Commitment contracts*

Philip Bond† Gustav Sigurdsson‡

February 16, 2017

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. Moreover, there are conditions under which preference reversal naturally arises. 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. Tarun Patel provided excellent research assistance.

†University of Washington.

‡Citigroup Global Markets.

1 Introduction

Preferences with hyperbolic time discounting, introduced by Strotz (1956), are widely used to model 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 future consumption needs. In such cases, an individual will be reluctant to commit to future consumption levels that are state-independent, and 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 that desired consumption at date 1 is negatively correlated with marginal utility (MU) at

[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 2005 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.

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.

Example 1: An individual is about to retire. At date 0, he has accumulated savings of $3\frac{1}{2}$, and anticipates a retirement spanning dates $1, 2, 3$. He also anticipates that at either date 1 or 2, he will have to make an essential home repair, at cost $\frac{1}{2}$. He has log preferences over consumption (minus any home repair expenditure) at each date. So to equalize MU across dates, self 0 wants consumption of $\frac{3}{2}, 1, 1$ at dates $1, 2, 3$ respectively if the home repair is needed at date 1, but consumption $1, \frac{3}{2}, 1$ if the home repair is needed at date 2. Self 0's problem is that, because of hyperbolic discounting, self 1 prefers the consumption stream

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

$\frac{3}{2}, 1, 1$ even if the home repair is not needed until date 2.[5]

Self 0 can overcome this problem as follows: he deposits $\frac{1}{2}$ into a liquid checking account, and uses 3 to buy an annuity that pays 1 at each of dates $1, 2, 3$. Importantly, the annuity contains an early redemption option: $\frac{2}{3}$ of the date 3 payment can be accessed at date 2, at a cost of $\frac{1}{6}$, so that annuity payments at the three dates become $1, \frac{3}{2}, \frac{1}{3}$.

To see why this arrangement achieves self 0's desired consumption, suppose self 1 consumes both the $\frac{1}{2}$ in the checking account and the 1 paid by the annuity, and consider self 2's behavior. If self 2 has to make the home repair then he exercises the annuity's early redemption option, consuming $\frac{3}{2}$ and leaving $\frac{1}{3}$ to self 3.[6] But if self 1 already made the home repair then self 2 finds early redemption too costly, despite his present-bias.[7]

Anticipating self 2's behavior, if self 1 does not need to pay for the home repair then he does not withdraw from the checking account, since he understands that if he does, self 2 will exercise the option of costly early redemption, hurting self 3.[8] Self 2 then liquidates the checking account at date 2 to cover the home repair. In contrast, if self 1 needs to pay for the home repair, then he liquidates the checking account to pay for it. Note that self 2 never exercises the early redemption option in equilibrium.

In essence, the conflict between selves 0 and 1 is resolved by the new conflict that the early redemption option creates between selves 1 and 2.

Example 2: In Example 1, the individual is able to commit to self 0's desired consumption plan. But suppose now that the expenditure shock takes a different form: either the individual must pay home repairs of $\frac{1}{4}$ at *both* dates 1 and 2, or else no home repair is necessary at either date. Consequently, self 0 wants to commit to consumption $\frac{5}{4}, \frac{5}{4}, 1$ if the home

[5]Formally, $\log\frac{3}{2} + \beta\log\left(1 - \frac{1}{2}\right) + \beta\log 1 = \log\left(\frac{3}{2\sqrt{2}}\right) > \log 1 + \beta\log\left(\frac{3}{2} - \frac{1}{2}\right) + \beta\log 1$.

[6]Formally, $\log\left(\frac{3}{2} - \frac{1}{2}\right) + \beta\log\frac{1}{3} = \log\left(\frac{1}{\sqrt{3}}\right) > \log\left(1 - \frac{1}{2}\right) + \beta\log 1$.

[7]Formally, $\log 1 + \beta\log 1 > \log\frac{3}{2} + \beta\log\frac{1}{3} = \log\frac{3}{2\sqrt{3}}$.

[8]Formally, $\log 1 + \beta\log\left(\frac{3}{2} - \frac{1}{2}\right) + \beta\log 1 > \log\frac{3}{2} + \beta\log\left(\frac{3}{2} - \frac{1}{2}\right) + \beta\log\frac{1}{3} = \log\left(\frac{3}{2\sqrt{3}}\right)$. Note that we compare only a full withdrawal of $\frac{1}{2}$ from the checking account with no withdrawal at all. Intermediate cases are covered by the analysis of Section 6.

repairs are needed, but to $\frac{7}{6}, \frac{7}{6}, \frac{7}{6}$ otherwise. Similar to Example 1, hyperbolic discounting causes self 1 to prefer $\frac{5}{4}, \frac{5}{4}, 1$ even if no repairs are needed, because this consumption stream offers higher date-1 consumption.[9]

The key to combining commitment and flexibility in Example 1 is that it is possible to offer self 2 an alternative consumption stream that he chooses if and only if self 1 overconsumed, and that also hurts self 1. But this is impossible in Example 2, as follows. First, the only consumption streams that raise self 2's utility while lowering self 1's utility are those that increase date 2 consumption but decrease date 3 consumption. But second, if self 2 prefers this stream when no repairs are needed—which is when self 1 has overconsumed in Example 2—then he also prefers such a stream when repairs are needed, since MU at date 2 is higher in this case. Consequently, it is impossible to impose a state-contingent punishment on self 1 for overconsumption. This point is formalized in Lemma 2 below.

Discussion: The key distinction between the two examples is that in Example 1 self 0's desired date 1 consumption 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 combine commitment and flexibility, and to attain exactly the outcome desired by self 0. In contrast, this is impossible if preference reversal is violated.

Closely related, commitment and flexibility can be combined in Example 1 because the realization of uncertainty at date 1 reveals something about self 2's preferences: specifically, if self 1 does not need to make the home repair then he knows that MU at date 2 will be high, and consequently, that self 2 will be prepared to punish him for overconsumption at date 1 by increasing date 2 consumption at the expense of date 3. If instead self 1 learns nothing about self 2's preferences, then commitment and flexibility cannot be combined.

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

Under preference reversal, the key qualitative feature of contracts that allow self 0 to combine commitment with flexibility is that they increase self 2's discretion relative to the discretion that he would have 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 self 2's discretion.

Examples 1 and 2 are consumption-savings problems with uncertainty over future essential expenditures. They are isomorphic to settings in which uncertainty is instead over future income. Consumption-savings problems with uncertainty over taste shocks are also closely related. Moreover, by switching to consumption of leisure instead of physical goods, Examples 1 and 2 can instead be interpreted as procrastination problems, a setting heavily studied in the literature (see footnote 1). Finally, Laibson (1996) considers uncertainty over investment returns as well as over income. Variation in rates of return directly induces correlation between self 0's desired date 1 consumption and MU at subsequent dates. We illustrate this with our final example, which concerns health "investments:"

Example 3: An individual's date 3 health is determined by how much he exercises at each of dates 1 and 2, denoted i_1 and i_2. However, the benefits of exercise are lower if he has a cold at date 1: his date 3 utility is $\log(i_1 + i_2)$ if he is cold-free, but only $\log\left(\frac{4}{5}i_1 + i_2\right)$ if he has a cold. He dislikes exercise: his utility at dates $t = 1, 2$ is $\log(1 - i_t)$. Hence self 0 would like to commit to an exercise regime of $i_1 = i_2 = \frac{1}{3}$ if he is cold-free at date 1, but to a regime of $i_1 = \frac{1}{4}$ and $i_2 = \frac{2}{5}$ if he has a cold. Self 0's problem is that self 1 prefers the latter exercise regime even when he is cold-free.[10]

Self 0 can overcome this problem as follows: he publicly announces (e.g., via a smart-

[10]Formally, $\log\left(1-\frac{1}{4}\right) + \beta\log\left(1-\frac{2}{5}\right) + \beta\log\left(\frac{1}{4}+\frac{2}{5}\right) = \log\left(\frac{3}{4}\sqrt{\frac{3}{5}\frac{13}{20}}\right) > \log\left(1-\frac{1}{3}\right) + \beta\log\left(1-\frac{1}{3}\right) + \beta\log\left(\frac{1}{3}+\frac{1}{3}\right) = \log\frac{4}{9}$.

phone app such as Strava) an exercise goal of $\frac{1}{3}$ at each of dates 1 and 2, with the understanding that if he exercises just $i_1 = \frac{1}{4}$ and so misses the date 1 goal, he can either largely "catch up" at date 2 by exercising $i_2 = \frac{2}{5}$ (note that $\frac{1}{3} + \frac{1}{3} \approx \frac{1}{4} + \frac{2}{5}$), or else he can quit and do no exercise at date 2, where quitting has a psychic cost of $\frac{1}{50}$ (e.g., the cost of missing the publicly announced goal).

To see why this achieves self 0's desired exercise regime, note that self 2 quits if and only if self 1 cheated by underexercising when cold-free.[11] The reason is that the health cost of quitting is lower when he was cold-free at date 1, since any exercise he did at that date was more effective in creating health "capital." Foreseeing this, self 1 meets his exercise goal of $\frac{1}{3}$ at date 1 when cold-free, since he does not want self 2 to quit.[12]

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."[13] 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 $\log \frac{3}{2} - \log 1$ 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

[11] Formally, $\log\left(1 - \frac{1}{50}\right) + \beta \log \frac{1}{4} > \log\left(1 - \frac{2}{5}\right) + \beta \log\left(\frac{1}{4} + \frac{2}{5}\right)$ and $\log\left(1 - \frac{2}{5}\right) + \beta \log\left(\frac{4}{5}\frac{1}{4} + \frac{2}{5}\right) > \log\left(1 - \frac{1}{50}\right) + \beta \log \frac{4}{5}\frac{1}{4}$.

[12] Formally, $\log\left(1 - \frac{1}{3}\right) + \beta \log\left(1 - \frac{1}{3}\right) + \beta \log\left(\frac{1}{3} + \frac{1}{3}\right) > \log\left(1 - \frac{1}{4}\right) + \beta \log\left(1 - \frac{1}{50}\right) + \beta \log \frac{1}{4}$.

[13] 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.

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.[14]

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 state-contingent punishment on self 1 for deviating—the key feature of our setting—because with two dates self 1 is effectively the only strategic agent.[15] 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, and characterize the contract that maximizes the profits of a monopolist counterparty facing a partially naïve agent (Section 7 below discusses partial naïveté). In particular, they characterize the combination of flat upfront fees and per-usage fees in the profit-maximizing contract.[16]

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

[14] 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 illustrated by Example 1, and is central to the design of contracts in our paper.

[15] Amador et al (2003) extend the analysis to three or more dates. They assume that shocks are independent across dates (see subsection 4.4 below), so that self 1 learns nothing about self 2's preferences.

[16] 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.

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. Krusell and Smith (2003) and Bernheim, Ray, and Yeltekin (2015) consider deterministic models in which an individual is infinitely lived, and show that equilibria exist in which the individual gains some commitment ability from the fact that deviations will cause future selves to punish him. Carrillo and Mariotti (2000) and Bénabou and Tirole (e.g., 2002, 2004) 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 $\theta_t \in \Theta_t$. At date 3, contemporaneous utility is $u_3(c_3)$. Note that $u_3(c_3)$ can be interpreted as a value function covering multiple future dates, i.e., the expected discounted future utility of an agent inheriting wealth c_3.[17] The contemporaneous utility functions u_t are strictly increasing and strictly concave in c_t. The total resources available to the agent across the three dates are W, and are state-independent.[18] This could either represent an initial endowment of the agent, or the present value of future income.

[17] As stated, the value function interpretation of u_3 depends on the distributions of any shocks from date 3 onwards being independent of θ_1 and θ_2. More generally, one could allow the distributions of future shocks to depend on θ_2, and write the date 3 value function as $u_3(c_3; \theta_2)$. In this case, one would then replace date 2 MU in our analysis with the ratio $\frac{u_2'(c_2;\theta_2)}{u_3'(c_3;\theta_2)}$.

[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.

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. 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). Finally, write $V^t = \sum_{s=t}^{3} u_s$ for utility under exponential discounting.

To summarize: the economy is defined by a state space $\Theta_1 \times \Theta_2$, a probability distribution over $\Theta_1 \times \Theta_2$, preferences $\{u_t\}$, endowment W and discount rate β.

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. Since our focus is on the effect of hyperbolic discounting on intertemporal efficiency, not its effect on insurance across states, we rule out transfers across states and impose the following resource constraint:

$$C_1(\theta_1) + C_2(\theta_1, \theta_2) + C_3(\theta_1, \theta_2) \leq W \text{ for all } (\theta_1, \theta_2) \in \Theta_1 \times \Theta_2. \tag{RC}$$

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 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)

[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.

constraints, which ensure that no self can gain by misrepresenting the state.[21] Self 2's IC constraints are: for all $\theta_1 \in \Theta_1$ and $\theta_2, \tilde{\theta}_2 \in \Theta_2$,

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

Self 1's IC constraints are: for all $\theta_1, \tilde{\theta}_1 \in \Theta_1$,

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

3.2 Benchmark: State θ_1 verifiable

Unverifiability of the state induces a potential trade-off between commitment and flexibility for self 0, 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

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

IC_1 is non-binding in (1) if self 0's utility is the same in (1) and the benchmark relaxed problem in which θ_1 is verifiable,

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

[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 $C\left(\theta_1, \theta_2, \theta_{21}\right)$. However, given that self 2's preferences are independent of state θ_1, the only way for a contract with $C\left(\theta_1, \theta_2, \theta_{21}\right) \neq C\left(\theta_1, \theta_2, \tilde{\theta}_{21}\right)$ to be incentive compatible is if self 2 is indifferent between $C\left(\theta_1, \theta_2, \theta_{21}\right)$ 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.

We focus on combining commitment and flexibility with respect to date 1 shocks because the analogous trade-off at date 2 corresponds to the two-date problem analyzed by Amador et al (2006), who establish that there is no way to fully resolve the date 2 conflict.

Definition 1 *"Commitment and flexibility can be combined" if and only if IC_1 is non-binding in (1).*

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]

Example, additive shocks: $u_t(c_t; \theta_t) = u(c_t - \theta_t)$ for $t = 1, 2$, where u has non-increasing absolute risk aversion (NIARA). This shock specification has a natural interpretation as either essential expenditure shocks, as in Examples 1 and 2, or as income shocks.

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

Assumption 1 *For $t = 1, 2$ and $\theta_t \neq \tilde{\theta}_t \in \Theta_t$, sign $\left(u_t'(c_t; \theta_t) - u_t'\left(c_t; \tilde{\theta}_t\right)\right)$ is independent of c_t.*

Given Assumption 1, 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 $\bar{\theta}_2$ and $\underline{\theta}_2$ for the maximal and minimal elements of Θ_2 under this ordering. Without loss, we assume $u_2'\left(\cdot; \tilde{\theta}_2\right) \neq u_2'(\cdot; \theta_2)$ if $\tilde{\theta}_2 \neq \theta_2$.

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):

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

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 and standard monotonicity result:[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 $\theta_1 \in \Theta_1$.*

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. 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

Our earlier examples illustrate that commitment and flexibility can be combined if states in which self 0 wants high date 1 consumption are followed by low date 2 MU, a condition

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

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

we term as preference reversal. Formalizing and generalizing:

Definition 2 *Preference reversal holds if for some C solving self 0's benchmark problem (2), then for any $\theta_1, \tilde{\theta}_1 \in \Theta_1$ with $C_1\left(\tilde{\theta}_1\right) > C_1(\theta_1)$, the distribution of θ_2 conditional on θ_1 strictly first-order stochastically dominates the distribution of θ_2 conditional on $\tilde{\theta}_1$.*

We establish that preference reversal[25] is necessary for the combination of commitment and flexibility, and give conditions under which it is also sufficient. We also characterize how preference reversal affects the characteristics of the contract maximizing self 0's utility (problem (1)). 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 assuming that the shocks θ_1 and θ_2 are perfectly correlated, and so θ_1 perfectly forecasts θ_2. Hence self 0's benchmark problem (2) has a solution in which consumption at all dates is determined solely by self 1's report, i.e., $C(\theta_1, \cdot)$ is constant, and self 2 has no discretion. We denote this *minimum discretion* solution to (2) by C^*.

Absent hyperbolic discounting (i.e., $\beta = 1$), self 1 and 2's preferences coincide, and so the minimum discretion contract C^* satisfies IC_1. More generally, let β^* be the cutoff hyperbolic discount rate for which C^* satisfies IC_1:

$$\beta^* = \inf\left\{\tilde{\beta} \in (0,1]\text{: } C^* \text{ satisfies } \text{IC}_1 \text{ for all } \beta \geq \tilde{\beta}\right\}. \quad (3)$$

Hence for $\beta \geq \beta^*$, self 0 can maximize utility by having self 1 make all decisions, with no

[25]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 $(\theta_1', \theta_2') \neq (\theta_1, \theta_2)$, $U^t(C(\theta_1, \theta_2); \theta_1, \theta_2) \geq U^t(x; \theta_1, \theta_2)$ and $U^t(C(\theta_1, \theta_2); \theta_1', \theta_2') < U^t(x; \theta_1', \theta_2')$ 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

discretion granted to self 2, and so commitment and flexibility are easily combined.

Our first result is Proposition 1, which establishes the point illustrated by our opening examples, namely that even when $\beta < \beta^*$ it may still be possible to combine commitment and flexibility, provided that preference reversal holds. Importantly, doing so requires departing from the minimum discretion contract C^*, and giving self 2 additional discretion.

Proposition 1 *Suppose $\beta < \beta^*$. Commitment and flexibility can be combined only if preference reversal holds. Moreover, if preference reversal holds, there exists $\hat{\beta} < \beta^*$ such that commitment and flexibility can be combined if $\beta \geq \hat{\beta}$.*

Given perfect correlation, the following corollary of Proposition 1 is immediate:

Corollary 1 *If preference reversal holds and $\beta \geq \hat{\beta}$ then self* 0 *attains the same utility as if both θ_1 and θ_2 were verifiable.*

We next discuss the economics behind Proposition 1. To aid exposition, we write $\bar{\theta}_1$ and $\underline{\theta}_1$ for the realizations of θ_1 such that $C_1^*\left(\bar{\theta}_1\right) > C_1^*\left(\underline{\theta}_1\right)$.[26] Consequently, under perfect correlation, preference reversal simplifies to the condition that $\bar{\theta}_1$ is deterministically followed by $\underline{\theta}_2$, while $\underline{\theta}_1$ is deterministically followed by $\bar{\theta}_2$.

To establish the necessity of preference reversal, suppose that it is violated, i.e., $\bar{\theta}_1$ is deterministically followed by $\bar{\theta}_2$, and $\underline{\theta}_1$ is deterministically followed by $\underline{\theta}_2$, as in Example 2. So IC_1 in state $\underline{\theta}_1$ for a contract C that delivers consumption C^* in equilibrium is

$$u_1\left(C_1^*\left(\underline{\theta}_1\right);\underline{\theta}_1\right) + \beta V^2\left(C^*\left(\underline{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right) \geq u_1\left(C_1^*\left(\bar{\theta}_1\right);\underline{\theta}_1\right) + \beta V^2\left(C\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right). \quad (4)$$

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 consumption state $\underline{\theta}_1$, which is followed

[26] Note that the case of C_1^* constant over Θ_1 is ruled out by the condition $\beta < \beta^*$, since C^* certainly satisfies IC_1 if C_1^* is constant. Also, note that $\bar{\theta}_1$ and $\underline{\theta}_1$ are defined in terms of date 1 consumption, while $\bar{\theta}_2$ and $\underline{\theta}_2$ are defined in terms of MU at date 2.

by $\underline{\theta}_2$ since preference reversal is violated. Because $\beta < \beta^*$, a necessary condition to satisfy (4) 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 low continuation utility in state $\underline{\theta}_2$. However, this is impossible, by the following result which is central to our analysis, and shows that self 2 cannot be induced to impose an effective punishment in date 2 states with low MU.

Lemma 2 *If C satisfies IC_2 then for any θ_1, θ_2 and $\tilde{\theta}_2 \geq \theta_2$,*

$$V^2\left(C\left(\theta_1, \theta_2\right); \theta_2\right) \geq V^2\left(C\left(\theta_1, \tilde{\theta}_2\right); \theta_2\right). \tag{5}$$

The proof of Lemma 2 is short, and we give it here. By Lemma 1, date 2 consumption is lower in the low MU state θ_2, i.e., $C_2\left(\theta_1, \theta_2\right) \leq C_2\left(\theta_1, \tilde{\theta}_2\right)$. So if $V^2\left(C\left(\theta_1, \theta_2\right); \theta_2\right) < V^2\left(C\left(\theta_1, \tilde{\theta}_2\right); \theta_2\right)$, hyperbolic discounting means that self 2 will not pick $C\left(\theta_1, \theta_2\right)$ in state θ_2, since the alternative $C\left(\theta_1, \tilde{\theta}_2\right)$ is both more attractive to a non-hyperbolic agent, and has greater immediate consumption, and so is certainly more attractive to a hyperbolic agent (formally, see Lemma A-1). Hence IC_2 can only hold if inequality (5) holds. Intuitively, the only way to get self 2 to punish self 1 is by having him consume heavily at date 2, but he is prepared to do this only if his MU at date 2 is relatively high.

To summarize, commitment and flexibility can be combined only if self 2 can be induced to punish self 1 after self 1 overconsumes in the low desired consumption state $\underline{\theta}_1$. Preference reversal is necessary for this because otherwise $\underline{\theta}_1$ is followed by a date 2 state with low MU, namely $\underline{\theta}_2$, and by Lemma 2, punishment is impossible in this state.

We next discuss the sufficiency of preference reversal for combining commitment with flexibility. To do so, suppose now that preference reversal holds, as in Example 1. So IC_1 in state $\underline{\theta}_1$ for a contract that implements consumption C^* on the equilibrium path is

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

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 consumption $C_1^*\left(\bar{\theta}_1\right)$ in the low consumption state $\underline{\theta}_1$, which is followed by $\bar{\theta}_2$ by preference reversal. As in the early redemption option of Example 1, one can set these contract terms to induce a conflict between selves 1 and 2, by raising $C_2\left(\bar{\theta}_1, \bar{\theta}_2\right)$ relative to $C_2^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$ and lowering $C_3\left(\bar{\theta}_1, \bar{\theta}_2\right)$ relative to $C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$. Precisely because of hyperbolic discounting, it is possible to do this so that $C\left(\bar{\theta}_1, \bar{\theta}_2\right)$ punishes self 1, i.e., $V^2\left(C\left(\bar{\theta}_1, \bar{\theta}_2\right); \bar{\theta}_2\right) < V^2\left(C^*\left(\bar{\theta}_1, \underline{\theta}_2\right); \bar{\theta}_2\right)$, and self 2 imposes the punishment, i.e., $U^2\left(C\left(\bar{\theta}_1, \bar{\theta}_2\right); \bar{\theta}_2\right) \geq U^2\left(C^*\left(\bar{\theta}_1, \underline{\theta}_2\right); \bar{\theta}_2\right)$.

One still needs to show that the above punishment can be made large enough to satisfy (6). In particular, the requirement that RC hold even off-equilibrium at $\left(\bar{\theta}_1, \bar{\theta}_2\right)$ precludes punishing self 1 by setting $C_3\left(\bar{\theta}_1, \bar{\theta}_2\right)$ very low, while setting $C_2\left(\bar{\theta}_1, \bar{\theta}_2\right)$ very high to satisfy IC_2. Proposition 1 deals with this complication by establishing a result for β close to β^*. In brief, for all $\beta < \beta^*$ 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,[27] Proposition 1 suggests that commitment and flexibility can be combined under preference reversal even when self 0 cannot directly delegate all decisions to self 0 using the minimum discretion contract C^*. Moreover, for some circumstances, the sufficiency half of Proposition 1 can be extended to cover *arbitrary* levels of hyperbolic discounting.

Proposition 2 *Suppose (I)* $\lim_{c \to 0} u_3(c) = -\infty$*, (II) there exists* κ *such that* $u_2\left(c; \bar{\theta}_2\right) \equiv u_3(c - \kappa)$ *for all* c*, and (III)* $C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right) \geq C_3^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$*. Commitment and flexibility can be combined when* $\beta < \beta^*$ *if and only if preference reversal holds.*

Condition (I) of Proposition 2 simply states that it is possible to impose arbitrarily

[27]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 $\beta \approx 0.9$. The length of a time period in these three studies is, respectively, a quarter, a year, and a week.

large utility punishments on self 1. It is satisfied by many utility functions, including, for example, those with constant relative risk aversion (CRRA) greater than 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 the next subsection.

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 *timing shocks*, in which a shock to MU occurs either at date 1 or 2. In Example 1, this arises from uncertainty of the timing of an essential expenditure; alternatively, self 0 may know that he will encounter an attractive consumption opportunity at either date 1 or 2, but not both (e.g., an out-of-town friend will visit at one of dates 1 and 2). Timing shocks generate preference reversal because they leave $C_1^* + C_2^*$ unchanged, and so high MU at date 2 is associated both with higher values of C_2^* and lower values of C_1^*.

A second leading case in which preference reversal holds is that of *one-period ahead shocks*, i.e., at date 1, self 1 learns about a shock that that affects MU at date 2. High date 2 MU raises C_2^* while reducing both C_1^* and C_3^*. In particular, the reduction in C_1^* means that preference reversal holds. 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. Note that condition (III) of Proposition 2 is satisfied by both one-period ahead shocks and timing shocks.

In the public finance interpretation of subsection 3.4, preference reversal holds if date 0 uncertainty centers on the efficacy 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 often arises naturally in investment problems, as illustrated by Example 3.

4.3 Implementation

The key feature of contracts that allow commitment and flexibility to be combined (when $\beta < \beta^*$) is that they grant self 2 discretion that would be unnecessary absent hyperbolic discounting. The specific form of discretion is that self 2 is able to increase date 2 consumption at the cost of decreasing date 3 consumption.

In practical terms, this discretion can be achieved by giving self 2 access to a partially illiquid asset. As Example 1 illustrates, if date 3 consumption comes from an annuity, but early redemption is possible, then this grants self 2 the discretion required. Similarly, if date 3 consumption stems from the use of a durable asset, but this asset can be sold (possibly at a discount) at date 2, then this again grants self 2 the discretion required.

The contracts used to combine commitment and flexibility in the proofs of Propositions 1 and 2 grant self 2 discretion only after high date 1 consumption. However, if condition (III) of Proposition 2 holds, then self 2 can be granted discretion unconditionally:

Lemma 3 *If $C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right) \geq C_3^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$ and commitment and flexibility can be combined, then they combined using a contract C in which self 2 has the option to increase date 2 consumption by some amount X_2 at the expense of decreasing date 3 consumption by some amount $X_3 \geq X_2$ regardless of self 1's consumption choice.*

In some leading cases, self 2's discretion arises even without self 0 making any explicit arrangements. This is the case in environments of the type illustrated by Example 1, in which self 0 allocates an initial endowment over future dates, knowing he will need to incur an essential expenditure at either date 1 or 2 (i.e., a timing shock). In this case, self 0's desired consumption has $C^*\left(\bar{\theta}_1, \underline{\theta}_2\right) = (c + x, c, c)$ and $C^*\left(\underline{\theta}_1, \bar{\theta}_2\right) = (c, c + x, c)$ for some $c, x > 0$. This corresponds to self 0 using $3c$ of his initial endowment to purchase either an annuity or durable assets that deliver c each period, and x to purchase a liquid asset. If the annuity/durable assets can be partially liquidated at date 2, then self 2's discretion arises out of the same arrangement that self 0 uses to spread consumption over time. Moreover, by Lemma 3 self 2 does not exercise this discretion if date 1 consumption is low.

A second class of problems in which self 2's discretion arises absent explicit arrangements is if the agent has a per-period income of w, but anticipates the possibility of a positive income shock of x at date 2, with self 1 learning whether this shock will occur (i.e., a one-period ahead shock). So self 0's desired consumption has $C^*\left(\bar{\theta}_1, \underline{\theta}_2\right) = \left(w + \frac{x}{3}, w - \frac{2x}{3}, w + \frac{x}{3}\right)$ and $C^*\left(\underline{\theta}_1, \bar{\theta}_2\right) = (w, w, w)$. Exactly these consumption streams arise if, when self 1 receives good news about date 2 income, he buys, on credit, a durable asset at price x that yields consumption of $\frac{x}{3}$ per period, and this purchase is made using a one-period debt contract that entails a repayment of x at date 2. And as above, the discretion that must be granted to self 2 to ensure that self 1 does not buy the durable asset on credit when he receives bad news about date 2 income arises almost automatically: once the durable asset is purchased, self 2 can then sell it, potentially at a discount.

For settings in which self 0 must divide an initial endowment across states in the face of one-period ahead shocks, a different but still simple implementation is feasible. First, self 0 deposits into illiquid certificates of deposit (CD) the minimum level of consumption he will need at date 1 (i.e., $C_1^*\left(\underline{\theta}_1\right)$) and the maximum level of consumption he will need at date 3 (i.e., $C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$), with the remainder of his initial endowment (i.e., $W - C_1^*\left(\underline{\theta}_1\right) - C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$) deposited into a date 2 CD. In addition, he arranges for a one-period line of credit of $C_1^*\left(\bar{\theta}_1\right) - C_1^*\left(\underline{\theta}_1\right)$ that may be drawn either at date 1 or 2. It is straightforward to check that if the credit line is drawn at date 1, the resulting consumption stream is $C^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$, while if it is drawn at date 2, the resulting consumption stream is $C^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$.[28]

[28] This calculation makes use of the fact that for one-period ahead shocks, $C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right) - C_3^*\left(\underline{\theta}_1, \bar{\theta}_2\right) = C_1^*\left(\bar{\theta}_1\right) - C_1^*\left(\underline{\theta}_1\right)$. More generally, this construction works for any shock specification if the credit line is made contingent on the date it is drawn, and in particular, is equal to $C_1^*\left(\bar{\theta}_1\right) - C_1^*\left(\underline{\theta}_1\right)$ if drawn at date 1, and to $C_3^*\left(\bar{\theta}_1, \underline{\theta}_2\right) - C_3^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$ if drawn at date 2.

Finally, the date 3 CD is only partially illiquid, and in particular can be accessed (at cost) at date 2, thereby giving self 2 the discretion to punish self 1 for overconsumption.

So to summarize, in many cases the partial illiquidity that gives self 2 the required discretion arises very naturally. It is also worth noting that this partial illiquidity plays a different role than illiquidity plays in commitment schemes discussed in the existing literature (see especially Strotz (1956) and Laibson (1997)). In the prior literature, illiquid assets are used by early selves to bind later selves to a particular consumption stream. In contrast, self 2's access to partially illiquid assets allows him to punish self 1 for overconsumption.[29]

Turning to the procrastination interpretation of our model, commitment and flexibility can be combined with a "suggestive" date 1 deadline that is costless to miss; but if the date 1 deadline is missed, then missing a date 2 deadline is costly, because it raises the total amount of work that must be completed by date 3. This is consistent with norms in many work environments. A more explicit example of such a scheme is a teacher-imposed rule on a permissible number of excused absences, in which absences beyond the excused number are costly to the student.[30] Likewise, some gyms impose a penalty charge on clients who accumulate too many unexcused absences from pre-paid group classes.

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 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 commitment and flexibility can be combined.

In contrast, under imperfect correlation, θ_2 may still be uncertain at date 1, and so self 2 is granted some discretion even in the benchmark problem (2). So here, we analyze

[29] Note that the date 1 and 2 CDs in the last example are fully illiquid, and play the same role as in the prior literature.

[30] The classroom example is consistent with the observation we make in Section 7 that partial naïveté may necessitate a benevolent principal imposing a contract on the individual.

whether it is useful to give self 2 discretion over and above this benchmark.

We must first deal with a complication: the definition of the minimum discretion contract C^* requires more care under imperfect correlation. The reason is that, in solving the benchmark problem (2), 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. As such, to show that the solution to (2) uniquely determines date 1 consumption, we must show the date 2 value function associated with resources $W - C_1(\theta_1)$ is concave. Moreover, because β enters IC_2, the date 2 value function depends on β, and so the solution to (2) depends on β.

The following result takes care of these details, and moreover, establishes continuity with respect to β: recall that the proof of Proposition 1 depends on the magnitude of self 1's temptation to overconsume being continuous as a function of β, and the results below make use of an analogous property.

Lemma 4 *Let C solve problem (2). Date 1 consumption $C_1(\theta_1)$ is uniquely determined and continuous in β. For any state (θ_1, θ_2) such that $\Pr(\theta_1, \theta_2) \neq 0$, utility $V^2(C(\theta_1, \theta_2); \theta_2)$ is likewise uniquely determined and continuous in β.*

Given Lemma 4, define the minimum discretion solution $C^*(\cdot; \beta)$ to the benchmark problem (2) in the same way as for the case of perfect correlation;[31] and then β^* by (3), as before. Moreover, whenever $C_1^*(\cdot; \beta^*)$ is non-constant in θ_1, define $\bar{\theta}_1$ and $\underline{\theta}_1$ as the elements of Θ_1 such that $C_1^*(\bar{\theta}_1; \beta^*) > C_1^*(\underline{\theta}_1; \beta^*)$.

The following very mild assumption, which holds generically in probabilities, is enough to ensure that $\bar{\theta}_1$ and $\underline{\theta}_1$ can be defined whenever θ_1 and θ_2 are non-independent:

[31] That is, for any θ_1 such that the date 2 state following θ_1 is deterministic, let $C^*(\cdot; \beta)$ be the solution to (2) in which $C(\theta_1, \cdot)$ is constant. Note that Lemma 4 ensures that $C_1^*(\theta_1; \beta)$ and $V^2(C^*(\theta_1, \theta_2; \beta); \theta_2)$ are uniquely determined. For (θ_1, θ_2) such that $\Pr(\theta_2|\theta_1) \in (0, 1)$, it is possible that there are multiple solutions to (2) that differ in consumption at dates 2 and 3. In such cases, let $C^*(\cdot; \beta)$ be any minimum discretion solution to (2). This indeterminacy does not affect any of the analysis below.

Assumption 4 *If θ_1 and θ_2 are non-independent then any solution to* $\max_{C \text{ s.t. } RC} E\left[U^0\left(C\left(\theta_1, \theta_2\right); \theta_1, \theta_2\right)\right]$ *has C_1 non-constant in θ_1.*

Lemma 5 *If θ_1 and θ_2 are non-independent then $C_1^*\left(\theta_1; \beta^*\right)$ is non-constant in θ_1.*

The following result then generalizes Proposition 1:

Proposition 3 *If the minimum discretion contract $C^*\left(\cdot; \beta\right)$ violates IC_1, commitment and flexibility can be combined only if preference reversal holds. Moreover, if (i) θ_1 and θ_2 are non-independent; (ii) the state following $\bar{\theta}_1$ is deterministic; and preference reversal holds for β in the neighborhood below β^*, then there exists $\hat{\beta} < \beta^*$ such that commitment and flexibility can be combined for all $\beta \geq \hat{\beta}$.*

As in Proposition 1, preference reversal is necessary to induce self 2 to punish self 1 for overconsumption. Moreover, under some conditions, preference reversal is sufficient as well as necessary. Condition (i) ensures that $\bar{\theta}_1, \underline{\theta}_1$ are well-defined. If condition (ii) is violated, the contract after $\bar{\theta}_1$ is fully determined by the solution to the benchmark problem (2), making it impossible for self 2 to punish self 1 for dishonestly reporting $\bar{\theta}_1$ in $\underline{\theta}_1$ without self 0 suffering some utility cost—regardless of whether or not preference reversal holds.

When condition (ii) fails, we accordingly turn instead to characterizing the solution to (1). Let $C^s\left(\cdot; \beta\right)$ be a solution to (1). We show that $C^s\left(\cdot; \beta\right)$ shares two key features with the contracts used to combine commitment and flexibility in Propositions 1-3.

Specifically, these results show that if $\bar{\theta}_1$ is deterministically 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 by giving self 2 discretion to consume more than $C_2\left(\bar{\theta}_1, \underline{\theta}_2\right)$. Proposition 4 establishes that both properties hold under preference reversal, even when uncertainty remains after the date 1 state $\bar{\theta}_1$. In contrast, if preference reversal is strictly violated—i.e., the distribution of θ_2 conditional on $\bar{\theta}_1$ strictly first-order stochastically dominates the distribution of θ_2 conditional on $\underline{\theta}_1$—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 high date 1 consumption is reduced rather than increased.

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

$$\hat{C}(\cdot,\cdot;\beta)=\arg\max_{C\text{ s.t. }IC_2,RC,C_1(\cdot)=C_1^s(\cdot;\beta)} E\left[U^0\left(C\left(\theta_1,\theta_2\right);\theta_1,\theta_2\right)\right].$$

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 $C^s(\cdot;\beta)$, i.e., $C_2^s\left(\bar{\theta}_1,\bar{\theta}_2\right)-C_2^s\left(\bar{\theta}_1,\underline{\theta}_2\right)$.

Proposition 4 *Suppose θ_1, θ_2 are non-independent, and for some θ_2 with* $\Pr\left(\bar{\theta}_1,\theta_2\right)\neq 0$,

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

(A) If preference reversal holds for β in the neighborhood below β^*, then there exists $\hat{\beta}<\beta^*$ such that if $\beta\in[\hat{\beta},\beta^*)$ and $C^*(\cdot;\beta)$ violates IC_1, the contract C^s increases discretion and strictly decreases $V^2\left(C^s\left(\bar{\theta}_1,\bar{\theta}_2;\beta\right);\bar{\theta}_2\right)$ relative to the benchmark $\hat{C}$.

(B) If preference reversal is strictly violated for β in the neighborhood below β^*, then there exists $\hat{\beta}<\beta^*$ such that if $\beta\in[\hat{\beta},\beta^*)$ and $C^*(\cdot;\beta)$ violates IC_1, the contract C^s decreases discretion and strictly increases $V^2\left(C^s\left(\bar{\theta}_1,\bar{\theta}_2;\beta\right);\bar{\theta}_2\right)$ relative to the benchmark $\hat{C}$.

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.s[32]

[32]If (7) does not hold, then for β close to β^* the contract C^s distorts date 1 consumption, but does not

The results on the continuation utility levels V^2 are intuitive. Self 0 must distort future consumption in order to prevent self 1 from overconsuming in state $\underline{\theta}_1$. Under preference reversal, 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 (formally, see Lemma A-3 in the appendix). 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. Because a higher continuation utility $V^2\left(C\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right)$ is associated with less distortion of consumption across dates 2 and 3, this corresponds 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 $C_2\left(\bar{\theta}_1,\bar{\theta}_2\right)$. 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]

Finally, we consider the case in which θ_1 and θ_2 are independently distributed (see Amador et al 2003). If $C_1^*\left(\cdot;\beta\right)$ is non-constant, preference reversal is violated, and so by Proposition 3 commitment and flexibility cannot be combined.[34] So as in Proposition 4, consumption must be distorted. But in this case, self 1 does not learn anything about the distribution of θ_2 from seeing θ_1, and so there is no reason for this distortion to take the form of changing self 2's discretion. Instead, and as noted in Amador et al (2003), it is

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.

[34]If $C_1^*\left(\cdot;\beta\right)$ is constant then IC_1 holds and commitment and flexibility are straightforwardly combined.

sufficient to simply destroy resources, while leaving self 2's discretion unperturbed.[35] The reason is that contract terms from date 2 onwards enter both the objective and IC_1 only through $E\left[V^2\left(C\left(\theta_1,\theta_2\right);\theta_2\right)\right]$, and any reduction in $E\left[V^2\left(C\left(\theta_1,\theta_2\right);\theta_2\right)\right]$ that is achieved via manipulating self 2's discretion can instead be achieved by destroying resources.

4.5 Continuum of states

We next allow θ_1 and θ_2 to take a continuum of different values: specifically, Θ_1 and Θ_2 are compact and convex subsets of $\Re$. We focus on the case in which the date 2 state following θ_1 is deterministic (subsection 4.1). We define the minimum discretion contract C^* that solves the benchmark problem (2) as before. Since all uncertainty is resolved at date 1, C^* removes all discretion from self 2, i.e., is constant in θ_2; and is independent of β.

Our existing definition of preference reversal is global. In addition, we say *local preference reversal* holds if, for every $\theta_1 \in \Theta_1$, there exists $\varepsilon > 0$ such that preference reversal holds over $(\theta_1-\varepsilon,\theta_1+\varepsilon)\cap\Theta_1$. Local preference reversal does not imply preference reversal.

Write $\phi\left(\theta_1\right)$ for the date 2 state that deterministically follows θ_1. We assume ϕ is differentiable, and utility $u_t\left(c_t;\theta_t\right)$ is twice differentiable in c_t and jointly differentiable in c_t and θ_t. Note that C^* violates IC_1 for any $\beta < 1$: because $C^*\left(\theta_1,\phi\left(\theta_1\right)\right)$ is continuous in θ_1, for any $\beta < 1$ self 1 gains by misreporting θ_1 so as to slightly increase 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.[36]

Our main result is the following analogue of Proposition 1:

Proposition 5 *If $\beta < 1$, commitment and flexibility can be combined only if local preference reversal holds. Moreover, if preference reversal holds, and* $\max_{\theta_1,\tilde{\theta}_1\in\Theta_1}\left|C^*\left(\theta_1,\phi\left(\theta_1\right)\right)-C^*\left(\tilde{\theta}_1,\phi\left(\tilde{\theta}_1\right)\right)\right|$

---

[35] Formally, there is a solution C^s to (1) such that, for all θ_1, there exists $k\left(\theta_1\right)\le W-C_1^s\left(\theta_1\right)$ such that $\left(C_2^s\left(\theta_1,\cdot;\beta\right),C_3^s\left(\theta_1,\cdot;\beta\right)\right)=\arg\max_{C \text{ s.t. } \text{IC}_2 \text{ and } C_2\left(\theta_1,\theta_2\right)+C_3\left(\theta_1,\theta_2\right)\le k\left(\theta_1\right)} E\left[U^0\left(C\left(\theta_1,\theta_2\right);\theta_1,\theta_2\right)\right]$.

[36] The proof of Proposition 5 formalizes this argument.

is sufficiently small, then commitment and flexibility can be combined.

The economic forces behind Proposition 5 are the same as for prior results. In particular, if local preference reversal does not hold, then combining commitment and flexibility would require punishing self 1 in some date 2 states with low MU, but by Lemma 2 this is impossible. In contrast, if preference reversal does hold, then self 1 must be punished in date 2 states with high MU, and this is potentially achievable.

The new element handled by Proposition 5 is that self 1 can overconsume to various degrees. Often, greater overconsumption necessitates a more severe punishment, i.e., holding fixed a misreport $\tilde{\theta}_1$, and denoting the true state by θ_1, the punishment needed is increasing in overconsumption $C_1^*\left(\tilde{\theta}_1\right) - C_1^*(\theta_1)$. The contract C must be designed so that after self 1 lies and reports $\tilde{\theta}_1$, self 2 picks the punishment appropriate to the true state θ_1.

We relegate most details to the appendix, and mention just a few key elements of the construction here. Consider a particular report by self 1, $\tilde{\theta}_1$. To combine commitment with flexibility, we clearly need to set $C\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right) = C^*\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right)$. Next, consider a particular realization of the true date 1 state, θ_1 say, with $C_1^*(\theta_1) < C_1^*\left(\tilde{\theta}_1\right)$. So the contract terms $C\left(\tilde{\theta}_1, \phi(\theta_1)\right)$ may need to impose a punishment on self 1 to deter him from overconsuming $C_1^*\left(\tilde{\theta}_1\right)$ in state θ_1. To do so, we design a contract such that $U^1\left(C\left(\tilde{\theta}_1, \phi(\theta_1)\right); \theta_1, \phi(\theta_1)\right) = U^1\left(C^*(\theta_1, \phi(\theta_1)); \theta_1, \phi(\theta_1)\right)$. As before, this is achieved by having $C\left(\tilde{\theta}_1, \phi(\theta_1)\right)$ offer greater date 2 consumption than $C^*\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right)$, but lower date 3 consumption, so that $V^2\left(C\left(\tilde{\theta}_1, \phi(\theta_1)\right); \phi(\theta_1)\right) < V^2\left(C^*\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right); \phi(\theta_1)\right)$. Also as before, such a construction is potentially possible because, by preference reversal, MU is higher in $\phi(\theta_1)$ than $\phi\left(\tilde{\theta}_1\right)$. Finally, to ensure that self 2 truthfully reports $\theta_2 = \phi(\theta_1)$, for any $\tilde{\theta}_1$ we design $C\left(\tilde{\theta}_1, \theta_2\right)$ so that C_2 is increasing in θ_2, and C_2 and C_3 satisfy the standard differential form of IC_2, namely

$$u_2'\left(C_2\left(\tilde{\theta}_1, \theta_2\right); \theta_2\right) \frac{\partial C_2\left(\tilde{\theta}_1, \theta_2\right)}{\partial \theta_2} + \beta u_3'\left(C_3\left(\tilde{\theta}_1, \theta_2\right)\right) \frac{\partial C_3\left(\tilde{\theta}_1, \theta_2\right)}{\partial \theta_2} = 0. \tag{8}$$

The condition on $\max_{\theta_1,\tilde{\theta}_1\in\Theta_1}\left|C^*(\theta_1,\phi(\theta_1))-C^*\left(\tilde{\theta}_1,\phi\left(\tilde{\theta}_1\right)\right)\right|$ in the sufficiency half of Proposition 5 plays the same role as the condition on β in Proposition 1, and ensures that the contract construction just described satisfies RC. The condition would be unnecessary if RC were allowed to be violated off the equilibrium path. Alternatively, the following result gives conditions that ensure that the contract described satisfies RC:

Proposition 6 *Suppose preference reversal holds, and in addition, (I)*

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

for some $n \geq 1$, function ζ, and CRRA utility function u, (II) $\frac{\partial\zeta(\phi^{-1}(\theta_2))}{\partial\theta_2} \in [0,\frac{1}{4})$, and (III) β is sufficiently close to $\beta^*=1$. Then commitment and flexibility can be combined.

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 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 $n^{-\gamma}(n+1)^{-\gamma}$, which is less than $\frac{1}{2}$ if $\gamma \geq 1$.

5 Investment problems

We next extend our analysis to cover investment problems, illustrated by Example 3. In an investment problem, investment rather than consumption is observable, and states θ_1 and θ_2 determine how investment affects subsequent "wealth." Investment problems very naturally deliver correlation between desired date 1 investment and marginal utilities at dates 2 and 3, since the return on investment affects wealth at later dates.

Let i_1 and i_2 denote investment at dates $t = 1, 2$. In addition, let $\ell \geq 0$ be an additional deadweight cost that a contract can stipulate at date 2. The following framework is general enough to encompass several applications: utilities at dates 1, 2, 3 are given by

$$u_1 \left(W - i_i \right); u_2 \left(\left(1 - \lambda \right) i_1 R \left(\theta_2 \right) + \lambda W - i_2 - \ell \right); u_3 \left(\lambda i_1 R \left(\theta_2 \right) + i_2 \right).$$

The parameter λ determines the maturity of the date 1 investment. On the one hand, $\lambda = 1$ corresponds to long-term date 1 investments, in which the return from date 1 investment is experienced at date 3. In particular, this case corresponds to the health investment setting of Example 3, in which both i_1 and i_2 directly affect long-term (i.e., date 3) health. On the other hand, $\lambda = 0$ corresponds to short-term date 1 investments, in which the return to date 1 investment arrives at date 2.

To map investment problems into our basic model, fix $\hat{\theta}_2 \in \Theta_2$, and define

$$c_1 = W - i_1; c_2 = (1 - \lambda) i_1 R \left(\hat{\theta}_2 \right) + \lambda W - i_2 - \ell; c_3 = \lambda i_1 R \left(\hat{\theta}_2 \right) + i_2.$$

So utilities at dates 1, 2, 3 are

$$u_1 \left(c_1 \right); u_2 \left(c_2 + \left(1 - \lambda \right) \left(W - c_1 \right) \left(R \left(\theta_2 \right) - R \left(\hat{\theta}_2 \right) \right) \right); u_3 \left(c_3 + \lambda \left(W - c_1 \right) \left(R \left(\theta_2 \right) - R \left(\hat{\theta}_2 \right) \right) \right).$$

Moreover, the constraint $\ell \geq 0$ is equivalent to

$$R \left(\hat{\theta}_2 \right) c_1 + c_2 + c_3 \leq \left(R \left(\hat{\theta}_2 \right) + \lambda \right) W,$$

which coincides with constraint RC in our basic model. Hence the only differences relative to our basic model are that θ_2 may affect MU at date 3 as well as at date 2, and moreover, this dependence is affected by c_1. However, it is straightforward to show that our main results still hold, with date 2 MU replaced with the ratio of date 2 MU to date 3 MU.

For conciseness, we focus on the case in which the return $R(\theta_2)$ is known to self 1 when he chooses i_1. This corresponds to perfect correlation of θ_1 and θ_2 in the basic model.

The return $R(\theta_2)$ affects desired date 1 consumption via both substitution and income effects. For long-term investment ($\lambda = 1$), preference reversal arises when the substitution effect dominates, so that lower $R(\theta_2)$ is associated both with higher desired c_1 and higher date 3 MU, and hence a lower ratio of date 2 MU to date 3 MU. This case is illustrated by Example 3.

For short-term investment ($\lambda = 0$), preference reversal instead arises when the income effect dominates, so that higher $R(\theta_2)$ is associated both with higher desired c_1 and lower date 2 MU. Here, the income effect dominates the substitution effect when the elasticity of intertemporal substitution (EIN) is below 1. Although there is a range of empirical estimates for the EIN, many estimates put the EIN substantially below 1 (e.g., Hall 1988).

6 Private savings

Thus far, we have assumed that the agent cannot 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,[37] 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). We define the minimum discretion

[37] See, for example, Kocherlakota (2004), Doepke and Townsend (2006), and He (2009).

contract C^* exactly as before, and to focus on the interesting case we assume that C^* violates an incentive constraint for at least some values of β. This implies that C_1^* must vary across states, and we define $\bar{\theta}_1$ and $\underline{\theta}_1$ as previously.

The possibility of private saving does not affect self 0's most-preferred consumption, but it places additional constraints on self 0's ability to actually attain this consumption. As such, it is immediate that preference reversal remains a necessary condition for combining commitment with flexibility, and so for the remainder of this section we assume that preference reversal holds, i.e., $\bar{\theta}_1$ is followed by $\underline{\theta}_2$ and $\underline{\theta}_1$ is followed by $\bar{\theta}_2$.

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 $C_1^*\left(\bar{\theta}_1\right) - C_1^*\left(\underline{\theta}_1\right)$ on to self 2. By doing so, self 1 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.

Suppose that the minimum discretion contract C^* is in place, and that self 1 overconsumes by reporting $\bar{\theta}_1$ in $\underline{\theta}_1$. A key object for the analysis of private savings is the savings level $s_1^*(\beta)$ at which self 1 is just indifferent between misreporting $\bar{\theta}_1$ in $\underline{\theta}_1$ and then privately saving $s_1^*(\beta)$; and truthfully reporting $\underline{\theta}_1$. Intuitively, $s_1^*(\beta)$ is the maximum amount of private savings that is relevant for the analysis.[38]

Our first result is that if commitment and flexibility cannot be combined using C^*, then they can be combined only if $\underline{\theta}_2 \leq \bar{\theta}_2 - s_1^*(\beta)$, a condition we term *strong preference reversal* (SPR), and that is stricter than preference reversal (simply $\underline{\theta}_2 < \bar{\theta}_2$).[39] The economics

[38] Formally, first define self 2's private savings choice as a function of self 1's private savings choice s_1,

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

Then $s_1^*(\beta)$ is defined by

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

where $s_1^*(\beta) = 0$ if the above set is empty.

[39] If commitment and flexibility cannot be combined using C^* then $s_1^*(\beta) > 0$, so SPR is indeed a stricter condition than preference reversal.

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 saves up to $s_1^*(\beta)$. His date 2 MU in this case is determined by $\bar{\theta}_2 - s_1^*(\beta)$, where $\bar{\theta}_2$ follows from preference reversal, which is reduced by private savings of $s_1^*(\beta)$. State-contingent punishment requires higher MU in the punishment state than in the non-punishment state, where date 2 MU is simply $\underline{\theta}_2$. This is exactly condition SPR.

Under conditions analogous to those of Proposition 1, SPR is sufficient as well as necessary. Specifically, define β^{*p} analogously to β^* as the hyperbolic discount rate such that if $\beta \geq \beta^{*p}$, allocating all decisions to self 1[40] delivers self 0's most preferred outcome, while this is not the case for β just below β^{*p}. If SPR holds strictly at $\beta = \beta^{*p}$, then commitment and flexibility can be combined for all β sufficiently close to β^{*p}.

Finally, we show that condition SPR is equivalent to condition (III) in Proposition 2, which as discussed in subsection 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 combine commitment with flexibility by writing a contract that induces self 2 to punish self 1 for overconsumption. 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

[40] Because private saving is possible, self 2 always has at least the discretion to save and pass resources from date 2 to date 3.

overconsume at date 1 without self 2 punishing him, because self 1 underestimates self 2's preference for date 2 consumption.

Table 1 presents numerical simulations that shed light on the impact of naïveté. All simulations use 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 $u_2\left(c_2;\bar{\theta}_2\right) = u\left(c_2 - \bar{\theta}_2\right)$ for a CRRA function u. Preference reversal holds here. For each level of risk aversion γ and shock size $\bar{\theta}_2$, in row (i) we report β^*, the hyperbolic discount rate at which it becomes impossible to combine commitment and flexibility simply by using the minimum discretion contract C^*. But—absent naïveté—from Proposition 2 we know that for any $\beta < \beta^*$ there exists a contract that nonetheless combines commitment and flexibility. Then, in each of rows (ii), (iii), (iv) we fix the degree of hyperbolic discounting at the empirically reasonable values of $\beta = 0.9, 0.8, 0.7$ respectively,[41] and calculate the maximum degree of naïveté for which commitment and flexibility can still be combined. For example, for $\gamma = 1$, $\bar{\theta}_2 = 0.05$ and $\beta = 0.9$, such a contract exists if self 1's belief $\tilde{\beta}$ about self 2's hyperbolic discount rate is 0.947 or less.

Panels (B) and (C) present the results for parallel analysis of timing shocks (B) and investment problems (C). For timing shocks, $\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right) = \Pr\left(\bar{\theta}_2|\underline{\theta}_1\right) = 1$, $u_1\left(c;\bar{\theta}_1\right) = u_2\left(c;\bar{\theta}_2\right) = u\left(c - \bar{\theta}_2\right)$ 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, we focus on long-term investments (see Example 3) with $\max_{\theta_2} R(\theta_2) = 1$. The table reports $\min_{\theta_2} R(\theta_2)$. For investment problems the case of risk-aversion $\gamma = 4$ is omitted, since in this case the substitution and income effects approximately offset each other and hence desired date 1 investment is approximately independent of θ_2.

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 combine commitment with flexibility by granting discretion to self 2. Moreover,

[41] See footnote 27 for estimates of β. Recall that the time period associated with these estimates is one week (the 0.9 estimate), one quarter (the 0.8 estimate) and one year (the 0.7 estimate).

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

Although commitment and flexibility can often be combined under partial naïveté, there is now the problem that self 0 may pick the wrong contract at date 0.

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 streams $C^*\left(\bar{\theta}_1, \underline{\theta}_2\right)$ and $C^*\left(\underline{\theta}_1, \bar{\theta}_2\right)$. First, note that the punishment component $C\left(\bar{\theta}_1, \bar{\theta}_2\right)$ 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 $U^0\left(C^*\left(\bar{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_1,\bar{\theta}_2\right)$.[42] 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.[43]

8 Conclusion

We characterize when an agent with hyperbolic discounting can 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 combine 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 savings instruments outside the contract.

In this paper, we focus on one particular form of time-inconsistent preferences, namely

[42]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.

[43]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 $\tilde{\beta}/\beta$ is constant. Again, we leave this interesting question for future research.

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 (2015), "Poverty and Self-Control," *Econometrica* 83, 1877-1911.

[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* $U^1(\tilde{c};\theta_1,\theta_2) \geq U^1(c;\theta_1,\theta_2)$*, 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 $\mathrm{IC}_2\left(\theta_1, \theta_2, \tilde{\theta}_2\right)$ 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 $\mathrm{IC}_2\left(\theta_1, \tilde{\theta}_2, \theta_2\right)$. **QED**

Lemma A-3 *Define $\bar{v}^2(k,\theta_2) = \max_{c_2} V^2(c_2, k - c_2; \theta_2)$, and $\bar{c}_2^k$ and $\underline{c}_2^k$ as 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*

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

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) consumption $(\underline{c}_2,\underline{c}_3)$ is uniquely defined, with $\underline{c}_2 + \underline{c}_3 = k$ and $\underline{c}_2$ strictly decreasing in v, and (C) if $\frac{u_2'(c;\underline{\theta}_2)}{u_2'(c;\bar{\theta}_2)} \neq \beta$ for some c, then $\bar{c}_2 - \underline{c}_2$ 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 $V^2\left(\underline{c}_2, \underline{c}_3; \underline{\theta}_2\right)$ unchanged while strictly increasing $U^2\left(\underline{c}_2, \underline{c}_3; \underline{\theta}_2\right)$, 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, Assumption 2 and the condition that $\frac{u_2'(c; \underline{\theta}_2)}{u_2'(c; \bar{\theta}_2)} \neq \beta$ for some c imply that $\bar{c}_2$ is uniquely defined by the condition $u_2'\left(c_2; \bar{\theta}_2\right) - \frac{1}{\beta} u_2'\left(c_2; \underline{\theta}_2\right) = 0$. Hence $\bar{c}_2$ is invariant to small changes in v.

Second, consider the case $\bar{c}_2 + \bar{c}_3 = k$. Observe 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 $u_2'\left(\bar{c}_2;\bar{\theta}_2\right) - \frac{1}{\beta}u_2'\left(\bar{c}_2;\underline{\theta}_2\right) < 0$ one can increase $V^2\left(\bar{c}_2,\bar{c}_3;\bar{\theta}_2\right)$ by reducing both $\bar{c}_2$ and $\bar{c}_2 + \bar{c}_3$ to leave $U^2\left(\bar{c}_2,\bar{c}_3;\underline{\theta}_2\right)$ unchanged (this uses $\underline{c}_2 < \bar{c}_2$ and (A-2) at equality). Assumption 2 and the condition that $\frac{u_2'(c;\underline{\theta}_2)}{u_2'(c;\bar{\theta}_2)} \neq \beta$ for some c then imply that $u_2'\left(c_2;\bar{\theta}_2\right) - \frac{1}{\beta}u_2'\left(c_2;\underline{\theta}_2\right) > 0$ for all $c_2 < \bar{c}_2$. Consider now a small upwards perturbation in v, to v^+ say; from (B), this is associated with a small decrease in $\underline{c}_2$. This tightens constraint (A-2) (again using $\underline{c}_2 < \bar{c}_2$ and (A-2) at equality). 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 $\bar{v}^2(k,\theta_2)$, $\bar{c}_2^k$, and $f(v;k,\beta)$ 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).*[44] *Moreover, f is strictly concave and differentiable in v at any value such that $f(v;k,\beta) < \bar{v}^2\left(k,\bar{\theta}_2\right)$.*

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

$$
\begin{aligned}
f(v;k,\beta) \quad &\equiv \quad \max_{\bar{u}_2,\bar{u}_3,\underline{u}_2,\underline{u}_3} \bar{u}_2 + \bar{u}_3 && \text{(A-4)}\\
&\text{s.t.} \quad \underline{u}_2 + \underline{u}_3 = v && \text{(A-5)}\\
&\text{and} \quad \underline{u}_2 + \beta\underline{u}_3 \geq u_2\left(u_2^{-1}\left(\bar{u}_2;\bar{\theta}_2\right);\underline{\theta}_2\right) + \beta\bar{u}_3 && \text{(A-6)}\\
&\text{and} \quad u_2^{-1}\left(\bar{u}_2;\bar{\theta}_2\right) + u_3^{-1}\left(\bar{u}_3\right) \leq k, u_2^{-1}\left(\underline{u}_2;\underline{\theta}_2\right) + u_3^{-1}\left(\underline{u}_3\right) \leq k. && \text{(A-7)}
\end{aligned}
$$

To establish concavity of f with respect to (v,k), it is sufficient to show that the constraint

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

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 strictly convex (they are inverses of strictly concave and increasing functions). Hence the constraint set is convex.

Strict concavity with respect to v follows straightforwardly, using the same arguments.

Finally, to establish differentiability of f at a point v_0 it is sufficient (given concavity) 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.[45] 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, $(\underline{c}_2, \underline{c}_3)$ 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 main text results (excluding Propositions 2 and 6)

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 $U^2(C(\theta_1,\theta_2);\theta_2) \geq U^2\left(C\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right)$, which imply $u_2\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\tilde{\theta}_2\right) - u_2\left(C_2(\theta_1,\theta_2);\tilde{\theta}_2\right) \geq u_2\left(C_2\left(\theta_1,\tilde{\theta}_2\right);\theta_2\right) - u_2(C_2(\theta_1,\theta_2);\theta_2)$, which by $\tilde{\theta}_2 > \theta_2$ implies $C_2\left(\theta_1,\tilde{\theta}_2\right) \geq C_2(\theta_1,\theta_2)$. **QED**

Proof of Proposition 1: Necessity is established in the main text. Here, we estab-

[45] Since f is concave, its one-sided derivatives both exist. Denote these one-sided derivatives by f^- and f^+, and note that, by concavity $f^- \geq f^+$. If the function g exists, then at v_0 we have $f^+ \geq g' \geq f^-$. Hence $f^- = f^+$, establishing differentiability.

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

Proof of Lemma 3: Since commitment and flexibility can be combined, either $\beta \geq \beta^*$, and so the result holds trivially; or else preference reversal holds, i.e., $\Pr(\bar{\theta}_1,\bar{\theta}_2) = \Pr(\underline{\theta}_1,\underline{\theta}_2) = 0$. Let C be a contract combining commitment and flexibility. By IC_2,

$$u_2(C_2(\bar{\theta}_1,\bar{\theta}_2);\bar{\theta}_2) - u_2(C_2(\bar{\theta}_1,\underline{\theta}_2);\bar{\theta}_2) \geq \beta u_3(C_3(\bar{\theta}_1,\underline{\theta}_2)) - \beta u_3(C_3(\bar{\theta}_1,\bar{\theta}_2)). \quad \text{(A-8)}$$

First, note that there exists C such that (A-8) holds with equality. To see this, note that Lemma 1 implies $C_2(\bar{\theta}_1,\bar{\theta}_2) \geq C_2(\bar{\theta}_1,\underline{\theta}_2)$ and $C_3(\bar{\theta}_1,\bar{\theta}_2) \leq C_3(\bar{\theta}_1,\underline{\theta}_2)$; and consequently, one can set (A-8) to equality by lowering $C_2(\bar{\theta}_1,\bar{\theta}_2)$, while continuing to satisfy all other IC constraints. Given this, assume directly that C satisfies (A-8) with equality.

Define $X_2 = C_2(\bar{\theta}_1,\bar{\theta}_2) - C_2(\bar{\theta}_1,\underline{\theta}_2)$ and $X_3 = C_3(\bar{\theta}_1,\underline{\theta}_2) - C_3(\bar{\theta}_1,\bar{\theta}_2)$. So (A-8) at equality is equivalent to

$$u_2\left(C_2\left(\bar{\theta}_1,\underline{\theta}_2\right)+X_2;\bar{\theta}_2\right)-u_2\left(C_2\left(\bar{\theta}_1,\underline{\theta}_2\right);\bar{\theta}_2\right)=\beta u_3\left(C_3\left(\bar{\theta}_1,\underline{\theta}_2\right)\right)-\beta u_3\left(C_3\left(\bar{\theta}_1,\underline{\theta}_2\right)-X_3\right).$$

By assumption $C_3\left(\bar{\theta}_1,\underline{\theta}_2\right)\geq C_3\left(\underline{\theta}_1,\bar{\theta}_2\right)$, and so $C_2\left(\bar{\theta}_1,\underline{\theta}_2\right)<C_2\left(\underline{\theta}_1,\bar{\theta}_2\right)$. So by concavity of u_2 and u_3 it follows that

$$u_2\left(C_2\left(\underline{\theta}_1,\bar{\theta}_2\right)+X_2;\bar{\theta}_2\right)-u_2\left(C_2\left(\underline{\theta}_1,\bar{\theta}_2\right);\bar{\theta}_2\right)<\beta u_3\left(C_3\left(\underline{\theta}_1,\bar{\theta}_2\right)\right)-\beta u_3\left(C_3\left(\underline{\theta}_1,\bar{\theta}_2\right)-X_3\right).$$

Hence under the contract C, self 2 does not wish to change consumption at dates 2 and 3 by $(X_2,-X_3)$ after self 1 truthfully reports $\underline{\theta}_1$. Finally, it is irrelevant whether or not self 2 would change consumption at dates 2 and 3 by $(X_2,-X_3)$ after self 1 misreports $\underline{\theta}_1$ in state $\bar{\theta}_1$, since by the same argument as in the proof of Proposition 1, self 1 never misreports in this way. **QED**

Proof of Lemma 4: Let $\bar{v}^2(k,\theta_2)$, $\bar{c}_2^k$, and $f(v;k,\beta)$ be as defined in Lemma A-3. For any θ_1 and β, define $g(k,\beta;\theta_1)$ by

$$g(k,\beta;\theta_1)\equiv\max_{v\in\left(V^2(0,0;\underline{\theta}_2),\bar{v}^2(k,\underline{\theta}_2)\right]}\Pr\left(\bar{\theta}_2|\theta_1\right)f(v;k,\beta)+\Pr\left(\underline{\theta}_2|\theta_1\right)v. \tag{A-9}$$

The maximizing choice of v in (A-9) 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.

Any solution C to (2) must have

$$C_1(\theta_1)\in\arg\max_{c_1}u_1(c_1;\theta_1)+g(W-c_1,\beta;\theta_1). \tag{A-10}$$

The uniqueness of $C_1(\theta_1)$ follows from the strict concavity of the objective in (A-10). Uniqueness of $V^2(C(\theta_1,\theta_2);\theta_2)$ for $\Pr(\theta_1,\theta_2)\neq 0$ then follows from the strict concavity of $f(v;k,\beta)$ in v (see Lemma A-4).

Because the objective in (A-10) is strictly concave in c_1, and continuous in β, it follows that $C_1(\theta_1)$ is continuous in β. For $\Pr(\theta_2|\theta_1) \in (0,1)$, the strict concavity of f in v and its continuity in k and β imply that $V^2(C(\theta_1,\theta_2);\theta_2)$ is continuous in β; while for $\Pr(\theta_2|\theta_1) = 1$ this is immediate. **QED**

Proof of Lemma 5: By the definition of β^* and Lemma 4, at least one constraint in IC_1 must hold with equality at $\beta = \beta^*$ and $C = C^*$, i.e., for some $\check{\theta}_1, \tilde{\theta}_1$, $E\left[U^1\left(C^*\left(\check{\theta}_1,\theta_2;\beta^*\right);\check{\theta}_1,\theta_2\right)|\check{\theta}_1\right] = E\left[U^1\left(C^*\left(\tilde{\theta}_1,\theta_2;\beta^*\right);\check{\theta}_1,\theta_2\right)|\check{\theta}_1\right]$. Suppose that, contrary to the claimed result, $C_1^*\left(\check{\theta}_1;\beta^*\right) = C_1^*\left(\tilde{\theta}_1;\beta^*\right)$. Hence

$$E\left[V^2\left(C^*\left(\check{\theta}_1,\theta_2;\beta^*\right);\theta_2\right)|\check{\theta}_1\right] = E\left[V^2\left(C^*\left(\tilde{\theta}_1,\theta_2;\beta^*\right);\theta_2\right)|\check{\theta}_1\right]. \tag{A-11}$$

There are two subcases. In the first subcase, $V^2\left(C^*\left(\check{\theta}_1,\theta_2;\beta^*\right);\theta_2\right) < \max_{c_2} V^2\left(c_2, W - C_1^*\left(\check{\theta}_1;\beta^*\right) - c_2;\theta_2\right)$ for some θ_2 such that $\Pr\left(\check{\theta}_1,\theta_2\right) \neq 0$. Note that this subcase can only arise if both $\Pr\left(\check{\theta}_1,\bar{\theta}_2\right) \neq 0$ and $\Pr\left(\check{\theta}_1,\underline{\theta}_2\right) \neq 0$. From strict concavity of the objective in (A-9) in the proof of Lemma 4, and the differentiability of f established in Lemma A-4, it follows that $V^2\left(C^*\left(\check{\theta}_1,\theta_2;\beta^*\right);\theta_2\right) \neq V^2\left(C^*\left(\tilde{\theta}_1,\theta_2;\beta^*\right);\theta_2\right)$ for any θ_2 with $\Pr\left(\tilde{\theta}_1,\theta_2\right) \neq 0$. It then follows from the strict concavity of (A-9) that $E\left[V^2\left(C^*\left(\check{\theta}_1,\theta_2;\beta^*\right);\theta_2\right)|\check{\theta}_1\right] > E\left[V^2\left(C^*\left(\tilde{\theta}_1,\theta_2;\beta^*\right);\theta_2\right)|\check{\theta}_1\right]$, contradicting (A-11).

In the second subcase, $V^2\left(C^*\left(\check{\theta}_1,\theta_2;\beta^*\right);\theta_2\right) = \max_{c_2} V^2\left(c_2, W - C_1^*\left(\check{\theta}_1;\beta^*\right) - c_2;\theta_2\right)$ for all θ_2 such that $\Pr\left(\check{\theta}_1,\theta_2\right) \neq 0$. It then follows that for any θ_2 such that $\Pr\left(\tilde{\theta}_1,\theta_2\right) \neq 0$, $V^2\left(C^*\left(\tilde{\theta}_1,\theta_2;\beta^*\right);\theta_2\right) = \max_{c_2} V^2\left(c_2, W - C_1^*\left(\tilde{\theta}_1;\beta^*\right) - c_2;\theta_2\right)$. (This is straightforward if the state following $\check{\theta}_1$ is non-deterministic. If instead $\check{\theta}_1$ is followed deterministically by some $\check{\theta}_2 \in \Theta_2$, the argument is as follows. Equality (A-11) implies that $V^2\left(C^*\left(\tilde{\theta}_1,\check{\theta}_2;\beta^*\right);\check{\theta}_2\right) = \max_{c_2} V^2\left(c_2, W - C_1^*\left(\tilde{\theta}_1;\beta^*\right) - c_2;\check{\theta}_2\right)$. Let $\tilde{\theta}_2 \neq \check{\theta}_2$, and note that we must have $\Pr\left(\tilde{\theta}_1,\tilde{\theta}_2\right) \neq 0$. We must have $V^2\left(C^*\left(\tilde{\theta}_1,\tilde{\theta}_2;\beta^*\right);\tilde{\theta}_2\right) = \max_{c_2} V^2\left(c_2, W - C_1^*\left(\tilde{\theta}_1;\beta^*\right) - c_2;\tilde{\theta}_2\right)$, since otherwise it is straightforward to construct a perturbation to C^* that increases self 0's utility, a contradiction.) It then further follows that $C^*(\cdot,\cdot;\beta^*)$ solves $\max_{C \text{ s.t. RC}} E\left[U^0\left(C\left(\theta_1,\theta_2\right);\theta_1,\theta_2\right)\right]$, since otherwise one can straightforwardly construct a perturbation that strictly increases self 0's utility while satisfying RC and IC_2, by either increasing date 1 consumption and decreasing date 3 consumption, or decreasing date 1 consumption and increasing date 2 consumption (this uses Lemma A-3(A)). Since $C_1^*\left(\check{\theta}_1;\beta^*\right) = C_1^*\left(\tilde{\theta}_1;\beta^*\right)$, this contradicts Assumption 4. **QED**

Proof of Proposition 3: *Necessity:* To establish the result, suppose that $C^*(\cdot;\beta)$ violates IC_1 and that preference reversal is violated, but that commitment and flexibility can nonetheless be combined. Note that since C^* is a solution to (2), for any $\check{\theta}_1,\tilde{\theta}_1$,

$$u_1\left(C_1^*\left(\check{\theta}_1;\beta\right);\check{\theta}_1\right)+E\left[V^2\left(C^*\left(\check{\theta}_1,\theta_2;\beta\right);\theta_2\right)|\check{\theta}_1\right] \geq u_1\left(C_1^*\left(\tilde{\theta}_1;\beta\right);\check{\theta}_1\right)+E\left[V^2\left(C^*\left(\tilde{\theta}_1,\theta_2;\beta\right);\theta_2\right)|\check{\theta}_1\right].$$

If $C_1^*\left(\check{\theta}_1;\beta\right) \geq C_1^*\left(\tilde{\theta}_1;\beta\right)$ then, by Lemma A-1, $C^*(\cdot;\beta)$ satisfies IC_1 in $\check{\theta}_1$. So by supposition, let $\check{\theta}_1,\tilde{\theta}_1$ be such that $C_1^*\left(\tilde{\theta}_1;\beta\right) > C_1^*\left(\check{\theta}_1;\beta\right)$ and $C^*(\cdot;\beta)$ violates IC_1 in $\check{\theta}_1$:

$$u_1\left(C_1^*\left(\check{\theta}_1;\beta\right);\check{\theta}_1\right)+\beta E\left[V^2\left(C^*\left(\check{\theta}_1,\theta_2;\beta\right);\theta_2\right)|\check{\theta}_1\right] < u_1\left(C_1^*\left(\tilde{\theta}_1;\beta\right);\check{\theta}_1\right)+\beta E\left[V^2\left(C^*\left(\tilde{\theta}_1,\theta_2;\beta\right);\theta_2\right)|\check{\theta}_1\right]. \tag{A-12}$$

The LHS of this inequality takes the same value in any solution to (2) (see Lemma 4). If both states $\underline{\theta}_2$ and $\bar{\theta}_2$ have positive probability after $\tilde{\theta}_1$, then the same is true of the RHS, and consequently, any solution to (2) violates IC_1, and commitment and flexibility cannot be combined (regardless of whether or not preference reversal holds), giving a contradiction.

So the only remaining case to consider is that in which $\tilde{\theta}_1$ is deterministically followed by $\bar{\theta}_2$: if $\tilde{\theta}_1$ were instead deterministically followed by $\underline{\theta}_2$, then $\Pr\left(\check{\theta}_1,\bar{\theta}_2\right) \neq 0$, and preference reversal holds. Lemma 2 implies $V^2\left(C\left(\tilde{\theta}_1,\underline{\theta}_2\right);\underline{\theta}_2\right) \geq V^2\left(C^*\left(\tilde{\theta}_1,\bar{\theta}_2\right);\underline{\theta}_2\right)$. By inequality (A-12), it follows that any solution to (2) violates IC_1, giving a contradiction and completing the proof.

Sufficiency: By the arguments above, if $C^*(\cdot;\beta)$ violates IC_1 for β in the neighborhood of β^*, then it violates IC_1 in state $\underline{\theta}_1$. Moreover, by Lemma 4, the quantity $E\left[U^1\left(C^*\left(\underline{\theta}_1,\theta_2;\beta\right);\underline{\theta}_1,\theta_2\right)|\underline{\theta}_1\right]-$

$E\left[U^1\left(C^*\left(\bar{\theta}_1,\theta_2;\beta\right);\underline{\theta}_1,\theta_2\right)|\underline{\theta}_1\right]$ is continuous in β. The remainder of the proof parallels that of Proposition 1. **QED**

Proof of Proposition 4: We prove part (A), i.e., $\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right) > \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)$. The proof of part (B) is parallel. The case $\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right) = 0$ is covered by Proposition 3; accordingly, we assume below that $\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right) \neq 0$. Let $f\left(v;k,\beta\right)$ be as defined in Lemma A-3.

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

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

From (A-13) and Lemma 4, it follows that for all (θ_1,θ_2) such that $\Pr\left(\theta_1,\theta_2\right) \neq 0$,

$$C_1^*\left(\theta_1;\beta\right) - C_1^s\left(\theta_1;\beta\right) \to 0 \quad \text{as} \quad \beta \to \beta^* \tag{A-14}$$

$$V^2\left(C^*\left(\theta_1,\theta_2;\beta\right);\theta_2\right) - V^2\left(C^s\left(\theta_1,\theta_2;\beta\right);\theta_2\right) \to 0 \quad \text{as} \quad \beta \to \beta^*. \tag{A-15}$$

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

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

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

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 $E\left[U^0\left(C^s\left(\theta_1,\theta_2;\beta\right);\theta_1,\theta_2\right)\right]$ 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 condition (7), and (A-14), for β in the neighborhood of β^*:

$$\begin{aligned} \Pr\left(\underline{\theta}_2|\bar{\theta}_1\right) + \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right) f'\left(V^2\left(C^*\left(\bar{\theta}_1, \underline{\theta}_2; \beta\right); \underline{\theta}_2\right); W - C_1^*\left(\bar{\theta}_1\right), \beta\right) &= 0 \quad \text{(A-17)} \\ \Pr\left(\underline{\theta}_2|\bar{\theta}_1\right) + \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right) f'\left(V^2\left(\hat{C}\left(\bar{\theta}_1, \underline{\theta}_2; \beta\right); \underline{\theta}_2\right); W - C_1^s\left(\bar{\theta}_1\right), \beta\right) &= 0. \quad \text{(A-18)} \end{aligned}$$

By supposition C^* violates IC_1. Specifically, this must happen in $\underline{\theta}_1$ (see proof of Proposition 3). Because $\Pr\left(\theta_2|\bar{\theta}_1\right) \neq 0$ for $\theta_2 = \underline{\theta}_2, \bar{\theta}_2$,

$$E\left[U^0\left(C^s\left(\theta_1, \theta_2; \beta\right); \theta_1, \theta_2\right)\right] < E\left[U^0\left(C^*\left(\theta_1, \theta_2; \beta\right); \theta_1, \theta_2\right)\right]. \quad \text{(A-19)}$$

Main step: We show, by contradiction, that for β in the neighborhood of β^*,

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

First, suppose that

$$\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right) + \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right) f'\left(v_{\underline{\theta}_2}; W - C_1^s\left(\bar{\theta}_1\right), \beta\right) = 0. \quad \text{(A-21)}$$

Consequently, an infinitesimal perturbation in $v_{\underline{\theta}_2}$ with a corresponding change to $v_{\bar{\theta}_2}$ to preserve (A-16) leaves $E\left[U^0\left(C^s\left(\theta_1, \theta_2; \beta\right); \theta_1, \theta_2\right)\right]$ unchanged and IC_2 intact while, since $\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right) \neq \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)$, strictly relaxing IC_1. By (A-19), one can use the newly-created slack in the constraints to a construct a perturbation that increases the objective $E\left[U^0\left(C^s\left(\theta_1, \theta_2; \beta\right); \theta_1, \theta_2\right)\right]$, giving a contradiction. Second, suppose instead that

$$\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right) + \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right) f'\left(v_{\underline{\theta}_2}; W - C_1^s\left(\bar{\theta}_1\right), \beta\right) > 0. \quad \text{(A-22)}$$

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

$$\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)\tilde{v}_{\underline{\theta}_2}+\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)f\left(\tilde{v}_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right)=\Pr\left(\underline{\theta}_2|\bar{\theta}_1\right)v_{\underline{\theta}_2}+\Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)f\left(v_{\underline{\theta}_2};W-C_1^s\left(\bar{\theta}_1\right),\beta\right). \tag{A-23}$$

Note that $\tilde{v}_{\underline{\theta}_2}$ is well-defined by the continuity and concavity of f (Lemma A-4), together with (A-14), (A-15), (A-17). From (A-23), and using $\Pr\left(\bar{\theta}_2|\underline{\theta}_1\right) > \Pr\left(\bar{\theta}_2|\bar{\theta}_1\right)$,

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

i.e.,

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

Hence switching from $v_{\underline{\theta}_2}$ to $\tilde{v}_{\underline{\theta}_2}$ strictly relaxes IC_1 while preserving both IC_2 and self 0's utility, again leading to a contradiction.

Completing the proof: Given the concavity of f, (A-18) and (A-20) imply that $v_{\underline{\theta}_2} > V^2\left(\hat{C}\left(\bar{\theta}_1,\underline{\theta}_2;\beta\right);\underline{\theta}_2\right)$. It follows that $v_{\bar{\theta}_2} < V^2\left(\hat{C}\left(\bar{\theta}_1,\bar{\theta}_2;\beta\right);\bar{\theta}_2\right)$, as claimed; and Lemma A-3 implies $C_2^s\left(\bar{\theta}_1,\bar{\theta}_2\right)-C_2^s\left(\bar{\theta}_1,\underline{\theta}_2\right)\geq\hat{C}_2\left(\bar{\theta}_1,\bar{\theta}_2\right)-\hat{C}_2\left(\bar{\theta}_1,\underline{\theta}_2\right)$. **QED**

Proof of Proposition 5: *Necessity:* Fix $\beta < 1$, and suppose C solves (2) but local preference reversal does not hold, i.e., there exists θ_1^0 such that, for any $\varepsilon > 0$, preference reversal does not hold over $\left(\theta_1^0-\varepsilon,\theta_1^0+\varepsilon\right)\cap\Theta_1$. Note that $U^1\left(C\left(\theta_1,\phi\left(\theta_1\right)\right);\theta_1,\phi\left(\theta_1\right)\right)=\left(1-\beta\right)u_1\left(C_1\left(\theta_1\right);\theta_1\right)+\beta V^1\left(C\left(\theta_1,\phi\left(\theta_1\right)\right);\theta_1,\phi\left(\theta_1\right)\right)$, and so, since C solves (2) and $C\left(\theta_1,\phi\left(\theta_1\right)\right)$ is differentiable in θ_1, for any $\hat{\theta}_1$, $\frac{\partial}{\partial\theta_1}U^1\left(C\left(\theta_1,\phi\left(\theta_1\right)\right);\hat{\theta}_1,\phi\left(\hat{\theta}_1\right)\right)\Big|_{\theta_1=\hat{\theta}_1}=\left(1-\beta\right)u_1\left(C_1\left(\hat{\theta}_1\right);\hat{\theta}_1\right)\frac{\partial}{\partial\theta_1}C_1\left(\theta_1\right)\Big|_{\theta_1=\hat{\theta}_1}$. Hence one can find $\check{\theta}_1,\tilde{\theta}_1$ close to θ_1^0 such that $C_1\left(\tilde{\theta}_1\right)>C_1\left(\check{\theta}_1\right)$, $\phi\left(\tilde{\theta}_1\right)\geq\phi\left(\check{\theta}_1\right)$, and $U^1\left(C\left(\tilde{\theta}_1,\phi\left(\tilde{\theta}_1\right)\right);\check{\theta}_1,\phi\left(\check{\theta}_1\right)\right)>U^1\left(C\left(\check{\theta}_1,\phi\left(\check{\theta}_1\right)\right);\check{\theta}_1,\phi\left(\check{\theta}_1\right)\right)$. By Lemma 2, if C solves (2) then $U^1\left(C\left(\tilde{\theta}_1,\phi\left(\check{\theta}_1\right)\right);\check{\theta}_1,\phi\left(\check{\theta}_1\right)\right)\geq U^1\left(C\left(\tilde{\theta}_1,\phi\left(\tilde{\theta}_1\right)\right);\check{\theta}_1,\phi\left(\check{\theta}_1\right)\right)$. But then $U^1\left(C\left(\tilde{\theta}_1,\phi\left(\check{\theta}_1\right)\right);\check{\theta}_1,\phi\left(\check{\theta}_1\right)\right)>U^1\left(C\left(\check{\theta}_1,\phi\left(\check{\theta}_1\right)\right);\check{\theta}_1,\phi\left(\check{\theta}_1\right)\right)$, contradicting

IC_1 and completing the proof.

Sufficiency: Because of preference reversal, ϕ is invertible.[46] For any $\theta_1 \in \Theta_1$, define the contract C by $C(\theta_1, \theta_2) = C^*(\theta_1, \phi(\theta_1))$ if $\theta_2 < \phi(\theta_1)$; while if $\theta_2 \geq \phi(\theta_1)$, define C by $C_1(\theta_1) = C_1^*(\theta_1)$ and the pair of differential equations (8) and

$$\begin{aligned}(1-\beta) u_2'(C_2(\theta_1,\theta_2);\theta_2) \frac{\partial C_2(\theta_1,\theta_2)}{\partial \theta_2} \quad = \quad & \max\left\{0, \frac{\partial}{\partial \theta_2} u_1\left(C_1(\theta_1);\phi^{-1}(\theta_2)\right) + \beta \frac{\partial}{\partial \theta_2} u_2(C_2(\theta_1,\theta_2);\theta_2)\right. \\ & \left. -\frac{d}{d\theta_2} U^1\left(C^*\left(\phi^{-1}(\theta_2),\theta_2\right);\phi^{-1}(\theta_2),\theta_2\right)\right\}, \end{aligned} \tag{A-24}$$

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-24) 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 to $\theta_1 = \phi^{-1}(\theta_2)$, for any $\tilde{\theta}_1$ and $\phi(\theta_1) \geq \phi\left(\tilde{\theta}_1\right)$,

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

Hence $IC_1\left(\theta_1, \tilde{\theta}_1\right)$ holds if $\phi(\theta_1) \geq \phi\left(\tilde{\theta}_1\right)$.

Next, we show that $IC_1\left(\theta_1, \tilde{\theta}_1\right)$ holds if $\phi(\theta_1) < \phi\left(\tilde{\theta}_1\right)$. By the construction of C, $C(\theta_1, \phi(\theta_1)) = C^*(\theta_1, \phi(\theta_1))$ and $C\left(\tilde{\theta}_1, \phi(\theta_1)\right) = C^*\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right)$. Preference re-

[46] Suppose to the contrary that $\phi(\theta_1) = \phi\left(\tilde{\theta}_1\right)$ for some $\theta_1 \neq \tilde{\theta}_1$, and that C solves (2). If $C_1(\theta_1) \neq C_1\left(\tilde{\theta}_1\right)$ then preference reversal is violated. If instead $C_1(\theta_1) = C_1\left(\tilde{\theta}_1\right)$ then (since $\phi(\theta_1) = \phi\left(\tilde{\theta}_1\right)$) $u_1'(\cdot;\theta_1) \equiv u_1'\left(\cdot;\tilde{\theta}_1\right)$ and $\theta_1 = \tilde{\theta}_1$.

versal implies $C_1^*(\theta_1) \geq C_1^*\left(\tilde{\theta}_1\right)$. Since C^* solves (2), $V^1(C^*(\theta_1, \phi(\theta_1)); \theta_1, \phi(\theta_1)) \geq V^1\left(C^*\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right); \theta_1, \phi(\theta_1)\right)$. Lemma A-1 implies $U^1(C^*(\theta_1, \phi(\theta_1)); \theta_1, \phi(\theta_1)) \geq U^1\left(C^*\left(\tilde{\theta}_1, \phi\left(\tilde{\theta}_1\right)\right); \theta_1,\right.$ establishing $IC_1\left(\theta_1, \tilde{\theta}_1\right)$.

Next, we show that IC$_2$ is satisfied. By construction, C satisfies (8), and C_2 is increasing in θ_2. Note that

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

Hence $U^2\left(C\left(\theta_1, \tilde{\theta}_2\right); \theta_2\right)$ is indeed maximized at $\tilde{\theta}_2 = \theta_2$.

Finally, RC is certainly satisfied for $\theta_2 \leq \phi(\theta_1)$. 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 Proposition 5 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 $\frac{\partial C_2(\theta_1,\theta_2)}{\partial \theta_2} + \frac{\partial C_3(\theta_1,\theta_2)}{\partial \theta_2} \leq 0$, implying that RC is satisfied for all $\theta_2 \in (\phi(\theta_1), \bar{\theta}_2]$. **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 $\beta = 0.9$	0.947	0.996	n/a	0.998	n/a	n/a
(iii)	Cutoff $\tilde{\beta}$ for $\beta = 0.8$	0.840	0.885	0.980	0.887	0.982	n/a
(iv)	Cutoff $\tilde{\beta}$ for $\beta = 0.7$	0.736	0.773	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 $\beta = 0.9$	0.947	0.996	n/a	0.998	n/a	n/a
(iii)	Cutoff $\tilde{\beta}$ for $\beta = 0.8$	0.841	0.885	0.980	0.887	0.984	n/a
(iv)	Cutoff $\tilde{\beta}$ for $\beta = 0.7$	0.736	0.773	0.857	0.775	0.861	n/a

Panel (C): Investment

		$\min_{\theta_2} R(\theta_2) = 0.8$		$\min_{\theta_2} R(\theta_2) = 0.6$	
	γ	1	2	1	2
(i)	β^*	0.901	0.892	0.786	0.757
(ii)	Cutoff $\tilde{\beta}$ for $\beta = 0.9$	0.979	n/a	n/a	n/a
(iii)	Cutoff $\tilde{\beta}$ for $\beta = 0.8$	0.875	0.972	n/a	n/a
(iv)	Cutoff $\tilde{\beta}$ for $\beta = 0.7$	0.772	0.860	0.772	0.946