta)$ violates IC_1,[31] constraint IC_1 binds in (1).

[31] The stipulation that $C^*(\cdot;\beta)$ violates IC_1 is required because, since C^* is a function of β, it is possible that C^* satisfies IC_1 for some subset of discount rate parameters below β^*.

(B, Preference Reversal) If $\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right) = 1$ *then there exists* $\hat{\beta} < \beta^*$ *such that for all* $\beta \geq \hat{\beta}$*, constraint* IC_1 *is non-binding in (1).*

Proposition 3 extends Proposition 1 beyond perfect correlation. Again, if high date 1 consumption is followed by high date 2 MU (no preference reversal), then if equilibrium consumption C^* cannot be implemented using the minimum level of self 2 discretion, it cannot be implemented by any level of discretion. In contrast, if high date 1 consumption is followed by low date 2 MU (preference reversal), then increasing self 2's discretion beyond the minimum amount is useful (at least provided hyperbolic discounting is not too severe).

The same economic forces drive Propositions 1 and 3. If the high consumption date 1 state $\bar{\theta}_1$ is followed by the high MU date 2 state $\bar{\theta}_2$, then $C\left(\bar{\theta}_1, \underline{\theta}_2\right)$ lies off the equilibrium path, and self 0 would like to punish self 1 for overconsumption in $\underline{\theta}_1$ by making continuation utility $V^2\left(C\left(\bar{\theta}_1, \underline{\theta}_2\right); \underline{\theta}_2\right)$ low. But by Lemma 2, this is impossible, leading to the conclusion of Part (A). If instead the high consumption date 1 state $\bar{\theta}_1$ is followed by the low MU date 2 state $\underline{\theta}_2$, then $C\left(\bar{\theta}_1, \bar{\theta}_2\right)$ lies off the equilibrium path, and self 0 would like to punish self 1 for overconsumption in $\underline{\theta}_1$ by making continuation utility $V^2\left(C\left(\bar{\theta}_1, \bar{\theta}_2\right); \bar{\theta}_2\right)$ low. This is potentially achievable, exactly as in the case of perfect correlation.

Second, we consider the case in which both $\left(\bar{\theta}_1, \underline{\theta}_2\right)$ and $\left(\bar{\theta}_1, \bar{\theta}_2\right)$ occur with strictly positive probability. Consequently, any adjustment to self 2's discretion after $C_1^*\left(\bar{\theta}_1; \beta\right)$ shifts consumption away from the solution to (1). It follows that if $C^*(\cdot; \beta)$ violates IC_1, then IC_1 certainly binds in (1), regardless of the preference reversal condition.

Given this, we turn instead to characterizing the contract that solves (1), which we denote by $C^s(\cdot; \beta)$. We show that $C^s(\cdot; \beta)$ shares two key features with the contracts used to reconcile commitment and flexibility in Propositions 1-3. Specifically, in these results, if preference reversal holds (state $\bar{\theta}_1$ always followed by $\underline{\theta}_2$, while state $\underline{\theta}_1$ at most sometimes followed by $\underline{\theta}_2$) then utility $V^2\left(C\left(\bar{\theta}_1, \bar{\theta}_2\right); \bar{\theta}_2\right)$ is set to a low level, in order to punish self 1 for overconsumption in state $\underline{\theta}_1$. Moreover, this low utility level is achieved via giving self 2 discretion to consume more than $C_2\left(\bar{\theta}_1, \underline{\theta}_2\right)$. Proposition 4 establishes that both properties still hold under the milder condition that $\underline{\theta}_2$ is simply more likely after state $\bar{\theta}_1$ than $\underline{\theta}_1$. In contrast, if instead $\underline{\theta}_2$ is less likely after state $\bar{\theta}_1$ than $\underline{\theta}_1$, then exactly the reverse is true: $V^2\left(C\left(\bar{\theta}_1, \bar{\theta}_2\right); \bar{\theta}_2\right)$ is set to a higher level than it would otherwise be, and self 2's discretion after higher date 1 consumption is reduced rather than increased.

To state these results formally, define $\hat{C}\left(\cdot,\cdot;\beta\right)$ as the consumption profile that maximizes self 0's utility in state θ_1, taking as given date 1 consumption $C_1^s\left(\cdot;\beta\right)$:

$$\hat{C}\left(\cdot,\cdot;\beta\right) = \arg \max_{C \text{ s.t. IC}_2,\text{ RC, } C_1(\cdot)=C_1^s(\cdot;\beta).} E\left[U^0\left(C\left(\theta_1, \theta_2\right)\right)\right].$$

The consumption difference $\hat{C}_2\left(\bar{\theta}_1, \bar{\theta}_2\right) - \hat{C}_2\left(\bar{\theta}_1, \underline{\theta}_2\right)$ is the benchmark level of discretion that self 0 would allocate to self 2 if he were not concerned about self 1's incentives, where date 1 consumption is fixed at the same level as the solution to (1), $C_1^s\left(\bar{\theta}_1;\beta\right)$. Our next result compares this benchmark level of discretion to the discretion that self 0 in fact grants to self 2 under $C^s\left(\cdot;\beta\right)$.

Proposition 4 *Suppose $C_1^*\left(\bar{\theta}_1;\beta^*\right) > C_1^*\left(\underline{\theta}_1;\beta^*\right)$ and, for some θ_2,*

$$V^2\left(C^*\left(\bar{\theta}_1, \theta_2;\beta^*\right);\theta_2\right) < \max_{c_2} V^2\left(c_2, W - C_1^*\left(\bar{\theta}_1;\beta^*\right) - c_2;\theta_2\right). \tag{7}$$

Then for all $\beta < \beta^*$ sufficiently close to β^*:

(A, No Preference Reversal) If $\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right) < \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)$ then $V^2\left(C^s\left(\bar{\theta}_1, \bar{\theta}_2;\beta\right);\bar{\theta}_2\right) > V^2\left(\hat{C}\left(\bar{\theta}_1, \bar{\theta}_2;\beta\right);\bar{\theta}_2\right)$ and $C_2^s\left(\bar{\theta}_1, \bar{\theta}_2;\beta\right) - C_2^s\left(\bar{\theta}_1, \underline{\theta}_2;\beta\right) \leq \hat{C}_2\left(\bar{\theta}_1, \bar{\theta}_2;\beta\right) - \hat{C}_2\left(\bar{\theta}_1, \underline{\theta}_2;\beta\right)$.

(B, Preference Reversal) If $\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right) > \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)$ then $V^2\left(C^s\left(\bar{\theta}_1, \bar{\theta}_2;\beta\right);\bar{\theta}_2\right) < V^2\left(\hat{C}\left(\bar{\theta}_1, \bar{\theta}_2;\beta\right);\bar{\theta}_2\right)$ and $C_2^s\left(\bar{\theta}_1, \bar{\theta}_2;\beta\right) - C_2^s\left(\bar{\theta}_1, \underline{\theta}_2;\beta\right) \geq \hat{C}_2\left(\bar{\theta}_1, \bar{\theta}_2;\beta\right) - \hat{C}_2\left(\bar{\theta}_1, \underline{\theta}_2;\beta\right)$.

Condition (7) says that given state $\bar{\theta}_1$ and date 1 consumption $C_1^*\left(\bar{\theta}_1\right)$, self 0's preferred division of $W - C_1^*\left(\bar{\theta}_1\right)$ across dates 2 and 3 violates IC_2. The condition is needed to ensure that the provision of incentives to self 1 entails a non-trivial trade-off between distorting date 1 consumption and distorting state-contingent date 2 consumption.[32]

The results on the continuation utility levels V^2 are intuitive. For $\beta < \beta^*$, self 0 must distort future consumption in order to prevent self 1 from overconsuming in state $\underline{\theta}_1$. Under the preference reversal condition, this is achieved by lowering $V^2\left(C^s\left(\bar{\theta}_1, \bar{\theta}_2; \beta\right); \bar{\theta}_2\right)$ while raising $V^2\left(C^s\left(\bar{\theta}_1, \underline{\theta}_2; \beta\right); \underline{\theta}_2\right)$, since doing so reduces self 1's expected continuation utility from falsely reporting state $\bar{\theta}_1$ in state $\underline{\theta}_1$, while doing as little damage as possible to continuation utility when self 1 truthfully reports $\bar{\theta}_1$.

The results on self 2's discretion follow from the results on continuation utility. The need to satisfy IC_2 leads to consumption $C\left(\bar{\theta}_1, \underline{\theta}_2\right)$ that allocates too many of the available resources $W - C_1^s\left(\bar{\theta}_1\right)$ to date 2, relative to the full information first-best (see part (B) of Lemma A-3 in the appendix). Because a higher continuation utility $V^2\left(C\left(\bar{\theta}_1, \underline{\theta}_2\right); \underline{\theta}_2\right)$ is closer to $\max_{c_2} V^2\left(c_2, W - C_1^s\left(\bar{\theta}_1\right) - c_2; \underline{\theta}_2\right)$, it is associated with less distortion of consumption across dates 2 and 3, corresponding to lower date 2 consumption. But lowering $C_2\left(\bar{\theta}_1, \underline{\theta}_2\right)$ in turn makes self 2 more tempted to falsely report $\bar{\theta}_2$, and hence $C\left(\bar{\theta}_1, \bar{\theta}_2\right)$ must be distorted more, which corresponds to raising $C_2\left(\bar{\theta}_1, \bar{\theta}_2\right)$ (see part (C) of Lemma A-3). Hence higher values of $V^2\left(C\left(\bar{\theta}_1, \underline{\theta}_2\right); \underline{\theta}_2\right)$ are associated with greater date 2 discretion, as measured by $C_2\left(\bar{\theta}_1, \bar{\theta}_2\right) - C_2\left(\bar{\theta}_1, \underline{\theta}_2\right)$.[33]

4.5 Continuum of states

We next allow both θ_1 and θ_2 to take a continuum of different values. We focus on the case of perfect correlation: as above, let $\phi\left(\theta_1\right)$ be the realization of θ_2 that deterministically follows θ_1.

[32] If (7) does not hold, then for β close to β^* the contract C^s distorts date 1 consumption, but does not distort consumption at dates 2 and 3 (conditional on date 1 consumption).

[33] Note that this discussion is all for the case in which β is moderate. For strong hyperbolic discounting (low β), self 0 removes all discretion from self 2, so that $C^s\left(\bar{\theta}_1, \theta_2; \beta\right)$is independent of state θ_2 (this extends Proposition 1 of Amador et al (2006) to the more general class of preferences covered here). In this case, the discretion results in Proposition 4 hold with equality.

Formally, let Θ_1 and Θ_2 be compact and convex subsets of $\Re$, with utility $u_t(c_t;\theta_t)$ jointly differentiable in c_t and θ_t for $t = 1, 2$. To avoid economically uninteresting mathematical complications, $C_1^*(\theta_1) \neq C_1^*\left(\tilde{\theta}_1\right)$ whenever $\theta_1 \neq \tilde{\theta}_1$, and accordingly we order Θ_1 so that $C_1^*(\theta_1)$ is strictly increasing in θ_1. Without loss, assume that Θ is such that $C^*(\theta_1, \phi(\theta_1))$ and ϕ are differentiable in θ_1. Moreover, it is straightforward to show that $\beta^* = 1$: because $C^*(\theta_1, \phi(\theta_1))$ is continuous in θ_1, for any $\beta < 1$ self 1 gains by slightly increasing date 1 consumption, since this generates a first-order gain for self 1, while introducing only second-order distortions in the intertemporal allocation of consumption.

Our main result is the following generalization of Proposition 1:

Proposition 5 *(A, No Preference Reversal) If ϕ is strictly increasing, IC_1 binds in (1) for $\beta < \beta^* = 1$.*

(B, Preference Reversal) If ϕ is strictly decreasing, then for any β, IC_1 is non-binding in (1) provided that $\max_{\theta_1,\tilde{\theta}_1\in\Theta_1}\left|C^*(\theta_1,\phi(\theta_1)) - C^*\left(\tilde{\theta}_1,\phi\left(\tilde{\theta}_1\right)\right)\right|$ is sufficiently small.

Proposition 5 extends the conclusion of prior results to the case of a continuous state space. Again, hyperbolic discounting affects outcomes if low date 1 consumption is associated with low date 2 MU (Part (A), no preference reversal), but not if it is associated with high date 2 MU (Part (B), preference reversal). The economic forces are the same as previously identified. In Part (A), it is impossible to punish self 1 for overconsuming, because the punishment needs to be inflicted in a date 2 state with low MU, which by Lemma 2 is impossible.

In contrast, in Part (B), such punishment is possible. The new complication relative to previous results is that now self 1 can overconsume to various degrees. In general, greater overconsumption necessitates a more severe punishment. In particular, a greater punishment is typically needed if self 1 lies and reports $\tilde{\theta}_1$ when the true state is θ_1 rather than the true state being $\theta_1' > \theta_1$. The challenge is to design the contract C so self 2 picks the punishment appropriate to the degree of overconsumption, i.e., picks different punishments in states θ_1' and θ_1.

The proof of Part (A) is almost immediate from the analysis of the binary case. The proof of Part (B) is constructive, and we sketch some of the elements here. A useful starting point is the following (standard) sufficient condition for IC_2:

Lemma 4 *$IC_2(\theta_1, \cdot, \cdot)$ is satisfied if*[34] *$C_2(\theta_1, \cdot)$ is increasing in θ_2 and for all $\theta_2 \in \Theta_2$,*

$$u_2'(C_2(\theta_1, \theta_2); \theta_2) \frac{\partial C_2(\theta_1, \theta_2)}{\partial \theta_2} = -\beta u_3'(C_3(\theta_1, \theta_2)) \frac{\partial C_3(\theta_1, \theta_2)}{\partial \theta_2}. \tag{8}$$

Next, consider any self 1 report $\tilde{\theta}_1$. Provided IC_2 is satisfied, then self 2 reports truthfully. Consequently, if self 2 reports $\tilde{\theta}_2 = \phi\left(\tilde{\theta}_1\right)$, he is confirming that self 1 reported truthfully. Accordingly, we set $C\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right) = C^*\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right)$. In contrast, if self 2 reports $\tilde{\theta}_2 > \phi\left(\tilde{\theta}_1\right)$, then since ϕ is decreasing he is reporting that self 1 reported too high a state, i.e., $\tilde{\theta}_1 > \phi^{-1}\left(\tilde{\theta}_2\right)$, meaning that self 1 overconsumed. To deter such overconsumption, we essentially[35] define $C\left(\tilde{\theta}_1, \theta_2\right)$ for $\theta_2 \geq \phi\left(\tilde{\theta}_1\right)$ to satisfy the differential equation (8) together with IC_1 at equality, i.e.,

$$U^1\left(C\left(\tilde{\theta}_1, \theta_2\right); \phi^{-1}(\theta_2), \theta_2\right) = U^1\left(C^*\left(\phi^{-1}(\theta_2), \theta_2\right); \phi^{-1}(\theta_2), \theta_2\right),$$

so that self 1 is indifferent between reporting $\tilde{\theta}_1$ and the true realization $\theta_1 = \phi^{-1}(\theta_2)$.

Finally, for part (B), the condition that $\max_{\theta_1, \tilde{\theta}_1 \in \Theta_1} \left| C^*(\theta_1, \phi(\theta_1)) - C^*\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right) \right|$ is not too large ensures that RC can be satisfied. In common with the role played by the assumption that β is close to β^* in Proposition 1, this condition is not needed if RC is allowed to be violated off the equilibrium path. Alternatively, the following result gives

[34] As an aside, it is worth noting that condition (8) is a *necessary* as well as a sufficient condition for IC_2 at any point at which $C(\theta_1, \theta_2)$ is differentiable with respect to θ_2; a proof is available upon request.

[35] Full details are in the proof of Lemma 5. The complication relative to the main text is that, by Lemma 1, $C_2\left(\tilde{\theta}_1, \theta_2\right)$ must be weakly increasing in θ_2.

conditions that guarantee the existence of a contract implementing C^* and satisfying RC.

Proposition 6 *Suppose the preference reversal condition $\phi' < 0$ is satisfied. If (I)*

$$\left(u_1\left(c_1;\theta_1\right), u_2\left(c_2;\theta_2\right), u_3\left(c_3\right)\right) = \left(u\left(c_1+\zeta\left(\theta_1\right)\right), u\left(c_2-\theta_2\right), \frac{1}{n}u\left(\frac{c_3}{n}\right)\right)$$

for some $n \geq 1$, function ζ, and CRRA utility function u, (II) $\frac{\partial\zeta\left(\phi^{-1}(\theta_2)\right)}{\partial\theta_2} \in [0, \frac{1}{4})$, and (III) β sufficiently close to $\beta^* = 1$, then IC_1 is non-binding in (1).

Condition (I) says that shocks are additive. The restriction on u_3 nests the value function interpretation noted earlier, i.e., if self 2 bequeaths c_3 to the future, then absent anticipation of further shocks, self 2 allocates consumption of c_3/n to each of n future dates. Condition (II) says that higher date 2 MU is associated not just with lower date 1 consumption (i.e., preference reversal), but also with weakly lower date 1 MU. In particular, (II) nests the case of one-period ahead additive shocks, i.e., ζ constant. The proof of Proposition 6 contains the explicit lower bound on β used in Condition (III): for one-period ahead shocks, and a relative risk aversion of γ, the lower bound is simply $n^{-\gamma}\left(n+1\right)^{-\gamma}$, which is less than $\frac{1}{2}$ if $\gamma \geq 1$.

5 Investment problems

We next extend our analysis to cover the following class of *investment problems*, illustrated by Example 3. An agent has wealth W at date 1, out of which he chooses consumption c_1. He invests the remainder and obtains $\left(W-c_1\right)R\left(\theta_2\right)$ at date 2. Neither $R\left(\theta_2\right)$ nor $\left(W-c_1\right)R\left(\theta_2\right)$ is observable by outsiders.

Self 0 can write contracts that specify the amount to invest at date 1, i.e., $W-c_1$, and also an amount to deposit into a savings account at date 2. We denote the date 2 deposit as c_3+d, which yields c_3 at date 3, and assume that $d \geq 0$, i.e., the recipient of the date

2 deposit does not lose money. So in this context, a contract specifies (c_1, d, c_3). The flow utilities at the three dates are, respectively, $u_1(c_1)$, $u_2((W - c_1) R(\theta_2) - d - c_3)$, $u_3(c_3)$.

To map this environment into our basic model, let $\bar{\theta}_2 = \arg\min_{\theta_2} R(\theta_2)$, and define

$$c_2 = (W - c_1) R\left(\bar{\theta}_2\right) - d - c_3. \tag{9}$$

Hence a contract can be written as specifying (c_1, c_2, c_3), with flow utilities $u_1(c_1)$,

$$u_2\left((W - c_1)\left(R(\theta_2) - R\left(\bar{\theta}_2\right)\right) + c_2\right), \tag{10}$$

and $u_3(c_3)$. The constraint $d \geq 0$ and equation (9) imply $R\left(\bar{\theta}_2\right) c_1 + c_2 + c_3 \leq R\left(\bar{\theta}_2\right) W$, which coincides with constraint RC in our basic model. Hence the only difference relative to our basic model is that date 2 MU depends on c_1 in addition to c_2 and θ_2. However, it is straightforward to verify that this has no impact on any of the analysis.

For conciseness, we focus here on the case in which the return $R(\theta_2)$ is known to self 1 when he decides how much to invest at date 1. This corresponds to the case in the benchmark model in which states θ_1 and θ_2 are perfectly correlated.

A higher return $R(\theta_2)$ increases the resources available to self 2, and so is associated with lower date 2 MU (see (10)). A higher return $R(\theta_2)$ also affects self 0's desired date 1 consumption, both positively, via the income effect, and negatively, via the substitution effect (i.e., the cost of date 1 consumption is higher). Hence, if the income effect dominates, the preference reversal condition holds, since the return $R(\theta_2)$ pushes desired date 1 consumption and date 2 MU in opposite directions. From Proposition 1, for hyperbolic discount rates not too far below β^* the IC_1 constraint is non-binding in (1), i.e., self 0 can reconcile commitment and flexibility. In contrast, these same results imply that if the substitution effect dominates then IC_1 binds in (1), i.e., commitment and flexibility cannot be fully reconciled.

As is well-known, if the EIN is below 1 then the income effect dominates the substitution effect in determining date 1 consumption's response to the return $R(\theta_2)$. Although there is a range of empirical estimates for the EIN, many estimates put the EIN substantially below 1 (see in particular Hall 1988).

6 Private savings

Thus far, we have assumed that the agent has no ability to save outside the contract. This assumption fits some applications well. For example, this is the case in procrastination problems where an agent's work is observable. This assumption also approximates the case in which private saving is possible, but only at a very disadvantageous interest rate.

In order to evaluate the consequences of relaxing this assumption, we have fully analyzed our environment for the case in which private savings are possible,[36] and the state space is binary with perfect correlation (i.e., the subsection 4.1 case). For conciseness, we summarize the results here: full details are contained in an earlier draft of the paper. Moreover, we focus here on the case of additive shocks (results for more general classes of shocks are likewise contained in an earlier draft).

The possibility of private saving does not affect self 0's most-preferred consumption, which we continue to denote by C^*; but it places additional constraints on self 0's ability to actually attain this consumption. Because private savings impose additional constraints, it is immediate that a necessary condition for consumption C^* to be implemented in equilibrium is that the preference reversal condition holds, i.e., $\bar{\theta}_1$ is followed by $\underline{\theta}_2$ and $\underline{\theta}_1$ and is followed by $\bar{\theta}_2$, and for the remainder of this section we assume this condition holds.

If self 1 falsely reports state $\bar{\theta}_1$ in $\underline{\theta}_1$, the possibility of private savings allows him to pass some of the extra consumption $C^*\left(\bar{\theta}_1\right) - C_1^*\left(\underline{\theta}_1\right)$ on to self 2. By doing so, self 1

[36]Recent papers in the growing literature on contracting with hidden savings include Kocherlakota (2004), Doepke and Townsend (2006), and He (2009).

reduces MU at date 2, and hence self 2's incentives to punish him for overconsumption. Moreover, the more self 1 privately saves, the lower are self 2's incentives to punish.

We denote by s_1^* the maximum amount that self 1 contemplates privately saving if he falsely reports state $\bar{\theta}_1$ in $\underline{\theta}_1$. Formally, we first define self 2's savings choice

$$\hat{s}_2(s_1;\beta) \equiv \arg\max_{s_2 \geq 0} u_2\left(s_1 + C_2^*\left(\bar{\theta}_1, \underline{\theta}_2\right) - s_2; \bar{\theta}_2\right) + \beta u_3\left(C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right) + s_2\right).$$

Note that $\beta \leq 1$ and $\bar{\theta}_2 > \underline{\theta}_2$ imply that $\hat{s}_2(s_1;\beta) = 0$ for all $s_1 \geq 0$ sufficiently small. We then define

$$\begin{aligned} s_1^*(\beta) = \sup\{s_1 \geq 0 \quad &: \quad U^1\left(C^*\left(\underline{\theta}_1, \bar{\theta}_2\right); \underline{\theta}_1, \bar{\theta}_2\right) < u_1\left(C_1^*\left(\bar{\theta}_1\right) - s_1; \underline{\theta}_1\right) \\ &+ \quad \beta u_2\left(s_1 + C_2^*\left(\bar{\theta}_1, \underline{\theta}_2\right) - \hat{s}_2(s_1;\beta); \bar{\theta}_2\right) + \beta u_3\left(C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right) + \hat{s}_2(s_1;\beta)\right)\}, \end{aligned}$$

where $s_1^*(\beta) = 0$ if the above set is empty. Since $\hat{s}_2(s_1, \beta) = 0$ in the neighborhood of $s_1 = 0$, it follows that $s_1^*(\beta) > 0$ whenever $\beta < \beta^*$, where β^* is defined exactly as before. Define $\beta^{*p} \geq \beta^*$ as the supremum value of β such that $s_1^*(\beta) > 0$. In words, for $\beta \geq \beta^{*p}$ consumption C^* can be implemented in equilibrium simply by allowing self 1 to choose between consumption profiles $C^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$ and $C^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$, while this is not generally the case for stronger hyperbolic discounting $\beta < \beta^{*p}$.

Our first result is that a necessary condition to implement C^* when $\beta < \beta^{*p}$ and $s_1^*(\beta) > 0$ is that

$$u_2'(x; \underline{\theta}_2) \leq u_2'\left(s_1^*(\beta) + x; \bar{\theta}_2\right). \qquad \text{(SPR)}$$

(Under additive shocks, condition SPR is independent of x.) Condition SPR is a strong preference reversal in the sense that it is a stricter requirement than preference reversal, which is simply $u_2'(\cdot; \underline{\theta}_2) < u_2'\left(\cdot; \bar{\theta}_2\right)$. The economics behind condition SPR is that self 2 must be induced to punish self 1 for choosing $C_1^*\left(\bar{\theta}_1\right)$ in state $\underline{\theta}_1$, even if self 1 privately saves up to s_1^*; but at the same time, self 2 needs to be deterred from selecting the punishment when self 1 chooses $C_1^*\left(\bar{\theta}_1\right)$ in state $\bar{\theta}_1$. Recalling that $\underline{\theta}_1$ is followed by $\bar{\theta}_2$, this combination is only possible if date 2 MU is higher in $\bar{\theta}_2$ than $\underline{\theta}_2$, even if MU in state $\bar{\theta}_2$ is lowered by self 1 privately saving s_1^*. This is exactly condition SPR.

Condition SPR is still just a necessary condition for consumption C^* to be implemented on the equilibrium path. However, our second result establishes the following analogue of part (B) of Proposition 1, namely that if condition SPR holds strictly, consumption C^* can be implemented on the equilibrium path for all hyperbolic discount rates sufficiently close to β^{*p}.[37] Hence in this case, commitment and flexibility can again be fully reconciled, even when the agent is able to save privately.

Finally, we show that condition SPR is equivalent to condition (III) in Proposition 2, which as discussed in 4.2 is satisfied by important classes of shocks.

7 Partial naïveté about future preferences

To analyze the impact of (partial) naïveté, we follow O'Donoghue and Rabin (2001) and let $\tilde{\beta} \geq \beta$ be self t's belief about the value of β that enters the preferences of future selves, i.e., self t's preferences are given by $U^t \equiv u_t + \beta \sum_{s=t+1}^{3} u_s$ while he believes the preferences of any future self $t' > t$ are given by $\tilde{U}^{t'} \equiv u_{t'} + \tilde{\beta} \sum_{s=t'+1}^{3} u_s$. Thus far, we have assumed that the agent is fully self-aware (sophisticated), i.e., $\tilde{\beta} = \beta$.

For discount rates $\beta < \beta^*$, under preference reversal and sophistication self 0 can reconcile commitment with flexibility by writing a contract that induces self 2 to punish self 1 for overconsumption, e.g., $C\left(\bar{\theta}_1, \bar{\theta}_2\right)$ in (6). Self 2 imposes the punishment because by doing so he increases date 2 consumption, which because of hyperbolic discounting he values heavily. The chief problem introduced by naïveté $\tilde{\beta} > \beta$ is that it may lead self 1 to believe that he can overconsume at date 1 without self 2 punishing him, because self 1 underestimates self 2's preference for date 2 consumption.

[37] For more general β, we derive an explicit sufficient condition.

In Table 1, we present numerical simulations that shed light on the impact of naïveté. In all simulations we assume CRRA preferences with risk aversion parameter γ. Panel (A) considers one-period ahead additive shocks, specifically, $u_1 \equiv u_3 \equiv u_2(\cdot;\underline{\theta}_2) = u$ and $u_2(c_2;\bar{\theta}_2) = u(c_2 - \bar{\theta}_2)$ for a CRRA function u, and θ_1 and θ_2 perfectly correlated. As discussed, the preference reversal condition holds here. For each combination of risk aversion γ and shock size $\bar{\theta}_2$, in row (i) we report β^*, the level of hyperbolic discounting at which self 0's desired consumption cannot be achieved by straightforward delegation to self 1. Absent naïveté, from Proposition 2 we know that commitment and flexibility can be fully reconciled for any $\beta < \beta^*$. In rows (ii) and (iii) we then fix the degree of hyperbolic discounting at the empirically reasonable values of $\beta = 0.8$ and $\beta = 0.7$,[38] and calculate the maximum degree of naïveté for which commitment and flexibility can still be reconciled, i.e., a contract that implements C^* exists. For example, for $\gamma = 1$, $\bar{\theta}_2 = 0.05$ and $\beta = 0.8$, such a contract exists provided that self 1's belief $\tilde{\beta}$ about self 2's hyperbolic discount rate is 0.84 or less.

Panels (B) and (C) present the results for parallel analysis of timing shocks (B) and investment problems (C). For timing shocks, $\Pr(\underline{\theta}_2|\bar{\theta}_1) = \Pr(\bar{\theta}_2|\underline{\theta}_1) = 1$, $u_1(c;\bar{\theta}_1) = u_2(c;\bar{\theta}_2) = u(c - \bar{\theta}_2)$ where $\bar{\theta}_2$ is as reported, and $u_1(\cdot;\underline{\theta}_1) \equiv u_2(\cdot;\underline{\theta}_2) \equiv u_3 \equiv u$. For investment problems, $R(\bar{\theta}_2)$=1, and the table reports $R(\underline{\theta}_2)$. Note that for investment problems the case of risk-aversion $\gamma = 1$ is omitted, since as noted above, preference reversal arises in investment problems only for $\gamma > 1$.

From Table 1, one can see, first, that for many parameterizations there is an empirically reasonable range of hyperbolic discount rates under which a fully sophisticated ($\tilde{\beta} = \beta$) agent can reconcile commitment with flexibility by granting discretion to self 2. Moreover, this conclusion continues to hold for at least moderate levels of naïveté: except for cases in which the shock is small and risk aversion is low, there is a reasonably large range of

[38] See footnote 29.

naïveté levels for which there exists a contract implementing equilibrium consumption C^*.

While the combination of commitment and flexibility is often possible when the agent is partially naïve, naïveté does significantly impact an agent's incentive to choose an appropriate contract. Under full sophistication, self 0 has every incentive to sign up for a contract that reconciles commitment and flexibility. In contrast, this is not the case when the agent is partially naïve.

There are two related issues here. First, as in Heidhues and Kőszegi (2010), an agent's naïveté means that self 0 may agree to a contract that increases date 1 consumption relative to C^*, but that distorts consumption at dates 2 and 3. In brief, self 0 finds the contract attractive because he incorrectly believes that he can increase both date 1 and total consumption by borrowing at a below-market rate; while the counterparty is happy to agree to the contract because he correctly understands that self 2 will choose repayment terms that correspond to the market rate.

Second, under partial naïveté self 0 may incorrectly believe that he does not have a commitment problem. In these circumstances, there is scope for a benevolent government to improve welfare (at least for self 0) by imposing a commitment contract.

However, it also important to note that while a government-mandated commitment contract can improve the welfare of a partially naïve agent, it can actually hurt a very naïve agent, relative to the alternative of simply allowing self 1 to choose freely between self 0's desired consumption paths $C^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$ and $C^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$. First, note that the punishment component of the contract must satisfy $V^2\left(C\left(\bar{\theta}_1, \bar{\theta}_2\right); \bar{\theta}_2\right) < V^2(C^*\left(\bar{\theta}_1, \underline{\theta}_2\right); \bar{\theta}_2)$, since otherwise the punishment would not deter self 1 from overconsuming in state θ. Consequently, at date 1 a completely naïve agent (i.e., $\tilde{\beta} = 1$) will claim the high consumption state $\bar{\theta}_1$ when the true state is $\underline{\theta}_1$, believing that self 2 will then report $\underline{\theta}_2$, delivering consumption $C^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$. However, after self 1 claims the high consumption state $\bar{\theta}_1$, self 2 in fact reports $\bar{\theta}_2$, delivering consumption $C\left(\bar{\theta}_1, \bar{\theta}_2\right)$, so that self 0's equilibrium utility in $\left(\underline{\theta}_1, \bar{\theta}_2\right)$ is $U^0\left(C\left(\bar{\theta}_1, \bar{\theta}_2\right); \underline{\theta}_1, \bar{\theta}_2\right)$. But because $V^2\left(C\left(\bar{\theta}_1, \bar{\theta}_2\right); \bar{\theta}_2\right) < V^2(C^*\left(\bar{\theta}_1, \underline{\theta}_2\right); \bar{\theta}_2)$, this is strictly less than the utility self 0 would get from a contract allowing self 1 to choose freely between $C^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$ and $C^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$, namely $U^0\left(C^*\left(\bar{\theta}_1, \underline{\theta}_2\right); \underline{\theta}_1, \bar{\theta}_2\right)$.[39] Consequently, although there is scope for government paternalism to improve welfare if the government has a reasonably precise estimate of the degree of naïveté, such paternalism is dangerous if agents are instead much more naïve than the government believes.[40]

8 Conclusion

Our analysis characterizes circumstances under which an agent with hyperbolic discounting is able to resolve the tension between commitment and flexibility. When this is possible, hyperbolic discounting has no impact on equilibrium consumption. The key condition we identify is preference reversal: high desired consumption at date 1 is associated with low MU at date 2. As discussed in Subsection 4.2 and Section 5, preference reversal arises naturally in a number of economic settings.

We have focused throughout on how unverifiability affects the individual's ability to reconcile commitment with flexibility. In doing so, we have abstracted from other possible impediments, such as a lack of exclusivity in contracting, a lack of commitment by contract counterparties, or other frictions in the contracting process. In this sense, our analysis provides an upper bound on an individual's ability to combine commitment with flexibility. However, our analysis of the impact of private savings, summarized in Section 6, deals with arguably the most important issue related to exclusivity, namely the possibility of using

[39]The argument here is closely related to Heidhues and Kőszegi (2010). Self 2 effectively borrows on expensive terms that self 1 naïvely believed he would not agree to.

[40]Eliaz and Spiegler (2006) analyze profit maximization by a monopolist who deals with a population of time-inconsistent individuals who differ in their degree of sophistication. The problem noted in the main text suggests that the parallel question of welfare maximization for a population of differentially sophisticated time-inconsistent individuals would also be interesting. We leave this topic for future research. Also related is the problem of designing a contract for a population of partially naïve agents who differ in the strength of their hyperbolic discounting, e.g., β varies across agents while $\hat{\beta}/\beta$ is constant. Again, we leave this interesting question for future research.

savings instruments outside the contract.

In this paper, we focus on one particular form of time-inconsistent preferences, namely the present-bias generated by hyperbolic discounting. However, our key insight—that time-inconsistent preferences not only generate commitment problems, but also allow their possible solution, since the preferences of later selves can be exploited to punish undesirable behavior by earlier selves—is more widely applicable. In particular, consider any source of time-inconsistent preferences that an individual is self-aware enough to anticipate. For example, an individual may understand today that, in the future, he will misinterpret the relevance of a small number of data points. Just as in the current setting, he can potentially commit to a course of action that avoids this bias, while at the same time maintaining flexibility to respond to shocks.

References

[1] Amador, M., I. Werning, and G.M. Angeletos (2003), "Commitment vs. Flexibility," *NBER Working Paper* 10151.

[2] Amador, M., I. Werning, and G.M. Angeletos (2006), "Commitment vs. Flexibility," *Econometrica* 74, 365-396.

[3] Ariely, D., and K. Wertenbroch (2002), "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," *Psychological Science* 13, 219-224.

[4] Augenblick, Ned, Muriel Niederle, and Charles Sprenger (2015), "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *Quarterly Journal of Economics* 130, 1067-1115.

[5] Bénabou, R., and J. Tirole (2002), "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics* 117, 871-915.

[6] Bénabou, R., and J. Tirole (2004), "Willpower and Personal Rules," *Journal of Political Economy* 112, 848-886.

[7] Bernheim, B.D., D. Ray, and S. Yeltekin (2013), "Self-Control, Saving, and the Low Asset Trap," working paper.

[8] Bolton, P., and Dewatripont, M. (2005), *Contract Theory*, MIT Press.

[9] Carrillo, J.D., and T. Mariotti (2000), "Strategic Ignorance as a Self-Disciplining Device," *Review of Economic Studies* 67, 529-544.

[10] Cole, H.L., and N.R. Kocherlakota (2001), "Efficient Allocations with Hidden Income and Hidden Storage," *Review of Economic Studies* 68, 523-542.

[11] DellaVigna, S., and U. Malmendier (2004), "Contract Design and Self-Control: Theory and Evidence," *Quarterly Journal of Economics* 119, 353-402.

[12] Doepke, M., and R.M. Townsend (2006), "Dynamic mechanism design with hidden income and hidden actions," *Journal of Economic Theory* 126, 235-285.

[13] Eliaz, K., and R. Spiegler (2006), "Contracting with Diversely Naive Agents," *Review of Economic Studies* 73, 689-714.

[14] Frederick, S., G. Loewenstein, and T. O'Donoghue (2002), "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature* 40, 351-401.

[15] Hall, Robert E., "Intertemporal Substitution in Consumption," *Journal of Political Economy 96*, 339-357.

[16] He, Zhiguo (2012), "Dynamic Compensation Contracts with Private Savings," *Review of Financial Studies* 25, 1494-1549.

[17] Heidhues, P. and Kőszegi, B. (2010), "Exploiting Naïvete about Self-Control in the Credit Market," *American Economic Review* 100, 2279-2303.

[18] Jackson, Thomas H. (1986), *The Logic and Limits of Bankruptcy Law*. Cambridge, MA: Harvard University Press.

[19] Kocherlakota, Naryana R. (2004), "Figuring out the impact of hidden savings on optimal unemployment insurance," *Review of Economic Dynamics* 7, 541-554.

[20] Krusell, P., and A.A. Smith (2003), "Consumption-Savings Decisions with Quasi-Geometric Discounting," *Econometrica* 71, 365-375.

[21] Laibson, D. (1996), "Hyperbolic Discount Functions, Undersaving and Savings Policy," *Working Paper*.

[22] Laibson, D. (1997), "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics* 112, 443-477.

[23] Laibson, D., A. Repetto, and J. Tobacman (1998), "Self-Control and Saving for Retirement," *Brookings Papers on Economic Activity* 57, 91-196.

[24] Laibson, D., A. Repetto, and J. Tobacman (2007), "Estimating Discount Functions with Consumption Choices over the Lifecycle," *Working Paper*.

[25] Luttmer, E.G.J, and T. Mariotti (2003), "Subjective Discounting in an Exchange Economy," *Journal of Political Economy* 111, 959-989.

[26] Maskin, E. (1999), "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies* 66, 23-38.

[27] Maskin, E. and T. Sjöström (2002), "Implementation Theory," in K. Arrow, A. Sen, and K. Suzumara, eds., *Handbook of Social Choice and Welfare*, vol. I. New York: Elsevier Science.

[28] Myerson, R. B. (1981), "Optimal Auction Design," *Mathematics of Operations Research* 6, 58-73.

[29] O'Donoghue, T., and M. Rabin (1999a), "Doing it Now or Later," *American Economic Review* 89, 103-124.

[30] O'Donoghue, T., and M. Rabin (1999b), "Incentives for Procrastinators," *Quarterly Journal of Economics* 114, 769-816.

[31] O'Donoghue, T., and M. Rabin (2001), "Choice and Procrastination," *Quarterly Journal of Economics* 116, 121-160.

[32] Palfrey, T. (2002), "Implementation Theory," in R. Aumann and S. Hart, eds., *Handbook of Game Theory with Economic Applications*, vol. III. New York: Elsevier Science.

[33] Serrano, R. (2004), "The Theory of Implementation of Social Choice Rules," *SIAM Review* 46, 377-414.

[34] Shui H. and L. Ausubel, (2005), "Time Inconsistency in the Credit Card Market," *Working Paper*.

[35] Skiba, P.M., and J. Tobacman (2008), "Payday Loans, Uncertainty, and Discounting: Explaining Patterns of Borrowing, Repayment, and Default," *Working Paper*.

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

Results omitted from main text

Lemma A-1 *If* $\tilde{c}_2 \geq c_2$ *and* $V^2(\tilde{c};\theta_2) \geq V^2(c;\theta_2)$*, then* $U^2(\tilde{c};\theta_2) \geq U^2(c;\theta_2)$*, with strict inequality if either* $\tilde{c}_2 > c_2$ *or* $V^2(\tilde{c};\theta_2) > V^2(c;\theta_2)$*. Likewise, if* $\tilde{c}_1 \geq c_1$ *and* $V^1(\tilde{c};\theta_1,\theta_2) \geq V^1(c;\theta_1,\theta_2)$*, then* $U^1(\tilde{c};\theta_1,\theta_2) \geq U^2(c;\theta_1,\theta_2)$*, with strict inequality if either* $\tilde{c}_1 > c_1$ *or* $V^1(\tilde{c};\theta_1,\theta_2) > V^1(c;\theta_1,\theta_2)$*.*

Proof of Lemma A-1: We prove the first statement; the second statement has a parallel proof. Rewriting $V^2(\tilde{c};\theta_2) \geq V^2(c;\theta_2)$ and $U^2(\tilde{c};\theta_2) \geq U^2(c;\theta_2)$ gives, respectively, $u_2(\tilde{c}_2;\theta_2) - u_2(c_2;\theta_2) \geq u_3(c_3) - u_3(\tilde{c}_3)$ and $u_2(\tilde{c}_2;\theta_2) - u_2(c_2;\theta_2) \geq \beta(u_3(c_3) - u_3(\tilde{c}_3))$. The result is then immediate. **QED**

Lemma A-2 *If C satisfies $IC_2\left(\theta_1,\theta_2,\tilde{\theta}_2\right)$ with equality and sign $\left(C_2(\theta_1,\theta_2) - C_2\left(\theta_1,\tilde{\theta}_2\right)\right) =$ sign $\left(\theta_2 - \tilde{\theta}_2\right)$ then C satisfies $IC_2\left(\theta_1,\tilde{\theta}_2,\theta_2\right)$, and does so strictly if $C_2(\theta_1,\theta_2) \neq C_2\left(\theta_1,\tilde{\theta}_2\right)$.*

Proof of Lemma A-2: Since $\text{IC}_2\left(\theta_1,\theta_2,\tilde{\theta}_2\right)$ holds with equality, $u_2(C_2(\theta_1,\theta_2);\theta_2) - u_2\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right) = \beta\left(u_3\left(C_3\left(\theta_1,\tilde{\theta}_2\right)\right) - u_3(C_3(\theta_1,\theta_2))\right)$. If either $C_2(\theta_1,\theta_2) \geq C_2\left(\theta_1,\tilde{\theta}_2\right)$ and $\theta_2 > \tilde{\theta}_2$, or $C_2(\theta_1,\theta_2) \leq C_2\left(\theta_1,\tilde{\theta}_2\right)$ and $\theta_2 < \tilde{\theta}_2$, then $u_2\left(C_2(\theta_1,\theta_2);\tilde{\theta}_2\right) - u_2\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\tilde{\theta}_2\right) \leq \beta\left(u_3\left(C_3\left(\theta_1,\tilde{\theta}_2\right)\right) - u_3(C_3(\theta_1,\theta_2))\right)$, which is equivalent to $\text{IC}_2\left(\theta_1,\tilde{\theta}_2,\theta_2\right)$.
QED

Lemma A-3 *Define $\bar{v}^2(k,\theta_2) = \max_{c_2} V^2(c_2,k-c_2;\theta_2)$ and $\bar{c}_2^k$, $\underline{c}_2^k$ as, respectively, the maximizers of $V^2\left(c_2,k-c_2;\bar{\theta}_2\right)$ and $V^2(c_2,k-c_2;\underline{\theta}_2)$. For $v \in \left(V^2(0,0;\underline{\theta}_2),\bar{v}^2(k,\underline{\theta}_2)\right]$, define*

$$
\begin{aligned}
f(v;k,\beta) \quad \equiv \quad & \max_{\bar{c}_2,\bar{c}_3,\underline{c}_2,\underline{c}_3} V^2\left(\bar{c}_2,\bar{c}_3;\bar{\theta}_2\right) && \text{(A-1)}\\
s.t.\quad & V^2(\underline{c}_2,\underline{c}_3;\underline{\theta}_2) = v &&\\
and\quad & U^2(\underline{c}_2,\underline{c}_3;\underline{\theta}_2) \geq U^2(\bar{c}_2,\bar{c}_3;\underline{\theta}_2) && \text{(A-2)}\\
and\quad & U^2\left(\bar{c}_2,\bar{c}_3;\bar{\theta}_2\right) \geq U^2\left(\underline{c}_2,\underline{c}_3;\bar{\theta}_2\right) && \text{(A-3)}\\
and\quad & \bar{c}_2 + \bar{c}_3 \leq k \text{ and } \underline{c}_2 + \underline{c}_3 \leq k. &&
\end{aligned}
$$

Fix k, β. For $v \geq V^2\left(\bar{c}_2^k, k-\bar{c}_2^k;\underline{\theta}_2\right)$, the following is true of the solution to problem (A-1): (A) The IC constraint (A-3) does not bind. If moreover $f(v;k,\beta) < \bar{v}^2\left(k,\bar{\theta}_2\right)$, then (B) the consumption allocation $(\underline{c}_2,\underline{c}_3)$ is uniquely defined, with $\underline{c}_2 + \underline{c}_3 = k$ and $\underline{c}_2$ strictly decreasing in v, and (C) $\bar{c}_2 - \underline{c}_2$ is non-negative, increasing in v, and strictly increasing if $\bar{c}_2 - \underline{c}_2 > 0$.

Proof of Lemma A-3: Note that $\bar{c}_2^k \geq \underline{c}_2^k$. To establish (A), let $(\bar{c}_2, \bar{c}_3, \underline{c}_2, \underline{c}_3)$ solve the relaxed version of problem (A-1) in which constraint (A-3) is not imposed. We show that (A-3) is nonetheless satisfied. The inequalities $v \geq V^2\left(\bar{c}_2^k, k - \bar{c}_2^k; \underline{\theta}_2\right)$ and $\bar{c}_2^k \geq \underline{c}_2^k$ together imply $\underline{c}_2 \leq \bar{c}_2^k$. If $\bar{c}_2 = \bar{c}_2^k$, it is immediate from Lemma A-1 that (A-3) holds. Accordingly, for the remainder of the proof of (A) we consider the case $\bar{c}_2 \neq \bar{c}_2^k$. Since $\bar{c}_2 \neq \bar{c}_2^k$, the IC constraint (A-2) must hold with equality. Similarly, the resource constraint $\underline{c}_2 + \underline{c}_3 \leq k$ must hold with equality, and $\underline{c}_2 \geq \underline{c}_2^k$, since if instead $\underline{c}_2 < \underline{c}_2^k$ it is possible to increase $\underline{c}_2$ and decrease $\underline{c}_3$ to leave $V^2(\underline{c}_2, \underline{c}_3; \underline{\theta}_2)$ unchanged while strictly increasing $U^2(\underline{c}_2, \underline{c}_3; \underline{\theta}_2)$, thereby relaxing the IC constraint (A-2).

It follows that $\bar{c}_2 \geq \underline{c}_2$, since if instead $\bar{c}_2 < \underline{c}_2$, then $V^2\left(\bar{c}_2, \bar{c}_3; \bar{\theta}_2\right) \leq V^2\left(\bar{c}_2, k - \bar{c}_2; \bar{\theta}_2\right) < V^2\left(\underline{c}_2, \underline{c}_3; \bar{\theta}_2\right)$, where the second inequality uses $\underline{c}_2 \leq \bar{c}_2^k$ and $\underline{c}_2 + \underline{c}_3 = k$. Since setting $(\bar{c}_2, \bar{c}_3) = (\underline{c}_2, \underline{c}_3)$ satisfies all constraints in the relaxed problem, this gives a contradiction. From $\bar{c}_2 \geq \underline{c}_2$ and the fact that the IC (A-2) holds with equality, Lemma A-2 implies that the IC (A-3) holds, completing the proof of (A).

Statement (B) is immediate from the observations above that $\underline{c}_2 \in \left[\underline{c}_2^k, \bar{c}_2^k\right]$ and $\underline{c}_2 + \underline{c}_3 = k$ if $\bar{c}_2 \neq \bar{c}_2^k$.

For statement (C), we have already shown that $\bar{c}_2 - \underline{c}_2 \geq 0$ if $\bar{c}_2 \neq \bar{c}_2^k$. It remains to establish that $\bar{c}_2 - \underline{c}_2$ is increasing. Since the IC constraint (A-2) holds with equality,

$$
\begin{aligned}
V^2\left(\bar{c}_2, \bar{c}_3; \bar{\theta}_2\right) &= V^2\left(\underline{c}_2, \underline{c}_3; \bar{\theta}_2\right) + u_2\left(\bar{c}_2; \bar{\theta}_2\right) - u_2\left(\underline{c}_2; \bar{\theta}_2\right) + u_3\left(\bar{c}_3\right) - u_3\left(\underline{c}_3\right) \\
&= V^2\left(\underline{c}_2, \underline{c}_3; \bar{\theta}_2\right) + u_2\left(\bar{c}_2; \bar{\theta}_2\right) - u_2\left(\underline{c}_2; \bar{\theta}_2\right) - \frac{1}{\beta}\left(u_2\left(\bar{c}_2; \underline{\theta}_2\right) - u_2\left(\underline{c}_2; \underline{\theta}_2\right)\right). \\
&= V^2\left(\underline{c}_2, \underline{c}_3; \bar{\theta}_2\right) + u_2\left(\bar{c}_2; \bar{\theta}_2\right) - \frac{1}{\beta} u_2\left(\bar{c}_2; \underline{\theta}_2\right) - \left(u_2\left(\underline{c}_2; \bar{\theta}_2\right) - \frac{1}{\beta} u_2\left(\underline{c}_2; \underline{\theta}_2\right)\right).
\end{aligned}
$$

Fix a value of v such that the solution to problem (A-1) has $\bar{c}_2 - \underline{c}_2 > 0$. To establish the result, we show that a small upwards perturbation in v strictly raises $\bar{c}_2 - \underline{c}_2$.

There are two cases. First, consider the case $\bar{c}_2 + \bar{c}_3 < k$. In this case, $\bar{c}_2$ must maximize

$u_2\left(\bar{c}_2;\bar{\theta}_2\right)-\frac{1}{\beta}u_2\left(\bar{c}_2;\underline{\theta}_2\right)$. Consequently, $\bar{c}_2$ is invariant to small changes in $\underline{c}_2$, and hence to small changes in v.

Second, consider the case $\bar{c}_2+\bar{c}_3=k$. Since the IC constraint (A-2) holds with equality, and $\bar{c}_2-\underline{c}_2>0$, it follows from concavity of $u_2\left(c_2;\underline{\theta}_2\right)+\beta u_3\left(k-c_2\right)$ that $u_2'\left(\bar{c}_2;\underline{\theta}_2\right)-\beta u_3'\left(k-\bar{c}_2\right)<0<u_2'\left(\underline{c}_2;\underline{\theta}_2\right)-\beta u_3'\left(k-\underline{c}_2\right)$. From this, it in turn follows that $u_2'\left(\bar{c}_2;\bar{\theta}_2\right)-\frac{1}{\beta}u_2'\left(\bar{c}_2;\underline{\theta}_2\right)\geq 0$, since if instead $u_2'\left(\bar{c}_2;\bar{\theta}_2\right)-\frac{1}{\beta}u_2'\left(\bar{c}_2;\underline{\theta}_2\right)<0$, it is possible to increase $V^2\left(\bar{c}_2,\bar{c}_3;\bar{\theta}_2\right)$ by reducing $\bar{c}_2$, and at the same time reducing $\bar{c}_3$ to preserve the IC constraint (A-2) at equality. Rearranging, we have $\frac{1}{\beta}\frac{u_2'\left(\bar{c}_2;\underline{\theta}_2\right)}{u_2'\left(\bar{c}_2;\bar{\theta}_2\right)}\leq 1$, which by Assumption 2 implies that $\frac{1}{\beta}\frac{u_2'\left(c_2;\underline{\theta}_2\right)}{u_2'\left(c_2;\bar{\theta}_2\right)}\leq 1$ for all $c_2\leq\bar{c}_2$, or equivalently, that $u_2'\left(c_2;\bar{\theta}_2\right)-\frac{1}{\beta}u_2'\left(c_2;\underline{\theta}_2\right)\geq 0$ for all $c_2\leq\bar{c}_2$. With these preliminaries, consider now a small upwards perturbation in v, to v^+ say. From above, this is associated with a small decrease in $\underline{c}_2$, which tightens the IC constraint (A-2). One way to keep (A-2) satisfied is to leave $\bar{c}_2$ unchanged and reduce $\bar{c}_3$. In contrast, any candidate solution in which $\bar{c}_2$ is lowered delivers a lower value of $V^2\left(\bar{c}_2,\bar{c}_3;\bar{\theta}_2\right)$. Hence the solution at v^+ certainly entails a weakly higher choice of $\bar{c}_2$, completing the proof. **QED**

Lemma A-4 *Let $\bar{v}^2\left(k,\theta_2\right)$, $\bar{c}_2^k$, and $f\left(v;k,\beta\right)$ be as defined in Lemma A-3. For $v\in\left[V^2\left(\bar{c}_2^k,k-\bar{c}_2^k;\underline{\theta}_2\right),\bar{v}^2\left(k,\underline{\theta}_2\right)\right]$, the function f is continuous in (v,k,β), and concave in (v,k).*[41] *Moreover, f is differentiable in v at any value such that $f\left(v;k,\beta\right)<\bar{v}^2\left(k,\bar{\theta}_2\right)$.*

Proof of Lemma A-4: Continuity follows from the Theorem of the Maximum. To establish concavity in v, we rewrite f using a change of variables and Lemma A-3(A):

$$f\left(v;k,\beta\right)\equiv\max_{\bar{u}_2,\bar{u}_3,\underline{u}_2,\underline{u}_3}\bar{u}_2+\bar{u}_3 \tag{A-4}$$

$$\text{s.t.}\quad \underline{u}_2+\underline{u}_3=v \tag{A-5}$$

$$\text{and}\quad \underline{u}_2+\beta\underline{u}_3\geq u_2\left(u_2^{-1}\left(\bar{u}_2;\bar{\theta}_2\right);\underline{\theta}_2\right)+\beta\bar{u}_3 \tag{A-6}$$

$$\text{and}\quad u_2^{-1}\left(\bar{u}_2;\bar{\theta}_2\right)+u_3^{-1}\left(\bar{u}_3\right)\leq k,\, u_2^{-1}\left(\underline{u}_2;\underline{\theta}_2\right)+u_3^{-1}\left(\underline{u}_3\right)\leq k. \tag{A-7}$$

[41] Amador et al (2003) state a related result for a continuous state space and multiplicative shocks.

To establish concavity of f with respect to (v,k), it is sufficient to show that the constraint set defined by (A-5), (A-6) and (A-7) is convex in $(\bar{u}_2, \bar{u}_3, \underline{u}_2, \underline{u}_3, v, k)$. Note first that

$$\frac{\partial}{\partial w} u_2\left(u_2^{-1}\left(w;\bar{\theta}_2\right);\underline{\theta}_2\right) = u_2'\left(u_2^{-1}\left(w;\bar{\theta}_2\right);\underline{\theta}_2\right)\left(u_2^{-1}\right)'\left(w;\bar{\theta}_2\right) = \frac{u_2'\left(u_2^{-1}\left(w;\bar{\theta}_2\right);\underline{\theta}_2\right)}{u_2'\left(u_2^{-1}\left(w;\bar{\theta}_2\right);\bar{\theta}_2\right)}.$$

By Assumption 2, the ratio $\frac{u_2'(c_2;\underline{\theta}_2)}{u_2'(c_2;\bar{\theta}_2)}$ is increasing in c_2, and hence $u_2\left(u_2^{-1}\left(w;\bar{\theta}_2\right);\underline{\theta}_2\right)$ is convex in w. Moreover, $u_2^{-1}(\cdot;\theta_2)$ and u_3^{-1} are certainly convex (they are inverses of concave strictly increasing functions). Hence the constraint set is convex.

Finally, we establish differentiability. Given concavity of f in v, to establish differentiability at a point v_0 it is sufficient to exhibit a differentiable function g such that $g(v) \leq f(v;k,\beta)$ in the neighborhood of v_0, with equality at v_0. Let $(\bar{c}_2, \bar{c}_3, \underline{c}_2, \underline{c}_3)$ be the solution to (A-1) at v_0. From Lemma A-3, it follows that $(\underline{c}_2, \underline{c}_3)$ is differentiable as a function of v. To construct the function g, there are two cases. First, if $u_2'(\bar{c}_2;\underline{\theta}_2) - \beta u_3'(\bar{c}_3) \neq 0$, define $g(v) = u_2\left(\bar{c}_2;\bar{\theta}_2\right) + u_3(\bar{c}_3)$ by perturbing $\bar{c}_2$ and $\bar{c}_3$ in equal but opposite directions to satisfy (A-2) with equality. Second, consider the case of $u_2'(\bar{c}_2;\underline{\theta}_2) - \beta u_3'(\bar{c}_3) = 0$. This case can only arise if $u_2'\left(\bar{c}_2;\bar{\theta}_2\right) - u_3'(\bar{c}_3) = 0$, which by the condition that $f(v;k,\beta) < \bar{v}^2\left(k,\bar{\theta}_2\right)$ implies that $\bar{c}_2 + \bar{c}_3 < k$. In this case, define $g(v) = u_2\left(\bar{c}_2;\bar{\theta}_2\right) + u_3(\bar{c}_3)$ by holding $\bar{c}_2$ fixed and perturbing $\bar{c}_3$ to leave (A-2) satisfied at equality. This completes the proof of differentiability. **QED**

Proofs of results stated in main text

Proof of Lemma 1: From IC_2, $U^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\tilde{\theta}_2\right) \geq U^2\left(C(\theta_1,\theta_2);\tilde{\theta}_2\right)$ and $U^2\left(C(\theta_1,\theta_2);\theta_2\right) \geq U^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right)$. Expanding, this implies $u_2\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\tilde{\theta}_2\right) - u_2\left(C_2(\theta_1,\theta_2);\tilde{\theta}_2\right) \geq u_2\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right) - u_2\left(C_2(\theta_1,\theta_2);\theta_2\right)$ which by $u_2'\left(\cdot;\tilde{\theta}_2\right) > u_2'(\cdot;\theta_2)$ implies $C_2\left(\theta_1,\tilde{\theta}_2\right) \geq C_2(\theta_1,\theta_2)$. **QED**

Proof of Proposition 1, Part (B): The proof is constructive. Define a contract C

by $C(\theta_1,\theta_2) = C^*(\theta_1,\theta_2)$ if $(\theta_1,\theta_2) \neq \left(\bar{\theta}_1,\bar{\theta}_2\right)$. It remains to define $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)$ and $C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)$. Note first that if $C\left(\bar{\theta}_1,\bar{\theta}_2\right)$ is set equal to $C^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$ then at $\beta=\beta^*$ constraint $\text{IC}_1\left(\underline{\theta}_1,\bar{\theta}_1\right)$ holds with equality. Moreover, $u_2'\left(C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right) = u_3'\left(C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)\right)$, so certainly $u_2'\left(C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right);\bar{\theta}_2\right) > u_3'\left(C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)\right)$. Choose $C\left(\bar{\theta}_1,\bar{\theta}_2\right)$ so that, at $\beta=\beta^*$, $U^2\left(C\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right) = U^2\left(C^*\left(\bar{\theta}_1,\underline{\theta}_1\right);\bar{\theta}_2\right)$ and $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right) > C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$ and $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right) + C_3\left(\bar{\theta}_1,\bar{\theta}_2\right) \leq C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right) + C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$. By Lemma A-1, we know $V^2\left(C\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right) < V^2\left(C^*\left(\bar{\theta}_1,\underline{\theta}_2\right);\bar{\theta}_2\right)$. So $\text{IC}_1\left(\underline{\theta}_1,\bar{\theta}_1\right)$ holds strictly at $\beta=\beta^*$, and moreover, by continuity there exists some interval $[\beta_1,\beta^*]$ such that $\text{IC}_1\left(\underline{\theta}_1,\bar{\theta}_1\right)$ holds for all β in this interval. Because $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right) > C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$, $\text{IC}_2\left(\bar{\theta}_1,\bar{\theta}_2,\underline{\theta}_2\right)$ is satisfied for all $\beta \leq \beta^*$. By Lemma A-2, $\text{IC}_2\left(\bar{\theta}_1,\underline{\theta}_2,\bar{\theta}_2\right)$ holds strictly at $\beta=\beta^*$, and so by continuity there exists some interval $[\beta_2,\beta^*]$ such that $\text{IC}_2\left(\bar{\theta}_1,\underline{\theta}_2,\bar{\theta}_2\right)$ holds over this interval. $\text{IC}_2(\underline{\theta}_1,\cdot,\cdot)$ holds trivially. Finally, Lemma A-1, the fact C^* solves (2), the fact that $C_1^*\left(\bar{\theta}_1\right) \geq C_1^*\left(\underline{\theta}_1\right)$ together imply that C satisfies $\text{IC}_1\left(\bar{\theta}_1,\underline{\theta}_1\right)$. Defining $\hat{\beta} = \max\{\beta_1,\beta_2\}$ completes the proof. **QED**

Proof of Lemma 2: Suppose to the contrary that $V^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\underline{\theta}_2\right) > V^2\left(C\left(\theta_1,\underline{\theta}_2\right);\underline{\theta}_2\right)$. By Lemma 1, $C_2\left(\theta_1,\tilde{\theta}_2\right) \geq C_2\left(\theta_1,\underline{\theta}_2\right)$. But then Lemma A-1 implies $U^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\underline{\theta}_2\right) > U^2\left(C\left(\theta_1,\underline{\theta}_2\right);\underline{\theta}_2\right)$, contradicting IC_2. **QED**

Proof of Proposition 2: The proof is constructive. Define a contract C by $C(\theta_1,\theta_2) = C^*(\theta_1,\theta_2)$ if $(\theta_1,\theta_2) \neq \left(\bar{\theta}_1,\bar{\theta}_2\right)$, and the remaining terms $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)$ and $C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)$ by

$$u_2\left(C_2\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right) = u_2\left(C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right);\bar{\theta}_2\right) + u_1\left(C_1^*\left(\bar{\theta}_1\right);\underline{\theta}_1\right) - u_1\left(C_1^*\left(\underline{\theta}_1\right);\underline{\theta}_1\right) \quad \text{(A-8)}$$

$$u_3\left(C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)\right) = u_3\left(C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)\right) - \frac{1}{\beta}\left(u_1\left(C_1^*\left(\bar{\theta}_1\right);\underline{\theta}_1\right) - u_1\left(C_1^*\left(\underline{\theta}_1\right);\underline{\theta}_1\right)\right) \quad \text{(A-9)}$$

Existence of $C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)$ satisfying (A-9) follows from Condition (I). To establish the existence of $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)$ satisfying (A-8), note that Condition (III) is equivalent to $C_1^*\left(\underline{\theta}_1\right) + C_2^*\left(\underline{\theta}_1,\bar{\theta}_2\right) \geq C_1^*\left(\bar{\theta}_1\right) + C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$, which implies

$$u_1\left(C_1^*\left(\underline{\theta}_1\right);\underline{\theta}_1\right) + u_2\left(C_2^*\left(\underline{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right) \geq u_1\left(C_1^*\left(\bar{\theta}_1\right);\underline{\theta}_1\right) + u_2\left(C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right);\bar{\theta}_2\right).$$

It follows that $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)$ is well-defined, and moreover, lies in the interval $\left[C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right),C_2^*\left(\underline{\theta}_1,\bar{\theta}_2\right)\right]$.

By construction, $\mathrm{IC}_2\left(\bar{\theta}_1,\bar{\theta}_2,\underline{\theta}_2\right)$ holds with equality. Hence, by Lemma A-2, $\mathrm{IC}_2\left(\bar{\theta}_1,\underline{\theta}_2,\bar{\theta}_2\right)$ holds also. $\mathrm{IC}_2(\underline{\theta}_1,\cdot,\cdot)$ holds trivially. $\mathrm{IC}_1\left(\bar{\theta}_1,\underline{\theta}_1\right)$ holds by the same argument as in the proof of Proposition 1. To see that $\mathrm{IC}_1\left(\underline{\theta}_1,\bar{\theta}_1\right)$ holds, observe first that self 1's utility U^1 from consumption c is related to V^1 by

$$U^1\left(c;\theta_1,\theta_2\right)=\left(1-\beta\right)u_1\left(c;\theta_1\right)+\beta V^1\left(c;\theta_1,\theta_2\right). \tag{A-10}$$

Consequently, the gain to self 1 in state $\left(\underline{\theta}_1,\bar{\theta}_2\right)$ of obtaining consumption $C^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$ instead of $C^*\left(\underline{\theta}_1,\bar{\theta}_2\right)$ is strictly less than

$$\left(1-\beta\right)\left(u_1\left(C_1^*\left(\bar{\theta}_1\right);\underline{\theta}_1\right)-u_1\left(C_1^*\left(\underline{\theta}_1\right);\underline{\theta}_1\right)\right), \tag{A-11}$$

since by definition $C^*\left(\underline{\theta}_1,\bar{\theta}_2\right)$ maximizes $V^1\left(\cdot;\underline{\theta}_1,\bar{\theta}_2\right)$. Second, $C\left(\bar{\theta}_1,\bar{\theta}_2\right)$ delivers self 1 utility in state $\left(\underline{\theta}_1,\bar{\theta}_2\right)$ that is exactly an amount (A-11) below $U^1\left(C^*\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_1,\bar{\theta}_2\right)$. Together, these two observations confirm that $\mathrm{IC}_1\left(\underline{\theta}_1,\bar{\theta}_1\right)$ holds.

It remains to show that $C\left(\bar{\theta}_1,\bar{\theta}_2\right)$ satisfies RC, or equivalently, $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)+C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)\leq C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)+C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$, which is equivalent to $u_3\left(C_3\left(\bar{\theta}_1,\bar{\theta}_2\right)\right)\leq u_3\left(C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)+C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)-C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)\right)$, which by (A-8) and (A-9), is in turn equivalent to

$$\begin{aligned}&u_3\left(C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)\right)-u_3\left(C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)+C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)-C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)\right)\\ \leq\;&\frac{1}{\beta}\left(u_2\left(C_2\left(\bar{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)-u_2\left(C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right);\bar{\theta}_2\right)\right).\end{aligned}$$

Since $\beta<1$ and $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)\geq C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)$, it is sufficient to establish this inequality at $\beta=1$, and given the concavity of u_2 and the earlier observation that $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)\leq C_2^*\left(\underline{\theta}_1,\bar{\theta}_2\right)$, it is in turn sufficient to establish

$$\begin{aligned} & u_3\left(C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)\right)-u_3\left(C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right)-\left(C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)-C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)\right)\right) \\ \leq\ & u_2\left(C_2^*\left(\underline{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)-u_2\left(C_2^*\left(\underline{\theta}_1,\bar{\theta}_2\right)-\left(C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)-C_2^*\left(\bar{\theta}_1,\underline{\theta}_2\right)\right);\bar{\theta}_2\right). \end{aligned}$$

Under Condition (II), this inequality is equivalent to $C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right) \geq C_2^*\left(\underline{\theta}_1,\bar{\theta}_2\right)-\kappa$, which is satisfied since the fact that C^* solves (2) implies $C_3^*\left(\underline{\theta}_1,\bar{\theta}_2\right)=C_2^*\left(\underline{\theta}_1,\bar{\theta}_2\right)-\kappa$, and from Condition (III), $C_3^*\left(\bar{\theta}_1,\underline{\theta}_2\right) \geq C_3^*\left(\underline{\theta}_1,\bar{\theta}_2\right)$. **QED**

Proof of Lemma 3: Let $\bar{v}^2\left(k,\theta_2\right)$, $\bar{c}_2^k$, and $f\left(v;k,\beta\right)$ be as defined in Lemma A-3. For any θ_1 and β, define $g\left(k,\beta;\theta_1\right)$ by

$$g\left(k,\beta;\theta_1\right) \equiv \max_{v\in\left(V^2\left(0,0;\underline{\theta}_2\right),\bar{v}^2\left(k,\underline{\theta}_2\right)\right]} \Pr\left(\bar{\theta}_2|\theta_1\right) f\left(v;k,\beta\right)+\Pr\left(\underline{\theta}_2|\theta_1\right) v. \tag{A-12}$$

Then

$$C_1^*\left(\theta_1;\beta\right)=\arg\max_{c_1} u_1\left(c_1;\theta_1\right)+g\left(W-c_1,\beta;\theta_1\right).$$

Note first that the maximizing choice of v in (A-12) satisfies $v \geq V^2\left(\bar{c}_2^k,k-\bar{c}_2^k;\underline{\theta}_2\right)$, since $f\left(V^2\left(\bar{c}_2^k,k-\bar{c}_2^k;\underline{\theta}_2\right);k,\beta\right)=\bar{v}^2\left(k,\bar{\theta}_2\right)$. From Lemma A-4, the function g is continuous in (k,β) and concave in k. Consequently, $C_1^*\left(\theta_1;\beta\right)$ is the maximizer of a strictly concave function, and hence is continuous as a function of β. Since $E\left[V^2\left(C_1^*\left(\theta_1,\theta_2;\beta\right);\theta_1,\theta_2\right)|\theta_1\right]=g\left(W-C_1^*\left(\theta_1;\beta\right),\beta;\theta_1\right)$, the result then follows from the continuity of g in (k,β). **QED**

Proof of Proposition 3, Part (B): Because $\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)=1$, consumption $C^*\left(\bar{\theta}_1,\bar{\theta}_2;\beta\right)=C^*\left(\bar{\theta}_1,\underline{\theta}_2;\beta\right)$, and moreover, is independent of β. Because C^* solves (2), and $C_1^*\left(\bar{\theta}_1;\beta^*\right) > C_1^*\left(\underline{\theta}_1;\beta^*\right)$, Lemma A-1 implies that C^* strictly satisfies $\mathrm{IC}_1\left(\bar{\theta}_1,\underline{\theta}_1\right)$ at $\beta=\beta^*$. The amount of slack in $\mathrm{IC}_1\left(\underline{\theta}_1,\bar{\theta}_1\right)$ under C^* equals

$$\begin{aligned} & u_1\left(C_1^*\left(\underline{\theta}_1;\beta\right);\underline{\theta}_1\right)+\beta\left(\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)V^2\left(C^*\left(\underline{\theta}_1,\bar{\theta}_2;\beta\right);\bar{\theta}_2\right)+\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)V^2\left(C^*\left(\underline{\theta}_1,\underline{\theta}_2;\beta\right);\underline{\theta}_2\right)\right) \\ -\ & u_1\left(C_1^*\left(\bar{\theta}_1;\beta\right);\underline{\theta}_1\right)-\beta\left(\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)V^2\left(C^*\left(\bar{\theta}_1,\underline{\theta}_2;\beta\right);\bar{\theta}_2\right)+\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)V^2\left(C^*\left(\bar{\theta}_1,\underline{\theta}_2;\beta\right);\underline{\theta}_2\right)\right). \end{aligned}$$

This quantity is strictly negative for β in the neighborhood below β^*, and by Lemma 3, is continuous in β, and hence equals 0 at $\beta = \beta^*$. After these preliminaries, the remainder of the proof is identical to that of Proposition 1. **QED**

Proof of Proposition 4: We prove the result for $\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right) < \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)$, i.e., $\frac{\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)}{\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)} < \frac{\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)}{\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)}$. The proof for $\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right) > \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)$ is parallel. Let $f\left(v;k,\beta\right)$ be as defined in Lemma A-3.

Preliminaries: By the Theorem of the Maximum, both $E\left[U^0\left(C^*\left(\theta_1,\theta_2;\beta\right)\right)\right]$ and $E\left[U^0\left(C^s\left(\theta_1,\theta_2;\beta\right)\right)\right]$ are continuous in β. By the definition of β, they coincide for $\beta \geq \beta^*$. Consequently,

$$E\left[U^0\left(C^*\left(\theta_1,\theta_2;\beta\right)\right)\right] - E\left[U^0\left(C^s\left(\theta_1,\theta_2;\beta\right)\right)\right] \to 0 \text{ as } \beta \to \beta^*. \qquad \text{(A-13)}$$

From (A-13), it follows that

$$C_1^*\left(\theta_1;\beta\right) - C_1^s\left(\theta_1;\beta\right) \to 0 \quad \text{as} \quad \beta \to \beta^* \text{ for } \theta_1 \in \Theta_1. \qquad \text{(A-14)}$$

$$C_t^*\left(\theta_1,\theta_2;\beta\right) - C_t^s\left(\theta_1,\theta_2;\beta\right) \to 0 \quad \text{as} \quad \beta \to \beta^* \text{ for } t = 2,3 \text{ and } \left(\theta_1,\theta_2\right) \in \Theta. \text{ (A-15)}$$

Write

$$v_{\theta_2} = V^2\left(C^s\left(\bar{\theta}_1,\theta_2;\beta\right);\theta_2\right) \text{ for } \theta_2 = \underline{\theta}_2,\bar{\theta}_2.$$

We next establish that for β in the neighborhood of β^*,

$$v_{\bar{\theta}_2} = f\left(v_{\underline{\theta}_2}; W - C_1^s\left(\bar{\theta}_1;\beta\right),\beta\right). \qquad \text{(A-16)}$$

Suppose to the contrary that $v_{\bar{\theta}_2} < f\left(v_{\underline{\theta}_2}; W - C_1^s\left(\bar{\theta}_1;\beta\right),\beta\right)$. Since $\frac{\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)}{\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)} \neq \frac{\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)}{\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)}$, there exists a perturbation that strictly increases $E\left[U^0\left(C^s\left(\theta_1,\theta_2;\beta\right)\right)\right]$ while preserving $\text{IC}_1\left(\underline{\theta}_1,\bar{\theta}_1\right)$ and IC_2. Moreover, from (A-14), (A-15) and Lemma A-1 it follows that $\text{IC}_1\left(\bar{\theta}_1,\underline{\theta}_1\right)$ is satisfied. But then the perturbation contradicts the definition of C^s.

By the definitions of f and $\hat{C}$, and by condition (7), for β in the neighborhood of β^*:

$$\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)+\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)f'\left(V^2\left(\hat{C}\left(\bar{\theta}_1,\underline{\theta}_2;\beta\right);\underline{\theta}_2\right);W-C_1^s\left(\bar{\theta}_1\right),\beta\right)=0. \tag{A-17}$$

Main step: We show, by contradiction, that

$$\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)+\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)f'\left(v_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right)>0. \tag{A-18}$$

First, suppose that

$$\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)+\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)f'\left(v_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right)=0. \tag{A-19}$$

Since f is concave (Lemma A-4), an infinitesimal reduction in $v_{\underline{\theta}_2}$ with a corresponding change in $v_{\bar{\theta}_2}$ to preserve (A-16) strictly relaxes the constraints in (1) without affecting $E\left[U^0\left(C^s\left(\theta_1,\theta_2;\beta\right)\right)\right]$. From (A-19) and $\beta<\beta^*$, we know $C_1^s\left(\cdot;\beta\right)\neq C_1^*\left(\cdot;\beta\right)$. Consequently, one can use the newly-created slack in the constraints to a construct a perturbation that increases the objective $E\left[U^0\left(C^s\left(\theta_1,\theta_2;\beta\right)\right)\right]$, giving a contradiction. Second, suppose instead that

$$\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)+\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)f'\left(v_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right)<0. \tag{A-20}$$

Let $\tilde{v}_{\underline{\theta}_2}<v_{\underline{\theta}_2}$ be such that

$$\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)\tilde{v}_{\underline{\theta}_2}+\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)f\left(\tilde{v}_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right)=\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)v_{\underline{\theta}_2}+\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)f\left(v_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right). \tag{A-21}$$

Note that $\tilde{v}_{\underline{\theta}_2}$ is well-defined by the continuity and concavity of f (Lemma A-4), together with (A-15). From (A-21),

$$\begin{aligned}\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)\frac{\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)}{\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)}\left(v_{\underline{\theta}_2}-\tilde{v}_{\underline{\theta}_2}\right) &> \Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)\left(v_{\underline{\theta}_2}-\tilde{v}_{\underline{\theta}_2}\right)\\ &= \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)\left(f\left(\tilde{v}_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right)-f\left(v_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right)\right),\end{aligned}$$

i.e.,

$$\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)v_{\underline{\theta}_2}+\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)f\left(v_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right)>\Pr\left(\underline{\theta}_2|\underline{\theta}_1\right)\tilde{v}_{\underline{\theta}_2}+\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right)f\left(\tilde{v}_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right).$$

Hence switching from $v_{\underline{\theta}_2}$ to $\tilde{v}_{\underline{\theta}_2}$ strictly relaxes the constraints in (1) while leaving the objective $E\left[U^0\left(C^s\left(\theta_1,\theta_2;\beta\right)\right)\right]$ unchanged, again leading to a contradiction.

Completing the proof: Given the concavity of f, (A-17) and (A-18) imply that $v_{\underline{\theta}_2} < V^2\left(\hat{C}\left(\bar{\theta}_1,\underline{\theta}_2;\beta\right);\underline{\theta}_2\right)$. Lemma A-3 then implies $C_2^s\left(\bar{\theta}_1,\bar{\theta}_2\right)-C_2^s\left(\bar{\theta}_1,\underline{\theta}_2\right)\leq\hat{C}_2\left(\bar{\theta}_1,\bar{\theta}_2\right)-\hat{C}_2\left(\bar{\theta}_1,\underline{\theta}_2\right)$, as claimed. **QED**

Proof of Lemma 4: Differentiating,

$$\frac{dU^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right)}{d\tilde{\theta}_2}=u_2'\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right)\frac{\partial C_2\left(\theta_1,\tilde{\theta}_2\right)}{\partial\tilde{\theta}_2}+\beta u_3'\left(C_3\left(\theta_1,\tilde{\theta}_2\right)\right)\frac{\partial C_3\left(\theta_1,\tilde{\theta}_2\right)}{\partial\tilde{\theta}_2}.$$

Substituting in (8),

$$\frac{dU^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right)}{d\tilde{\theta}_2}=\left(u_2'\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right)-u_2'\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\tilde{\theta}_2\right)\right)\frac{\partial C_2\left(\theta_1,\tilde{\theta}_2\right)}{\partial\tilde{\theta}_2}.$$

Since $C_2\left(\theta_1,\tilde{\theta}_2\right)$ is increasing in $\tilde{\theta}_2$, $sign\left(\frac{dU^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right)}{d\tilde{\theta}_2}\right)=-sign\left(\tilde{\theta}_2-\theta_2\right)$, implying the result. **QED**

Proof of Proposition 5: For Part (A), fix $\beta<1$, and choose θ_1 and $\tilde{\theta}_1>\theta_1$ such that C^* violates $\text{IC}_1\left(\theta_1,\tilde{\theta}_1\right)$. Since ϕ is increasing, $\phi\left(\tilde{\theta}_1\right)>\phi\left(\theta_1\right)$. From Proposition 1, IC_1 would bind in (1) if $\Theta=\left\{\theta_1,\tilde{\theta}_1\right\}\times\left\{\phi\left(\theta_1\right),\phi\left(\tilde{\theta}_1\right)\right\}$. *A fortiori*, IC_1 binds in (1).

The proof of Part (B) is constructive. For any $\theta_1\in\Theta_1$, define the contract C by

$C(\theta_1,\theta_2) = C^*(\theta_1,\phi(\theta_1))$ if $\phi^{-1}(\theta_2) > \theta_1$; while for $\phi^{-1}(\theta_2) \leq \theta_1$, or equivalently $\theta_2 \geq \phi(\theta_1)$, define C by $C_1(\theta_1) = C_1^*(\theta_1)$ and the pair of differential equations (8) and

$$(1-\beta)\,u_2'(C_2(\theta_1,\theta_2);\theta_2)\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2} = \max\left\{0, \frac{\partial}{\partial\theta_2}u_1\left(C_1(\theta_1);\phi^{-1}(\theta_2)\right) + \frac{\partial}{\partial\theta_2}u_2(C_2(\theta_1,\theta_2);\theta_2) - \frac{d}{d\theta_2}U^1\left(C^*\left(\phi^{-1}(\theta_2),\theta_2\right);\phi^{-1}(\theta_2),\theta_2\right)\right\}, \tag{A-22}$$

subject to the boundary condition that $C(\theta_1,\theta_2) = C^*(\theta_1,\phi(\theta_1))$ at $\theta_2 = \phi(\theta_1)$.

The differential equations (8) and (A-22) imply that, for any $\tilde{\theta}_1$ and $\theta_2 \geq \phi\left(\tilde{\theta}_1\right)$,

$$\frac{d}{d\theta_2}U^1\left(C\left(\tilde{\theta}_1,\theta_2\right);\phi^{-1}(\theta_2),\theta_2\right) \leq \frac{d}{d\theta_2}U^1\left(C^*\left(\phi^{-1}(\theta_2),\theta_2\right);\phi^{-1}(\theta_2),\theta_2\right).$$

Given the boundary condition, it follows that, for any $\tilde{\theta}_1$ and $\theta_2 \geq \phi\left(\tilde{\theta}_1\right)$,

$$U^1\left(C^*\left(\tilde{\theta}_1,\theta_2\right);\phi^{-1}(\theta_2),\theta_2\right) \leq U^1\left(C^*\left(\phi^{-1}(\theta_2),\theta_2\right);\phi^{-1}(\theta_2),\theta_2\right).$$

Changing variables $\theta_1 = \phi^{-1}(\theta_2)$, for any $\tilde{\theta}_1$ and $\theta_1 \leq \tilde{\theta}_1$,

$$U^1\left(C^*\left(\tilde{\theta}_1,\phi(\theta_1)\right);\theta_1,\phi(\theta_1)\right) \leq U^1\left(C^*(\theta_1,\phi(\theta_1));\theta_1,\phi(\theta_1)\right).$$

Hence $IC_1\left(\theta_1,\tilde{\theta}_1\right)$ holds for any θ_1 and $\tilde{\theta}_1 > \theta_1$.

Next, we show that $IC_1\left(\theta_1,\tilde{\theta}_1\right)$ holds for any θ_1 and $\tilde{\theta}_1 < \theta_1$. In this case, $C\left(\tilde{\theta}_1,\phi(\theta_1)\right) = C^*\left(\tilde{\theta}_1,\phi\left(\tilde{\theta}_1\right)\right)$ since $\phi^{-1}(\phi(\theta_1)) > \tilde{\theta}_1$, while $C(\theta_1,\phi(\theta_1)) = C^*(\theta_1,\phi(\theta_1))$. Hence we must show that $U^1(C^*(\theta_1,\phi(\theta_1));\theta_1,\phi(\theta_1)) \geq U^1\left(C^*\left(\tilde{\theta}_1,\phi\left(\tilde{\theta}_1\right)\right);\theta_1,\phi(\theta_1)\right)$. By the ordering of Θ_1, $C_1^*\left(\tilde{\theta}_1\right) \leq C_1^*(\theta_1)$, and by the optimality of C^*, $V^1(C^*(\theta_1,\phi(\theta_1));\theta_1,\phi(\theta_1)) \geq V^1\left(C^*\left(\tilde{\theta}_1,\phi\left(\tilde{\theta}_1\right)\right);\theta_1,\phi(\theta_1)\right)$. By the obvious analogue of Lemma A-1 for self 1's preferences, the result then follows.

Finally, RC is certainly satisfied for $\phi^{-1}(\theta_2) \geq \theta_1$, or equivalently, $\theta_2 \leq \phi(\theta_1)$. For

$\theta_2 > \phi(\theta_1)$, observe that, by (8),

$$\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2} + \frac{\partial C_3(\theta_1,\theta_2)}{\partial\theta_2} = \left(1 - \frac{1}{\beta}\frac{u_2'(C_2(\theta_1,\theta_2);\theta_2)}{u_3'(C_3(\theta_1,\theta_2))}\right)\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}.$$

At $\theta_2 = \phi(\theta_1)$, $u_2'(C_2(\theta_1,\theta_2);\theta_2) = u_3'(C_3(\theta_1,\theta_2))$ by the definition of C^*, so the term $1 - \frac{1}{\beta}\frac{u_2'(C_2(\theta_1,\theta_2);\theta_2)}{u_3'(C_3(\theta_1,\theta_2))}$ is strictly negative. The condition stated in Part (B) ensures that this expression remains negative for all $\theta_2 \in \left(\phi(\theta_1), \bar{\theta}_2\right)$. By construction, $\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2} \geq 0$. Hence $\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2} + \frac{\partial C_3(\theta_1,\theta_2)}{\partial\theta_2} \leq 0$, which implies RC is satisfied for all $\theta_2 \in (\phi(\theta_1), \bar{\theta}_2]$.
QED

Proof of Proposition 6: To establish the result, we show that Conditions (I)-(III) imply that, for all $\theta_2 \geq \phi(\theta_1)$, the contract C constructed in the proof of Proposition 5 satisfies

$$\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2} + \frac{\partial C_3(\theta_1,\theta_2)}{\partial\theta_2} \leq 0. \tag{A-23}$$

Since $C(\theta_1, \phi(\theta_1)) = C^*(\theta_1, \phi(\theta_1))$, this is sufficient to establish that C satisfies RC.

Preliminaries: Let γ be the relative risk aversion associated with u. By Condition (III), assume that

$$\beta > n^{-\gamma}\left(\left(\frac{1+(n+1)\frac{\partial\zeta\left(\phi^{-1}(\theta_2)\right)}{\partial\theta_2}}{n+2}\right)^{-1} - 1\right)^{-\gamma}. \tag{A-24}$$

(Note that by Condition (II), the RHS of (A-24) is strictly below 1.) Next, differentiation yields

$$\begin{aligned}\frac{d}{d\theta_2}U^1\left(C^*\left(\phi^{-1}(\theta_2),\theta_2\right);\phi^{-1}(\theta_2),\theta_2\right) &= (1-\beta)\frac{\partial}{\partial\theta_2}C_1^*\left(\phi^{-1}(\theta_2)\right)u_1'\left(C_1^*\left(\phi^{-1}(\theta_2)\right);\phi^{-1}(\theta_2)\right)\\ &+ \frac{\partial\phi^{-1}(\theta_2)}{\partial\theta_2}\frac{\partial}{\partial\theta_1}u_1\left(C_1^*\left(\phi^{-1}(\theta_2)\right);\theta_1\right)\Big|_{\theta_1=\phi^{-1}(\theta_2)}\\ &+ \beta\frac{\partial}{\partial\theta_2}u_2\left(C_2^*\left(\phi^{-1}(\theta_2),\theta_2\right);\theta_2\right).\end{aligned}$$

Substituting into (A-22) of the proof of Proposition 5,

$$
\begin{aligned}
& (1-\beta)\, u_2'\left(C_2(\theta_1,\theta_2);\theta_2\right) \frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2} \\
= \;& \frac{\partial}{\partial\theta_2} u_1\left(C_1(\theta_1);\phi^{-1}(\theta_2)\right) + \beta \frac{\partial}{\partial\theta_2} u_2\left(C_2(\theta_1,\theta_2);\theta_2\right) \\
- \;& (1-\beta)\frac{\partial}{\partial\theta_2} C_1^*\left(\phi^{-1}(\theta_2)\right) u_1'\left(C_1^*\left(\phi^{-1}(\theta_2)\right);\phi^{-1}(\theta_2)\right) \\
- \;& \frac{\partial\phi^{-1}(\theta_2)}{\partial\theta_2}\frac{\partial}{\partial\theta_1}\, u_1\left(C_1^*\left(\phi^{-1}(\theta_2)\right);\theta_1\right)\Big|_{\theta_1=\phi^{-1}(\theta_2)} - \beta\frac{\partial}{\partial\theta_2} u_2\left(C_2^*\left(\phi^{-1}(\theta_2),\theta_2\right);\theta_2\right) \\
= \;& \frac{\partial\zeta\left(\phi^{-1}(\theta_2)\right)}{\partial\theta_2} u'\left(C_1(\theta_1)+\zeta\left(\phi^{-1}(\theta_2)\right)\right) - \beta u'\left(C_2(\theta_1,\theta_2)-\theta_2\right) \\
- \;& (1-\beta)\frac{\partial}{\partial\theta_2} C_1^*\left(\phi^{-1}(\theta_2)\right) u'\left(C_1^*\left(\phi^{-1}(\theta_2)\right)+\zeta\left(\phi^{-1}(\theta_2)\right)\right) \\
- \;& \frac{\partial\phi^{-1}(\theta_2)}{\partial\theta_2}\zeta'\left(\phi^{-1}(\theta_2)\right) u'\left(C_1^*\left(\phi^{-1}(\theta_2)\right)+\zeta\left(\phi^{-1}(\theta_2)\right)\right) + \beta u'\left(C_2^*\left(\phi^{-1}(\theta_2),\theta_2\right)-\theta_2\right),
\end{aligned}
$$

where the second inequality follows from Condition (I). By the preference reversal condition, $\phi^{-1}(\theta_2) < \theta_1$ for any $\theta_2 > \phi(\theta_1)$, and hence $C_1^*\left(\phi^{-1}(\theta_2)\right) < C_1(\theta_1)$ for any $\theta_2 > \phi(\theta_1)$. By Condition (II), $\frac{\partial\zeta\left(\phi^{-1}(\theta_2)\right)}{\partial\theta_2} \geq 0$, and so (making use of the optimality of C^*)

$$
\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2} \leq -\frac{\beta}{1-\beta} - \frac{u'\left(C_2^*\left(\phi^{-1}(\theta_2),\theta_2\right)-\theta_2\right)}{u'\left(C_2(\theta_1,\theta_2)-\theta_2\right)}\left(\frac{\partial}{\partial\theta_2}C_1^*\left(\phi^{-1}(\theta_2)\right) - \frac{\beta}{1-\beta}\right). \tag{A-25}
$$

Moreover, given Condition (I) it is straightforward to solve explicitly for C^*:

$$
\begin{aligned}
C_1^*\left(\phi^{-1}(\theta_2)\right) &= \frac{1}{n+2}\left(W-\theta_2+\zeta\left(\phi^{-1}(\theta_2)\right)\right) - \zeta\left(\phi^{-1}(\theta_2)\right) \\
C_2^*\left(\phi^{-1}(\theta_2),\theta_2\right) &= \frac{1}{n+2}\left(W-\theta_2+\zeta\left(\phi^{-1}(\theta_2)\right)\right) + \theta_2 \\
C_3^*\left(\phi^{-1}(\theta_2),\theta_2\right) &= \frac{n}{n+2}\left(W-\theta_2+\zeta\left(\phi^{-1}(\theta_2)\right)\right).
\end{aligned}
$$

Hence (using Condition (II))

$$\frac{\partial}{\partial\theta_2}\left(C_1^*\left(\phi^{-1}\left(\theta_2\right)\right)+C_2^*\left(\phi^{-1}\left(\theta_2\right),\theta_2\right)\right)=\frac{n}{n+2}\left(1-\frac{\partial}{\partial\theta_2}\zeta\left(\phi^{-1}\left(\theta_2\right)\right)\right)>0. \qquad \text{(A-26)}$$

Claim 1: $C_2\left(\theta_1,\theta_2\right)\leq C_2^*\left(\phi^{-1}\left(\theta_2\right),\theta_2\right)$ for $\theta_2\geq\phi\left(\theta_1\right)$.

Proof of Claim 1: At $\theta_2=\phi\left(\theta_1\right)$, $C_2\left(\theta_1,\theta_2\right)=C_2^*\left(\phi^{-1}\left(\theta_2\right),\theta_2\right)$, and so from (A-25) and (A-26),

$$\frac{\partial C_2\left(\theta_1,\theta_2\right)}{\partial\theta_2}\leq-\frac{\partial}{\partial\theta_2}C_1^*\left(\phi^{-1}\left(\theta_2\right)\right)<\frac{\partial}{\partial\theta_2}C_2^*\left(\phi^{-1}\left(\theta_2\right),\theta_2\right).$$

Hence the claimed inequality holds strictly for θ_2 close to $\phi\left(\theta_1\right)$. Suppose now that the claim is not true. By continuity, it follows that there is some $\theta_2>\phi\left(\theta_1\right)$ such that $C_2\left(\theta_1,\theta_2\right)=C_2^*\left(\phi^{-1}\left(\theta_2\right),\theta_2\right)$ and $\frac{\partial}{\partial\theta_2}\left(C_2\left(\theta_1,\theta_2\right)-C_2^*\left(\phi^{-1}\left(\theta_2\right),\theta_2\right)\right)\geq 0$. But then (A-25) and and (A-26) again imply $\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}<\frac{\partial}{\partial\theta_2}C_2^*\left(\phi^{-1}\left(\theta_2\right),\theta_2\right)$. The contradiction completes the proof of the claim.

Claim 2: $\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}\leq-\frac{\partial}{\partial\theta_2}C_1^*\left(\phi^{-1}\left(\theta_2\right)\right)$.

Proof of Claim 2: Since $\frac{\partial}{\partial\theta_2}C_1^*\left(\phi^{-1}\left(\theta_2\right)\right)<0$, it follows from Claim 1 that $\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}\leq-\frac{\beta}{1-\beta}-\left(\frac{\partial}{\partial\theta_2}C_1^*\left(\phi^{-1}\left(\theta_2\right)\right)-\frac{\beta}{1-\beta}\right)$.

Claim 3: $u_2'\left(C_2\left(\theta_1,\theta_2\right);\theta_2\right)-\beta u_3'\left(C_3\left(\theta_1,\theta_2\right)\right)\geq 0$ for $\theta_2\geq\phi\left(\theta_1\right)$.

Proof of Claim 3: Since $C\left(\theta_1,\phi\left(\theta_1\right)\right)=C^*\left(\theta_1,\phi\left(\theta_1\right)\right)$, the claimed inequality holds strictly at $\theta_2=\phi\left(\theta_1\right)$. Suppose now that the claim is not true. By continuity, it follows that there is some $\theta_2>\phi\left(\theta_1\right)$ such that $u_2'\left(C_2\left(\theta_1,\theta_2\right);\theta_2\right)=\beta u_3'\left(C\left(\theta_1,\theta_2\right)\right)$ and $\frac{\partial}{\partial\theta_2}\left(u_2'\left(C_2\left(\theta_1,\theta_2\right);\theta_2\right)-\beta u_3'\left(C_3\left(\theta_1,\theta_2\right)\right)\right)\leq 0$, i.e.,

$$\left(\frac{\partial C_2\left(\theta_1,\theta_2\right)}{\partial\theta_2}-1\right)u_2''\left(C_2\left(\theta_1,\theta_2\right);\theta_2\right)-\beta\frac{\partial C_3\left(\theta_1,\theta_2\right)}{\partial\theta_2}u_2''\left(C_3\left(\theta_1,\theta_2\right)\right)\leq 0.$$

Substituting (8) into this inequality delivers

$$\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}\left(1+\frac{\frac{u_3''(C_3(\theta_1,\theta_2))}{u_3'(C_3(\theta_1,\theta_2))}}{\frac{u_2''(C_2(\theta_1,\theta_2);\theta_2)}{u_2'(C_2(\theta_1,\theta_2);\theta_2)}}\right)\geq 1.$$

Under Condition (I), this inequality is in turn equivalent to

$$\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}\left(1+\frac{\frac{1}{n}\gamma\left(\frac{1}{n}C_3(\theta_1,\theta_2)\right)^{-1}}{\gamma\left(C_2(\theta_1,\theta_2)-\theta_2\right)^{-1}}\right)\geq 1.$$

Since $u'(C_2(\theta_1,\theta_2)-\theta_2)=\beta u'\left(\frac{1}{n}C_3(\theta_1,\theta_2)\right)$, it follows that $\frac{n^{\gamma}C_3(\theta_1,\theta_2)^{-\gamma}}{(C_2(\theta_1,\theta_2)-\theta_2)^{-\gamma}}=\beta^{-1}$, and hence the above inequality is equivalent to

$$\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}\left(1+\frac{1}{n}\beta^{-\frac{1}{\gamma}}\right)\geq 1.$$

Inequality (A-24) and Claim 2 imply

$$\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}\left(1+\frac{1}{n}\beta^{-\frac{1}{\gamma}}\right)<\left(-\frac{\partial C_1^*\left(\phi^{-1}(\theta_2)\right)}{\partial\theta_2}\right)\left(\frac{1+(n+1)\frac{\partial\zeta\left(\phi^{-1}(\theta_2)\right)}{\partial\theta_2}}{n+2}\right)^{-1}=1.$$

The contradiction completes the proof of Claim 3.

Completing the proof: Finally, (A-23) follows from the fact that (8) and Claim 3 combine to give

$$\begin{aligned}0 &= u_2'(C_2(\theta_1,\theta_2);\theta_2)\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}+\beta u_3'(C_3(\theta_1,\theta_2))\frac{\partial C_3(\theta_1,\theta_2)}{\partial\theta_2}\\ &\geq \beta u_3'(C_3(\theta_1,\theta_2))\frac{\partial C_2(\theta_1,\theta_2)}{\partial\theta_2}+\beta u_3'(C_3(\theta_1,\theta_2))\frac{\partial C_3(\theta_1,\theta_2)}{\partial\theta_2}.\end{aligned}$$

QED

Table 1: Results of numerical simulations, as described in Section 7

Panel (A): One-period ahead shocks

		$\bar{\theta}_2 = 0.05$			$\bar{\theta}_2 = 0.1$		
	γ	1	2	4	1	2	4
(i)	β^*	0.951	0.905	0.822	0.903	0.818	0.679
(ii)	Cutoff $\tilde{\beta}$ for $\beta = 0.8$	0.840	0.885	0.980	0.887	0.982	n/a
(iii)	Cutoff $\tilde{\beta}$ for $\beta = 0.7$	0.736	0.774	0.858	0.776	0.861	n/a

Panel (B): Timing shocks

		$\bar{\theta}_2 = 0.05$			$\bar{\theta}_2 = 0.1$		
	γ	1	2	4	1	2	4
(i)	β^*	0.950	0.903	0.816	0.902	0.813	0.660
(ii)	Cutoff $\tilde{\beta}$ for $\beta = 0.8$	0.841	0.885	0.980	0.887	0.984	n/a
(iii)	Cutoff $\tilde{\beta}$ for $\beta = 0.7$	0.736	0.773	0.857	0.775	0.861	n/a

Panel (C): Investment

		$R(\underline{\theta}_2) = 1.025$		$R(\underline{\theta}_2) = 1.05$	
	γ	2	4	2	4
(i)	β^*	0.858	0.801	0.747	0.657
(ii)	Cutoff $\tilde{\beta}$ for $\beta = 0.8$	0.884	0.977	n/a	n/a
(iii)	Cutoff $\tilde{\beta}$ for $\beta = 0.7$	0.773	0.854	0.854	n/a