Long-Term Contracting with Time-Inconsistent Agents[*]

Daniel Gottlieb[†] and Xingtan Zhang[‡]

September 23, 2020

Abstract

We study contracts between naive present-biased consumers and risk-neutral firms. We show that the welfare loss from present bias vanishes as the contracting horizon grows. This is true both when bargaining power is on the consumers' and on the firms' side, when consumers cannot commit to long-term contracts, and when firms do not know the consumers' naiveté. However, the welfare loss from present bias does not vanish when firms do not know the consumers' present bias or when they cannot offer exclusive contracts.

JEL: D81, D86, D91, I18

[*]We thank Alessandro Lizzeri (editor), four anonymous referees, David Dillenberger, Drew Fudenberg, Faruk Gul, Michael Grubb, Botond Kőszegi, Shengwu Li, George Mailath, Pietro Ortoleva, Wolfgang Pesendorfer, Philipp Strack, Dmitry Taubinsky, Jeremy Tobacman, Juuso Toikka, Edward Van Wesep, Leeat Yariv, Muhamet Yildiz, John Zhu, and seminar participants at Harvard-MIT, LSE, Princeton, UC Berkeley, University of Colorado Boulder, Wharton, ESSET Gerzensee, SITE Psychology and Economics, the 2019 AEAs, and the World Congress of the Econometric Society for comments. Gottlieb is grateful for the hospitality of the Economics Department at Harvard University, where part of this research was conducted. Zhang acknowledges financial support from the Capital Markets Program at University of Colorado Boulder.

[†]London School of Economics, d.gottlieb@lse.ac.uk.

[‡]University of Colorado Boulder, xingtan.zhang@colorado.edu.

1 Introduction

A vast literature in behavioral economics studies markets with present-biased consumers who underestimate their bias ("partial naiveté"). An important finding from this literature is that the equilibrium is inefficient and regulation that accounts for internalities can increase welfare.[1] Models in this literature generally have only three periods, which is the minimum needed for present bias to play a role. But this is an unrealistic assumption since, in these models, periods are thought to be very short, typically no more than a day (O'Donoghue and Rabin, 2015).

Our paper considers a general contracting model with partially naive present-biased consumers and an arbitrary number of periods. We show that to exploit consumer naiveté firms offer contracts with two options at each point in time: a "baseline" that provides high consumption in the future in exchange for low consumption today, and an "alternative" option that provides greater consumption today but less consumption in the future. In every period, consumers think they will pick the baseline but pick the alternative option instead, effectively postponing the reduction in consumption to the next period. As a result, the equilibrium has smooth consumption in all but the last period. Because the relative weight on the last period shrinks as the contracting horizon grows, the consumption on the equilibrium path of present-biased consumers converges to the path that maximizes their long-run preferences. Therefore, the welfare loss from present bias vanishes as the contracting horizon grows.

This result can be interpreted in two ways. If one takes the models as currently formulated as good approximations of reality, then it suggests that there is no role for regulation that corrects for present bias as long as contractual relationships are long enough. If instead, one believes that inefficiency is prevalent in markets with present-biased consumers, then something must be missing from how these markets are often modeled. To understand which features can prevent efficiency in markets with present-biased consumers, we extend

[1] See, for example, Gruber and Kőszegi (2001); O'Donoghue and Rabin (2003); DellaVigna and Malmendier (2004); Heidhues and Kőszegi (2010).

the model in several directions.

First, we drop the assumption that consumers can commit to long-term contracts. Such "one-sided commitment" is prevalent in many markets (e.g., life insurance, long-term care, annuities, mortgages, and car loans).[2] For a fixed horizon, removing commitment power can help present-biased consumers, who are tempted to overborrow. Since firms would not lend to consumers who can walk away from contracts, removing commitment power limits their ability to borrow, increasing welfare when consumers are sufficiently present biased.

We then generalize the vanishing inefficiency result for settings with one-sided commitment, showing that the equilibrium converges to the path that maximizes the consumer's long-run preferences subject to non-lapsing constraints. Because the equilibrium with commitment maximizes long-run preferences without these constraints, removing commitment power cannot improve welfare when the contracting horizon is long enough.

Second, we show that the vanishing inefficiency result also holds when firms, rather than consumers, have the bargaining power. However, the equilibrium converges to a different point on the Pareto frontier. In the limit, consumers are worse off (according to their long-run preferences) than with their outside option and firms obtain higher profits than if they were facing consumers who maximized their long-run preferences. Nevertheless, the outcome is efficient and only distributional concerns would justify interventions.

Third, we study how private information affects the equilibrium. When firms do not know the consumer's naiveté, the equilibrium remains unchanged and, therefore, is still asymptotically efficient. Adding a sophisticated type also does not affect the equilibrium contract of naive consumers, so the inefficiency from their contracts still vanishes as the horizon grows. However, when firms do not know the consumer's present bias, they face an adverse selection problem that leads them to provide insufficient savings to more present-biased types, and the equilibrium is no longer asymptotically efficient.

Fourth, we assume that consumers can simultaneously sign contracts with multiple

[2] In long-term insurance markets – such as life insurance, long-term care insurance, or annuities – policyholders are allowed to cancel their policies at all times, but firms cannot drop them. In mortgages and other credit markets, borrowers can prepay their debt, but debtors cannot force them to repay before the contract is due.

firms, so contracts are not exclusive. This is the case, for example, with credit cards and many other loans. We show that contract non-exclusivity prevents any commitment from being provided, so the equilibrium allocation remains inefficient no matter how long the horizon is.

Fifth, we consider more general forms of present bias (beyond quasi-hyperbolic discounting). Since quasi-hyperbolic discounting only distorts decisions involving the current period, one might think that with more general discount functions, the efficiency result would break. We characterize the equilibrium for general present-biased discounting and show that this conjecture is incorrect.

And sixth, we study contracting over effort and show that, in that case, the equilibrium path is determined by the agent's short-run discount, so the equilibrium remains inefficient regardless of the horizon. With effort, firms offer a baseline requiring all effort to be exerted in the immediate future. On the equilibrium path, the agent keeps postponing some effort one period into the future according to their short-run discount factor, which is lower than their long-run discount factor due to present bias.

The main message of our paper is that the inefficiency of markets with present-biased consumers crucially depends on the length of the relationship, on how easy it is to contract with multiple firms, and on the information that firms have. When contracts are exclusive and firms know the consumer's present bias, the equilibrium is approximately efficient as long as the contracting horizon is long enough.

The paper is organized as follows. In Section 2, we present the basic model. In Section 3, we discuss several extensions and their implications for welfare. Subsection 3.1 drops the assumption that consumers can commit to long-term contracts, 3.2 assumes that the bargaining power is on the firm's rather than the consumer's side, 3.3 introduces heterogeneity, 3.4 assumes that contracts are not exclusive, 3.5 relaxes the assumption of quasi-hyperbolic discounting, and 3.6 considers contracting over effort. Then, Section 4 concludes.

Related Literature. Our paper fits into a recent literature on contracting with behavioral agents, summarized in Kőszegi (2014) and Grubb (2015). The basic model in Section

2 builds on the credit card model of Heidhues and Kőszegi (2010) by considering more than two consumption periods and allowing for uncertainty. Our paper is also related to a literature that studies commitment contracts with time-inconsistent agents (c.f. Amador et al. (2006); Halac and Yared (2014); Galberti (2015); Bond and Sigurdsson (2017)).[3]

Subsection 3.1 is related to a literature on dynamic risk-sharing with one-sided commitment. Several papers show that front-loaded payment schedules help mitigate a consumer's lack of commitment power. For example, Hendel and Lizzeri (2003) theoretically and empirically examine how life insurers mitigate reclassification risk by offering front-loaded policies. More recently, Handel et al. (2017) and Atal et al. (2018) show that front-loaded long-term health insurance contracts can produce substantial welfare gains by insuring policyholders against reclassification risk. The main difference between these models and ours is that consumers in our model are dynamically inconsistent.

2 Basic Model

There is one consumer (agent) and a finite number of firms. Time is discrete and finite. To allow for arbitrary non-stationary settings, we model the stochastic environment as follows. There is a finite state space $\mathbb{S}_t$ for each $t \in \mathbb{N}$. The agent earns income $w(s_t)$ at state s_t. Let $p(s_t|s_\tau)$ denote the probability of reaching state s_t conditional on state s_τ. We say that state s_t follows state s_τ if $p(s_t|s_\tau) > 0$. A state specifies all previously realized uncertainty, so a state cannot follow two different states. We consider the T-period truncation of this setting; that is, an environment with state spaces $\mathbb{S}_t$ and conditional probabilities $p(\cdot|\cdot)$ up

[3]There are two key differences between our paper and this literature. First, this literature considers sophisticated agents, whereas our primary focus is on partially naive agents. Second, we study a different incentive aspect. This literature studies the trade-off between commitment and flexibility (agents have commitment power but, because they face an unverifiable taste shock, they value the flexibility to adjust to different taste shocks), whereas, in Section 3.1, we study the agent's incentive to lapse and re-contract with other firms. Our paper is also related to Bisin et al. (2015), who study the interaction between government policy and private commitments by present-biased voters, Heidhues and Strack (2019) who characterize stopping behavior by time-inconsistent agents, and to Harris and Laibson (2001) and Cao and Werning (2018), who study the Markov equilibria in infinite-horizon problems with sophisticated consumers and show there can be multiple non-smooth equilibria. Multiplicity and non-smoothness do not arise in our setting because our model has a finite (albeit arbitrary) horizon.

to period T, at which point the game ends.

Without loss of generality, we assume that no uncertainty is realized before the initial period: $\mathbb{S}_1 = \{\emptyset\}$. Let $E[\cdot|s_t]$ denote the expectation operator conditional on state s_t and let $E[\cdot]$ denote the unconditional (time-1) expectation. By taking degenerate distributions, our framework allows for deterministic income paths. Also, since the probabilities of reaching future states may depend on the current state, our framework also allows for persistent shocks, which is important to encompass environments with reclassification risk.

Firms are risk neutral and can freely save or borrow at the interest rate $R \geq 1$, so that each firm maximizes its expected discounted profits. The expected profits at state s_τ of a firm that collects state-dependent payments $\{\pi(s_t)\}_{t\geq\tau}$ are

$$E\left[\sum_{t\geq\tau} \frac{\pi(s_t)}{R^{t-\tau}} \;\middle|\; s_\tau\right].$$

The agent has quasi-hyperbolic discounting and needs a firm to transfer consumption across states.[4] At state s_τ, the agent evaluates the state-dependent consumption $\{c(s_t)\}_{t\geq\tau}$ according to

$$u(c(s_\tau)) + \beta E\left[\sum_{t>\tau} \delta^{t-s} u(c(s_t)) \;\middle|\; s_\tau\right], \tag{1}$$

where $\beta \in (0,1)$, $\delta \in (0,1]$ and $u : \mathbb{R}_+ \to \mathbb{R}$ is continuous, strictly increasing, strictly concave, and twice continuously differentiable in the interior of its domain.[5] We are interested in time-inconsistent consumers who underestimate their bias – i.e., they are *partially naive* as defined by O'Donoghue and Rabin (1999). Such a consumer believes that, in all future periods, he will behave like someone with time-consistency parameter $\hat{\beta} \in (\beta, 1]$. For brevity, we refer to a partially naive time-inconsistent consumer simply as a *time-*

[4]The assumption that the agent cannot perfectly transfer resources across states without a firm is not innocuous (see Subsection 3.4). As shown by Augenblick et al. (2015), it may also be related to why lab experiments often fail to find evidence of present bias. This assumption is probably better suited for contracts involving larger amounts (such as mortgages or life insurance policies) which cannot be easily financed through other means.

[5]Continuity rules out utility functions that are unbounded from below, which is important for equilibrium existence. One can incorporate unbounded utility functions by assuming that the consumer has a minimum (subsistence) level of consumption and renormalizing that level to zero.

inconsistent consumer.[6] As a benchmark, we also consider the case of time-consistent consumers ($\hat{\beta} = \beta = 1$). Following most of the literature, we take the agent's long-run preferences as the relevant ones in our welfare calculations.[7] Therefore, consumers maximize welfare in the time-consistent benchmark but not when they are time-inconsistent.

For now, we assume that firms know the consumer's preferences. We also assume that the consumer has all bargaining power and that all parties can commit to long-term contracts, so the consumer makes a take-it-or-leave-it offer of a contract in the first period, which is honored until the game ends. These three assumptions are dropped in Section 3.

When an individual is not time consistent, his ranking of consumption streams depends on when the stream is evaluated. As usual, we model the behavior of such an agent by treating his decision in each period as if it was decided by a different "self." Because the consumer is naive, each self may mispredict how his future selves will choose. We consider *perception-perfect equilibria* (O'Donoghue and Rabin, 1999, 2001), which require that: (i) each self picks an optimal strategy given its prediction of how future selves will behave, and (ii) predictions about the behavior of future selves must be consistent with how a future self with time-consistency parameter $\hat{\beta}$ would choose.[8]

2.1 Time-Consistent Consumers

As a benchmark, we first consider a time-consistent consumer. Because parties can commit to long-term contracts, the equilibrium consumption maximizes the agent's utility in period 1,

$$E\left[\sum_{t=1}^{T} \delta^{t-1} u\left(c\left(s_t\right)\right)\right], \tag{2}$$

[6]Online Appendix D considers sophisticates, who fully understand their time inconsistency ($\hat{\beta} = \beta$).

[7]See, e.g., DellaVigna and Malmendier (2004); O'Donoghue and Rabin (1999, 2001).

[8]See Online Appendix E for a formal definition. As we show there, there is no loss of generality in restricting attention to pure strategies. Our game-theoretical equilibrium concept coincides with the non-strategic competitive equilibrium of Heidhues and Kőszegi (2010). We formulate the model as a game because it can be more straightforwardly generalized to settings with one-sided commitment (Subsection 3.1), monopoly (3.2), private information (3.3), and contracting over effort (3.6).

subject to the zero-profits constraint,

$$\sum_{t=1}^{T} E\left[\frac{w\left(s_t\right)-c\left(s_t\right)}{R^{t-1}}\right]=0. \tag{3}$$

Indeed, no firm would accept a contract with negative expected profits. If profits were positive, the agent would benefit by offering a contract with slightly higher consumption. Because the objective function in (2) is strictly concave and (3) is a linear constraint, there is a unique solution. So, any equilibrium of the game provides the same consumption, which solves the program above. Let W_T^C denote the equilibrium welfare of the time-consistent consumer, which evaluates the objective (2) at the equilibrium consumption.

2.2 Time-Inconsistent Consumers

Before presenting a general analysis of equilibrium with time-inconsistent consumers, we start with a simple illustrative example. There are $T=3$ periods and there is no uncertainty. The net interest rate is zero ($R=1$) and the agent's total income equals one ($w_1+w_2+w_3=1$). The agent is fully naive ($\hat{\beta}=1$), has discount factors $\beta=\frac{1}{2}$ and $\delta=1$, and utility function $u(c)=\sqrt{c}$.

Based on Subsection 2.1, one may think that the equilibrium consumption maximizes the agent's utility in period 1 subject to zero profits:

$$\max_{c_1,c_2,c_3} \sqrt{c_1}+\beta\left(\sqrt{c_2}+\sqrt{c_3}\right) \tag{4}$$

subject to

$$c_1+c_2+c_3=1. \tag{5}$$

We will show that this is not the case. To see this, first note that the solution to this program is $c_1=\frac{2}{3}$, $c_2=c_3=\frac{1}{6}$, which gives the agent a utility of $\sqrt{3/2}$.

Suppose a firm decides to offer a contract that gives $c_1=\frac{8}{27}$ in the first period and allows the agent to pick between two different options in the second period: a *baseline* and an *alternative* option. The baseline provides as little consumption as possible in the second period in exchange for a high consumption in the future: $c_2(B) = 0$ and $c_3(B) = \frac{50}{27}$. The alternative option offers a smoother path: $c_2(A) = \frac{8}{27}$ and $c_3(A) = \frac{7}{81}$.

Since the agent thinks that his future selves are perfectly patient, he believes that he will pick the baseline option:

$$\sqrt{c_2(B)} + \sqrt{c_3(B)} \approx 1.36 > 0.84 \approx \sqrt{c_2(A)} + \sqrt{c_3(A)},$$

which gives him the same perceived utility as in the solution of program (4)-(5):

$$\sqrt{c_1} + \beta\left[\sqrt{c_2(B)} + \sqrt{c_3(B)}\right] = \sqrt{3/2}.$$

Therefore, the agent accepts to switch to this new contract.

However, in period 2, the agent picks the alternative option instead of the baseline, since

$$\sqrt{c_2(A)} + \beta\sqrt{c_3(A)} \approx 0.69 > 0.68 \approx \sqrt{c_2(B)} + \beta\sqrt{c_3(B)}.$$

And because the agent ends up with the alternative option, the firm makes a profit of

$$1 - c_1 - c_2(A) - c_3(A) \approx 0.32.$$

That is, a flexible contract allows the firm to exploit the agent's incorrect beliefs and make positive profits. The firm exploits the difference in beliefs by offering a baseline option with very low consumption in period 2 ($c_2(B) = 0$) in exchange for a large future consumption ($c_3(B) = \frac{50}{27}$). And while the agent thinks that he will choose this baseline option, he ends up switching to the alternative option, which has a lower NPV but a much higher immediate consumption ($c_2(A) = \frac{8}{27},\ c_3(A) = \frac{7}{81}$).

Having shown that we cannot have an equilibrium with inflexible contracts in this simple example, we now characterize the equilibrium in the general case.

2.2.1 Equilibrium

Any contract that is accepted with positive probability must maximize the consumer's utility in period 1 subject to two types of constraints: zero profits, which is the same as before, and incentive constraints, which are due to consumer naiveté.

Because the consumer mispredicts his future preferences, he may disagree with the firm about the actions that his future selves will take. So we need to distinguish between what the consumer thinks he will choose and what firms think that the consumer will choose (which we interpret as the correct beliefs). This disagreement gives rise to two sets of incentive constraints. Following Heidhues and Kőszegi (2010), we refer to them as *perceived choice (PC)* and *incentive compatibility (IC)* constraints.

PC requires the consumer to believe that his future selves will choose the actions that maximize his perceived utility. IC requires firms to believe that the consumer's future selves will choose the actions that maximize the consumer's true utility. The option that the consumer thinks that his future selves will choose is called the *baseline* option (B). The option that firms think that the consumer's future selves will choose is called the *alternative* option (A). In principle, these options can coincide, in which case the consumer and the firms agree on which actions will be chosen. But we will show that, in equilibrium, these options are always different.[9]

A time-t option history h^t is a list of options chosen by the consumer up to time t: $h^1 = \emptyset$, $h^2 \in \{(A), (B)\}$, $h^3 \in \{(A, A), (A, B), (B, A), (B, B)\}$, etc. Since there are no actions after the last period, there is no space for disagreement at $t = T$, so that $h^T = h^{T-1}$. Figure 1 depicts the option histories when there are four periods.

A *consumption vector* specifies the agent's consumption in all states that happen with positive probability for all option histories:[10]

$$\mathbf{c} \equiv \{(c(s_1), c(s_2, h^2), c(s_3, h^3), \cdots, c(s_T, h^T)) : \; p(s_2|s_1)p(s_3|s_2)\cdots p(s_T|s_{T-1}) > 0\}.$$

A *consumption path* specifies the consumption that happens with positive probability using

[9]See Yildiz (2003) for another game in which different prior beliefs may cause players to disagree about which actions will be taken in equilibrium.

[10]Since a state of the world encodes all previously realized uncertainty, the distribution over future states conditional on s_t can only have full support in the trivial case in which no uncertainty is realized until state s_t. We omit the subscript T from the consumption vector $\mathbf{c}$ to simplify notation.

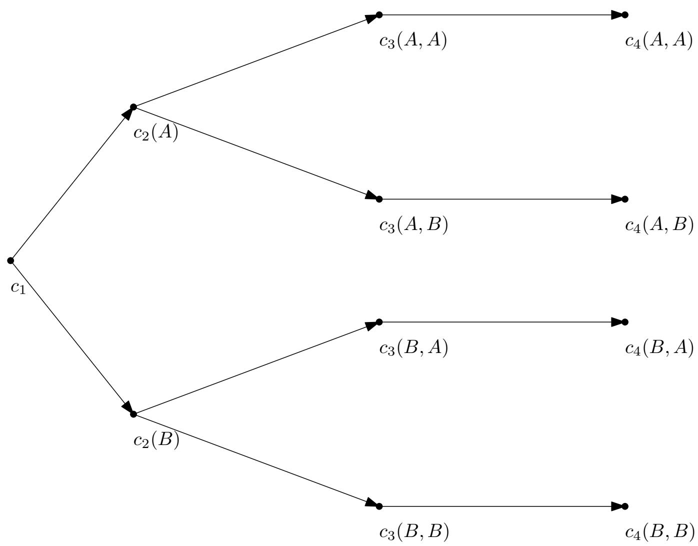

Figure 1: Option histories with $T = 4$ and no uncertainty. The consumer initially believes he will choose the baseline in each node, which would give him consumption $(c_1, c_2(B), c_3(B, B), c_4(B, B))$. In each node, he switches to the alternative. In period 2, he receives $c_2(A)$ and thinks that he will get $(c_3(A, B), c_4(A, B))$. Then, he switches to the alternative again in period 3, obtaining $(c_3(A, A), c_4(A, A))$. The consumption vector consists of consumption in all nodes, whereas the consumption path is $(c_1, c_2(A), c_3(A, A), c_4(A, A))$. With uncertainty, consumption also depends on the state of the world.

correct beliefs about the options that the consumer chooses:

$$\mathbf{c}^E \equiv \{(c(s_1), c(s_2, A), c(s_3, A, A), \cdots, c(s_T, A, ..., A)) : \ p(s_2|s_1)p(s_3|s_2)\cdots p(s_T|s_{T-1}) > 0\}.$$

Note that, unlike the consumption vector, the consumption path only includes outcomes conditional on the consumer repeatedly picking option A.

The *equilibrium program* (P) is:

$$\max_{\{c(s_t,h^t)\}} u(c(s_1)) + \beta E\left[\sum_{t=2}^{T} \delta^{t-1} u(c(s_t, B, B, ..., B))\right], \tag{6}$$

subject to

$$\sum_{t=1}^{T} E\left[\frac{w\left(s_t\right) - c\left(s_t, A, A, ..., A\right)}{R^{t-1}}\right] = 0, \tag{Zero Profits}$$

$$\begin{aligned} &u(c(s_\tau, (h^{\tau-1}, B))) + \hat{\beta} E\left[\textstyle\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t, (h^{\tau-1}, B, B, ..., B)\right)\right)\Big|\ s_\tau\right] \\ &\geq u(c(s_\tau, (h^{\tau-1}, A))) + \hat{\beta} E\left[\textstyle\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t, (h^{\tau-1}, A, B, ..., B)\right)\right)\Big|\ s_\tau\right], \end{aligned} \tag{PC}$$

and

$$u(c(s_\tau,(h^{\tau-1},A))) + \beta E\left[\textstyle\sum_{t>\tau}\delta^{t-\tau}u\left(c\left(s_t,(h^{\tau-1},A,B,...,B)\right)\right)\middle|\ s_\tau\right]$$
$$\geq u(c(s_\tau,(h^{\tau-1},B))) + \beta E\left[\textstyle\sum_{t>\tau}\delta^{t-\tau}u\left(c\left(s_t,(h^{\tau-1},B,B,...,B)\right)\right)\middle|\ s_\tau\right]. \tag{IC}$$

The following lemma establishes that the equilibrium program (P) characterizes the equilibrium consumption vector:

Lemma 1. ***c*** *is the consumption vector in a perception-perfect equilibrium if and only if it solves program (P).*

2.2.2 Auxiliary Program

Consider a dynamically consistent agent who differs from the one described in Subsection 2.1 in that he discounts consumption in the last period by an additional factor β. The equilibrium consumption for this agent solves the following *auxiliary program*:

$$\max_{\{c(s_t)\}} E\left[\sum_{t=1}^{T-1}\delta^{t-1}u\left(c(s_t)\right) + \beta\delta^{T-1}u\left(c(s_T)\right)\right], \tag{7}$$

subject to the zero-profits constraint (3). The auxiliary program has a unique solution since the objective function is strictly concave and the constraint is linear.

The following lemma establishes that the consumption path for time-inconsistent agents coincides with the solution of the auxiliary program:

Lemma 2. $\boldsymbol{c}^E$ *is the consumption path in a perception-perfect equilibrium if and only if it solves the auxiliary program.*

The auxiliary program highlights that underweighting consumption in the last period is the only distortion from time inconsistency in this model. To illustrate the lemma, suppose there are four periods and income is constant (w). Because there is no uncertainty, we can omit the state of the world from histories. The equilibrium contract solves:

$$\max_{\{c(h^t)\}} u(c_1) + \beta[\delta u(c_2(B)) + \delta^2 u(c_3(B,B)) + \delta^3 u(c_4(B,B))], \tag{8}$$

subject to

$$c_1 + \frac{c_2(A)}{R} + \frac{c_3(A,A)}{R^2} + \frac{c_4(A,A)}{R^3} = w\left(1 + \frac{1}{R} + \frac{1}{R^2} + \frac{1}{R^3}\right), \tag{9}$$

$$u(c_2(B)) + \hat{\beta}[\delta u(c_3(B,B)) + \delta^2 u(c_4(B,B))] \geq u(c_2(A)) + \hat{\beta}[\delta u(c_3(A)) + \delta^2 u(c_3(A,B))], \tag{10}$$

$$u(c_2(A)) + \beta[\delta u(c_3(A,B)) + \delta^2 u(c_4(A,B))] \geq u(c_2(B)) + \beta[\delta u(c_3(B,B)) + \delta^2 u(c_4(B,B))], \tag{11}$$

$$u(c_3(A,B)) + \hat{\beta}\delta u(c_4(A,B)) \geq u(c_3(A,A)) + \hat{\beta}\delta u(c_4(A,A)), \tag{12}$$

$$u(c_3(A,A)) + \beta\delta u(c_4(A,A)) \geq u(c_3(A,B)) + \beta\delta u(c_4(A,B)), \tag{13}$$

where (9) is the zero profits constraint, (10) and (12) are the PC constraints, and (11) and (13) are the IC constraints in periods 2 and 3, respectively.

First, note that the IC constraints (11) and (13) must both bind. Otherwise, it would be possible to achieve a higher utility by increasing $c_4(B,B)$ and $c_4(A,B)$. Substitute these two binding constraints in the objective to eliminate $c_4(B,B)$ and $c_4(A,B)$:

$$u(c_1) + \delta u(c_2(A)) + \delta^2 u(c_3(A,A)) + \beta\delta^3 u(c_4(A,A)) - (1-\beta)[\delta u(c_2(B)) + \delta^2 u(c_3(A,B))]. \tag{14}$$

Use the binding ICs again to rewrite the PC constraints (10) and (12) as monotonicity conditions:

$$c_2(A) \geq c_2(B), \tag{15}$$

and

$$c_3(A,A) \geq c_3(A,B). \tag{16}$$

So the program reduces to the maximization of (14) subject to zero profits (9) and the monotonicity conditions (15) and (16). Because the objective function (14) is decreasing in $c_2(B)$ and $c_3(A,B)$, the solution entails

$$c_2(B) = c_3(A,B) = 0. \tag{17}$$

Equation (17) implies that the monotonicity conditions (15) and (16) automatically hold. Substitute $c_2(B) = c_3(A, B) = 0$ in the objective function (14) to obtain the auxiliary program:

$$\max_{\{c_1, c_2(A), c_3(A,A), c_4(A,A)\}} u(c_1) + \delta u(c_2(A)) + \delta^2 u(c_3(A, A)) + \beta\delta^3 u(c_4(A, A)). \tag{18}$$

subject to the zero-profit condition (9).

To understand why the baseline offers the lowest consumption possible in period 2, consider the following perturbation argument. Starting from any interior $c_2(B)$, lower $u(c_2(B))$ by $\epsilon > 0$ and raise $u(c_4(B, B))$ by $\frac{\epsilon}{\beta\delta^2}$. This perturbation keeps the IC (11) unchanged but increases the objective (8) by $\delta\beta\left(\frac{1}{\beta} - 1\right)\epsilon > 0$. This is because the objective function evaluates consumption from the perspective of self 1, whereas the IC (11) evaluates it according to self 2's preferences. Since the agent is present biased, shifting consumption into the future while keeping his self-2 utility constant increases his utility from the perspective of self 1. Therefore, the solution of the program shifts as much consumption as possible to the last period: $c_2(B) = 0$.

To summarize, the consumer thinks he will pick the baseline option. Since he underestimates the present bias of his time-2 self, the baseline provides as little consumption as possible in period 2 in return for higher future consumption. The firm accepts to offer higher consumption in the future because it knows that the consumer's future selves will not pick the baseline. Instead, the firm offers an alternative option that induces the consumer's future self to switch.

With more than 4 periods, the consumer initially thinks that he will follow the option path $(B, B, ..., B)$. The firm designs the options to induce him to keep switching to the alternative. Because each future self is more present biased than the consumer anticipates, the cheapest way to induce switching from $(B, B, ..., B)$ to $(A, B, ..., B)$ is to postpone the reduction in consumption from period 2 to period 3. The switching decision is made by the period-2 self, which discounts the future starting at period 3 by an additional β. Then, in period 3, the alternative option needs to again induce the agent to switch from

$(A, B, B, ..., B)$ to $(A, A, B, ..., B)$, which is done by postponing the reduction in consumption into period 4. That decision is made by the period-3 self, which discounts the future starting at period 4 by an additional β. This argument proceeds until we reach period $T - 1$, when the consumer can no longer push the reduction in consumption further into the future. Since there is only one period left, consumer and firms can no longer disagree about when the reduction in consumption will happen. Thus, on the equilibrium path $(A, A..., A)$, the factor β only applies to the last period. Intuitively, each self postpones the reduction in consumption one period into the future (by choosing the alternative option rather than the baseline) until the last period.

Recall that we used the binding ICs – equations (11) and (13) – and equation (17) to eliminate the baseline options from the auxiliary program. Substituting the solution of the auxiliary program back in these equations, we can recover the baseline options. Since neither the auxiliary program nor these equations depend on the consumer's naiveté parameter, it follows that, in equilibrium, both the baseline and the alternative options are not functions of $\hat{\beta}$.

Corollary 1. *There exists a perception-perfect equilibrium that does not depend on the consumer's naiveté $\hat{\beta} \in (\beta, 1]$. Moreover, any perception-perfect equilibrium has the same consumption path, which is continuous in $\beta \in (0, 1]$.*

2.2.3 Vanishing Inefficiency

We now use Lemma 2 to obtain our main result. Let W_T^I denote the equilibrium welfare of the time-inconsistent consumer, which evaluates the consumption path according to the agent's long-run preferences (2), and recall that W_T^C is the welfare of a time-consistent consumer. Since the time-consistent consumer maximizes welfare, the welfare loss from dynamic inconsistency is $W_T^C - W_T^I \geq 0$.

Theorem 1. *Suppose u is bounded and $\delta < 1$. Then, $lim_{T \nearrow +\infty} \left(W_T^C - W_T^I\right) = 0$.*

The theorem states that the welfare loss from dynamic inconsistency converges to zero as the contracting length grows. The assumption that u is bounded and $\delta < 1$ ensures that the discounted welfare converges.[11]

Recall that the only inefficiency from time inconsistency is underweighting the last period. Because the effect of the last period vanishes as the number of periods grows, the solution of the auxiliary program converges to the equilibrium consumption with time-consistent consumers as $T \nearrow +\infty$. So, even though the time-inconsistent consumer does not maximize his welfare function and has incorrect beliefs, in any equilibrium, he obtains approximately the maximum welfare possible if the number of periods is large. In fact, it is precisely the fact that the agent has incorrect beliefs that causes the inefficiency to vanish, since the optimal way to exploit the agent's naiveté is to offer contracts that postpone the reduction in consumption until the last period. In Online Appendix D, we show that the inefficiency does not vanish when the consumer is sophisticated.

3 Extensions and Limitations

3.1 One-Sided Commitment

We now assume that consumers cannot commit to long-term contracts. As argued in the introduction, this type of one-sided commitment is common in many markets, including life insurance, long-term care, annuities, mortgages, and car loans. Moreover, regulations that allow consumers to terminate agreements at will are often motivated by an attempt to protect them. However, in standard models, removing a rational consumer's commitment power can only hurt the consumer.

Our setting is a natural candidate for studying the effect of regulating commitment power, because committing to future actions and lapsing on previous agreements are inherently intertemporal decisions, and present bias is the most well-studied bias in intertemporal decision-making. Moreover, there is evidence that present bias is an important feature

[11] Along with $\delta < 1$, assuming bounded utility is sufficient to ensure that discounted welfare converges, but can be substantially relaxed. When the discounted welfare does not converge, similar results can be obtained using other criteria (see Subsection 3.5).

in some credit markets where regulation prevents consumers from being able to commit to long-term contracts, such as in mortgage or credit card markets.[12]

We keep the assumption that the consumer has all bargaining power and model one-sided commitment as follows. The consumer offers a contract in each period. If a firm has accepted a contract, the consumer decides whether to keep it or replace it with a new one. If multiple firms accept a contract, the consumer picks each of them with some positive probability.[13]

3.1.1 Benchmark: Time-Consistent Consumers

Consider first the benchmark case of a time-consistent consumer. With one-sided commitment, the consumer can freely switch to a new contract ("lapse") as long as he can find a firm willing to provide such contract. To obtain the equilibrium consumption, there is no loss of generality in restricting attention to contracts in which the consumer never lapses.[14] Therefore the equilibrium consumption must satisfy *non-lapsing constraints*, which require that the consumer's outside option at any state does not exceed the value from keeping the original contract.[15]

[12]See Schlafmann (2016), Ghent (2011) and Atlas et al. (2017) for evidence of present bias in mortgage markets, and Meier and Sprenger (2010) for credit card markets. Our equilibrium pattern of repeatedly postponing repayments is broadly consistent with the findings from Carter et al. (2017) on payday loans. See also Sulka (2020), who calibrates a model of retirement saving and finds that present bias lowers the representative household's pension wealth by 10%.

[13]In this formulation, the only cost of walking away from a contract are the forgone payments made into that contract. This assumption is appropriate to insurance settings, where policyholders are allowed to drop coverage at no additional cost by stopping to pay their premiums (see Hendel and Lizzeri 2003; Handel et al. 2017). In some markets (such as mortgages or car loans), the cost also includes the loss of the collateral, whereas in others there are also reputational costs. With present-biased preferences, the timing of these costs matters. For example, with immediate costs, a naive consumer may think he will incur the cost and leave, but end up deciding not to. With delayed costs, a naive consumer may think he will not walk away from a contract, but end up doing so.

[14]To see this, consider an equilibrium in which the consumer lapses in some state of the world, replacing the original contract with a contract from another firm. Since the other firm cannot lose money by offering this new contract, the old firm could have accepted a contract that substituted the terms of the old contract from this period on with the terms of the new contract, and the consumer would have remained with the old firm.

[15]When a non-lapsing constraint binds, there are also equilibria in which the consumer lapses and recontracts with another firm. These equilibria are equivalent to the one with no lapsing in the sense that the consumer obtains the same consumption and all firms make the same profits.

The outside option at state s_τ is defined by the recursion:

$$V^C(s_\tau) \equiv \max_{\{c(s_t)\}} u\left(c(s_\tau)\right) + E\left[\left.\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t\right)\right)\right| s_\tau\right], \tag{19}$$

subject to

$$\sum_{t=\tau}^{T} E\left[\left.\frac{w\left(s_t\right) - c\left(s_t\right)}{R^{t-\tau}}\right| s_\tau\right] = 0, \tag{20}$$

and

$$u\left(c(s_{\tilde{\tau}})\right) + E\left[\left.\sum_{t>\tilde{\tau}} \delta^{t-\tilde{\tau}} u\left(c\left(s_t\right)\right)\right| s_{\tilde{\tau}}\right] \geq V^C(s_{\tilde{\tau}}),\ \ \forall s_{\tilde{\tau}} \text{ with } p(s_{\tilde{\tau}}|s_\tau) > 0. \tag{21}$$

Equation (20) is the zero-profits condition and (21) is the non-lapsing condition of the new contract. The equilibrium consumption with one-sided commitment solves this program at period 1.

While one can use program (19)-(21) to obtain the equilibrium consumption by backward induction, there is an easier approach when the consumer is time consistent. Consider, instead, the program that replaces non-lapsing constraints by the requirement that, at each point in time, the expected future income cannot exceed the expected future consumption:

$$\sum_{t \geq \tau} E\left[\left.\frac{w\left(s_t\right) - c\left(s_t\right)}{R^{t-\tau}}\right| s_\tau\right] \leq 0,\ \forall s_\tau. \tag{22}$$

We refer to (22) as *front-loading constraints*. Contracts satisfying (20) and (22) cannot have negative accumulated profits at any time, so the consumer initially overpays the firm and is repaid later. This overpayment discourages the consumer from switching contracts.

In general, front-loading constraints are weaker than non-lapsing constraints: if the continuation contract gave positive expected profits at some state, a consumer would be able to increase his utility by replacing it with another contract that gives zero profits. When consumers are dynamically consistent, however, maximizing (19) subject to either (20) and (21) or (20) and (22) gives the same solutions.[16] Since (19) is a strictly concave function

[16]To see this, suppose a solution to the program subject to (20) and (22) did not satisfy the non-lapsing constraints (21). Then, there would exist a continuation contract that gives zero profits while increasing the

and (20) and (22) are linear constraints, consumption is the same in any equilibrium of the game.

3.1.2 Time-Inconsistent Consumers

We now turn to the more interesting case of a time-inconsistent consumer. As with time-consistent consumers, the equilibrium with one-sided commitment must satisfy non-lapsing constraints. Yet, because the parties may disagree about which options will be chosen, we need to distinguish between non-lapsing constraints according to the beliefs of the consumer and the beliefs of firms. Equilibrium requires both of them to hold. To write down these constraints, we first define the outside options recursively.

The *actual outside option* at state s_τ given option history h^τ is the highest utility that the consumer can obtain at that state:

$$V^I(s_\tau) \equiv \max_{\{c(s_t,h^t_\tau)\}_{t\geq\tau}} u\left(c(s_\tau,h^\tau_\tau)\right) + \beta E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c(s_t,(h^\tau_\tau,B,...,B))\right)\middle| s_\tau\right],$$

subject to (PC), (IC), the zero-profits constraint

$$E\left[\sum_{t\geq\tau} \frac{w(s_t)-c(s_t,(h^\tau_\tau,A,A,...,A))}{R^{t-\tau}}\middle| s_\tau\right] = 0, \tag{23}$$

and the non-lapsing constraints

$$u\left(c_\tau(s_{\tilde\tau},(h^{\tilde\tau-1}_\tau,A))\right) + \beta E\left[\sum_{t>\tilde\tau} \delta^{t-\tilde\tau} u\left(c(s_t,(h^{\tilde\tau-1}_\tau,A,B,...,B))\right)\middle| s_{\tilde\tau}\right] \geq V^I(s_{\tilde\tau}), \tag{24}$$

$$u\left(c_\tau(s_{\tilde\tau},(h^{\tilde\tau-1}_\tau,B))\right) + \hat\beta E\left[\sum_{t>\tilde\tau} \delta^{t-\tilde\tau} u\left(c(s_t,(h^{\tilde\tau-1}_\tau,B,B,...,B))\right)\middle| s_{\tilde\tau}\right] \geq \hat V^I(s_{\tilde\tau}), \tag{25}$$

for all $s_{\tilde\tau}$ following s_τ and all $h^{\tilde\tau}_\tau$ that are continuation histories of h^τ, where $\hat V$ is the "perceived outside option," which we define next.

The *perceived outside option* at state s_τ given option history h^τ is the highest utility

consumer's continuation utility. Substituting the original continuation contract by this new one would then increase the consumer's utility while giving non-negative profits at $t = 1$, contradicting the optimality of the original contract.

that the consumer believes he will be able to achieve at that state:

$$\hat{V}^I(s_\tau) \equiv \max_{\{c(s_t,h_\tau^t)\}_{t\geq\tau}} u\left(c_\tau(s_\tau,h_\tau^\tau)\right) + \hat{\beta}E\left[\sum_{t>\tau} \delta^{t-\tau}u\left(c(s_t,(h_\tau^\tau,B,B,...,B))\right)\middle| s_\tau\right],$$

subject to (PC), (IC), zero profits (23), and the non-lapsing constraints (24)-(25).

The *equilibrium program with one-sided commitment* (P') adds the non-lapsing constraints to the program with two-sided commitment (P):

$$\max_{c(s_t,h^t)} u(c(s_1)) + \beta E\left[\sum_{t=1}^{T} \delta^{t-1}u(c(s_t,(B,B,...,B)))\right],$$

subject to (Zero Profits), (PC), (IC), the actual non-lapsing constraints,

$$u\Big(c(s_\tau,(h^{\tau-1},A))\Big) + \beta E\left[\sum_{t>\tau} \delta^{t-\tau}u\Big(c(s_t,(h^{\tau-1},A,B,...,B))\Big)\middle| s_\tau\right] \geq V^I(s_\tau), \ \ \forall s_\tau, \tag{NL}$$

and the perceived non-lapsing constraints,

$$u\Big(c(s_\tau,(h^{\tau-1},B))\Big) + \hat{\beta} E\left[\sum_{t>\tau} \delta^{t-\tau}u\Big(c(s_t,(h^{\tau-1},B,B,...,B))\Big)\middle| s_\tau\right] \geq \hat{V}^I(s_\tau), \ \ \forall s_\tau. \tag{PNL}$$

It may seem counter-intuitive to impose perceived non-lapsing constraints, since they are associated with histories that are off the equilibrium path. But equilibrium requires the consumer to pick his optimal actions given how he thinks his future selves will behave. If (PNL) did not hold, the consumer would not expect his future selves to choose the baseline option, and so the solution of program (P') would not be an equilibrium consumption.

Lemma 3. ***c*** *is the consumption vector of a perception-perfect equilibrium of the one-sided commitment model if and only if it solves program (P').*

3.1.3 Auxiliary Program and Vanishing Inefficiency

We now show that, as in the model with two-sided commitment, the consumption path coincides with the equilibrium with a dynamically consistent agent who discounts the last period by an additional factor β. Since this agent is dynamically consistent, we can replace the non-lapsing constraints by front-loading constraints (as in Subsection 2.1). Therefore, the *auxiliary program with one-sided commitment* maximizes (7) subject to the zero-profits (3) and front-loading (22) constraints. This program has a unique solution since the objective function is strictly concave and the constraints are linear.

Lemma 4. *$\boldsymbol{c}^E$ is the consumption path in a perception-perfect equilibrium of the one-sided commitment model if and only if it solves the auxiliary program with one-sided commitment.*

The proof of the lemma is similar to the proof of Lemma 2 with one additional step, in which we verify that the perceived non-lapsing constraints do not bind and can therefore be ignored. This step is needed for us to remove baseline consumption from the equilibrium program. The fact that PNL does not bind follows from two observations. First, PNL depends on the individual's perceived time-consistency parameter $\hat{\beta}$, while non-lapsing constraints depend on the true β. So, holding a consumption stream constant, NL implies PNL. Second, because the individual saves more in the baseline than in the alternative option, whenever future consumption in the alternative option is large enough to ensure that the agent does not lapse, future consumption in the baseline option must also be large enough to prevent lapsing.[17]

As in the case of two-sided commitment, Lemma 4 implies that the welfare loss from dynamic inconsistency vanishes as the contracting length grows. Let ${W'}_T^C$ and ${W'}_T^I$ denote the equilibrium welfare of time-consistent (Subsection 3.1.1) and time-inconsistent (3.1.2) consumers, respectively.

Theorem 2. *Suppose u is bounded and $\delta < 1$. Then, $lim_{T \nearrow +\infty} \left({W'}_T^C - {W'}_T^I \right) = 0.$*

To understand the welfare effect of removing commitment power, we now compare the equilibrium with one- and two-sided commitment. Recall that a time-inconsistent agent gets the same consumption path as a dynamically consistent agent who under-weights the

[17]In the Supplementary Appendix, we also generalize Corollary 1 to the one-sided commitment model.

last period by an additional factor β. Commitment power allows him to smooth consumption in the first $T-1$ periods while leaving too little consumption for the last period. This underconsumption in the last period is large when the consumer is sufficiently time inconsistent (β is low), in which case the last period consumption is close to zero. So, if underconsuming in the last period hurts the agent enough and β is low, the agent is better off without commitment.[18]

As the contracting length grows, however, the relative weight of the last period shrinks making it increasingly difficult for the time-inconsistent consumer to obtain higher welfare without commitment. In fact, it follows from Theorems 1 and 2 that removing commitment power cannot increase welfare if the contracting length is large enough. To see this, recall that a time-consistent consumer with commitment power maximizes welfare subject to zero profits. Removing commitment power is equivalent to introducing front-loading constraints, so his welfare with one-sided commitment cannot be higher. But, since the welfare of time-inconsistent consumers converges to the welfare of the time-consistent consumer, the same must be true for time-inconsistent consumers when T is large.

To summarize, while removing commitment power can help present-biased consumers when the contracting horizon is small, it cannot help when the horizon is large.

3.2 Market Power

Our previous analysis assumes that the consumer had all the bargaining power. We now consider the case in which the bargaining power is on the firm's side. Since the firm can commit to a long-term contract, there is no loss of generality in assuming that the firm makes a take-it-or-leave it offer to the consumer at time 1. If the consumer rejects the firm's offer, he gets $\underline{\mathbf{c}} \equiv \{\underline{c}(s_t) : s_t \in \mathbb{S}_t, t = 1, ..., T\}$. This is the consumer's outside option, that is, his best way to smooth consumption if he turns down the firm's offer.[19] The

[18]See the supplementary appendix for a formal statement and proof.

[19]Specifying an outside option in terms of consumption rather than a reservation utility makes the comparison between the allocations of individuals with different utility functions (in this case, time-inconsistent and time-consistent individuals) more transparent.

consumer accepts the firm's offer if and only if his perceived utility from the contract is weakly higher than the perceived utility from the outside option.

We assume that, for all T, the expected NPV of the agent's income is bounded and exceeds the NPV of his outside option:

$$E \sum_{t=1}^{T} \frac{\underline{c}(s_t)}{R^{t-1}} \leq E \sum_{t=1}^{T} \frac{w(s_t)}{R^{t-1}} < K, \tag{26}$$

for some $K > 0$. The inequality on the left in (26) ensures that the firm can find a profitable contract that the agent accepts, whereas the one on the right ensures that the expected NPV of income converges as the horizon T goes to infinity. We also assume that the outside option provides positive consumption in at least one state after the first period:[20]

$$\underline{c}(s_t) > 0 \text{ for some } s_t \text{ with } t > 1 \text{ and } p(s_t|\emptyset) > 0. \tag{27}$$

Conditions (26) and (27) ensure that there exists a non-trivial solution to the firm's program.

Let $\hat{W}_T^C(\underline{\mathbf{c}})$ and $\hat{W}_T^I(\underline{\mathbf{c}})$ denote the equilibrium welfare of time-consistent and time-inconsistent consumers, respectively, and let $\Pi_T^C(\underline{\mathbf{c}})$ and $\Pi_T^I(\underline{\mathbf{c}})$ denote the firm's equilibrium profit when the consumer is time consistent and time inconsistent. Picking different outside options for a time-consistent consumer (who maximizes welfare), we obtain different points on the Pareto frontier. Formally, let $\mathcal{P}_T \equiv \{(\hat{W}_T^C(\underline{\mathbf{c}}'), \Pi_T^C(\underline{\mathbf{c}}'))\}_{\forall \underline{\mathbf{c}}'}$ denote the set of consumer welfare and firm profits on the Pareto frontier. Let $\mathcal{P}_\infty \equiv \lim_{T \nearrow +\infty} \mathcal{P}_T$ denote the limit of the Pareto frontier as the contracting horizon grows.

The theorem below presents the vanishing inefficiency result when the bargaining power is on the firm's side:

Theorem 3. *Suppose u is bounded and $\delta < 1$. Then:*

1. $\lim_{T \nearrow +\infty} (\hat{W}_T^I(\underline{\boldsymbol{c}}), \Pi_T^I(\underline{\boldsymbol{c}})) \in \mathcal{P}_\infty$,

[20] If condition (27) fails, the outside option provides all consumption in the first period, which gives the same utility to naive and time-consistent consumers. In that case, in equilibrium, the agent will get the same perceived consumption as his outside option in all but the last period. Then, as the horizon grows, the equilibrium welfare of naive and time-consistent consumers will converge to the same point. That is, when the agent's outside option only pays in the first period, firms are be unable to "exploit" consumers (as will be shown in Theorem 3).

2. $\lim_{T \nearrow +\infty} \hat{W}_T^I(\underline{\boldsymbol{c}}) < \lim_{T \nearrow +\infty} E\left[\sum_{t \geq 1} \delta^{t-1} u(\underline{c}(s_t))\right]$, *and*

3. $\lim_{T \nearrow +\infty} \Pi_T^I(\underline{\boldsymbol{c}}) > \lim_{T \nearrow +\infty} \Pi_T^C(\underline{\boldsymbol{c}})$.

Part 1 states that the equilibrium converges to an efficient allocation. However, unlike when the bargaining power is on the consumer's side, the limit is different from the equilibrium of a time-consistent consumer with the same outside option. The consumer's welfare converges to a point *below his opportunity cost*, which is the equilibrium welfare of a time-consistent consumer (Part 2), whereas the firm's profits converge to a point above what it would get if the consumer were time consistent (Part 3). Since the consumer has incorrect beliefs about his future choices, he overestimates the value of the contract. The firm then "exploits" the consumer's naiveté, giving him a welfare below his outside option and obtaining greater profits than it would get if the consumer maximized his long-run preferences.

3.3 Consumer Heterogeneity

In Section 2, we assumed that firms know the consumer's preferences. We now briefly describe what happens if firms do not know either the consumer's naiveté or time-consistency parameters. For a formal analysis, see Online Appendix A.

3.3.1 Unknown Naiveté

When firms do not know the consumer's naiveté, all consumers get their full-information contracts in any equilibrium. Suppose first that all consumers are naive and recall that the naive consumer's full-information contract does not depend on his naiveté (Corollary 1). Then, all consumers receive the same contract and their welfare loss vanishes as the contracting horizon grows (as in Theorem 1).

Next, suppose there are also sophisticated consumers. Although naive and sophisticated consumers have different full-information contracts, neither of them prefers someone else's contract. When faced with the naive consumer's full-information contract, a sophisticated consumer understands that his future selves will pick the alternative options. Therefore, he prefers his full-information contract, which does not give flexibility to his future selves. Because naive consumers think they will stick with the baseline, they have a higher perceived utility from their own full-information contract. Then, the welfare loss from time inconsistency vanishes for all naive consumers as the contracting horizon grows (but not for sophisticated ones).

3.3.2 Unknown Time Inconsistency

Suppose now that the firm does not know the consumer's time-consistency parameter $\beta \in (0, \hat{\beta}]$. Note that the model allows some consumers to be sophisticated (when $\hat{\beta} < 1$) or time consistent ($\hat{\beta} = 1$). When the contracting horizon is large enough, there is no equilibrium in which multiple consumers get their full-information contracts. If multiple consumers received their full-information contracts, the more time-consistent of them would pick another consumer's full-information contract and choose the baseline rather than the alternative option. The firm offering this contract would lose money, since baseline options are unprofitable decoys not meant to be chosen on the equilibrium path.

In the only equilibrium that survives the D1 criterion, consumers get the "least costly separating allocation." In this allocation, only the most time-consistent consumer gets his full-information contract when the horizon is large. All other contracts are distorted to prevent deviations by more time-consistent consumers, making them save less than with full information. Moreover, this informational inefficiency does not vanish as the contracting horizon grows. In particular, when there is a time-consistent consumer ($\hat{\beta} = \beta = 1$), he is the only one who receives an efficient contract. When there is a sophisticated consumer ($\hat{\beta} = \beta < 1$), no consumer gets an asymptotically efficient contract (since the welfare loss of sophisticates does not vanish).

To summarize, with heterogeneous *naiveté* parameters, consumers receive their full-information contract in any equilibrium, and the welfare loss from time inconsistency vanishes as the horizon grows. On the other hand, with heterogeneous *time-consistency* parameters, only the most time-consistent consumer receives his full-information contract. Because firms worry that consumers will pick the contracts designed for more present-biased consumers and stick with the baseline, less time-consistent types get an allocation with insufficient savings, and this informational inefficiency persists as the horizon grows.

3.4 Non-Exclusive Contracts

In Section 2, we assumed that the consumer cannot simultaneously contract with multiple firms (i.e., contracts are exclusive). While, in Subsection 3.1 we allowed consumers to drop contracts and recontract with other firms, we kept the exclusivity assumption. This is a reasonable assumption in markets such as auto, property, or health insurance, where the consumer cannot be reimbursed by multiple insurers for the same loss. In some other markets, such as credit cards or life insurance, consumers are able to simultaneously contract with multiple firms – i.e., contracts are not exclusive. We now describe the equilibrium with non-exclusive contracts (see Online Appendix B for all formal results).

With non-exclusivity, firms cannot add unprofitable options that naive consumers think they will choose but end up not choosing. If offered the contract from Section 2, for example, the consumer would borrow from another firm to finance the fees specified in the baseline. Therefore, equilibrium contracts must make zero profits both along the equilibrium path and the consumer's perceived path. Moreover, consumers are able to undo any previous commitment.

Formally, we establish an equivalence between this model and a consumption-savings problem with no illiquid assets. The consumer is endowed with the expected PDV of income at the start. The only asset available is a risk-free bond that pays a gross return R. The consumer chooses how much to save in each period. The perception-perfect equilibrium is obtained by backward induction taking into account how much each self thinks that his future selves will save. This consumption-savings problem corresponds to a benchmark case in which there are no commitment devices: savings in each period are decided by the individual's current self given his incorrect beliefs about his future choices.[21]

The equilibrium with non-exclusive contracts is inefficient both because firms do not provide commitment devices and because the individual has mistaken beliefs about his future choices. It is straightforward to construct uniform bounds on consumer welfare in this model, so the vanishing inefficiency result fails to hold.

3.5 General Discounting Functions

So far, we have assumed that consumers have quasi-hyperbolic discounting, which is the most standard model of present bias. With quasi-hyperbolic discounting, there is a stark distinction between now and the future. While, in each period, the individual has a higher preference for consumption in that period, he agrees with his previous selves on how to compare consumption in any two periods in the future. Psychologists have proposed a different model of present bias, based on a hyperbolic functional form. With hyperbolic discounting, the individual becomes gradually more impatient as a period approaches.

In this subsection, we consider the model from Section 2 with general present-biased preferences. To simplify exposition, we assume that the consumer has a constant income w in each period. At time τ, the agent evaluates a consumption stream $\{c_t\}_{t\geq\tau}$ according to the separable representation:

$$\sum_{t=\tau}^{T} D_{t-\tau}u(c_t), \tag{28}$$

where the discount factor $D_t > 0$ is strictly decreasing in t and $D_0 = 1$.

It is well known that preferences represented by (28) are time consistent if and only if $D_t = D_1^t$ (i.e., discounting is exponential). We assume, instead, that preferences are *present biased*:

$$\frac{D_{t+2}}{D_{t+1}} \geq \frac{D_{t+1}}{D_t}, \tag{29}$$

for all $t = 0, 1, ...$, with strict inequality for at least one t. Present bias means that the in-

[21]As pointed out by Laibson (1997) and Bisin et al. (2015), investing in illiquid assets is a way for individuals to commit to a path of consumption.

dividual becomes (weakly) more impatient as a period approaches, with at least one period in which she becomes strictly more impatient.[22] With quasi-hyperbolic discounting, (29) holds as an equality in all but the first period, whereas the inequality is strict in all periods with hyperbolic discounting.

Let $(\hat{D}_1, \hat{D}_2, \cdots)$ denote the agent's beliefs about the discount factor of his future selves, where as before we normalize $\hat{D}_0 = 1$. A *sophisticated* agent knows that his future selves will have the same discount factor as he has: $\hat{D}_t = D_t$ for all t. A *naive* consumer underestimates the present bias of his future selves:

$$\frac{\hat{D}_{t+1}}{\hat{D}_t} \geq \frac{D_{t+1}}{D_t}$$

for all $t = 0, 1, ...$ with strict inequality for at least one t. Note that, with quasi-hyperbolic discounting, this inequality becomes the usual condition: $\hat{\beta} > \beta$.

The next proposition generalizes Lemma 2, characterizing the consumption path as the solution to a simpler auxiliary program:

Proposition 1. *The consumption path of a naive agent coincides with the equilibrium of an agent with utility function:*

$$u(c_1) + \frac{D_{T-1}}{D_{T-2}} u(c_2) + \frac{D_{T-1}}{D_{T-3}} u(c_3) + \cdots + \frac{D_{T-1}}{D_1} u(c_{T-1}) + D_{T-1} u(c_T). \tag{30}$$

To understand the coefficients in equation (30), consider the case of four periods and no uncertainty as in Figure 1. For simplicity (and without loss of generality), let $u(0) = 0$. The IC constraints are:

$$u(c_2(A)) + D_1 u(c_3(A,B)) + D_2 u(c_4(A,B)) \geq u(c_2(B)) + D_1 u(c_3(B,B)) + D_2 u(c_4(B,B)),$$
$$u(c_3(A,A)) + D_1 u(c_4(A,A)) \geq u(c_3(A,B)) + D_1 u(c_4(A,B)).$$

As in the sketch of the proof of Lemma 2, the baseline option shifts all intermediate con-

[22] See Prelec (1989) and DeJarnette et al. (2020). Note that (29) holds as an equality for all t if and only if the discount function is exponential, so preferences are time consistent.

sumption to the last period:

$$c_2(B) = c_3(B, B) = c_3(A, B) = 0. \tag{31}$$

Intuitively, because of present bias, self 1 benefits by shifting consumption to the future at a rate that keeps his future selves indifferent (thereby preserving incentive compatibility).

Both ICs must bind (otherwise, the agent would benefit from increasing his baseline consumption). Substituting (31) in the binding ICs and reorganizing terms to eliminate $c_4(A, B)$, we obtain:

$$\frac{u(c_2(A))}{D_2} + \frac{u(c_3(A, A))}{D_1} + u(c_4(A, A)) = u(c_4(B, B)). \tag{32}$$

Equation (32) specifies the period-2 utility needed to convince the agent to switch from the baseline. Since self 2 discounts consumption in the last period by D_2, for each increase in the last-period perceived utility $u(c_4(B, B))$, the firm must either raise $u(c_2(A))$ by $\frac{1}{D_2}$, $u(c_3(A, A))$ by $\frac{1}{D_1}$, or $u(c_4(A, A))$ by 1. But because the time-1 self discounts last-period utility by D_3, we must multiply these terms by D_3 to obtain the consumer's perceived utility at the time of contracting as a function of the alternative options chosen by his future selves. More generally, the coefficients in the objective function (30) are the ratio between how the period-1 self and the period-t self value consumption in the last period. These discount rates appear because the contract needs to convince each future self to switch to the alternative, rather than sticking to the baseline and waiting until the last period to consume.

Comparing the coefficients in (30) with the agent's discount factor D_t, we find that the equilibrium program assigns a lower weight on first-period consumption than the agent's utility function. The need to convince future selves to switch requires offering alternative options that take the consumer's future preferences into account. Since a sophisticated consumer would maximize self-1's utility subject to zero profits, it follows that a naive consumer always saves more than a sophisticate in the first period.[23]

[23] See the supplementary appendix for a formal proof.

Proposition 1 sheds light on how the shape of the discount function determines the equilibrium. With quasi-hyperbolic discounting, (7) and (30) coincide. The example below considers a generalization of quasi-hyperbolic discounting, which allows the immediate future to be treated differently than a more distant future:

Example 1. Consider a consumer with discount function $D_1 = \beta\delta$ and $D_t = \beta\gamma\delta^t$ for $t \geq 2$, where $0 < \beta \leq \gamma \leq 1$ and $\delta \in (0,1)$. The consumer believes that in future periods, he will behave as someone with parameters $\hat{\beta} \geq \beta$ and $\hat{\gamma} \geq \gamma$, with at least one of these two inequalities strict (so the consumer is naive). Note that this model reduces to quasi-hyperbolic discounting when $\gamma = \hat{\gamma} = 1$.

From Proposition 1, the equilibrium consumption path maximizes:

$$\sum_{t=1}^{T-2} \delta^{t-1}u(c_t) + \gamma\delta^{T-2}u(c_{T-1}) + \beta\gamma\delta^{T-1}u(c_T),$$

subject to zero profits. Here, present bias leads to the underweighting of consumption in the last two periods, with the last period being more distorted than the previous one. By the same argument as in Theorem 1, the inefficiency from present bias vanishes as the contracting horizon grows.

As the previous example illustrates, the vanishing inefficiency result generalizes to smoother versions of quasi-hyperbolic discounting. We now turn to the case of *hyperbolic discounting*:

$$D_t = \frac{1}{1+kt}, \tag{33}$$

where $k \geq 0$ is the "time-inconsistency parameter." The agent is time consistent if $k = 0$ and time inconsistent if $k > 0$. A naive agent has time-inconsistency parameter $k > 0$ but believes that, in the future, he will behave as an agent with time-inconsistency parameter $\hat{k} \in [0, k)$.

To understand the key difference between hyperbolic and quasi-hyperbolic discounting, consider how an individual discounts between periods t and $t+1$. With quasi-hyperbolic discounting, one util in $t+1$ periods is worth the same as $\frac{D_{t+1}}{D_t} = \delta$ utils in t periods for any $t > 0$. With hyperbolic discounting, one util in $t+1$ periods is worth $\frac{D_{t+1}}{D_t} = \frac{1+kt}{1+k(t+1)}$ in t periods, which is increasing and converges to 1 as t goes to infinity. That is, with hyperbolic discounting, the cost of waiting an additional period decreases as periods get further into the future. Moreover, waiting an additional period is virtually costless as long as that period happens far enough in the future.

Richer discount factors, such as hyperbolic, introduce two complications that are absent in exponential or quasi-hyperbolic discounting. First, the welfare criterion is not obvious when all selves disagree about how to discount the future. Second, the sum of discounted utility typically does not converge as the horizon grows. For example, with hyperbolic discounting, the discounted sum of any constant per-period utility $u(\bar{c}) > 0$ diverges:

$$\lim_{\tau \nearrow +\infty} \sum_{t=0}^{\tau} \frac{1}{1+k \cdot t} u(\bar{c}) = +\infty.$$

To deal with these issues, we adopt the limit-of-means criterion. Let $W_T^H(c) \equiv u(c_1) + u(c_2) + \cdots + u(c_T)$ denote the welfare from consumption stream $c = (c_1, ..., c_T)$. This welfare function corresponds to the utility of an individual evaluating outcomes from a sufficiently distant past.[24]

From Proposition 1, the naive agent's consumption path maximizes

$$u(c_1) + \frac{1+(T-2)k}{1+(T-1)k} u(c_2) + \frac{1+(T-3)k}{1+(T-1)k} u(c_3) + \cdots + \frac{1}{1+(T-1)k} u(c_T), \qquad (34)$$

subject to zero profits. Since, for any fixed k, $\lim_{T \nearrow \infty} \frac{D_{T-1}}{D_{T-t}} = \lim_{T \nearrow \infty} \frac{1+(T-t)k}{1+(T-1)k} = 1$, the objective function (30) converges to the undiscounted sum.

Let $c^H \equiv (c_1^H, \cdots, c_T^H)$ denote the consumption path of a naive agent, and let $c^* \equiv (c_1^*, \cdots, c_T^*)$ denote the equilibrium consumption for time-consistent agents ($k = 0$). Note

[24] With quasi-hyperbolic discounting, the only disagreement between selves concerns immediate consumption, which a person disapproves of at every other moment in the person's life. Because of this, most of the literature takes the agent's long-run preferences as the relevant welfare criterion (see footnote 7). We do not claim that, with hyperbolic discounting, undiscounted payoffs are the only normatively justifiable welfare measure. A common alternative approach is to treat each self as a different person and adopt a Pareto optimality criterion. Unfortunately, the incompleteness of the Pareto relation often makes this criterion excessively weak, especially when disagreement between selves is more pronounced.

that time-consistent agents maximize the welfare function W_T^H. We write $W_T^* = W_T^H(c^*)$ and $W_T^{H,I} = W_T^H(c^H)$ for the equilibrium welfare of time-consistent and time-inconsistent agents, respectively.

We make the following technical assumption:

Assumption 1. The utility function satisfies $\limsup_{\xi \nearrow +\infty} \frac{(u'(\xi)\log(u'(\xi)))^2}{|u''(\xi)|} < +\infty$.

Assumption 1 says that $u''(\cdot)$ does not go to zero too quickly as consumption grows. It is satisfied, for example, by any CARA preference.[25] The proposition below presents the convergence result for consumers with hyperbolic discounting:

Proposition 2. *Suppose Assumption 1 holds, u is bounded, and $\lim_{c \searrow 0} u'(c) = +\infty$. Then,*

$$\lim_{T \nearrow +\infty} \frac{W_T^{H,I} - W_T^*}{T} = 0.$$

3.6 Effort

We have previously focused on contracts over consumption. In many applications, parties contract over effort rather than money. This subsection shows that the predictions of the model are remarkably different when contracting over effort.

Consider an agent who needs to complete a task in at most $T \geq 3$ periods. In each period $t \in \{1, ..., T\}$, the agent exerts effort e_t at a cost of $C(e_t)$. The cost of effort $C : \mathbb{R}_+ \to \mathbb{R}_+$ is increasing, continuous, convex, and satisfies $C(0) = 0$. Completing the task requires a total effort of at least $E_T > 0$, so the agent faces the constraint:

$$\sum_{t=1}^{T} e_t \geq E_T. \tag{35}$$

For concreteness, we can think of effort as hours worked and E_T as the time it takes to complete the task. We hold the task fixed for simplicity here, although it is straightforward to add an ex-ante stage in which the task is chosen, making E_T endogenous.

[25] Assumption 1 holds for any CRRA utility with risk aversion greater than 1 and, more generally, for any HARA utility, $u(c) = \frac{\gamma}{1-\gamma}\left(\frac{a}{\gamma}c + b\right)^{1-\gamma}$, with $\gamma > 1$.

To highlight the role of the discount function, we consider the general setting from Subsection 3.5. As in that subsection, we assume that the firm knows the agent's preferences, that the agent has all bargaining power, and that parties can commit to long-term contracts.

Firms accept an effort plan as long as they believe that the task will be completed. A sophisticated agent chooses an effort plan $(e_1, ..., e_T)$ that minimizes his discounted cost:

$$\sum_{t=1}^{T} D_{t-1}C(e_t)$$

subject to (35). The next proposition characterizes the effort path of a naive agent:

Proposition 3. *The effort path of a naive agent coincides with the equilibrium of a time-consistent agent with discount factor* $\tilde{D}_t = D_1^t$.

Note that the time-consistent agent in the auxiliary program in Proposition 3 has exponential discounting. Since he uses the one-period rate of a present-biased discount function, (29) implies that he acts more impatiently than a sophisticated agent, who uses his true discount factor $\{D_t\}$. In particular, with quasi-hyperbolic discounting, the auxiliary program has discount factor $\tilde{D}_t = (\beta\delta)^t < \beta\delta^t$. Therefore, not only does the inefficiency from present bias persist, but the naive agent acts even less patiently than a sophisticated agent. Instead of maximizing his long-run preferences, the agent perpetually chooses according to his short-run preferences.

Recall from Proposition 1 that in the consumption model, firms exploit naive agents by offering a baseline option that postpones consumption until the last period. At the time of contracting, self 1 uses the long-run discount factor D_{T-1} to decide how much to consume now and how much to leave to the last period. Then, each of his future selves deviates from the baseline, effectively bringing some of this future consumption to the present. Since the t-period's self discounts last-period utility by D_{T-t}, the auxiliary program assigns weight $\frac{D_{T-1}}{D_{T-t}}$ to actual consumption in period $t < T$. In particular, this weight equals δ^{t-1} when the agent has quasi-hyperbolic discounting.

With effort, the way to exploit naiveté is to offer a baseline option that requires zero effort after period 2. When designing the contract, self 1 uses the one-period discount factor D_1 to decide how much effort to exert immediately and how much to leave to period 2. Then, each of his future selves deviates from the baseline, effectively postponing some effort into the following period, while thinking that he will not exert any effort afterwards. Since, each self decides how much effort to do immediately and how much to leave to the following period, along the equilibrium path, effort costs are discounted using the one-period discount rate: $\tilde{D}_t = D_1^t$.

More generally, with pleasant tasks (such as consumption), naive agents are exploited by offering a baseline that postpones these tasks far into the future, making them act more in line with their long-run discount factors. With unpleasant tasks (such as effort), they are exploited by offering a baseline that concentrates all future effort in the next period, making them act according to their short-run discount factors. This explains why the equilibrium converges to the one that maximizes their long-run preferences in the former but maximizes their short-run preferences in the latter.[26]

4 Conclusion

In this paper, we study contracting between firms and present-biased consumers. Our main result is that the welfare loss from present bias vanishes as the contracting horizon grows.

Recall that the equilibrium does not depend on the consumer's naiveté parameter as long as the consumer is partially naive ($\hat{\beta} > \beta$) but jumps discontinuously as the consumer becomes sophisticated ($\hat{\beta} = \beta$). But for sophisticated agents, the welfare loss from present bias does not vanish as the horizon grows. Therefore, when T is large, the consumer's welfare jumps downwards at $\hat{\beta} = \beta$, meaning that the consumer is harmed by learning his true present bias. This finding contrasts with a general intuition that educating behavioral individuals about their biases would increase their welfare.

[26] This observation is related to Kőszegi (2005) and Gottlieb (2008), who find that, with 3 periods and non-exclusive contracts, there is an asymmetry in the market's ability to provide commitment for goods with immediate benefits and delayed costs and those with immediate costs and delayed benefits.

Learning one's present bias in our model is hard. First, starting from any belief about the time-consistency parameter (captured by the naiveté parameter $\hat{\beta} > \beta$), local updates do not affect the equilibrium choices or payoffs.[27] Second, because consumers who know their true time-consistency parameter get a lower welfare, there are no incentives to learn.

Note also that under the conditions of Theorem 1, adding restrictions to the space of contracts cannot benefit consumers if the contracting horizon is long. Under those conditions, enforcing long-term contracts may be enough to ensure efficiency. We studied one particular restriction: removing commitment power from consumers. Other examples of regulations that try to protect consumers include limiting the fees that firms can charge and allowing consumers to pull out of a contract under certain conditions.

The vanishing inefficiency result generalizes to when the bargaining power is on the firm's side or when firms do not know the consumer's naiveté. However, it breaks down when firms do not know the consumer's present bias or when they cannot offer exclusive contracts. When firms do not know the consumer's present bias, firms face an adverse selection problem and the equilibrium consumption corresponds to the least costly separating allocation, in which all but the most time-consistent type undersave. When they cannot offer exclusive contracts, the equilibrium fails to provide any commitment devices. Finally, when contracting over unpleasant rather than pleasant tasks (such as effort instead of consumption), the equilibrium path caters to naive agents' short-run rather than long-run preferences, exacerbating instead of dissipating the inefficiency from present bias.

[27]The consumption path is discontinuous in the agent's naiveté, with "almost sophisticated" agents getting the same consumption as any other naive agent (which is bounded away from the consumption of a sophisticated agent). However, the consumption path is continuous in the agent's time-consistency, so "almost time-consistent" agents get approximately the same consumption as a time-consistent agent.

References

Amador, M., I. Werning, and G.-M. Angeletos (2006). Commitment vs. flexibility. *Econometrica 74*(2), 365–396.

Atal, J. P., H. Fang, M. Karlsson, and N. R. Ziebarth (2018). Long-term health insurance: Theory meets evidence. *Working Paper*.

Atlas, S. A., E. J. Johnson, and J. W. Payne (2017). Time preferences and mortgage choice. *Journal of Marketing Research 54*(3), 415–429.

Augenblick, N., M. Niederle, and C. Sprenger (2015). Working over time: Dynamic inconsistency in real effort tasks. *Quarterly Journal of Economics 130*(3), 1067–1115.

Banks, J. S. and J. Sobel (1987). Equilibrium selection in signaling games. *Econometrica 55*(3), 647–661.

Bisin, A., A. Lizzeri, and L. Yariv (2015). Government policy with time inconsistent voters. *American Economic Review 105*(6), 1711–37.

Bond, P. and G. Sigurdsson (2017). Commitment contracts. *Review of Economic Studies 85*(1), 194–222.

Cao, D. and I. Werning (2018). Saving and dissaving with hyperbolic discounting. *Econometrica 86*(3), 805–857.

Carter, S. P., K. Liu, P. M. Skiba, and J. Sydnor (2017). Time to repay or time to delay? the effect of having more time before a payday loan is due. *Working Paper*.

DeJarnette, P., D. Dillenberger, D. Gottlieb, and P. Ortoleva (2020). Time lotteries and stochastic impatience. *Econometrica 88*(2), 619–656.

DellaVigna, S. and U. Malmendier (2004). Contract design and self-control: Theory and evidence. *Quarterly Journal of Economics 119*(2), 353–402.

Galberti, S. (2015). Commitment, flexibility, and optimal screening of time inconsistency. *Econometrica 83*(4), 1425–1465.

Ghent, A. (2011). Subprime mortgages, mortgage choice, and hyperbolic discounting. *Working paper, Zicklin School of Business*.

Gottlieb, D. (2008). Competition over time-inconsistent consumers. *Journal of Public Economic Theory 10*(4), 673–684.

Grubb, M. D. (2015). Overconfident consumers in the marketplace. *Journal of Economic Perspectives 29*(4), 9–35.

Gruber, J. and B. Kőszegi (2001). Is addiction “rational”? theory and evidence. *Quarterly Journal of Economics 116*(4), 1261–1303.

Halac, M. and P. Yared (2014). Fiscal rules and discretion under persistent shocks. *Econometrica 82*(5), 1557–1614.

Handel, B., I. Hendel, and M. D. Whinston (2017). The welfare effects of long-term health insurance contracts. *Working paper*.

Harris, C. and D. Laibson (2001). Dynamic choices of hyperbolic consumers. *Econometrica 69*(4), 935–957.

Heidhues, P. and B. Kőszegi (2010). Exploiting naivete about self-control in the credit market. *American Economic Review 100*(5), 2279–2303.

Heidhues, P. and P. Strack (2019). Identifying procrastination from the timing of choices. *Working Paper*.

Hendel, I. and A. Lizzeri (2003). The role of commitment in dynamic contracts: Evidence from life insurance. *Quarterly Journal of Economics 118*(1), 299–327.

Kőszegi, B. (2005). On the feasibility of market solutions to self-control problems. *Swedish Economic Policy Review 12*, 71–94.

Kőszegi, B. (2014). Behavioral contract theory. *Journal of Economic Literature 52*(4), 1075–1118.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics 112*(2), 443–477.

Meier, S. and C. Sprenger (2010). Present-biased preferences and credit card borrowing. *American Economic Journal: Applied Economics 2*(1), 193–210.

O’Donoghue, T. and M. Rabin (1999). Doing it now or later. *American Economic Review 89*(1), 103–124.

O'Donoghue, T. and M. Rabin (2001). Choice and procrastination. *Quarterly Journal of Economics 116*(1), 121–160.

O'Donoghue, T. and M. Rabin (2003). Studying optimal paternalism, illustrated by a model of sin taxes. *American Economic Review 93*(2), 186–191.

O'Donoghue, T. and M. Rabin (2015). Present bias: Lessons learned and to be learned. *American Economic Review 105*(5), 273–79.

Prelec, D. (1989). Decreasing impatience: Definition and consequences. *Harvard Business School Working Paper no. 90-015*.

Schlafmann, K. (2016). Housing, mortgages, and self control. *Working paper, Stockholm University*.

Sulka, T. (2020). Exploitative contracting in a life cycle savings model. *Working Paper*.

Yildiz, M. (2003). Bargaining without a common prior: An immediate agreement theorem. *Econometrica 71*(3), 793–811.

Appendix: Selected Proofs

Proof of Theorems 1 and 2. We establish the result for one-sided commitment (Theorem 2). The two-sided commitment case (Theorem 1) is similar and, therefore, omitted.

For each β, let $V_T^A(\beta)$ denote the maximum value attained by the solution of the auxiliary program with one-sided commitment. Notice that the feasible set is independent of β. When $\beta = 1$, the auxiliary program becomes the time-consistent agent's program, so that $V_T^A(1) = {W'}_T^C$. Note that $\lim_{T\nearrow\infty}({W'}_T^I - V_T^A(\beta)) = \lim_{T\nearrow\infty}(1-\beta)E\delta^{T-1}u(c(s_T)) = 0$. Since the objective function is linear in β, it follows from the Envelope Theorem that $\frac{\partial V_T^A(\beta)}{\partial\beta} = E\delta^{T-1}u\left(c(s_T)\right) \geq \delta^{T-1}u\left(0\right)$. Applying Lagrange's Mean Value Theorem gives

$$V_T^A(1) - V_T^A(\beta) = \frac{\partial V_T^A}{\partial\beta}(\beta')\cdot(1-\beta) \geq \delta^{T-1}u(0)(1-\beta),$$

$$V_T^A(\beta) - V_T^A(0) = \frac{\partial V_T^A}{\partial\beta}(\beta'')\cdot\beta \geq \delta^{T-1}u(0)\beta,$$

where $\beta' \in (\beta, 1)$ and $\beta'' \in (0, \beta)$. Taking T to infinity leads to:

$$\liminf_{T \nearrow \infty} (V_T^A(1) - V_T^A(\beta)) \geq 0, \ \liminf_{T \nearrow \infty} (V_T^A(\beta) - V_T^A(0)) \geq 0.$$

To obtain the theorem, it suffices to show that:

$$\limsup_{T \nearrow +\infty} [V_T^A(1) - V_T^A(0)] \leq 0.$$

If this is true, we obtain $\lim_{T \nearrow \infty} (V_T^A(1) - V_T^A(\beta)) = 0$. It then follows that

$$\lim_{T \nearrow +\infty} ({W'}_T^C - {W'}_T^I) = \lim_{T \nearrow \infty} (V_T^A(1) - V_T^A(\beta)) + \lim_{T \nearrow \infty} (V_T^A(\beta) - {W'}_T^I) = 0.$$

Consider the auxiliary program with one-sided commitment when $\beta = 0$, which attains maximum value $V_T^A(0)$. Let $\mathbf{c^0} \equiv \{c^0(s_t) : s_t \in S_t(s_1), 1 \leq t \leq T\}$ denote a solution to this program. Since the objective function does not depend on $c(s_T)$ when $\beta = 0$, the solution has the lowest possible value for $c(s_T)$ that still satisfies the constraints: $\mathbf{c}^0(s_T) = w(s_T)$. Substituting this equality back, we obtain the same program that determines the consumption of a time-consistent agent with a contracting horizon consisting of the first $(T-1)$ periods.

Let c^C denote the equilibrium consumption of a time-consistent agent. Since c^C is in the feasible set, income cannot exceed consumption for any last-period state: $c^C(s_T) \geq w(s_T)$. Therefore, by revealed preference ($V_T^A(0)$ maximizes expected utility in the first $T-1$ periods and uses weakly higher resources), we must have

$$\begin{aligned} V_T^A(0) &= E \textstyle\sum_{t=1}^{T-1} \delta^{t-1} u\left(c^0(s_t)\right) \\ &\geq E \textstyle\sum_{t=1}^{T-1} \delta^{t-1} u\left(c^C(s_t)\right) \\ &= V_T^A(1) - \delta^{T-1} E u\left(c^C(s_T)\right), \end{aligned}$$

where the first line uses the definition of $V_T^A(0)$, the second line uses revealed preference, and the third line uses the definition of $V_T^A(1)$. Since $\delta < 1$ and u is bounded, we have

$$\lim_{T \nearrow +\infty} \delta^{T-1} E u\left(c^C(s_T)\right) = 0,$$

which establishes that $\limsup_{T\nearrow+\infty}[V_T^A(1)-V_T^A(0)]\leq 0$. $\square$

Proof of Theorem 3. Let $\mathcal{P}_\infty \equiv \lim_{T\nearrow+\infty}\mathcal{P}_T$. We will use the following result, which is proved in the supplementary appendix:

Claim 1. *Suppose u is bounded and $\delta < 1$. Then, $\mathcal{P}_\infty$ exists.*

For any outside option $\underline{\mathbf{c}}$, the equilibrium profit with time-consistent consumers solves:

$$\Pi_T^C(\underline{\mathbf{c}}) := \max_{\{c(s_t)\}} E\sum_{t=1}^{T}\frac{w(s_t)-c(s_t)}{R^{t-1}},$$

subject to $E\sum_{t=1}^{T}\delta^{t-1}u(c(s_t)) \geq E\sum_{t=1}^{T}\delta^{t-1}u(\underline{c}(s_t))$. With time-inconsistent consumers, the equilibrium profit is determined by:

$$\Pi_T^I(\underline{\mathbf{c}}) := \max_{\{c(\cdot)\}} E\sum_{t=1}^{T}\frac{w(s_t)-c(s_t,A,\cdots,A)}{R^{t-1}},$$

subject to (IC), (PC), and

$$u(c(s_1))+\beta E\sum_{t>1}\delta^{t-1}u(c(s_t,B,\cdots,B)) \geq \underline{U},$$

where $\underline{U} \equiv u(\underline{c}(s_1))+\beta E\sum_{t>1}\delta^{t-1}u(\underline{c}(s_t))$. Consider the following auxiliary program:

$$\Pi_{T,\beta} \equiv \max_{\{c(s_t)\}} E\sum_{t=1}^{T}\frac{w(s_t)-c(s_t)}{R^{t-1}}, \tag{A1}$$

subject to

$$E\left[\sum_{t=1}^{T-1}\delta^{t-1}u(c(s_t))+\beta\delta^{T-1}u(c(s_T))\right] \geq \underline{U}'. \tag{A2}$$

We claim that the equilibrium consumption for time-inconsistent consumers must solve the auxiliary program with $\underline{U}'=\underline{U}+(1-\beta)u(0)(\delta+\cdots+\delta^{T-2})$. To establish this result, it is helpful to work with the dual program for the time-inconsistent agents:

$$\max_{\{c_t(\cdot)\}} u(c(s_1))+\beta E\sum_{t=2}^{T}\delta^{t-1}u(c(s_t,B,\cdots,B)),$$

subject to (IC), (PC), and

$$E \sum_{t=1}^{T} \frac{w(s_t) - c(s_t, A, \cdots, A)}{R^{t-1}} \leq V_T^I(\underline{\mathbf{c}}).$$

Note that $V_T^I(\underline{\mathbf{c}})$ is the maximum profit to the firm when the consumer gets utility $\underline{U}$. We can now follow the same steps as in the proof of Lemma 2 to simplify (IC) and (PC), obtaining the following program:

$$\max_{\{c(s_t)\}} E \left[\sum_{t=1}^{T-1} \delta^{t-1} u(c(s_t)) + \beta \delta^{T-1} u(c(s_T)) \right] + (1-\beta) u(0)(\delta + \cdots + \delta^{T-2}),$$

subject to $E \sum_{t=1}^{T} \frac{w(s_t) - c(s_t)}{R^{t-1}} \leq \Pi_T^I(\underline{\mathbf{c}})$.

Note that this is the equilibrium program of a dynamically consistent agent who discounts the last period by an extra β. Therefore, the consumption path solves the auxiliary program with

$$\underline{U}' = \underline{U} + (1-\beta) u(0)(\delta + \cdots + \delta^{T-2}). \tag{A3}$$

We now obtain the convergence result. Since that the participation constraints must be binding both in the auxiliary program and in the program for time-consistent consumers, we must have:

$$E \left[\sum_{t=1}^{T-1} \delta^{t-1} u(c^A(s_t, \underline{U}')) + \beta \delta^{T-1} u(c^A(s_T, \underline{U}')) \right] = \underline{U}',$$

where $c^A := (c^A(s_1, \underline{U}'), \cdots, c^A(s_T, \underline{U}'))$ denotes the equilibrium consumption in the auxiliary program. Omitting the dependence of c^A on $\underline{U}'$ for notational simplicity, we have:

$$\begin{aligned}
\hat{W}_T^I(\underline{\mathbf{c}}) &= E \sum_{t=1}^{T} \delta^{t-1} u(c^A(s_t)) \\
&= E \left[\sum_{t=1}^{T-1} \delta^{t-1} u(c^A(s_t)) + \beta \delta^{T-1} u(c^A(s_T)) + (1-\beta) \delta^{T-1} u(c^A(s_T)) \right] \\
&= \underline{U}' + E \left[(1-\beta) \delta^{T-1} u(c^A(s_T)) \right], \qquad \text{(A4)}
\end{aligned}$$

where the first line uses the definition of $\hat{W}_T^I(\underline{\mathbf{c}})$, the second line comes from algebraic manipulations, and the third line comes from (A2). Choose an outside option, $\underline{\mathbf{c}}'$, for a time-consistent consumer such that the time-consistent consumer's utility is given by $\underline{U}'$:

$$\hat{W}_T^C(\underline{\mathbf{c}}') = E\sum_{t=1}^{T}\delta^{t-1}u(\underline{c}'(s_t)) = \underline{U}'. \tag{A5}$$

The existence of $\underline{\mathbf{c}}'$ is guaranteed because we can pick $\underline{c}_1' = \underline{c}_1$ and $u(\underline{c}'(s_t)) = \beta u(\underline{c}(s_t)) + (1-\beta)u(0), \forall t \geq 2$. Moreover, since $\underline{c}'(s_t) \leq \underline{c}(s_t)$, condition (26) holds for $\underline{\mathbf{c}}'$.

Combining (A4) and (A5), we obtain:

$$\hat{W}_T^I(\underline{\mathbf{c}}) = \hat{W}_T^C(\underline{\mathbf{c}}') + E\left[(1-\beta)\delta^{T-1}u(c^A(s_T))\right],$$

and, since u bounded and $\delta < 1$, it follows that

$$\lim_{T\nearrow\infty}|\hat{W}_T^C(\underline{\mathbf{c}}') - \hat{W}_T^I(\underline{\mathbf{c}})| = \lim_{T\nearrow\infty}|E\left[(1-\beta)\delta^{T-1}u(c^A(s_T))\right]| = 0.$$

We now turn to the firm's profit in the program (A1). Let λ denote the Lagrangian multiplier with the constraint (A2). The first-order condition gives

$$\lambda\delta^{t-1}u'(c^A(s_t)) = \frac{1}{R^{t-1}}, \forall t = 1, \cdots, T-1,$$

and

$$\lambda\beta\delta^{T-1}u'(c^A(s_T)) = \frac{1}{R^{T-1}}.$$

Note that

$$\Pi_{T,\beta} = E\sum_{t=1}^{T}\frac{w(s_t)-c(s_t)}{R^{t-1}} + \lambda E\left[\sum_{t=1}^{T-1}\delta^{t-1}u(c(s_t)) + \beta\delta^{T-1}u(c(s_T)) - \underline{U}'\right].$$

Differentiating with respect to β leads to

$$\frac{\partial\Pi_{T,\beta}}{\partial\beta} = E\left[\lambda\delta^{T-1}u(c^A(s_T))\right] \geq \lambda\delta^{T-1}u(0),$$

where the inequality uses $c^A(s_T) \geq 0$. Applying Lagrange's Mean Value Theorem gives:

$$\Pi_{T,1} - \Pi_{T,\beta} = \frac{\partial \Pi_{T,\beta}}{\partial \beta}|_{\beta=\beta'}(1-\beta) \geq \lambda \delta^{T-1} u(0)(1-\beta),$$
$$\Pi_{T,\beta} - \Pi_{T,0} = \frac{\partial \Pi_{T,\beta}}{\partial \beta}|_{\beta=\beta''}(\beta - 0) \geq \lambda \delta^{T-1} u(0)\beta,$$

where $\beta' \in (\beta, 1), \beta'' \in (0, \beta)$. Note that λ is bounded because $\lambda = \frac{1}{u'(c^A(s_1))} \leq \frac{1}{u'(K)}$, where the inequality uses $c^A(s_1) \leq K$ from Condition (26). As $T \nearrow \infty$, we obtain

$$\liminf_{T \nearrow \infty}(\Pi_{T,1} - \Pi_{T,\beta}) \geq 0, \ \liminf_{T \nearrow \infty}(\Pi_{T,\beta} - \Pi_{T,0}) \geq 0.$$

Note that $\Pi_{T,1} = \Pi_T^C(\underline{\mathbf{c}}')$ and $\Pi_{T,\beta} = \Pi_T^I(\underline{\mathbf{c}})$. To show that $\lim_{T \nearrow \infty}(\Pi_{T,1} - \Pi_{T,\beta}) = 0$, it is sufficient to show that $\limsup_{T \nearrow \infty}(\Pi_{T,1} - \Pi_{T,0}) = 0$. The program for $\Pi_{T,0}$ is:

$$\Pi_{T,0} = \max_{\{c_t\}} E \sum_{t=1}^{T} \frac{w(s_t) - c(s_t)}{R^{t-1}},$$

subject to $E \sum_{t=1}^{T-1} \delta^{t-1} u(c(s_t)) \geq \underline{U}'$.

Since $c(s_T)$ does not appear in the constraint, $c(s_T) = 0$. The program reduces to:

$$\Pi_{T,0} = \max_{\{c_t\}} E \left[\sum_{t=1}^{T-1} \frac{w(s_t) - c(s_t)}{R^{t-1}} \right] + E \left[\frac{w(s_T)}{R^{T-1}} \right],$$

subject to $E \sum_{t=1}^{T-1} \delta^{t-1} u(c(s_t)) \geq \underline{U}'$. Note that $\Pi_{T,0}$ maximizes the objective from the $(T-1)$-period program with $\beta = 1$ (plus the constant $E\left[\frac{w(s_T)}{R^{T-1}}\right]$). By a revealed preference argument,

$$\Pi_{T,0} \geq \Pi_{T-1,1} + E \left[\frac{w(s_T)}{R^{T-1}} \right]. \tag{A6}$$

Since the NPV of income is assumed to be bounded (Condition (26)), it follows that $\lim_{T \nearrow \infty} E\left[\frac{w(s_T)}{R^{T-1}}\right] = 0$. Then,

$$\limsup_{T \nearrow \infty}(\Pi_{T,1} - \Pi_{T,0}) \leq \limsup_{T \nearrow \infty} \left(\Pi_{T,1} - \Pi_{T-1,1} - E \left[\frac{w(s_T)}{R^{T-1}} \right] \right)$$
$$= \limsup_{T \nearrow \infty}(\Pi_{T,1} - \Pi_{T-1,1}) = 0,$$

where the first step comes from (A6), the second step comes from $\lim_{T\nearrow\infty} E\left[\frac{w(s_T)}{R^{T-1}}\right] = 0$, and the last step is because the limit of $\Pi_{T,1}$ exists (Claim 1). Since $\Pi_{T,1}$ is the problem for time consistent consumers with the outside option $\underline{\mathbf{c}}'$ and $\Pi_{T,\beta}$ is the problem for time inconsistent consumers with the outside option $\underline{\mathbf{c}}$, it follows that $\lim_{T\nearrow\infty} |\Pi_T^C(\underline{\mathbf{c}}') - \Pi_T^I(\underline{\mathbf{c}})| = 0$.

To recap, we have shown that there exists an outside option $\underline{\mathbf{c}}'$ such that

$$\lim_{T\nearrow\infty} |\hat{W}_T^C(\underline{\mathbf{c}}') - \hat{W}_T^I(\underline{\mathbf{c}})| = 0, \lim_{T\nearrow\infty} |\Pi_T^C(\underline{\mathbf{c}}') - \Pi_T^I(\underline{\mathbf{c}})| = 0.$$

This establishes that $(\hat{W}_T^I(\underline{\mathbf{c}}), \Pi_T^I(\underline{\mathbf{c}}))$ converges to a point on the Pareto frontier for time-consistent consumers (first part of the theorem): $\lim_{T\nearrow\infty}(\hat{W}_T^I(\underline{\mathbf{c}}), \Pi_T^I(\underline{\mathbf{c}})) \in \mathcal{P}_\infty$.

We now prove the second part of the theorem. Recall that a time-inconsistent consumer's welfare in the equilibrium $\lim_{T\nearrow\infty} \hat{W}_T^I(\underline{\mathbf{c}}) = \lim_{T\nearrow\infty} \underline{U}'$, where

$$\begin{aligned}
\lim_{T\nearrow\infty} \underline{U}' &= \lim_{T\nearrow\infty} \left[u(\underline{c}_1) + \beta E\left[\sum_{t\geq 2} \delta^{t-1} u(\underline{c}(s_t))\right] + (1-\beta)(\delta + \cdots + \delta^{T-2})u(0)\right] \\
&\leq \lim_{T\nearrow\infty} E\left[\sum_{t\geq 1} \delta^{t-1} u(\underline{c}(s_t))\right],
\end{aligned}$$

where the equation comes from (A3) and the inequality comes from

$$\beta u(\underline{c}(s_t)) + (1-\beta)u(0) \leq u(\underline{c}(s_t)), \forall t.$$

The inequality is strict for at least one $\underline{c}_t(s_t)$ because of condition (27). So a time-inconsistent consume is worse off contracting with the firm in that he receives lower welfare in the equilibrium than the welfare he would've received by consuming the outside option. The third part follows from the definition of the Pareto frontier and the second part of the theorem.

□

Proof of Proposition 1. A naive agent's equilibrium consumption vector solves

$$\max_{c(\cdot)} \sum_{t=1}^{T} D_{t-1} u(c_t(B, \cdots, B)),$$

subject to the zero-profit condition, and

$$\sum_{t=\tau}^{T} D_{t-\tau} u(c_t(A, B, \cdots, B)) \geq \sum_{t=\tau}^{T} D_{t-\tau} u(c_t(B, B, \cdots, B)), \forall \tau \geq 2, \tag{A7}$$

$$\sum_{t=\tau}^{T} \hat{D}_{t-\tau} u(c_t(B, B, \cdots, B)) \geq \sum_{t=\tau}^{T} \hat{D}_{t-\tau} u(c_t(A, B, \cdots, B)), \forall \tau \geq 2. \tag{A8}$$

As in the proof of Lemma 2, there exists a solution in which the baseline option shifts all intermediate consumption to the last period:

$$c_2(B) = c_3(B, B) = \cdots = c_{T-1}(B, \cdots, B) = 0. \tag{A9}$$

Otherwise, we could lower $u(c_t(B, \cdots, B))$ by $\frac{\epsilon}{D_{t-2}}$ and increase $u(c_T(B, \cdots, B))$ by $\frac{\epsilon}{D_{T-2}}$, giving the naive agent a weakly higher utility since $-\frac{D_{t-1}}{D_{t-2}} + \frac{D_{T-1}}{D_{T-2}} \geq 0$. If $\frac{D_{t-1}}{D_{t-2}} = \frac{D_{T-1}}{D_{T-2}}$, the naive agent is indifferent between the original contract and this new one, and both contracts have the same consumption path. Since we focus on the consumption path, we can without loss of generality work with (A9). Using the same argument, self $t \geq 2$'s perceived consumption stream also features shifting all intermediate consumption to the last period: $c_{t+1}(\underbrace{A, \cdots, A}_{t-1}, B) = \cdots = c_{T-1}(\underbrace{A, \cdots, A}_{t-1}, \underbrace{B, \cdots, B}_{T-t-1}) = 0$.

Substituting in the binding IC constraints (A7), we obtain the objective function in our auxiliary program (up to a constant): $\sum_{t=1}^{T} \frac{D_{T-1}}{D_{T-t}} u(c_t(A, \cdots, A))$. □

Proof of Proposition 3. The effort path of the naive agent solves:

$$\min_{\mathbf{e}} \sum_{t=1}^{T} D_{t-1} C(e_t(B, B, ..., B))$$

subject to

$$\sum_{t=1}^{T} e_t(A, ..., A) = E_T, \tag{A10}$$

$$\sum_{t=\tau}^{T} D_{t-\tau} C(e_t(A, B, \cdots, B)) \leq \sum_{t=\tau}^{T} D_{t-\tau} C(e_t(B, B, \cdots, B)), \forall \tau \geq 2, \tag{IC}$$

$$\sum_{t=\tau}^{T} \hat{D}_{t-\tau} C(e_t(B, B, \cdots, B)) \leq \sum_{t=\tau}^{T} \hat{D}_{t-\tau} C(e_t(A, B, \cdots, B)), \forall \tau \geq 2. \tag{PC}$$

The (IC) at time $\tau = 2$ must be binding (otherwise, we could decrease $e_T(B, \cdots, B)$, improving the objective). We claim that $e_t(B, \cdots, B) = 0$ for $t \geq 3$. Suppose otherwise that $e_t(B, \cdots, B) > 0$ for some t. Consider a perturbation that shifts effort in the baseline from period t to period 2: $C(e_2(B)) + \epsilon$ and $C(e_t(B, \cdots, B)) - \frac{\epsilon}{D_{t-2}}$. This perturbation keeps (IC) unchanged, relaxes (PC), and lowers the objective function by $\left(D_1 - \frac{D_{t-1}}{D_{t-2}}\right) \epsilon \leq 0$. By the same argument, self $t \geq 2$'s perceived effort path also shifts all future effort to period $t+1$:

$$e_{t+2}(\underbrace{A, \cdots, A}_{t-1}, B, B) = \cdots = e_{T-1}(\underbrace{A, \cdots, A}_{t-1}, \underbrace{B, \cdots, B}_{T-t-1}) = e_T(\underbrace{A, \cdots, A}_{t-1}, \underbrace{B, \cdots, B}_{T-t-1}) = 0.$$

Substituting in the objective function, we obtain:

$$\begin{aligned}
&\sum_{t=1}^{T} D_{t-1} C(e_t(B, B, ..., B)) \\
&= C(e_1) + D_1 C(e_2(B)) \\
&= C(e_1) + D_1 \left(\sum_{t=2}^{T} D_{t-\tau} C(e_t(A, B, \cdots, B)) \right) \\
&= C(e_1) + D_1 C(e_2(A)) + D_1^2 C(e_3(A, B)) \\
&= C(e_1) + D_1 C(e_2(A)) + D_1^2 \left(\sum_{t=3}^{T} D_{t-\tau} C(e_t(A, B, \cdots, B)) \right) \\
&= \cdots = \sum_{t=1}^{T} D_1^{t-1} C(e_t(A, \cdots, A)),
\end{aligned}$$

where the first and third equations come from $C(0) = 0$, and the second and forth equations come from the binding ICs for selves 2 and 3, respectively, the fifth and sixth equations come from iterating the same procedure as in the first four equations.

It is straightforward to check that (PC) are slack and can be ignored, so the effort path of a naive agent is the same as with a time-consistent agent with discount factor $\tilde{D}_t = D_1^t$. $\square$

Online Appendix

A Consumer Heterogeneity

In this appendix, we present the results summarized in Subsection 3.3. To simplify the exposition, we assume that the consumer has a constant deterministic income w in each period. The setting is the same as in Section 2, except that firms do not know the consumer's "type" (either $\hat{\beta}$ or β). Instead, they have some prior distribution over the consumer's possible types.

This is a dynamic game with incomplete information, where the consumer's contract can signal his type to the firm. After seeing the contract offered by the consumer, the firm must update its beliefs about the consumer's type. We therefore incorporate the standard consistency condition from perfect Bayesian equilibrium, which requires the firm's beliefs to be consistent with Bayes' rule on histories that are reached with positive probability.

A pure strategy for type θ of the time-1 self is a consumption vector ("contract") $\mathbf{c}(\theta)$. A pure strategy for the firm is a mapping d from the space of possible contracts to $\{0, 1\}$, which specifies whether the firm accepts ($d = 1$) or rejects ($d = 0$) each contract offered by the time-1 self.

A time-t history describes all actions by the consumer and all uncertainty realized until period t: $\mathtt{h}_t = (\mathbf{c}, h^t, s_t)$, where $h^t \in \{A, B\}^{t-1}$ is an option history (as defined in Section 2). Let $\mathcal{H}_t$ denote the set of all possible time-t histories. A pure strategy for type θ of self $t \in \{2, ..., T-1\}$ is a mapping from the history of previous actions and realized uncertainty to an action $\sigma_t : \mathcal{H}_t \to \{A, B\}$.[28]

We can now generalize the definition of perception-perfect equilibrium (see Appendix E) to incorporate imperfect information. A "perception-perfect equilibrium with Bayesian rationality on the firm side" (henceforth *equilibrium*) is a contract and a pair of strategies

[28] As in the model without heterogeneity, there is no loss of generality in assuming that there is only one option in period T. Therefore, we do not need to include choices by self T.

for each consumer type,

$$(\mathbf{c}(\theta), \sigma_2(\theta), ..., \sigma_{T-1}(\theta)) \text{ and } (\hat{\sigma}_2(\theta), ..., \hat{\sigma}_{T-1}(\theta)),$$

and a strategy d for the firm such that:

- For each θ, $\mathbf{c}(\theta)$ maximizes the expected experienced utility (E1) of the time-1 self of type θ under the assumption that each self $r > 1$ uses strategy $\hat{\sigma}_r(\theta)$ and the firm uses strategy d.
- For each θ, history $\mathtt{h}_t$, and $t > 1$, the strategy $\sigma_t(\theta)$ maximizes the expected experienced utility (E1) of the time-t self of type θ conditional on $\mathtt{h}_t$ under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r(\theta)$.
- For each θ, history $\mathtt{h}_t$, and $t > 1$, the strategy $\hat{\sigma}_t(\theta)$ maximizes the expected perceived utility (E2) of the time-t self of type θ conditional on $\mathtt{h}_t$ under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r(\theta)$.
- For each $\mathbf{c}$, the acceptance decision $d(\mathbf{c})$ maximizes the firm's expected discounted profits given the firm's beliefs about the consumer's type θ under the assumption that each type of the consumer uses strategy $\sigma_t(\theta)$ in all periods $t > 1$; and
- For any contract offered by some type $\tilde{\mathbf{c}} \in \mathbf{c}(\Theta)$, the firm's beliefs about the consumer's type θ are derived by Bayes' rule.

A.1 Unknown Naiveté

In this subsection, we show that the results from Section 2 remain unchanged when the firm does not know the consumer's naiveté parameter $\hat{\beta}$. Suppose the firm has a prior distribution with full support over the non-degenerate type space $\Theta \subseteq (\beta, 1]$. Note that the type space may be discrete or continuous.

We refer to the equilibrium consumption in the model in which the firm knows the consumer's type (Section 2) as the "full-information contract." We now show that any equilibrium of this game has complete pooling at the full-information contract:

Proposition 4. *Suppose the firm does not know the consumer's naiveté parameter. There exists an equilibrium. Moreover, in any equilibrium, there is complete pooling at the full-information contract.*

The key intuition for Proposition 4 is that the equilibrium contract with full information does not depend on $\hat{\beta}$ (Corollary 1). Since that contract maximizes each type's perceived utility and gives zero profits, there are no beliefs by the firm about the consumer's type that would allow the consumer to obtain a higher perceived utility while not giving negative profits to the firm.

Proposition 4 implies that, as in Theorem 1, the welfare loss from dynamic inconsistency vanishes as the contracting horizon grows.

A.2 Adding Sophisticated Consumers

In the previous subsection, we assumed that, while the firm did not know the consumer's naiveté parameter $\hat{\beta}$, it still knew that the consumer was (at least partially) naive, so that $\hat{\beta} > \beta$ for all types. We now introduce a sophisticated consumer type into the analysis. Formally, consider the type space $\Theta \subseteq [\beta, 1]$ and suppose that the support of the firm's beliefs about the consumer's naiveté parameter includes both the sophisticated type ($\hat{\beta} = \beta$) and at least one naive type ($\hat{\beta} > \beta$). Note that, as before, the type space can be discrete or continuous.

The proposition below shows that in the equilibrium of this game, all naive types get the full-information contract:

Proposition 5. *Suppose the support of the firm's beliefs includes both the sophisticated type and at least one naive type. Then, in any equilibrium, all types get their full-information contract.*

The intuition is as follows. When given the naive consumers' full-information contract, the sophisticated consumer understands that his future selves will pick the alternative, rather than the baseline option. So he prefers to offer his own full-information contract, which prevents his future selves from deviating. However, because naive consumers believe they will pick the baseline option, they have a higher perceived utility from their own full-information contract. And because they are each offered their full-information contracts, there are no beliefs that firms can have about consumer types that would justify them offering any other contract. Then, as in the model with full information, the inefficiency from time inconsistency vanishes for all naive consumers as the contracting horizon grows (but not for the sophisticated consumer).

A.3 Unknown Time-Consistency Parameter

Suppose now that the firm does not know the consumer's time-consistency parameter β. The firm has a prior distribution with full support over the non-degenerate type space $\Theta \subseteq (0, \hat{\beta}]$. When $\hat{\beta} \in \Theta$, the model allows for both a sophisticated time-inconsistent type ($\hat{\beta} = \beta < 1$) and a time-consistent type ($\hat{\beta} = \beta = 1$). To avoid situations in which consuming all resources in the first period maximizes welfare (which would coincide with a present-biased consumer's choice), we assume that $\lim_{c \searrow 0} u'(c) = +\infty$ in this subsection.

We first show that, unlike when the consumer's naiveté is not known, there is no equilibrium in which multiple types get their full-information contracts:

Lemma 5. *There exists $\overline{T}$ such that for all $T > \overline{T}$, any equilibrium has at most one type offering his full information contract.*

The intuition of Lemma 5 is as follows. If more than one type offered their full-information contracts, the more time-consistent one would pick the full-information contract of the less time-consistent type and choose the baseline rather than the alternative option. The firm offering this contract would lose money, since the baseline was an unprofitable decoy option not meant to be chosen on the equilibrium path.

Having shown that the full information contracts cannot be supported in equilibrium, we now turn to the characterization of the equilibrium.[29] As is common in signaling games,

[29] Although this appendix assumes that the consumer has the bargaining power, which leads to a signaling

if we do not impose restrictions on beliefs that firms can have off the equilibrium path, there are many equilibria. We adopt the D1 criterion (Banks and Sobel, 1987) to deal with this multiplicity issue. For simplicity, we assume that there are only two types: β_L with probability q and β_H with probability $1-q$, where $0 < \beta_L < \beta_H \leq \hat{\beta} \leq 1$. Note that when $\beta_H = \hat{\beta} = 1$, the model has one time-inconsistent type and one time-consistent type. It is straightforward to extend our results to any finite number of types.

We will show that the allocation in the unique equilibrium that survives D1 corresponds to the "least-costly separating allocation." In this allocation, the high type gets the full-information consumption, whereas the low type gets the allocation that maximizes his perceived utility among those leaving zero profits to the firm and ensuring that the high type does not wish to deviate. That is, the allocation of the high type solves program (6), whereas the allocation of the low type solves the following program:

$$\max_{\{c(s_t)\}} u(c_1) + \beta_L E\left[\sum_{t=2}^{T} \delta^{t-1} u(c_t(B, B, ..., B))\right], \tag{A1}$$

subject to zero-profit condition, (IC), (PC), and the constraint that requires type β_H to prefers his full-information contract over the contract for type β_L, which choosing the baseline options:

$$u(c_1)+\beta_H E\left[\sum_{t=2}^{T} \delta^{t-1} u(c_t(B, B, ..., B))\right] \leq u(c_1^H)+\beta_H E\left[\sum_{t=2}^{T} \delta^{t-1} u(c_t^H(B, B, ..., B))\right]. \tag{A2}$$

From the previous lemma, it follows that (A2) must bind when T is large enough, so the equivalence to the auxiliary program no longer holds for the low type.

The next lemma formally shows that the high type gets his full-information contract in any equilibrium that survives D1 (i.e., there is no distortion at the top):

Lemma 6. *In any equilibrium that survives D1, type β_H gets his full-information contract.*

We next show that there is an equilibrium that survives D1 in which consumers get the

model, a similar result can be shown when the bargaining power is on the firm side. In that case, the model becomes one of screening. It can be shown that for generic income paths the full-information contract is not incentive compatible as the contract length grows.

least-costly separating allocation. Furthermore, we show that the equilibrium that survives D1 is unique.

Lemma 7. *There exists an equilibrium that survives D1 in which consumers get the least-costly separating allocation. Furthermore, in any equilibrium surviving D1, the consumers get the least costly separating allocation.*

Let W_T^L denote the equilibrium welfare of type β_L, and recall that W_T^C is the equilibrium welfare the time-consistent consumer (who maximizes welfare). The proposition below establishes that in equilibrium, the less time-consistent type gets a contract in which he consumes more than the full-information amount in the first period, thereby under-saving for the future. Moreover, this informational distortion does not vanish as the contracting length grows, so his equilibrium allocation does not converge to the Pareto frontier:

Proposition 6. *Suppose* $\lim_{c\searrow 0} u'(c) = +\infty$. *In any equilibrium satisfying D1:*

- *There exists* $\bar{T}$ *such that, for all* $T > \bar{T}$, *type* β_L *consumes more in the first period than in his full-information contract.*
- *The welfare loss is uniformly bounded away from zero as the contracting horizon grows:* $\liminf_{T\nearrow\infty}(W_T^C - W_T^L) > 0$.

B Non-Exclusive Contracts

This appendix considers the model in which contracts are not exclusive, so consumers can, at any point in time, sign a new contract with another firm. As in the model with one-sided commitment, to characterize the equilibrium consumption, there is no loss of generality in restricting attention to equilibria in which the consumer never contracts with another firm.[30]

[30]We assume that contracting is costless. If the cost of contracting with another firm is large enough, we return to the baseline model in which consumers can commit to long-term contracts. More generally, one can envision situations in which the cost of contracting is positive but not too large, so consumers only have partial commitment.

When contracts are not exclusive, firms cannot add unprofitable baseline options that naive consumers think they will choose but end up not choosing. If they offered such a contract, the consumer would stick to the baseline and readjust consumption by contracting with another firm. Therefore, any equilibrium contract must make zero profits both along the consumer's perceived path and the equilibrium (i.e., firm's perceived) path. In fact, our next lemma shows that, starting at any history, the expected PDV of future consumption must be the same in all option histories:

Lemma 8. *Suppose contracts are not exclusive. For any (s_t, h^t), the expected present discounted value of consumption in any option history path following h^t must be the same.*

The proof is in the supplementary appendix, but its intuition is straightforward. With non-exclusive contracts, the consumer can always smooth consumption by contracting with a new firm. Therefore, he would always pick the option path with the highest PDV of consumption.

Consider an (auxiliary) consumption-savings problem, in which the consumer is endowed with the expected PDV of income $\{w(s_t)\}$ in period 1. The only asset available is a risk-free bond that pays a gross return R, and the consumer can freely save or borrow. Since the consumer is time-inconsistent and naive, each period's self decides how much to consume and underestimates the present bias of future selves. As before, we focus on perception-perfect equilibria of this game.

To obtain the equilibrium consumption, we need to specify both how much the agent thinks his future selves will consumer and how much they actually consume. Let a_1 denote the asset available to the agent at time 1: $a_1 \equiv E \sum_{t=1}^{T} \frac{w(s_t)}{R^{t-1}}$. The agent, who has a_t asset at time t and believes he will choose consumption in periods $s > t$ according to $\hat{c}(a_s)$, believes that in period t, he will consume:

$$\hat{c}_t(a_t) \in \arg\max_{\tilde{c}} u(\tilde{c}) + \hat{\beta} \sum_{s>t} \delta^{s-t} u(\hat{c}_s(a_s)), \tag{B1}$$

subject to

$$\tilde{c} + \sum_{s>t} \frac{\hat{c}_s(a_s)}{R^{s-t}} \leq a_t, \tag{B2}$$

$$a_{t+1} = R(a_t - \tilde{c}), \tag{B3}$$

$$a_{s+1} = R(a_s - c_s(a_s)) \text{ for all } s > t. \tag{B4}$$

However, in period t, he chooses to consume:

$$c_t(a_t) \in \arg\max_{\tilde{c}} u(\tilde{c}) + \beta \sum_{s>t} \delta^{s-t} u(\hat{c}_s(a_s)), \tag{B5}$$

subject to (B2), (B3), and (B4).

The next proposition establishes the equivalence between non-exclusive contracts and the consumption-savings problem.

Proposition 7. *The problem with non-exclusive contracts is equivalent to the consumption-saving problem. In particular, the consumption paths in the two problems are the same.*

Lastly, we show that the welfare loss in the consumption-savings problem does not vanish as the contracting horizon goes to infinity. Note that if in the welfare-maximizing allocation the agent consumes all resources in the first period, leaving zero consumption in all future periods, there is no scope for contracting with other firms after the first period. Then, there is no welfare loss from non-exclusive contracting. To rule out this uninteresting case, we proceed as in Subsection A.3 and assume that $\lim_{c\searrow 0} u'(c) = +\infty$.

Proposition 8. *Suppose contracts are not exclusive, u is bounded, $\delta < 1$, and $\lim_{c\searrow 0} u'(c) = +\infty$. The welfare loss from time inconsistency is uniformly bounded away from zero as the contracting horizon T goes to infinity.*

C Effort

This appendix formally presents the analysis from Subsection 3.6, where we considered the effort model. As in the consumption model, when the agent is naive, contracts involve two options in each period: a baseline effort (B) that the agent thinks his future selves will pick and an alternative effort (A) that they end up picking. As before, let h^t denote the options chosen by the agent up to time t.

The effort path of the naive agent solves:

$$\min_{\mathbf{e}} \sum_{t=1}^{T} D_{t-1} C(e_t(B, B, ..., B)$$

subject to

$$\sum_{t=1}^{T} e_t(A, ..., A) = E_T, \tag{C1}$$

$$\sum_{t=\tau}^{T} D_{t-\tau} C(e_t(A, B, \cdots, B)) \leq \sum_{t=\tau}^{T} D_{t-\tau} C(e_t(B, B, \cdots, B)), \forall \tau \geq 2, \tag{C2}$$

$$\sum_{t=\tau}^{T} \hat{D}_{t-\tau} C(e_t(B, B, \cdots, B)) \leq \sum_{t=\tau}^{T} \hat{D}_{t-\tau} C(e_t(A, B, \cdots, B)), \forall \tau \geq 2. \tag{C3}$$

That is, the agent minimizes his perceived discounted cost subject to the task-completion constraint (C1), IC (C2), and PC (C3). This program is analogous to the one in the proof of Proposition 1, except that the zero profits constraint is replaced by the task completion constraint (C1) and the agent minimizes his discounted effort costs rather than maximizes his discounted utility. As before, PC requires the agent to believe that his future selves pick B, whereas IC requires them to switch to A instead.

Since the firm and the agent disagree on the options that the agent will pick, they have different beliefs about the total effort that will be exerted on the equilibrium path. The firm accepts a contract as long as it believes that the agent will complete the task, regardless of what the agent believes. Therefore, as with the zero profits constraint in Proposition 1, the task-completion constraint (C1) only needs to hold according to the firm's beliefs.

We solve this program in the proof of Proposition 3 in the supplementary appendix. Here we illustrate it by solving the case of three periods and quasi-hyperbolic discounting, as we did in the text for Lemma 2. This illustration helps clarify the difference between contracting over consumption and over effort.

The equilibrium program becomes:

$$\min_{\mathbf{e}} C(e_1) + \beta \left[\delta C\left(e_2\left(B\right)\right) + \delta^2 C\left(e_3\left(B\right)\right)\right]$$

subject to

$$e_1\left(A\right) + e_2\left(A\right) + e_3\left(A\right) = E_3, \tag{C4}$$

$$C\left(e_2\left(B\right)\right) + \hat{\beta}\delta C\left(e_3\left(B\right)\right) \leq C\left(e_2\left(A\right)\right) + \hat{\beta}\delta C\left(e_3\left(A\right)\right) \tag{PC}$$

$$C\left(e_2\left(A\right)\right) + \beta\delta C\left(e_3\left(A\right)\right) \leq C\left(e_2\left(B\right)\right) + \beta\delta C\left(e_3\left(B\right)\right) \tag{IC}$$

Note first that (IC) must bind. Otherwise, we could reduce the perceived cost in the objective function by reducing $e_3(B)$. Since (IC) binds, (PC) can be written as a monotonicity constraint:

$$e_3\left(B\right) \leq e_3\left(A\right). \tag{C5}$$

In words, because agents under-estimate their present bias, they think they will leave less effort for the last period than they end up leaving. We ignore this monotonicity constraint for now and verify that it holds later.

For each $\epsilon > 0$ small, consider a perturbation to the baseline efforts $\tilde{e}_2\left(B\right)$ and $\tilde{e}_3\left(B\right)$ that shifts effort from period 2 to period 3 according to self 2's preferences:

$$C\left(\tilde{e}_2\left(B\right)\right) = C\left(e_2\left(B\right)\right) + \epsilon,\ C\left(\tilde{e}_3\left(B\right)\right) = C\left(e_3\left(B\right)\right) - \frac{\epsilon}{\beta\delta}.$$

By construction, this perturbation preserves IC. Moreover, since the objective function evaluates costs from the perspective of self 1, this perturbation improves the objective. Thus, to minimize costs, the solution leaves as little effort as possible to the last period:

$$e_3(B) = 0. \tag{C6}$$

It follows directly from (C6) that the monotonicity condition (C5) holds. Substituting back in IC, we obtain:

$$C\left(e_2\left(B\right)\right) = C\left(e_2\left(A\right)\right) + \beta\delta C\left(e_3\left(A\right)\right). \tag{C7}$$

Substituting (C6) and (C7) in the objective function, we obtain

$$C(e_1) + \beta\delta C\left(e_2\left(A\right)\right) + \left(\beta\delta\right)^2 C\left(e_3\left(A\right)\right),$$

which is the cost of a time-consistent agent with discount factor $\beta\delta$.

Note how the argument above differs from the one in Lemma 2. The way to exploit naiveté in the consumption model is to postpone consumption in the baseline from period 2 to period 3. So, when deciding whether to consume now or to leave resources for the future, self 1 decides according to his long-run discount rate $\beta\delta^2$. Then, self 2 deviates from B to A, effectively bringing some consumption from $c_3(B)$ to period 2 (and reducing the consumption left to self 3). Since self 2 discounts period-3 consumption by $\beta\delta$, the rate between $u(c_1)$ and $u(c_2(A))$ is $\frac{\beta\delta^2}{\beta\delta} = \delta$, as shown in the auxiliary program in Lemma 2.

On the other hand, the way to exploit naiveté in the effort model is to require all effort in period 2, leaving zero effort for the future. Thus, self 1 decides how much effort to leave to the future according to his 1-period discount $\beta\delta$. Then, self 2 deviates from the baseline, leaving some effort for period 3. He decides how much to leave for period 3 also based on his 1-period discount $\beta\delta$. Therefore, the rate between $u(c_1)$ and $u(c_3(A))$ is $(\beta\delta)^2$.

D Sophisticated Agents

In Section 2, we characterized the equilibrium with either time-consistent or (partially) naive present-biased consumers. We now consider the case of sophisticated consumers, who correctly predict their future preferences ($\hat{\beta} = \beta$. We are interested in the asymptotic welfare of sophisticated consumers as the contracting horizon grows.

Recall that the welfare function does not discount future periods by the additional term β. Therefore, if consuming all resources in the first period maximizes welfare, the sophisticated agent must also consume all the resources in the first period. In this case, the equilibrium of the sophisticated agent trivially maximizes welfare. To rule out this uninteresting case, we proceed as in Subsection A.3 and assume that $\lim_{c \searrow 0} u'(c) = +\infty$. We will show that, when this is the case, the welfare loss from present bias of sophisticated agents does not vanish.

A sophisticated agent evaluates future consumption according to (1) with $\hat{\beta} = \beta$. Therefore, he fully understands that his future selves will behave like someone with the same time-consistency parameter as his current self. As with time-consistent consumers, since a sophisticated consumer agrees with the firm about his future preferences, there is no need to allow for options in the contract. Therefore, there is no loss of generality in restricting contracts to be vectors of state-dependent consumption. Because parties can commit to long-term contracts, any contract that is accepted by a firm must maximize the utility of the period-1 self subject to the zero-profits constraint. The equilibrium contract solves the following program:

$$\max_{\{c(s_t)\}} u(c(s_1)) + \beta E\left[\sum_{t=2}^{T} \delta^{t-1} u\left(c\left(s_t\right)\right)\right], \tag{D1}$$

subject to the zero-profits constraint,

$$\sum_{t=1}^{T} E\left[\frac{w\left(s_t\right) - c\left(s_t\right)}{R^{t-1}}\right] = 0. \tag{D2}$$

Let W_T^S denote the equilibrium equilibrium welfare of the sophisticated consumer, which evaluates the consumption path according to the agent's long-run preferences (2), and recall that W_T^C is the welfare in the benchmark case of a time-consistent consumer. Since the time-consistent consumer maximizes welfare, the welfare loss from dynamic inconsistency cannot be negative:

$$W_T^C - W_T^S \geq 0.$$

We now show that unlike with partially naive agents, the consumption path of a sophisticate does not converge to the welfare-maximizing path as the contracting horizon grows. Therefore, the previous inequality is strict:

Proposition 9. *Suppose u is bounded, $\delta < 1$, and $\lim_{c \searrow 0} u'(c) = +\infty$. Then, the welfare loss of a sophisticated consumer is uniformly bounded away from zero:*

$$\liminf_{T \nearrow +\infty} \left(W_T^C - W_T^S \right) > 0.$$

Note that, in our model, the individual can consume in all periods, including when contracts are signed. If, instead, contracting occurred before consumption (say, at period 0), sophisticated consumers would commit to the ex-ante optimal contract – see, DellaVigna and Malmendier (2004); Heidhues and Kőszegi (2010). The inefficiency with naive consumers, as well as the asymptotic efficiency result, remains unchanged if we add a contracting period with no consumption.

E Equilibrium Definition and Mixed Strategies

In this appendix, we present a formal definition of perception-perfect equilibria and show that the results in the paper generalize to mixed strategy equilibria.

For each state-dependent consumption $\{c(s_t)\}_{t \geq \tau}$, let

$$u(c(s_\tau)) + \beta E \left[\sum_{t > \tau} \delta^{t-\tau} u(c(s_t)) \;\middle|\; s_\tau \right] \tag{E1}$$

denote self τ's "experienced utility," and let

$$u(c(s_\tau)) + \hat{\beta} E \left[\sum_{t > \tau} \delta^{t-\tau} u(c(s_t)) \;\middle|\; s_\tau \right] \tag{E2}$$

denote self τ's "perceived utility."

As described in the text, it is without loss of generality to focus on contracts that offer a baseline (B) and an alternative (A) options in each period $t = 2, ...T-1$. This is no longer true with mixed strategies. We generalize the definitions to allow for arbitrary message spaces when we consider mixed strategy equilibria.

Recall that we adopted the convention that each state describes all previous realization of uncertainty. A time-t history describes all actions by the consumer and all uncertainty realized until period t: $\mathtt{h}_t = (\mathbf{c}, h^t, s_t)$, where $\mathbf{c}$ is a contract offered by the time-1 self, $h^t \in \{A, B\}^{t-1}$ is an option history (which lists the options taken by all previous selves as defined in Section 2), and s_t is a state of the world at time t. Let $\mathcal{H}_t$ denote the set of all possible time-t histories.

A pure strategy for the time-1 self is a consumption vector $\mathbf{c}$. A pure strategy for the firm is a mapping d from the space of possible consumption vectors to $\{0, 1\}$ specifying whether the firm accepts ($d = 1$) or rejects ($d = 0$) each consumption vector offered by the time-1 self. A pure strategy for self $t \in \{2, ..., T-1\}$ is a mapping from the time-t history to an option, that is, $\sigma_t : \mathcal{H}_t \to \{A, B\}$.[31]

Before stating the equilibrium definition, we need to specify each player's payoffs. We start with the firm, which has correct beliefs. Let $\Pi(\mathbf{c}, \sigma_2, ..., \sigma_{T-1}, \hat{d})$ denote the firm's expected profits from accepting ($\hat{d} = 1$) or rejecting ($\hat{d} = 0$) the consumption vector $\mathbf{c}$ when selves $t > 1$ of the consumer play σ_t.

Since the consumer has incorrect beliefs about his future preferences, we need to distinguish between the actions that the consumer thinks he will choose and the actions that he ends up choosing. The agent's *perceived utility* determines what he thinks he will choose in the future, whereas the agent's *experienced utility* determines what he will end up choosing (see equations E1 and E2 in Appendix E).

- Let $U_1(\mathbf{c}, \sigma_2, ..., \sigma_{T-1}, d)$ denote the time-1 self's expected experienced utility from offering contract $\mathbf{c}$ if each future self $r > 1$ plays strategy σ_r and the firm plays d.
- For $t > 1$, let $U_t(\sigma_t, ..., \sigma_{T-1} | \mathtt{h}_t)$ denote the expected experienced utility of the time-t self conditional on history $\mathtt{h}_t$ when each self $r > t$ plays strategy σ_r.
- Let $\hat{U}_t(\sigma_t, ..., \sigma_{T-1} | \mathtt{h}_t)$ denote the expected perceived utility of the time-t self conditional on history $\mathtt{h}_t$ when each self $r > t$ plays strategy σ_r.

[31] It is without loss of generality to focus on consumption vectors that do not offer any options to the time-T self, since it would be a dominant strategy for him to pick the one with the highest consumption (both according to the experienced and the perceived utility).

Definition 1. A perception-perfect equilibrium is a consumption vector $\mathbf{c}$, a pair of strategies $(\sigma_2, ..., \sigma_{T-1})$ and $(\hat{\sigma}_2, ..., \hat{\sigma}_{T-1})$ and an acceptance decision d such that:

- $\mathbf{c}$ maximizes self 1's expected experienced utility under the assumption that his future selves use strategy $\hat{\sigma}_r$ and the firm uses strategy d:

$$U_1(\mathbf{c}, \hat{\sigma}_2, ..., \hat{\sigma}_{T-1}, d) \geq U_1(\mathbf{c}', \hat{\sigma}_2, ..., \hat{\sigma}_{T-1}, d), \;\; \forall \mathbf{c}'.$$

- For all $t > 1$ and all $\mathtt{h}_t$, $\sigma_t(\mathtt{h}_t)$ maximizes self-t's expected experienced utility under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$:

$$U_t(\sigma_t, \hat{\sigma}_{t+1}, ..., \hat{\sigma}_{T-1}|\mathtt{h}_t) \geq U_t(\sigma_t', \hat{\sigma}_{t+1}, ..., \hat{\sigma}_{T-1}|\mathtt{h}_t), \;\; \forall \sigma_t'.$$

- For all $t > 1$ and all $\mathtt{h}_t$, $\hat{\sigma}_t(\mathtt{h}_t)$ maximizes the consumer's time-t expected perceived utility under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$:

$$\hat{U}_t(\hat{\sigma}_t, \hat{\sigma}_{t+1} ..., \hat{\sigma}_{T-1}|\mathtt{h}_t) \geq \hat{U}_t(\sigma_t', \hat{\sigma}_{t+1} ..., \hat{\sigma}_{T-1}|\mathtt{h}_t), \;\; \forall \sigma_t'.$$

- For all $\mathbf{c}$, $d(\mathbf{c})$ maximizes the firm's expected discounted profits under the assumption that the consumer uses strategies σ_t for all t:

$$\Pi(\mathbf{c}, \sigma_2, ..., \sigma_{T-1}, d) \geq \Pi(\mathbf{c}, \sigma_2, ..., \sigma_{T-1}, d'), \;\; \forall d'.$$

We now generalize the definition of perception-perfect equilibrium to allow for mixed strategies. In this case, we need to work with more general message spaces, since the restriction to two possible messages is no longer without loss of generality.

Let M_2 be a non-empty, compact space of possible messages in period 2 with generic element m_2. For each $t \in \{3, ..., T-1\}$, let $M_t(m_2, ..., m_{t-1})$ be a non-empty, compact space of possible messages in period t conditional on previous messages $(m_2, ..., m_{t-1})$. Let $\mathbf{M}$ denote the space of all possible messages in all periods.[32] A consumption vector (or "contract") $\mathbf{c}$ specifies, for each period, the consumption conditional on all messages up to t and all realized uncertainty: $c(m_2, ..., m_{t-1}, m_t, s_t)$. A period-$t$ history

[32]That is, $\mathbf{M} \equiv \{M_1, ..., M_{T-1}(m_1, ..., m_{T-2}) \;:\; m_1 \in M_1, ..., (m_1, ..., m_{T-1}) \in M_1 \times ... \times M_{T-1}(m_1, ..., m_{T-2})\}$.

$\mathtt{h}_t = (\mathbf{c}, \mathbf{M}, m_2, ..., m_{t-1}, s_t)$ consists of a consumption vector and a message space offered at time 1, the messages sent in all previous periods, and the state of the world at t describing all realized uncertainty.

With a slight abuse notation, we now allow σ to be a mixed strategy as well. A mixed strategy for the time-1 self σ_1 is a distribution over (compact and non-empty) message spaces and contracts. A mixed strategy for the firm σ_{firm} is a distribution over acceptance decisions for each contract and message space offered by the time-1 self. A mixed strategy for self $t \in \{2, ..., T-1\}$ specifies, for each period-t history, a distribution over messages: $\sigma_t(\mathbf{c}, m_2, ..., m_{t-1}, s_t) \in \Delta(M_t(m_2, ..., m_{t-1}))$.

We now extend the payoffs to allow for mixed strategies:

- Let $\Pi(\sigma_1, ..., \sigma_{T-1}, \sigma_{\text{firm}})$ denote the firm's expected profits from playing σ_{firm} when each consumer self plays strategy σ_1.
- Let $U_1(\sigma_1, ..., \sigma_{T-1}, \sigma_{\text{firm}})$ denote the expected experienced utility (E1) of the time-1 self from playing strategy σ_1 if each future self $r > 1$ plays σ_r and the firm plays σ_{firm}.
- For $t > 1$ let $U_t(\sigma_t, ..., \sigma_{T-1}|\mathtt{h}_t)$ denote the expected experienced utility (E1) of the time-t self conditional on history $\mathtt{h}_t$ when each self $r > t$ plays strategy σ_r.
- For $t > 1$ let $\hat{U}_t(\sigma_t, ..., \sigma_{T-1}|\mathtt{h}_t)$ denote the expected perceived utility (E2) of the time-t self conditional on history $\mathtt{h}_t$ when each self $r > t$ plays strategy σ_r.

We can now state the equilibrium definition:

Definition 2. A perception-perfect equilibrium in mixed strategies is a pair of strategies for the consumer $(\mathbf{c}, \sigma_2, ..., \sigma_{T-1})$ and $(\hat{\sigma}_2, ..., \hat{\sigma}_{T-1})$ and a strategy for the firm σ_{firm} such that:

- σ_1 maximizes self 1's expected experienced utility under the assumption that his future selves use strategy $\hat{\sigma}_r$ and the firm uses strategy σ_{firm}:

$$U_1(\sigma_1, \hat{\sigma}_2, ..., \hat{\sigma}_{T-1}, \sigma_{\text{firm}}) \geq U_1((\mathbf{c}', \mathbf{M}'), \hat{\sigma}_2, ..., \hat{\sigma}_{T-1}, \sigma_{\text{firm}}), \;\; \forall \mathbf{c}', \mathbf{M}'.$$

- For all $t > 1$ and all $\mathtt{h}_t$, $\sigma_t(\mathtt{h}_t)$ maximizes self-t's expected experienced utility under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$:

$$U_t(\sigma_t, \hat{\sigma}_{t+1}, ..., \hat{\sigma}_{T-1}|\mathtt{h}_t) \geq U_t(m_t, \hat{\sigma}_{t+1}, ..., \hat{\sigma}_{T-1}|\mathtt{h}_t), \ \ \forall m_t.$$

- For all $t > 1$ and all $\mathtt{h}_t$, $\hat{\sigma}_t(\mathtt{h}_t)$ maximizes the consumer's time-t expected perceived utility under the assumption that selves $r > t$ use strategy $\hat{\sigma}_r$:

$$\hat{U}_t(\hat{\sigma}_t, \hat{\sigma}_{t+1}..., \hat{\sigma}_{T-1}|\mathtt{h}_t) \geq \hat{U}_t(m_t, \hat{\sigma}_{t+1}..., \hat{\sigma}_{T-1}|\mathtt{h}_t), \ \ \forall m_t.$$

- For all $\mathbf{c}$, $\sigma_{\text{firm}}(\mathbf{c})$ maximizes the firm's expected discounted profits under the assumption that the consumer uses strategies σ_t in periods $t > 1$:

$$\Pi(\mathbf{c}, \mathbf{M}, \sigma_2, ..., \sigma_{T-1}, \sigma_{\text{firm}}) \geq \Pi(\mathbf{c}, \mathbf{M}, \sigma_2, ..., \sigma_{T-1}, d'), \ \forall d' = 0, 1.$$

We can now establish that our restriction to pure strategies in the text was without loss of generality. The *equilibrium program* is:

$$\max_{\{c(s_t,h^t)\}} u(c(s_1)) + \beta E\left[\sum_{t=2}^{T} \delta^{t-1} u(c(s_t, \hat{\sigma}_2, \hat{\sigma}_3, ..., \hat{\sigma}_{T-1}))\right], \tag{E3}$$

subject to

$$\sum_{t=1}^{T} E\left[\frac{w\left(s_t\right) - c\left(s_t, \sigma_2, \sigma_3, ..., \sigma_{T-1}\right)}{R^{t-1}}\right] = 0, \tag{Zero Profits}$$

$$u(c(s_\tau, \left(h^{\tau-1}, \hat{m}_\tau\right))) + \hat{\beta} E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t, \left(h^{\tau-1}, \hat{m}_\tau, \hat{\sigma}_{\tau+1}, ..., \hat{\sigma}_{T-1}\right)\right)\right) \middle| \ s_\tau\right] \tag{PC}$$

$$\geq u(c(s_\tau, \left(h^{\tau-1}, m'_\tau\right)))+\hat{\beta} E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t, \left(h^{\tau-1}, m'_\tau, \hat{\sigma}_{\tau+1}, ..., \hat{\sigma}_{T-1}\right)\right)\right) \middle| \ s_\tau\right], \forall \hat{m}_\tau \in supp(\hat{\sigma}_\tau), m'_\tau \in M_\tau$$

and

$$u(c(s_\tau, \left(h^{\tau-1}, m_\tau\right))) + \beta E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t, \left(h^{\tau-1}, m_\tau, \hat{\sigma}_{\tau+1}, ..., \hat{\sigma}_{T-1}\right)\right)\right) \middle| \ s_\tau\right] \tag{IC}$$

$$\geq u(c(s_\tau, \left(h^{\tau-1}, m'_\tau\right)))+\beta E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t, \left(h^{\tau-1}, m'_\tau, \hat{\sigma}_{\tau+1}, ..., \hat{\sigma}_{T-1}\right)\right)\right) \middle| \ s_\tau\right], \forall m_\tau \in supp(\sigma_\tau), m'_\tau \in M_\tau$$

The next lemma establishes that the consumption path for time-inconsistent agents still coincides with the solution of the auxiliary program when we allow for mixed strategies:

Lemma 9. *In any perception-perfect equilibrium in mixed strategies, the consumption path solves the auxiliary program* (7)*.*

The proof is in the supplementary appendix. To illustrate it, consider the case in which $T = 3$ and the consumer has income w in each period. Suppose self 1 believes that self 2 will pick option B_1 with probability θ_1 and B_2 with probability θ_2. (PC) states that

$$u(c_2(B_1)) + \hat{\beta}\delta u(c_3(B_1)) = u(c_2(B_2)) + \hat{\beta}\delta u(c_3(B_2)) \geq u(c_2(A)) + \hat{\beta}\delta u(c_3(A)),$$

whereas (IC) requires

$$u(c_2(A)) + \beta\delta u(c_3(A)) \geq u(c_2(B_1)) + \beta\delta u(c_3(B_1)),$$
$$u(c_2(A)) + \beta\delta u(c_3(A)) \geq u(c_2(B_2)) + \beta\delta u(c_3(B_2)).$$

So self 1's perceived utility is

$$u(c_1) + \beta\delta(\theta_1 u(c_2(B_1)) + \theta_2 u(c_2(B_2))) + \beta\delta^2(\theta_1 u(c_3(B_1)) + \theta_2 u(c_3(B_2))).$$

We claim that both ICs must bind. First, note that at least one of them must bind (otherwise, we can raise $c_3(B_1)$ and $c_3(B_2)$ without affecting any other constraints). Suppose the IC associated with B_1 binds but not the one associated with B_2:

$$u(c_2(B_1)) + \beta\delta u(c_3(B_1)) = u(c_2(A)) + \beta\delta u(c_3(A)) > u(c_2(B_2)) + \beta\delta u(c_3(B_2)).$$

Recall that self 1 perceives that self 2 will mix between option B_1 and B_2, so self 1 believes that self 2 must be indifferent:

$$u(c_2(B_1)) + \hat{\beta}\delta u(c_3(B_1)) = u(c_2(B_2)) + \hat{\beta}\delta u(c_3(B_2)).$$

It follows that $c_2(B_1) > c_2(B_2), c_3(B_1) < c_3(B_2)$. Together with the perceived-choice constraints, it implies that

$$u(c_2(B_1)) + \delta u(c_3(B_1)) < u(c_2(B_2)) + \delta u(c_3(B_2)).$$

Consider an alternative contract by setting the consumption associated with option B_1 equal to the consumption associated with option B_2. This contract strictly increases self 1's perceived utility, a contradiction to the optimality of the original contract. So both IC constraints are binding. Then, $c_2(B_1) = c_2(B_2)$ and $c_3(B_1) = c_3(B_2)$. Similar to the proof of Lemma 2, $c_2(B_1) = c_2(B_2) = 0$. Substituting them back to the objective function leads to the auxiliary program

$$u(c_1) + \delta u(c_2(A)) + \beta\delta^2 u(c_3(A)).$$

On the other hand, if the alternative options have more than one options, i.e., the program is now

$$u(c_1) + \delta[\theta_1 u(c_2(A_1)) + \theta_2 u(c_2(A_2))] + \beta\delta^2[\theta_1 u(c_3(A_1)) + \theta_2 u(c_3(A_2))],$$

subject to the zero-profit condition

$$c_1 + \frac{\theta_1 c_2(A_1) + \theta_2 c_2(A_2)}{R} + \frac{\theta_1 c_3(A_1) + \theta_3 c_2(A_2)}{R^2} = w + \frac{w}{R} + \frac{w}{R^2}.$$

If options A_1 and A_2 are different, by Jensen's inequality and the strict concavity of $u(\cdot)$, merging these two options A_1 and A_2 would strictly increase self 1's payoff.

Supplementary Appendix (Not For Publication)

Omitted Proofs

Proof of Lemmas 1 and 3. We consider the model with one-sided commitment (Lemma 3). The proof of the two-sided commitment case (Lemma 1), which follows similar steps but is simpler, is omitted.

Suppose the period-t self of the consumer offers a contract $\mathcal{C}'_t$. Specifically, a contract at time t, $\mathcal{C}'_t$, specifies consumption on each possible state in each future time $\tau \geq t$. Denote the set of possible states by $K_{t,\tau}$, in which the first subscript corresponds to the time in which the contract is offered and the second subscript corresponds to the decision-making time τ. The contract specifies consumption for each different income states, so the contracting space is generally greater than the space of income states. In addition, perception-perfect equilibrium imposes no restrictions on $K_{t,\tau}$, i.e., $K_{t,\tau}$ can be arbitrary. To keep analysis tractable, we assume that $K_{t,\tau}$ has a product structure and only depends on decision making time τ. Otherwise, we can always add more states that are never reached so that it has a product structure and the resulting equilibrium is outcome-equivalent to the original equilibrium. Specifically, we write $K_{t,\tau} = \mathbb{S}_\tau \times H_\tau$, in which H_τ consists of all the possible income-independent messages/actions that the agent can send at time τ. The income-independent messages can be arbitrary. One of the reasons that an income-independent message can arise is from the consumer's different beliefs. Since we allow any contracts, we cannot impose what types of income-independent messages the consumer can send. For simplicity, we call H_τ the income-independent history. Without loss of generality, $H_1 = \emptyset$. Denote h_t a generic element in H_t. We call h_t an income-independent message. Denote $H_\tau(h_t)$ the states that can be reached at time τ from an earlier history $h_t \in H_t$ for $\tau > t$.

Fix a contract, we next write down the agent's strategy profile. Consider an agent who makes a decision at time τ. Suppose the income-independent messages that has been reached is $h_{\tau-1}$, which is an element in $H_{\tau-1}$. At time τ, the agent learns the income state, i.e., s_τ is realized. The agent needs to decide which message $a_\tau \in \Delta(H_\tau(h_{\tau-1}))$ to send, where $\Delta(\cdot)$ represents the set of lotteries. If there is one-sided commitment, the agent also needs to decide whether he will lapse or not, in which case, the strategy can be summarized by a pair (d_τ, a_τ), where $d_\tau \in \Delta(\{0,1\})$. If $d_\tau = 1$ with probability 1, then the agent stays, otherwise the contract is lapsed with a positive probability.

As described in the body of the paper, the perception-perfect equilibrium is solved by treating the agent's decisions in each period as if it were taken by a different player (i.e., a different "self"). The main claim is that for any perception-perfect equilibrium, the consumption vector must solve the program (P').

For the ease of exposition, we say that two perception-perfect equilibria are *equivalent* if all selves of the consumers have same actual and perceived consumption. We will establish the result through two separate claims:

Claim 2. *Fix a perception-perfect equilibrium. There exists an equivalent perception-perfect equilibrium in which the agent never lapses ($d_\tau = 1, \forall \tau$).*

Proof. Consider a perception-perfect equilibrium in which the agent lapses in some period $d_\tau = 0$ with a positive probability, replacing it with a contract $\mathcal{C}''_\tau$ from another firm. Since the other firm cannot lose money by offering this new contract, the old firm could have accepted a contract that substituted the terms of the old contract from this period on with the terms of the new contract, and the agent would have accepted to remain with the old firm. The constructed new contracts together with the agent's optimal decision forms a perception-perfect equilibrium that is equivalent to the original one. □

Claim 3. *Fix a perception-perfect equilibrium. There is an equivalent perception-perfect equilibrium that offers two options following any history: $\#|H_t(h_{t-1})| \leq 2$, for all $h_{t-1} \in H_{t-1},\ t \geq 2$.*

Proof. From the previous claim, we can restrict attention to equilibria in which the agent never lapses. Suppose $t_1 < t_2 < t_3$. Note that self t_1's prediction about self t_3's decision coincides with self t_2's prediction about self t_3's decision. Restricting $H_t(h_{t-1})$ to two messages – one that the agent will choose and another one that the agent thinks that he will choose – does not affect the actual consumption or the perceived consumption. Put differently, if $H_t(h_{t-1})$ has at least three messages, then there is at least one of them that the agent never sends and the agent never believes other selves would send. Therefore, we can restrict the income-independent message space to be at most two: one that the agent actually choose, and one that the agent thought he would choose. $\square$

Given these two claims, a contract offered by self t, $\mathcal{C}'_t$, must maximize the agent's utility subject to the zero profits, incentive compatibility, perceived choice, and non-lapsing constraints, concluding the proof of Lemma 3.

$\square$

Proof of Lemmas 2 and 4. In the text, we presented the proof for the case with two-sided commitment when there is no uncertainty and $T = 4$. Here, we consider the model with one-sided commitment case (Lemma 4), still assuming no uncertainty and $T = 4$. The proof for stochastic income and arbitrary T is presented in the supplementary appendix.

There are two ICs:

$$u(c_2(A)) + \beta[\delta u(c_3(A,B)) + \delta^2 u(c_4(A,B))] \geq u(c_2(B)) + \beta[\delta u(c_3(B,B)) + \delta^2 u(c_4(B,B))], \tag{sA1}$$

$$u(c_3(A,A)) + \beta\delta u(c_4(A,A)) \geq u(c_3(A,B)) + \beta\delta u(c_4(A,B)). \tag{sA2}$$

First, note that (sA1) must bind at an optimum (otherwise, we can raise $c_4(B,B)$, giving the agent a higher utility). Substitute the binding (sA1) in the objective to eliminate $c_4(B,B)$:

$$u(c_1) + \delta u(c_2(A)) + \beta[\delta^2 u(c_3(A,B)) + \delta^3 u(c_4(A,B))] + (\beta - 1)\delta u(c_2(B)).$$

Similarly, (sA2) must bind (otherwise, we can raise $c_4(A,B)$, increasing the agent's utility).

Use the binding (sA2) to rewrite the objective as:

$$u(c_1)+\delta u(c_2(A))+\delta^2 u(c_3(A,A))+\beta\delta^3 u(c_4(A,A))-(1-\beta)[\delta u(c_2(B))+\delta^2 u(c_3(A,B))].$$

Since $\beta < 1$, we should pick $c_2(B)$ and $c_3(AB)$ as small as possible subject to the constraints. Substituting $c_2(B) = c_3(AB) = 0$ back in this expression concludes the proof of Lemma 2. For the proof of Lemma 4, it remains to be verified that the non-lapsing constraints imply perceived non-lapsing constraints if we set $c_2(B) = c_3(A,B) = 0$.

Let $\hat{\mathbf{c}}$ denote a solution to the perceived outside option program, and let $\hat{V}_2^I = u(\hat{c}_2) + \hat{\beta}(\delta u(\hat{c}_3(B)) + \delta^2 u(\hat{c}_4(B)))$. We will use binding ICs constraints to obtain a lower bound on the perceived payoff of keeping the contract and show that is greater than the perceived outside option $\hat{V}_2^I$. We first use the the binding IC for self 2 to rewrite the perceived payoff:

$$\begin{aligned}
&u(c_2(B)) + \hat{\beta}(\delta u(c_3(BB)) + \delta^2 u(c_4(BB)) \\
&= u(0) + \frac{\hat{\beta}}{\beta}\beta(\delta u(c_3(B,B)) + \delta^2 u(c_4(B,B)) \\
&= u(0) + \frac{\hat{\beta}}{\beta}\left[u(c_2(A)) + \beta(\delta u(c_3(A,B)) + \delta^2 u(c_4(AB))) - u(0)\right] \\
&= \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \frac{\hat{\beta}}{\beta}\left[u(c_2(A)) + \beta(\delta u(c_3(A,B)) + \delta^2 u(c_4(AB)))\right],
\end{aligned}$$

where the first equality follows from $c_2(B) = 0$ and the second uses the binding IC constraint (sA1). From the non-lapsing constraint at time 2, we know that $u(c_2(A)) + \beta(\delta u(c_3(A,B)) + \delta^2 u(c_4(AB) \geq V_2^I$, giving a lower bound to the perceived payoff.

$$u(c_2(B)) + \hat{\beta}(\delta u(c_3(BB)) + \delta^2 u(c_4(BB)) \geq \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \frac{\hat{\beta}}{\beta} V_2^I.$$

Since V_2^I is the best possible outside option at time 2, in particular, it is greater than or equal to the utility provided by the contract $\hat{c}$, implying

$$\begin{aligned}&u(c_2(B)) + \hat{\beta}(\delta u(c_3(BB)) + \delta^2 u(c_4(BB))\\&\geq \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \frac{\hat{\beta}}{\beta}\left[u(\hat{c}_2) + \beta(\delta u(\hat{c}_3(B)) + \delta^2 u(\hat{c}_4(B))\right].\end{aligned}$$

Rearranging,

$$\begin{aligned}&u(c_2(B)) + \hat{\beta}(\delta u(c_3(BB)) + \delta^2 u(c_4(BB))\\&= \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \left(\frac{\hat{\beta}}{\beta} - 1\right) u(\hat{c}_2) + \left[u(\hat{c}_2) + \hat{\beta}(\delta u(\hat{c}_3(B)) + \delta^2 u(\hat{c}_4(B))\right]\\&\geq u(\hat{c}_2) + \hat{\beta}(\delta u(\hat{c}_3(B)) + \delta^2 u(\hat{c}_4(B))) = \hat{V}_2^I,\end{aligned}$$

where the inequality comes from $\hat{c}_2 \geq 0$ and $\hat{\beta} \geq \beta$ and the last line comes from the definition of $\hat{V}_2^I$. This shows that the perceived non-lapsing constraints hold.

We next verify that all the perceived choice constraints hold. Notice that

$$\begin{aligned}&u(c_3(A,B)) + \hat{\beta}\delta u(c_4(A,B)) = u(0) + \hat{\beta}\delta u(c_4(A,B))\\&= \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \frac{\hat{\beta}}{\beta}\left(u(c_3(A,A)) + \beta\delta u(c_4(A,A))\right)\\&= \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \left(\frac{\hat{\beta}}{\beta} - 1\right) u(c_3(A,A)) + u(c_3(A,A)) + \hat{\beta}\delta u(c_4(A,A))\\&\geq u(c_3(A,A)) + \hat{\beta}\delta u(c_4(A,A)),\end{aligned} \tag{sA3}$$

where the first line uses $u(c_3(A,B)) = 0$, the second line uses the self 3's binding IC constraint (sA2), the third line comes algebraic manipulations, and the last line uses $\hat{\beta} > \beta$

and $c_3(A,A) \geq 0$. Similarly,

$$
\begin{aligned}
&u(c_2(B)) + \hat{\beta}[\delta u(c_3(B,B)) + \delta^2 u(c_4(B,B))] \\
&= u(0) + \hat{\beta}[\delta u(c_3(B,B)) + \delta^2 u(c_4(B,B))] \\
&= (1-\frac{\hat{\beta}}{\beta})u(0) + \frac{\hat{\beta}}{\beta}\left[u(c_2(A)) + \beta[\delta u(c_3(A,B)) + \delta^2 u(c_4(A,B))]\right] \\
&\geq u(c_2(A)) + \hat{\beta}[\delta u(c_3(A,B)) + \delta^2 u(c_4(A,B))], \qquad \text{(sA4)}
\end{aligned}
$$

where the first line uses $c_2(B) = 0$, the second line uses the self 2's binding IC constraint (sA1), and the last line uses $\hat{\beta} > \beta$ and $c_2(A) \geq 0$. So the perceived choice constraints hold.

So far, we have shown that $c_2(B) = c_3(AB) = 0$ under the equilibrium contract. We also showed that we can disregard the perceived choice constraints and perceived non-lapsing constraints. Recall that c_t^E denotes the consumption on the equilibrium path at time t. Substituting the binding ICs, the non-lapsing constraints on the equilibrium path can be simplified to $u(c_t^E) + \delta u(c_{t+1}^E) + \cdots + \beta\delta^{t-4} u(c_4^E) \geq V_t^I$.

Therefore, the original program reduces to the auxiliary program:

$$
\max_{(c_1,c_2,c_3,c_4)} u(c_1) + \delta u(c_2) + \delta^2 u(c_3) + \beta\delta^3 u(c_4), \qquad \text{(sA5)}
$$

subject to

$$
\sum_{t=1}^{4} \frac{c_t}{R^{t-1}} = \sum_{t=1}^{4} \frac{w}{R^{t-1}}, \qquad \text{(sA6)}
$$

$$
u(c_t) + \delta u(c_{t+1}) + \cdots + \beta\delta^{4-t} u(c_4) \geq V_t^I, \forall 2 \leq t \leq 4. \qquad \text{(sA7)}
$$

□

Proof of Corollary 1. We can focus on the auxiliary program. Let $x(s_t) \equiv u(c(s_t))$ denote the agent's utility from the consumption he gets in state s_t. We study the dual program:

$$\max_{\{x(s_t)\}} \sum_{t=1}^{T} \sum_{s_t \in \mathbb{S}_t} p(s_t|s_1) \frac{w(s_t) - u^{-1}(x(s_t))}{R^{t-1}}, \tag{sA8}$$

subject to

$$\sum_{t=1}^{T-1} \sum_{s_t \in \mathbb{S}_t} \delta^{t-1} p(s_t|s_1) x(s_t) + \beta \sum_{s_T \in \mathbb{S}_T} \delta^{T-1} p(s_T|s_1) x(s_T) \geq \underline{u}. \tag{sA9}$$

This program corresponds to the maximization of a strictly concave function over a convex set, so that, by the Theorem of the Maximum, the solution is unique. Moreover, the consumption path is continuous in $\beta \in (0, 1]$. Finally, the program does not involve $\hat{\beta}$, so the consumption path is not a function of the consumer's naiveté.

Once we pin down the unique consumption path, the baseline options are either zero or determined by the binding IC constraints, which do not depend on $\hat{\beta}$ (see the proof of Lemma 2). So the equilibrium consumption vector is not a function of the consumer's naiveté.

□

Proof of Claim 1 (from the proof of Theorem 3). First, the time consistent agent's welfare is exactly given by the outside option,

$$\hat{W}_T^C(\underline{\mathbf{c}}) = E \sum_{t=1}^{T} \delta^{t-1} u(\underline{c}(s_t)).$$

The limit of $\hat{W}_T^C(\underline{\mathbf{c}})$ exists by the root test:

$$\limsup_{T \nearrow \infty} \sqrt[T]{\delta^{T-1} |u(\underline{c}(s_T))|} \leq \delta < 1.$$

Second, we show the limit of $\Pi_T^C(\underline{\mathbf{c}})$ exists using the Cauchy convergence criterion. Specifically, we claim that for sufficiently large T,

$$E\frac{w(s_T) - \underline{c}(s_T)}{R^{T-1}} \le \Pi_T^C(\underline{\mathbf{c}}) - \Pi_{T-1}^C(\underline{\mathbf{c}}) \le E\frac{w(s_T)}{R^{T-1}} + \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)}. \tag{sA10}$$

The claim follows from a revealed-preference argument. Suppose $(c_1', \cdots, c'(s_{T-1}))$ solves the program $\Pi_{T-1}^C(\underline{\mathbf{c}})$. Then $(c_1', \cdots, c'(s_{T-1}), \underline{c}(s_T))$ is in the feasible set of the program $\Pi_T^C(\underline{\mathbf{c}})$. By the revealed-preference argument, it immediately follows that

$$\Pi_T^C(\underline{\mathbf{c}}) \ge E\frac{w(s_T) - \underline{c}(s_T)}{R^{T-1}} + \Pi_{T-1}^C(\underline{\mathbf{c}}).$$

Suppose $(c_1^*, \cdots, c^*(s_T))$ solves the program $\Pi_T^C(\underline{\mathbf{c}})$. We show that

$$\left(c_1^* + \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)}, c^*(s_2), \cdots, c^*(s_{T-1})\right) \tag{sA11}$$

is in the feasible set of the program $\Pi_{T-1}^C(\underline{\mathbf{c}})$. To see that, note that from the Lagrange's Mean Value Theorem, it follows that

$$u\left(c_1^* + \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)}\right) - u(c_1^*) = u'(\xi)\frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)}, \tag{sA12}$$

where $\xi \in (c_1^*, c_1^* + \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)})$. For sufficiently large T, $c_1^* + \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)} < K$. So $u'(\xi) \ge u'(K)$. Going back to equation (sA12) leads to

$$u\left(c_1^* + \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)}\right) - u(c_1^*) \ge 2\delta^{T-1}\max\{|u(\cdot)|\}. \tag{sA13}$$

Then,

$$
\begin{aligned}
& u\left(c_1^* + \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)}\right) + E\sum_{t=2}^{T-1}\delta^{t-1}u(c^*(s_t)) \\
& \geq 2\delta^{T-1}\max\{|u(\cdot)|\} + E\sum_{t=1}^{T-1}\delta^{t-1}u(c^*(s_t)) \\
& \geq 2\delta^{T-1}\max\{|u(\cdot)|\} + E\sum_{t=1}^{T}\delta^{t-1}u(\underline{c}_t) - E\delta^{T-1}u(c^*(s_T)) \\
& = E\sum_{t=1}^{T-1}\delta^{t-1}u(\underline{c}_t) + \left(\delta^{T-1}\max\{|u(\cdot)|\} - E\delta^{T-1}u(c^*(s_T))\right) + \left(\delta^{T-1}\max\{|u(\cdot)|\} + E\delta^{T-1}u(\underline{c}_T)\right) \\
& \geq E\sum_{t=1}^{T-1}\delta^{t-1}u(\underline{c}_t),
\end{aligned}
$$

where the first inequality comes from (sA13), the second comes from noting that $(c_1^*, \cdots, c^*(s_T))$ solves program $\Pi_T^C(\underline{\mathbf{c}})$, the equality comes from algebraic manipulations, and the last step uses the boundedness of u. So we have shown that (sA11) is in the feasible set of $\Pi_{T-1}^C(\underline{c})$. A revealed-preference argument implies that

$$
\Pi_{T-1}^C(\underline{\mathbf{c}}) \geq E\sum_{t=1}^{T-1}\frac{w(s_t) - c^*(s_t)}{R^{t-1}} - \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)}.
$$

Recall that $\Pi_T^C(\underline{\mathbf{c}}) = E\sum_{t=1}^{T}\frac{w(s_t)-c^*(s_t)}{R^{t-1}}$. Substituting it back to the previous inequality, we obtain

$$
\Pi_{T-1}^C(\underline{\mathbf{c}}) \geq \Pi_T^C(\underline{\mathbf{c}}) - E\frac{w(s_T) - c^*(s_T)}{R^{T-1}} - \frac{2\delta^{T-1}\max\{|u(\cdot)|\}}{u'(K)},
$$

establishing the right-hand-side of (sA10) because of $c^*(s_T) \geq 0$. Since $E\sum_{t=1}^{T}\frac{w(s_T)-\underline{c}(s_T)}{R^{T-1}}$ exists and $\delta < 1$, for $\forall\epsilon$, we can find T_0 such that $\forall T_1, T_2 > T_0$, $|\Pi_{T_1}^C(\underline{\mathbf{c}}) - \Pi_{T_2}^C(\underline{\mathbf{c}})| < \epsilon$. This establishes that $\{\Pi_T^C(\underline{\mathbf{c}})\}$ satisfies the Cauchy convergence criterion, therefore the limit exists. □

Proof of Proposition 2. Since c^* maximizes the welfare function $W_T^H(c)$, it immediately follows that $W_T^{H,I} \leq W_T^*, \forall T$. Thus,

$$\limsup_{T \nearrow \infty} \frac{W_T^{H,I} - W_T^*}{T} \leq 0. \tag{sA14}$$

Denote $d_{t,T} = \frac{D_{T-1}}{D_{T-t}}, \forall t = 1, \cdots, T$. The objective function in the naive agent's auxiliary program becomes

$$\sum_{t=1}^{T} d_{t,T} u(c_t). \tag{sA15}$$

It follows that

$$\begin{aligned} W_T^{H,I} = \sum_{t=1}^{T} u(c_t^H) &= \sum_{t=1}^{T} [d_{t,T} u(c_t^H) + (1 - d_{t,T}) u(c_t^H)] \\ &\geq \sum_{t=1}^{T} [d_{t,T} u(c_t^*) + (1 - d_{t,T}) u(c_t^H)], \end{aligned}$$

where the first line comes from the definition and algebraic manipulations and the last step comes from the fact that c^H maximizes (sA15) and that c^* is feasible. Rearranging,

$$\begin{aligned} W_T^{H,I} &\geq \sum_{t=1}^{T} \left[u(c_t^*) + (1 - d_{t,T})[u(c_t^H) - u(c_t^*)] \right] \\ &= \sum_{t=1}^{T} \left[u(c_t^*) + \frac{k(t-1)}{1 + k(T-1)}[u(c_t^H) - u(c_t^*)] \right] \\ &= W_T^* + \sum_{t=1}^{T} \frac{k(t-1)}{1 + k(T-1)}[u(c_t^H) - u(c_t^*)], \end{aligned} \tag{sA16}$$

where the first line comes from algebraic manipulations, the second line uses the definition of $d_{t,T}$, and the last line comes from the definition of W_T^*.

We next show a series of lemmas to bound the second term. Let λ^H denote the Lagrangian multiplier from the zero-profit condition in the naive agent's program, and let λ^* denote the Lagrangian multiplier from the time-consistent agent's program. Note that the solution must be interior solution since $\lim_{c \searrow 0} u'(0) = +\infty$.

Lemma 10. *There exist* $\underline{\lambda}, \overline{\lambda} \in (0, +\infty)$ *such that*

$$\underline{\lambda} \leq \min(\lambda^H, \lambda^*) \leq \max(\lambda^H, \lambda^*) \leq \overline{\lambda}.$$

Proof. From the first-order-condition, we know that

$$\lambda^H = u'(c_1^H), \lambda^* = u'(c_1^*).$$

Note that the first period consumption must be between 0 and $\sum_{t=1}^{\infty} \frac{w}{R^{t-1}} = \frac{w}{1-R}$. The lemma follows immediately by letting $\overline{\lambda} = u'(0)$ and $\underline{\lambda} = u'\left(\frac{w}{1-R}\right)$. □

Lemma 11. *There exists a constant* $A > 0$ *such that* $|t(u(c_t^H) - u(c_t^*))| < A, \forall t, \forall T$.

Proof. From the first-order-condition, we know that

$$\frac{\lambda^H d_{t,T}}{R^{t-1}} = u'(c_t^H), \frac{\lambda^*}{R^{t-1}} = u'(c_t^*)$$

Denote $g(\cdot) = (u')^{-1}(\cdot)$. Inverting above equations to solve for c_t^H and c_t^*,

$$c_t^H = g\left(\frac{\lambda^H d_{t,T}}{R^{t-1}}\right), c_t^* = g\left(\frac{\lambda^*}{R^{t-1}}\right).$$

Note that $\frac{du(g(x))}{dx} = \frac{x}{u''(g(x))}$. Applying Lagrangian Mean Value Theorem, there exists η, where $\frac{\min(\lambda^*, \lambda^H d_{t,T})}{R^{t-1}} \leq \eta \leq \frac{\max(\lambda^*, \lambda^H d_{t,T})}{R^{t-1}}$, such that

$$|t(u(c_t^H) - u(c_t^*))| = t\left|u\left(g\left(\frac{\lambda^H d_{t,T}}{R^{t-1}}\right)\right) - u\left(g\left(\frac{\lambda^*}{R^{t-1}}\right)\right)\right| \tag{sA17}$$

$$= t\left|\frac{\eta}{u''(g(\eta))}\left(\frac{\lambda^H d_{t,T}}{R^{t-1}} - \frac{\lambda^*}{R^{t-1}}\right)\right|. \tag{sA18}$$

Using a change of variable $x = \frac{1}{R^{t-1}}$, then

$$x\underline{\lambda} d_{t,T} \leq x\min(\lambda^*, \lambda^H d_{t,T}) \leq \eta \leq x\max(\lambda^*, \lambda^H d_{t,T}) \leq x\overline{\lambda}.$$

So $x \geq \frac{\eta}{\overline{\lambda}}$. Note that $d_{t,T} = \frac{1+k(T-t)}{1+k(T-1)} \geq \frac{1}{1+k(t-1)}$. So,

$$x \leq \frac{\eta}{\underline{\lambda} d_{t,T}} \leq \frac{\eta(1+k(t-1))}{\underline{\lambda}} = \frac{\eta(1-k\frac{\log(x)}{\log R})}{\underline{\lambda}} \leq \frac{\eta(1-k\frac{\log(\eta)-\log(\overline{\lambda})}{\log R})}{\underline{\lambda}}. \tag{sA19}$$

We can rewrite (sA18) as

$$\begin{aligned}
|t(u(c_t^H) - u(c_t^*))| &\leq \left(-\frac{\log x}{\log R} + 1\right) \left|\frac{\eta}{u''(g(\eta))}\right| \frac{2\overline{\lambda}}{R^{t-1}} \\
&= \left(-\frac{\log x}{\log R} + 1\right) \left|\frac{\eta}{u''(g(\eta))}\right| 2\overline{\lambda} x \\
&\leq 2\frac{\overline{\lambda}}{\underline{\lambda}} \left(-\frac{\log \eta - \log(\overline{\lambda})}{\log R} + 1\right) \left(1 - k\frac{\log \eta - \log \overline{\lambda}}{\log R}\right) \frac{\eta^2}{|u''(g(\eta))|} \\
&\leq constant * \frac{(\log \eta)^2 \eta^2}{|u''(g(\eta))|},
\end{aligned}$$

where the first line uses $t = -\frac{\log x}{\log R} + 1$ and Lemma 10, the second line uses $x = \frac{1}{R^{t-1}}$, the third uses (sA19), and the last line collects the first-order terms. Let $\xi = g(\eta)$ and use Assumption 1, so there exists $A > 0$ such that $|t(u(c_t^H) - u(c_t^*))| < A$. □

Lemma 12. $\sum_{t=1}^{T} \frac{1}{t} \geq \log(T)$ *for any* $T \geq 1$.

Proof. Note that $\log(t+1) - \log(t) = \int_t^{t+1} \frac{1}{\theta} d\theta \leq \frac{1}{t}$. Sum over t from 1 to $(T-1)$ to obtain: $\log(T) \leq \sum_{t=1}^{T-1} \frac{1}{t} \leq \sum_{t=1}^{T} \frac{1}{t}$. □

Lemma 13. *There exists a constant* $A' > 0$ *such that*

$$\sum_{t=1}^{T} \frac{k(t-1)}{1+k(T-1)} |u(c_t^H) - u(c_t^*)| < A', \forall T$$

Proof. Using Lemma 11, it follows that

$$\begin{aligned}\sum_{t=1}^{T}\frac{k(t-1)}{1+k(T-1)}[u(c_t^H)-u(c_t^*)] &\leq \sum_{t=1}^{T}\frac{k(t-1)}{1+k(T-1)}\frac{A}{t}\\ &\leq \sum_{t=1}^{T}\frac{k}{1+k(T-1)}A+\sum_{t=1}^{T}\frac{-k}{1+k(T-1)}A\frac{1}{t}\\ &\leq \frac{kAT}{1+k(T-1)}+\frac{-k}{1+k(T-1)}A\log(T),\end{aligned}$$

where the first line comes from the lemma 11, the second line comes from algebraic manipulations, and the last line comes from $k \geq 0$ and lemma 12. Note that as $T \nearrow \infty$, the first term converges to A, and the second term converges to 0. So there exists a constant $A' > 0$ such that

$$\sum_{t=1}^{T}\frac{k(t-1)}{1+k(T-1)}[u(c_t^H)-u(c_t^*)] < A', \forall T.$$

□

Returning to (sA16), we have

$$\liminf_{T\nearrow\infty}\frac{W_T^{H,I}-W_T^*}{T} \geq -\liminf_{T\nearrow\infty}\frac{A'}{T} = 0. \tag{sA20}$$

Together with (sA14), it implies that $\lim_{T\nearrow\infty}\frac{W_T^{H,I}-W_T^*}{T}$ exists, and

$$\lim_{T\nearrow\infty}\frac{W_T^{H,I}-W_T^*}{T} = 0.$$

□

Proof of Proposition 4. It is easy to construct off-path beliefs that support the full-information allocation as an equilibrium. We need to show that no other allocation can be supported as an equilibrium. Suppose there exists a type $\hat{\beta}_0$ that does not pick the full information contract in equilibrium. There are two possibilities: (i) $\hat{\beta}_0$ is separated in equilibrium (i.e., no other type picks the same contract at $\hat{\beta}_0$), or (ii) $\hat{\beta}_0$ is pooled in equilibrium (i.e. there exists another type that picks the same contract as $\hat{\beta}_0$).

Consider case (i) first. Since $\hat{\beta}_0$ is the only type picking its contract, that contract must satisfy IC, PC, and zero profits. Recall that the full information contract is the unique contract that maximizes self 1's perceived utility subject to IC, PC, and zero profits. Consider a deviation in which type $\hat{\beta}_0$ offers the full-information contract in all histories except in period 1, where it offers a slightly lower consumption than with full information. Note that lowering c_1 does not affect IC and PC and, by taking c_1 arbitrarily close to the full-information consumption, we ensure that the consumer gets a strictly higher perceived utility while leaving strictly positive profits to the firm, contradicting the assumption that the original allocation was part of an equilibrium.

Next consider case (ii), so there are at least two types pooled at a contract different from the full information contract. If the firm breaks even on each consumer, then by the same argument as before, all consumers would strictly benefit from deviating to offering the full information contract (with a slightly lower c_1), which also gives strictly positive profits for the firm. If instead there is a cross subsidy between types, a type that is providing a positive profit can strictly benefit from deviating to the full information contract (with a slightly lower c_1). Moreover, by taking c_1 close enough to one in the full-information contract (which maximizes the perceived utility and leaves zero profits), we ensure that deviation is profitable. $\square$

Proof of Proposition 5. Suppose we have an equilibrium in which at least one naive type does not pick the full-information contract. Using the same argument as in Proposition 4, that type cannot be separated or pooled with other naive types only. Therefore, the only remaining case is one where at least one naive type pools with the sophisticated type.

But note that the contract that a sophisticated type would offer under full information offers a fixed consumption in each period (no alternative options), maximizing his perceived utility at time 1 under the zero profits constraint. Therefore, he must be cross subsidized in order to choose another contract (i.e., the firm must make strictly negative profits from serving him). But since the firm would not accept a contract that makes negative profits, this means that the firm must make strictly positive profits on some naive type that is pooling with the sophisticated type. But then this naive type would strictly profit from deviating to full information contract, which maximizes his perceived utility at time 1 subject to the zero profits constraint. □

Proof of Lemma 5. Without loss of generality, we use the following normalization $u(0) = 0$ in our analysis below. The proof follows by contradiction. Suppose there is an equilibrium in which two types, β_L and $\beta_H > \beta_L$, offer their full information contracts, $\mathcal{C}^L$ and $\mathcal{C}^H$. We show that these contracts cannot be part of an equilibrium when T is large enough since type β_H would deviate and pick $\mathcal{C}^L$, leaving the firm with negative profits.

Note that, by the binding IC constraint for type β_L, if type β_H picks $\mathcal{C}^L$, he ends up choosing B rather than A. In this case, the firm offering $\mathcal{C}^L$ makes negative profits. To see this, suppose instead that the firm makes a non-negative profit from this contract. But this would mean that the non-flexible contract that gives only the baseline consumption would also solve type β_L's program, which contradicts Corollary 1.

Type β_H's perceived utility from $\mathcal{C}^L$ equals:

$$\begin{aligned}
&u(c_1^L) + \beta_H \sum_{t=2}^{T} \delta^{t-1} u(c_t^L(B, \cdots, B)) \\
&= \frac{\beta_L - \beta_H}{\beta_L} u(c_1^L) + \frac{\beta_H}{\beta_L} u(c_1^L) + \beta_H \sum_{t=2}^{T} \delta^{t-1} u(c_t^L(B, \cdots, B)) \\
&= \frac{\beta_L - \beta_H}{\beta_L} u(c_1^L) + \frac{\beta_H}{\beta_L} \left[u(c_1^L) + \beta_L \sum_{t=2}^{T} \delta^{t-1} u(c_t^L(B, \cdots, B)) \right] \\
&= \frac{\beta_L - \beta_H}{\beta_L} u(c_1^L) + \frac{\beta_H}{\beta_L} \left[\sum_{t=1}^{T-1} \delta^{t-1} u(c_t^L(A, \cdots, A)) + \beta_L \delta^{T-1} u(c_T^L(A, \cdots, A)) \right] \\
&= \frac{\beta_L - \beta_H}{\beta_L} u(c_1^L) + \frac{\beta_H}{\beta_L} \sum_{t=1}^{T-1} \delta^{t-1} u(c_t^L(A, \cdots, A)) + \beta_H \delta^{T-1} u(c_T^L(A, \cdots, A)) \\
&= u(c_1^L) + \frac{\beta_H}{\beta_L} \sum_{t=2}^{T-1} \delta^{t-1} u(c_t^L(A, \cdots, A)) + \beta_H \delta^{T-1} u(c_T^L(A, \cdots, A))
\end{aligned}$$

$$> \sum_{t=1}^{T-1} \delta^{t-1} u(c_t^H(A, \cdots, A)) + \beta_H \delta^{T-1} u(c_T^H(A, \cdots, A))$$

$$= u(c_1^H) + \beta_H \left[\sum_{t=2}^{T} \delta^{t-1} u(c_t^H(B, \cdots, B)) \right].$$

where the second, third, fifth, and sixth lines follow from algebraic manipulations, the fourth line substitutes the binding (IC) for the low type, and the last line uses the binding (IC) for the high type. The strict inequality on the seventh line uses the following facts: $\beta_H > \beta_L$, $u_t(A, A..., A) \geq 0$ with strict inequality for at least one t, and, from Theorem 1, the welfare of time-inconsistent consumers converge to the welfare of time-consistent consumers

$$\lim_{T \nearrow \infty} \sum_{t=1}^{T-1} \delta^{t-1} u(c_t^L) + \beta_H \delta^{T-1} u(c_T^L) = \lim_{T \nearrow \infty} \sum_{t=1}^{T-1} \delta^{t-1} u(c_t^H) + \beta_H \delta^{T-1} u(c_T^H).$$

Therefore, for T sufficiently large, the β_H consumer would have an incentive to deviate and choose β_L consumer's full-information contract while taking the baseline option. $\square$

Proof of Lemma 6. We argue by contradiction. Fix an equilibrium in which β_H does not get his full-information contract. First, suppose that firms make non-negative profits from β_H. Suppose β_H deviates and offers his full-information contract. By a single-crossing argument, if type β_L got β_H's full-information contract, he would always choose option A, so the firm would break even on both types under type β_H's full-information contract. Since the full-information contract maximizes β_H's perceived utility among those that make zero profits, β_H has an incentive to deviate to it.

Suppose, instead, that the firm makes strictly negative profits on type β_H. Then firm optimality requires that both types pool on the same contract $\mathcal{C}$ and the firm makes strictly positive profits on type β_L. To generate different profits, these two types must be getting different allocations on the equilibrium path.

We will construct a deviation contract $\mathcal{C}(\epsilon)$ such that whenever the β_H consumer weakly benefits from the deviation, the β_L consumer strictly benefits from the deviation. By D1 criteria, we should assign zero weight to the type β_H and all the weight to the type β_L consumer. Given that firms make positive profits from the β_L consumer's equilibrium contract, firms would charge a price such that the β_L consumer are better off with $\mathcal{C}(\epsilon)$ than the contract $\mathcal{C}$, a contradiction. Since both types have the same naiveté parameter $\hat{\beta}$, they both believe they will choose the same options. Let $(\hat{c}_2, \cdots, \hat{c}_T)$ denote their perceived consumption stream. Construct a perturbation of the equilibrium contract, $\mathcal{C}(\epsilon)$, by decreasing the last-period perceived consumption by ϵ and adjusting the other options so that (IC) and (PC) hold for both types. Upon observing contract $\mathcal{C}(\epsilon)$, the firm must assign full weight to type β_L. This is because whenever β_H benefits from deviating to this contract (i.e., when the firm's price is lower than $-\beta_H \delta^{T-1} \frac{u'(c_T)}{u'(c_1)} \epsilon$, β_L also benefits from this deviation (i.e., when the firm's price is lower than $-\beta_L \delta^{T-1} \frac{u'(c_T)}{u'(c_1)} \epsilon$. Therefore, this candidate equilibrium does not satisfy D1. □

Proof of Lemma 7. To show that the proposed equilibrium survives D1, we show that if β_L can benefit from a deviation to $\mathcal{C}'$, then β_H strictly benefits from the deviation as well.

Recall that they have the same perceived time-consistency parameter $\hat{\beta}$, so their perceived consumption from the contract $\mathcal{C}'$ are the same, denoted as $(c'_1, c'_2, \cdots, c'_T)$. Suppose β_L can benefit from the deviation:

$$u(\bar{c}_1) + \beta_L \sum_{t=2}^{T} \delta^{t-1} u(\bar{c}_t(B, \cdots, B)) < u(c'_1) + \beta_L \sum_{t=2}^{T} \delta^{t-1} u(c'_t). \tag{sA21}$$

By a single-crossing argument, since $\bar{c}$ solves β_L's program (and therefore his IC must bind), type β_H's IC cannot hold:

$$u(c_1^H) + \beta_H \sum_{t=2}^{T} \delta^{t-1} u(c_t^H(B, \cdots, B)) < u(c'_1) + \beta_H \sum_{t=2}^{T} \delta^{t-1} u(c'_t), \tag{sA22}$$

so β_H also benefits from this deviation. According to D1, we must assign zero weight on β_L and full weight on β_H. Because β_H gets his full-information contract in any equilibrium satisfying D1 (Lemma 6), we must have

$$u(c_1^H)+\beta_H\sum_{t=2}^{T}\delta^{t-1}u(c_t^H(B,\cdots,B))\geq u(c_1')+\beta_H\sum_{t=2}^{T}\delta^{t-1}u(c_t'), \tag{sA23}$$

a contradiction to (sA22). So we have

$$u(\bar{c}_1)+\beta_L\sum_{t=2}^{T}\delta^{t-1}u(\bar{c}_t(B,\cdots,B))\geq u(c_1')+\beta_L\sum_{t=2}^{T}\delta^{t-1}u(c_t'), \tag{sA24}$$

showing that β_L does not have a profitable deviation and the proposed equilibrium survives D1.

Next, we show that in any equilibrium satisfying D1, the consumption path corresponds to the least costly separating allocation. Suppose there exists another equilibrium that survives D1. As we showed above, β_H gets his full-information contract. Let $\mathcal{C}'$ denote β_L's equilibrium contract, and the contract is different from the leastly costly separation allocation (A1). Suppose β_L deviates and offers a contract that coincides with the solution to (A1) except that it reduces consumption in the first period by a small $\epsilon > 0$. By the IC constraint, the β_H consumer is strictly worse off by choosing this new contract instead of his full-information contract. By D1, firms must assign full weight to β_L. By choosing ϵ small enough, β_L strictly benefits from the deviation. $\square$

Proof of Proposition 6. From the previous lemmas, the equilibrium is given by the least-costly separation. The equilibrium-path consumption for the low type solves the following program:

$$\max u(c_1)+l(c_1),$$

subject to

$$u(c_1)+\frac{\beta_H}{\beta_L}l(c_1)=V^H, \tag{sA25}$$

where $l(\cdot)$ is defined as

$$l(c_1) = \max_{(c_2, c_3, \cdots, c_T)} \sum_{t=2}^{T-1} \delta^{t-1} u(c_t) + \beta \delta^{T-1} u(c_T).$$

subject to

$$\sum_{t=2}^{T} \frac{c_t}{R^{t-1}} = \sum_{t=1}^{T} \frac{w}{R^{t-1}} - c_1.$$

Note that by substituting (sA25) to the objective function, maximizing $u(c_1) + l(c_1)$ is equivalent to maximizing c_1. If β_L's full-information contract cannot be sustained in a equilibrium (as must be the case if T is large), it means that

$$u(c_1^L) + \frac{\beta_H}{\beta_L} l(c_1^L) > V^H. \tag{sA26}$$

Evaluating c_1 at $\sum_{t=1}^{T} \frac{w}{R^{t-1}}$ implies that the

$$u\left(\sum_{t=1}^{T} \frac{w}{R^{t-1}}\right) + \frac{\beta_H}{\beta_L} \delta l \left(\sum_{t=1}^{T} \frac{w}{R^{t-1}}\right) = u\left(\sum_{t=1}^{T} \frac{w}{R^{t-1}}\right) < V^H. \tag{sA27}$$

By the intermediate value theorem, it follows that the maximal root of (sA25) must be greater than the first period consumption in the full-information contract: $\bar{c}_1 > c_1^L$. This completes the first part of the proposition.

We next show that the welfare loss must be bounded below away from 0. We argue by contradiction. Suppose there exists a subsequence $\{T_n : n \in \mathbb{N}\}$ such that $\lim_{n \nearrow \infty}(W_{T_n}^C - W_{T_n}^L) = 0$. To be clear that our variables now depends on T_n, we write variables as a function of T_n. Note that

$$\begin{aligned}
\lim_{n\nearrow\infty}(W^C_{T_n} - W^L_{T_n}) &= \lim_{n\nearrow\infty}(V^H(T_n) - W^L_{T_n}) \\
&= \lim_{n\nearrow\infty}\big(V^H(T_n) - u(\overline{c}_1(T_n)) - l(\overline{c}_1(T_n))\big) \\
&= \lim_{n\nearrow\infty}\left(V^H(T_n) - u(\overline{c}_1(T_n)) - (V^H(T_n) - u(\overline{c}_1(T_n)))\frac{\beta_L}{\beta_H}\right) \\
&= \lim_{n\nearrow\infty}\big(V^H(T_n) - u(\overline{c}_1(T_n))\big)\left(1 - \frac{\beta_L}{\beta_H}\right),
\end{aligned}$$

where the first equality comes from the vanishing inefficiency result for the β_H consumer, the second equality comes from the definition of $l(c_1)$ and $(1-\beta)\delta^{T-1}u(\overline{c}(s_T)) \to 0$, the third equality comes from (sA25), and the fourth equality comes from algebraic manipulations.

It follows that $\lim_{n\nearrow\infty}\big(V^H(T_n) - u(\overline{c}_1(T_n))\big) = 0$. From the vanishing inefficiency result for the β_H consumer, it implies that

$$\lim_{n\nearrow\infty}\big(W^L_{T_n} - u(\overline{c}_1(T_n))\big) = \lim_{n\nearrow\infty}\big(W^C_{T_n} - u(\overline{c}_1(T_n))\big) = \lim_{n\nearrow\infty}\big(V^H(T_n) - u(\overline{c}_1(T_n))\big) = 0. \tag{sA28}$$

Note that $W^L_{T_n} = \sum_{t=1}^{T_n}\delta^{t-1}u(\overline{c}_t(T_n))$. We obtain

$$\lim_{n\nearrow\infty}\sum_{t=2}^{T_n}\delta^{t-1}u(\overline{c}_t(T_n)) = 0.$$

Recall that we normalize $u(0) = 0$, so $u(c) \geq 0, \forall c \geq 0$. We must have $\lim_{n\nearrow\infty}\overline{c}_t(T_n) = 0, \forall t$. By the zero-profits condition, the β_L consumer consumes everything in the first period in the limit: $\lim_{n\nearrow\infty}\overline{c}_1(T_n) = \sum_{t=1}^{\infty}\frac{w}{R^{t-1}}$. This consumption stream cannot achieve the first-best welfare (i.e., W^C_T), as shifting some consumption to future periods can strictly improves welfare since $\lim_{c\searrow 0}u'(c) = +\infty$. Specifically, fix a small $\epsilon_0 > 0$, it is straightforward to show that

$$u\left(\sum_{t=1}^{\infty}\frac{w}{R^{t-1}}\right) < u\left(\sum_{t=1}^{\infty}\frac{w}{R^{t-1}} - \epsilon_0\right) + \delta u(R\epsilon_0) \leq \lim_{T\nearrow\infty} W_T^C.$$

This is a contradiction to (sA28) that $\lim_{n\nearrow\infty}\left(W_{T_n}^C - u(\overline{c}_1(T_n))\right) = 0$.

So the welfare loss does not vanish as the contracting horizon grows: $\liminf_{T\nearrow\infty}(W_T^C - W_T^L) > 0$. □

Proof of Lemma 8. To prove the lemma, we argue by contradiction. There exist two option history paths of consumption stream starting with (s_t, h^t) that have different expected present discounted values. Without loss of generality, assume that one path, denoted as $\hat{\mathbf{c}}$, has a higher expected present value than the other path, denoted as $\tilde{\mathbf{c}}$.

We note that since $\mathbf{c}$ is the equilibrium consumption vector, it must satisfy the *no additional contracting constraints*. Given that the present value of $\hat{\mathbf{c}}$ is higher than the present value of $\tilde{\mathbf{c}}$. There are two possibilities, either $\tilde{\mathbf{c}}$ starts with the baseline option or $\tilde{\mathbf{c}}$ starts with the alternative option. In either case, we show that the no additional contracting constraints would be violated. First, suppose $\tilde{\mathbf{c}}$ starts with the baseline option. In this case, the baseline option would not be the consumer's perceived consumption, because the consumer perceives that he has an incentive to recontract with another firm, who can give the consumer slightly higher consumption in the baseline option. Specifically, consider another contract $\mathbf{c}'$, which has the same term as $\mathbf{c}$ except that we increase ϵ in the consumption in the baseline option of $\tilde{\mathbf{c}}$. Similarly, if $\tilde{\mathbf{c}}$ starts with the alternative option, the consumer can recontract with another firm, who gives him slightly higher consumption in the alternative option. □

Proof of Proposition 7. Note first that uncertainty over states plays no role in the program with non-exclusive contracts. Starting from any allocation in which consumption within a period is random, the agent increase his perceived utility by signing a contract with another firm to smooth consumption in that period. So we can without loss of generality substitute each period's income by its expected value. Using Lemma 8, we find that the program with non-exclusive contracts becomes identical to the consumption-savings problem. $\square$

Proof of Proposition 8. Consider a problem with a sophisticated consumer who has the commitment power and whose time-consistency parameter is $\frac{\beta}{\hat{\beta}}$. Without loss of generality, we assume that there is no uncertainty. Recall that a_1 is the PDV of income. The sophisticate's program is

$$c^S = \max_{\{c(\cdot)\}} u(c_1) + \frac{\beta}{\hat{\beta}} \sum_{t=2}^{T} \delta^{t-1} u(c_t),$$

subject to

$$\sum_{t=1}^{T} \frac{c_t}{R^{t-1}} = a_1.$$

We claim that the welfare in the above program is an upper bound of the welfare in the consumption-savings problem for the naive consumer. Let $\bar{c}_1$ denote the first-period consumption in the consumption-savings problem. We will show that the naive agent consumes strictly more than the sophisticated agent: $\bar{c}_1 \geq c_1^S$.

The proof proceeds through four lemmas. The first one adapts arguments from Harris and Laibson (2001).

Lemma 14. *The perceived consumption functions* $(\hat{c}_2(\cdot), \cdots, \hat{c}_T(\cdot))$ *satisfy:*

$$(\delta R)^{t-1} u'(\hat{c}_t) = (\delta R)^t u'(\hat{c}_{t+1}) \left[1 - \hat{c}'_{t+1}(a_{t+1}) + \hat{\beta} \hat{c}'_{t+1}(a_{t+1}))\right], \forall 1 < t < T,$$

where $a_{t+1} = R(a_t - \hat{c}_t(a_t))$.

Proof. The proof follows by induction, starting at period $T-1$. The last period consumption is $c_T(a_T) = a_T$. Consumption in the penultimate period is:

$$\hat{c}_{T-1}(a_{T-1}) = \arg\max_{\tilde{c}} \{u(\tilde{c}) + \hat{\beta} \delta u(\hat{c}_T(a_T)) \text{ subject to } a_T = R(a_{T-1} - \tilde{c})\}.$$

Since $\lim_{c \searrow 0} u'(c) = +\infty$, the unique solution must be interior and satisfy the FOC:

$$u'(\hat{c}_{T-1}) = \hat{\beta}\delta R u'(\hat{c}_T).$$

Since $\hat{c}'_T(a_T) = 1$, the statement in the lemma holds for $t = T - 1$.

Moving to the induction step, suppose the statement holds for $\tau < T$ and recall that:

$$\hat{c}_\tau(a_\tau) = \arg\max_{\tilde{c}}\{u(\tilde{c}) + \hat{\beta}\sum_{t=\tau+1}^{T}\delta^{t-\tau}u(\hat{c}_t(a_t)) \text{ subject to (B2), (B3), and (B4)}\}.$$

The unique solution must be interior and satisfy the FOC:

$$u'(\hat{c}_\tau) + \hat{\beta}\sum_{t=\tau+1}^{T}\delta^{t-\tau}u'(\hat{c}_t(a_t))\frac{\partial \hat{c}_t(a_t)}{\partial \hat{c}_\tau} = 0.$$

Substitute

$$\begin{aligned}\frac{\partial \hat{c}_t(a_t)}{\partial \hat{c}_\tau} &= \hat{c}'_t(a_t)\frac{\partial a_t}{\partial \hat{c}_\tau} = \hat{c}'_t(a_t)\frac{\partial a_t}{\partial a_{t-1}}\cdots\frac{\partial a_{\tau+1}}{\partial \hat{c}_\tau} \\ &= -R^{t-\tau}\hat{c}'_t(a_t)(1-\hat{c}'_{t-1}(a_{t-1}))\cdots(1-\hat{c}'_{\tau+1}(a_{\tau+1})),\end{aligned}$$

to rewrite the FOC as:

$$u'(\hat{c}_\tau) = \hat{\beta}\sum_{t=\tau+1}^{T}(\delta R)^{t-\tau}u'(\hat{c}_t(a_t))\hat{c}'_t(a_t)(1-\hat{c}'_{t-1}(a_{t-1}))\cdots(1-\hat{c}'_{\tau+1}(a_{\tau+1})).$$

The FOC at $\tau + 1$ is:

$$u'(\hat{c}_{\tau+1}) = \hat{\beta}\sum_{t=\tau+2}^{T}(\delta R)^{t-\tau-1}u'(\hat{c}_t(a_t))\hat{c}'_t(a_t)(1-\hat{c}'_{t-1}(a_{t-1}))\cdots(1-\hat{c}'_{\tau+2}(a_{\tau+2})).$$

Multiply both sides by $\delta R(1-\hat{c}'_{\tau+1}(a_{\tau+1}))$ and substitute back in the equation for $u'(\hat{c}_\tau)$

to verify that the statement in the lemma also holds for $t = \tau$:

$$u'(\hat{c}_\tau) = \hat{\beta}(\delta R)u'(\hat{c}_{\tau+1})\hat{c}'_{\tau+1}(a_{\tau+1}) + (\delta R)u'(\hat{c}_{\tau+1})[1 - \hat{c}'_{\tau+1}(a_{\tau+1})].$$

$\square$

Lemma 15. *The first-period consumption* $\overline{c}_1$ *satisfies:*

$$u'(\overline{c}_1) = \delta R u'(\hat{c}_2)\frac{\beta}{\hat{\beta}}\left[1 - \hat{c}'_2(a_2) + \hat{\beta}\hat{c}'_2(a_2)\right],$$

where $a_2 = R(a_1 - \overline{c}_1)$.

Proof. Similar to the proof of last lemma,

$$u'(\overline{c}_1) = \beta \sum_{t=2}^{T} (\delta R)^{t-1} u'(\hat{c}_t(a_t))\hat{c}'_t(a_t)(1 - \hat{c}'_{t-1}(a_{t-1})) \cdots (1 - \hat{c}'_2(a_2)).$$

The FOC at $t = 2$ gives to

$$u'(\hat{c}_2) = \hat{\beta} \sum_{t=3}^{T} (\delta R)^{t-2} u'(\hat{c}_t(a_t))\hat{c}'_t(a_t)(1 - \hat{c}'_{t-1}(a_{t-1})) \cdots (1 - \hat{c}'_3(a_3)).$$

Multiply by $\delta R(1 - \hat{c}'_2(a_2))$ on both sides, and substitute back to the equation for $u'(\overline{c}_1)$, then we obtain

$$u'(\overline{c}_1) = \beta \delta R u'(\hat{c}_2)\hat{c}'_2(a_2) + \frac{\beta}{\hat{\beta}}\delta R u'(\hat{c}_2)(1 - \hat{c}'_2(a_2)).$$

$\square$

Lemma 16. $(\delta R)^{t-1}u'(\hat{c}_t(a_t)) \geq \frac{\hat{\beta}}{\beta}u'(\overline{c}_1)$ *for all* $t > 1$.

Proof. It is straightforward to see that $\hat{c}'_t(a_t) \in [0, 1], \forall t > 1$. It follows that $1 - \hat{c}'_t(a_t) +$

$\hat{\beta}\hat{c}'_t(a_t) \in [\hat{\beta}, 1]$. From the previous two lemmas, we have

$$(\delta R)^{t-1} u'(\hat{c}_t(a_t)) \leq (\delta R)^t u'(\hat{c}_{t+1}(a_{t+1})), \forall 1 < t < T.$$
$$u'(\overline{c}_1) \leq \frac{\beta}{\hat{\beta}} \delta R u'(\hat{c}_2).$$

It immediately follows that $(\delta R)^{t-1} u'(\hat{c}_t(a_t)) \geq \frac{\hat{\beta}}{\beta} u'(\overline{c}_1), \forall t > 1$. □

Lemma 17. *The naive agent consumes weakly more than the sophisticated agent in the first period:* $\overline{c}_1 \geq c_1^S$.

Proof. We argue by contradiction. Suppose $\overline{c}_1 < c_1^S$. Then $u'(\overline{c}_1) > u'(c_1^S)$. From the FOC of the sophisticate's problem, we know that

$$u'(c_1^S) = \frac{\beta}{\hat{\beta}} (\delta R)^{t-1} u'(c_t^S).$$

Together with the previous lemma, we obtain

$$\frac{\beta}{\hat{\beta}} (\delta R)^{t-1} u'(\hat{c}_t(a_t)) > \frac{\beta}{\hat{\beta}} (\delta R)^{t-1} u'(c_t^S).$$

Thus, $\hat{c}_t(a_t) < c_t^S$, which is a contradiction because of the zero-profits condition

$$\overline{c}_1 + \sum_{t=2}^{T} \frac{\hat{c}_t(a_t)}{R^{t-1}} = a_1 = \sum_{t=1}^{T} \frac{c_t^S}{R^{t-1}}.$$

□

We are now ready to show the proposition. Let $\overline{\mathbf{c}} = (\overline{c}_1, \cdots, \overline{c}_T)$ denote the naive consumer's equilibrium allocation. Since $\overline{\mathbf{c}}$ also satisfies the zero-profit condition, a revealed-preference argument applied to the sophisticate's program gives:

$$u(c_1^S) + \frac{\beta}{\hat{\beta}} \sum_{t=2}^{T} \delta^{t-1} u(c_t^S) \geq u(\overline{c}_1) + \frac{\beta}{\hat{\beta}} \sum_{t=2}^{T} \delta^{t-1} u(\overline{c}_t). \tag{sA29}$$

The naive consumer's welfare is:

$$
\begin{aligned}
\sum_{t=1}^{T}\delta^{t-1}u(\bar{c}_t) &= \left(1-\frac{\hat{\beta}}{\beta}\right)u(\bar{c}_1)+\frac{\hat{\beta}}{\beta}u(\bar{c}_1)+\sum_{t=2}^{T}\delta^{t-1}u(\bar{c}_t)\\
&= \left(1-\frac{\hat{\beta}}{\beta}\right)u(\bar{c}_1)+\frac{\hat{\beta}}{\beta}\left[u(\bar{c}_1)+\frac{\beta}{\hat{\beta}}\sum_{t=2}^{T}\delta^{t-1}u(\bar{c}_t)\right]\\
&\leq \left(1-\frac{\hat{\beta}}{\beta}\right)u(\bar{c}_1)+\frac{\hat{\beta}}{\beta}\left[u(c_1^S)+\frac{\beta}{\hat{\beta}}\sum_{t=2}^{T}\delta^{t-1}u(c_t^S)\right]\\
&= \left(1-\frac{\hat{\beta}}{\beta}\right)u(\bar{c}_1)+\frac{\hat{\beta}}{\beta}u(c_1^S)+\sum_{t=2}^{T}\delta^{t-1}u(c_t^S)\\
&\leq u(c_1^S)+\sum_{t=2}^{T}\delta^{t-1}u(c_t^S),
\end{aligned}
$$

where equalities come from algebraic manipulation, the first inequality comes from (sA29), and the last inequality comes from the previous lemma $\bar{c}_1 \geq c_1^S$. So the naive consumer's welfare is bounded above by the sophisticate's welfare, which does not converge to the time-consistent consumer's welfare (Proposition 9), establishing the result.

□

Proof of Proposition 9. We argue by contradiction. Suppose instead that

$$\liminf_{T\nearrow+\infty}\left(W_T^C - W_T^S\right)=0,$$

so that there exists a subsequence $\{T_n : n\in\mathbb{N}\}$ with $\lim_{n\nearrow+\infty}\left(W_{T_n}^C - W_{T_n}^S\right)=0$.

Let $c^S = (c_1^S(T),\cdots,c^S(s_T,T))$ denote the equilibrium consumption for the sophisticated agent in the (truncation of the) model with T periods. Let $c^C = (c_1^C(T),\cdots,c^C(s_T,T))$ denote the equilibrium consumption for the time-consistent agent in the (truncation of the) model with T periods. Passing to subsequences, we can assume both limits $\lim_{n\nearrow\infty} c_1^S(T_n)$ and $\lim_{n\nearrow\infty} c_1^C(T_n)$ exist.[33]

[33] That is, there exists a subsequence $\{T_{n_m}\}$ of $\{T_n\}$ such that the limit of $c_1^S(T_{n_m})$ exists. Similarly, consider the sequence $\{c_1^C(T_{n_{m_k}})\}$. Again, pick a subsequence $\{T_{n_{m_o}}\}$ of $\{T_{n_m}\}$ such that the limit $c_1^C(T_{n_{m_k}})$ exists. For notational simplicity, and with no loss of generality, we can replace the original sequence $\{T_n\}$

We first claim that the sophisticate consumes strictly more in the first period than the time-consistent consumer in the limit: $\lim_{n\nearrow\infty} c_1^S(T_n) > \lim_{n\nearrow\infty} c_1^C(T_n)$. Suppose instead that $\lim_{n\nearrow\infty} c_1^S(T_n) \le \lim_{n\nearrow\infty} c_1^C(T_n)$. The FOCs of the time-consistent consumer's program give:

$$u'(c_1^C(T_n)) = (\delta R)^{t-1} E u'(c^C(s_t, T_n)), \quad u'(c_1^S(T_n)) = \beta(\delta R)^{t-1} E u'(c^S(s_t, T_n)).$$

We claim that $\liminf_{n\nearrow\infty}(c^S(s_t, T_n) - c^C(s_t, T_n)) < 0, \forall t > 1$. Otherwise, there exists $t > 1$ and $\liminf_{n\nearrow\infty}(c^S(s_t, T_n) - c^C(s_t, T_n)) \ge 0$. Passing to subsequences, we can assume that $\lim_{n\nearrow\infty} c_t^S(s_t, T_n)$ and $\lim_{n\nearrow\infty} c_t^C(s_t, T_n)$ exist.

It follows that

$$\begin{aligned}
\lim_{n\nearrow\infty} (\delta R)^{t-1} E u'(c^S(s_t, T_n)) &> \lim_{n\nearrow\infty} \beta(\delta R)^{t-1} E u'(c^S(s_t, T_n)) \\
&= \lim_{n\nearrow\infty} u'(c_1^S(T_n)) \\
&\ge \lim_{n\nearrow\infty} u'(c_1^C(T_n)) \\
&= \lim_{n\nearrow\infty} (\delta R)^{t-1} E u'(c^C(s_t, T_n)) \\
&\ge \lim_{n\nearrow\infty} (\delta R)^{t-1} E u'(c^S(s_t, T_n)),
\end{aligned}$$

where the first inequality is strict because of $\beta < 1$, the second equation comes from the sophisticate's FOC, the third comes from $\lim_{n\nearrow\infty} c_1^S(T_n) \le \lim_{n\nearrow\infty} c_1^C(T_n)$, the fourth comes from the time-consistent consumer's FOC, the last comes from $\liminf_{n\nearrow\infty}(c^S(s_t, T_n) - c^C(s_t, T_n)) \ge 0$. This is a contradiction. So $\liminf_{n\nearrow\infty}(c^S(s_t, T_n) - c^C(s_t, T_n)) < 0, \forall t > 1$. But then it violates the zero-profit condition since

$$0 = \liminf_{n\nearrow\infty} \sum_{t=1}^{T_n} \frac{E(c^S(s_t, T_n) - c^C(s_t, T_n))}{R^{t-1}} < 0.$$

What we have shown now is that in the first period the sophisticate consume strictly more

with this last subsequence $\{T_{n_{m_o}}\}$.

than the time consistent consumer in the limit: $\lim_{n \nearrow \infty} c_1^S(T_n) > \lim_{n \nearrow \infty} c_1^C(T_n)$.

Define $l_T(c_1)$ as

$$l_T(c_1) = \max_{c(s_2),\cdots,c(s_T)} \sum_{t=2}^{T} \delta^{t-1} u(c(s_t)),$$

subject to $E \sum_{t=2}^{T} \frac{c(s_t)}{R^{t-1}} = E \sum_{t=1}^{T} \frac{w(s_t)}{R^{t-1}} - c_1$. We claim that $l_T''(c_1) < 0$. Let λ_l denote the Lagrangian on the zero-profit constraint.

$$l_T(c_1) = \sum_{t=2}^{T} \delta^{t-1} u(c(s_t)) + \lambda_l \left(E \sum_{t=1}^{T} \frac{w(s_t)}{R^{t-1}} - c_1 - E \sum_{t=2}^{T} \frac{c(s_t)}{R^{t-1}} \right).$$

Taking derivative with respect to c_1: $l_T'(c_1) = -\lambda_l$. Then, $l_T''(c_1) = -\lambda_l'$.

Taking derivative on the both sides of FOC, $\delta^{t-1} u'(c_t) = \frac{\lambda_l}{R^{t-1}}$, with respect to c_1:

$$\frac{\partial c_t}{\partial c_1} = \frac{\lambda_l'}{(\delta R)^{t-1} u''(c_t)}.$$

Taking derivative with respect to c_1 on the zero-profit condition:

$$-1 = \sum_{t=2}^{T} \frac{\frac{\partial c_t}{\partial c_1}}{R^{t-1}} = \sum_{t=2}^{T} \frac{\lambda_l'}{(\delta R)^{t-1} u''(c_t)}.$$

Thus, $\lambda_l' > 0$ because $u'' < 0$. So $l_T''(c_1) = -\lambda_l' < 0$.

It implies that $u(c_1) + l_T(c_1)$ is a concave function of c_1 for any T. Since u is bounded and $\delta < 1$, we can use dominated convergence theorem. Taking limit of T to infinity, $\limsup_{T \nearrow \infty} u''(c_1) + l_T''(c_1) \leq \limsup_{T \nearrow \infty} u''(c_1) < 0$, since we assume strict concavity of u. So $\lim_{T \nearrow \infty} [u(c_1) + l_T(c_1)]$ is a strict concave function of c_1. Together with our first claim that $\lim_{n \nearrow \infty} (c_1^S(T_n) - c_1^C(T_n)) > 0$ and the fact that $c_1^C(T)$ maximizes $u(c_1) + l_T(c_1)$, it follows that

$$\lim_{n \nearrow +\infty} \left(u(c_1^C(T_n)) + l_T(c_1^C(T_n)) - u(c_1^S(T_n)) - l_T(c_1^S(T_n)) \right) > 0,$$

i.e. $\lim_{n \nearrow +\infty} \left(W_{T_n}^C - W_{T_n}^S \right) > 0$, a contradiction. So the welfare loss for sophisticated agents is bounded below away from 0. $\square$

Proof of Lemma 9. We first show that the (IC) constraints for self-2 must be binding for all $m_2 \in supp(\sigma_2)$ and $m_2' \in supp(\hat{\sigma}_2)$. We note that the (IC) must be binding for at least one m_2', because otherwise we can increase consumption on the perceived path and increase the self 1's payoff. Now suppose there exists $m_2 \in supp(\sigma_2)$ and $m_2'' \in supp(\hat{\sigma}_2)$ such that the corresponding (IC) is slack. In this case, we show that from self 1's perspective, the perceived path m_2'' gives a higher payoff than the perceived path m_2' (using a coefficient of 1). To see that, notice that from (PC) constraint, the perceived self-2 is indifferent between m_2' and m_2'' (using a coefficient of $\hat{\beta}$), but self-2 strictly prefers m_2' over m_2'' (using a coefficient of β). By the single crossing property, it implies that m_2'' gives a strictly higher payoff than m_2' in calculating self 1's perceived payoff (using a coefficient of 1). Then, replacing terms in options m_2' with terms in options m_2'' would not affect any constraints, but it would increase self 1's perceived payoff, a contradiction to the optimality of the original contract.

Next we show that $c_2(s_2, \hat{m}_2) = 0, \forall \hat{m}_2 \in supp(\hat{\sigma}_2)$. Otherwise, consider a perturbation in which lowers $u(c_2(s_2, \hat{m}_2))$ by $\beta\epsilon$ and increases $u(c_T(s_T, \hat{m}_2, \hat{\sigma}_3, \cdots, \hat{\sigma}_{T-1}))$ by ϵ. This perturbation preserves the IC constraints and maintains all other constraints, but increases self 1's perceived payoff.

Substituting the binding IC constraint into the objective function, we obtain (up to a constant):

$$\max_{\{c(s_t,h^t)\}} u(c(s_1)) + \delta E u(c(s_2, \sigma_2)) + \beta E\left[\sum_{t=3}^{T} \delta^{t-1} u(c(s_t, \sigma_2, \hat{\sigma}_3, ..., \hat{\sigma}_{T-1}))\right].$$

Repeating the same analysis, we have a new program (up to a constant):

$$\max_{\{c(s_t,h^t)\}} E\delta^{t-1} u(c(s_t, \sigma_2, \cdots, \sigma_{T-1})) + \beta E\left[\delta^{t-1} u(c(s_T, \sigma_2, \sigma_3, ..., \sigma_{T-1}))\right].$$

subject to the zero-profit condition.

Our final step is showing that the equilibrium path $\sigma_2, \sigma_3, ..., \sigma_{T-1}$ involves only one option. This is because of Jensen's inequality and the strict concavity of $u(\cdot)$. If there are multiple options in the σ_τ, then merging those options can strictly increase self 1's perceived payoff. This completes the proof. □

Proof of Lemma 2 for General Income Distributions and Arbitrary T

This appendix establishes the equivalence between the naive agent's program and the auxiliary program for general income distributions and arbitrary T. As in the text, we consider the one-sided commitment case. With two-sided commitment, one can ignore the non-lapsing constraints in the proof below.

Recall that the naive agent's program is

$$\max_{c(s_t,h^t)} u(c(s_1)) + \beta E\left[\sum_{t=2}^{T} \delta^{t-1} u(c(s_t,(B,B,...,B)))\right],$$

subject to

$$\sum_{t=1}^{T} E\left[\frac{w(s_t) - c(s_t,(A,A,...,A))}{R^{t-1}}\right] = 0, \tag{Zero Profits}$$

$$u(c(s_\tau,(h^{\tau-1},B))) + \hat{\beta} E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t,\left(h^{\tau-1},B,B,...,B\right)\right)\right)\bigg|\ s_\tau\right] \tag{PCC}$$

$$\geq u(c(s_\tau,(h^{\tau-1},A))) + \hat{\beta} E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t,\left(h^{\tau-1},A,B,...,B\right)\right)\right)\bigg|\ s_\tau\right],$$

and

$$u(c(s_\tau,(h^{\tau-1},A))) + \beta E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t,\left(h^{\tau-1},A,B,...,B\right)\right)\right)\bigg|\ s_\tau\right] \tag{IC}$$

$$\geq u(c(s_\tau,(h^{\tau-1},B))) + \beta E\left[\sum_{t>\tau} \delta^{t-\tau} u\left(c\left(s_t,\left(h^{\tau-1},B,B,...,B\right)\right)\right)\bigg|\ s_\tau\right],$$

and non-lapsing constraints:

$$u\left(c(s_\tau,(h^{\tau-1},A))\right)+\beta E\left[\sum_{t>\tau}\delta^{t-\tau}u\left(c(s_t,(h^{\tau-1},A,B,...,B))\right)\middle|s_\tau\right]\geq V(s_\tau),\ \ \forall s_\tau, \tag{NL}$$

and

$$u\left(c(s_\tau,(h^{\tau-1},B))\right)+\hat{\beta} E\left[\sum_{t>\tau}\delta^{t-\tau}u\left(c(s_t,(h^{\tau-1},B,B,...,B))\right)\middle|s_\tau\right]\geq \hat{V}(s_\tau),\ \ \forall s_\tau. \tag{PNL}$$

We first note that the incentive compatibility constraints (IC) must be binding on the equilibrium path, because otherwise we can increase $c(s_T,h^\tau,B,B,...,B)$ without affecting all other constraints while weakly increase the agent's perceived utility. Given incentive constraints are binding, we can simplify (PC) as

$$u(c(s_\tau,(h^{\tau-1},B)))\leq u(c(s_\tau,(h^{\tau-1},A))). \tag{sA30}$$

Substituting the binding IC constraints in the objective gives

$$E\sum_{t=1}^{T-1}\delta^{t-1}u(c(s_t,A,\cdots,A))+\beta\delta^{T-1}u(c(s_T,A,\cdots,A))+(\beta-1)\delta^{t-1}u(c(s_t,A,\cdots,A,B)).$$

Since $\beta<1$, we want to choose $c(s_t,A,\cdots,A,B)$ as small as possible (subject to the constraints). We now show that under the optimal contract, $c(s_t,A,\cdots,A,B)=0$. We need to verify that setting $c(s_t,A,\cdots,A,B)=0$ would not violate all other constraints. First, PC holds because (sA30) holds.

We then verify that PNL holds if NL holds. Suppose $\{\hat{c}(s_t,h^t_\tau):t\geq\tau\}$ solves the perceived outside option program $\hat{V}^I(s_\tau)$. So we have

$$\hat{V}^I(s_\tau)=u(\hat{c}(s_\tau,h^\tau_\tau))+\hat{\beta}E[\delta^{t-\tau}u(\hat{c}(s_t,(h^\tau_\tau,B,\cdots,B)))|s_\tau]. \tag{sA31}$$

We next verify the perceived non-lapsing constraint at $(s_\tau,(h^{\tau-1},B))=(s_\tau,(A,\cdots,A,B))$.

Other perceived non-lapsing constraints can be similarly verified. Note that

$$
\begin{aligned}
& u(c(s_\tau, (h^{\tau-1}, B))) + \hat{\beta} E\left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, B, \cdots, B))) \middle| s_\tau\right] \\
&= \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0)) + \frac{\hat{\beta}}{\beta} u(c(s_\tau, (h^{\tau-1}, B))) + \hat{\beta} E\left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, B, \cdots, B))) \middle| s_\tau\right] \quad \text{(sA32)} \\
&= \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \frac{\hat{\beta}}{\beta} \left(u(c(s_\tau, (h^{\tau-1}, A))) + \beta E\left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, A, B, \cdots, B))) \middle| s_\tau\right]\right) \quad \text{(sA33)} \\
&\geq \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \frac{\hat{\beta}}{\beta} V^I(s_\tau) \quad \text{(sA34)} \\
&\geq \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \frac{\hat{\beta}}{\beta} \left(u(\hat{c}(s_\tau, h_\tau^\tau)) + \beta E\left[\sum_{t>\tau} \delta^{t-\tau} u(\hat{c}(s_t, h_\tau^\tau, B, \cdots, B)) \middle| s_\tau\right]\right) \quad \text{(sA35)} \\
&= \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + \frac{\hat{\beta}}{\beta} u(\hat{c}(s_\tau, h_\tau^\tau)) + \hat{\beta} E\left[\sum_{t>\tau} \delta^{t-\tau} u(\hat{c}(s_t, h_\tau^\tau, B, \cdots, B)) \middle| s_\tau\right] \quad \text{(sA36)} \\
&= \left(1 - \frac{\hat{\beta}}{\beta}\right) u(0) + (\frac{\hat{\beta}}{\beta} - 1) u(\hat{c}(s_\tau, h_\tau^\tau)) + \hat{V}^I(s_\tau) \quad \text{(sA37)} \\
&\geq \hat{V}^I(s_\tau), \quad \text{(sA38)}
\end{aligned}
$$

where (sA32) follows from $c(s_\tau, (h^{\tau-1}, B)) = 0$, (sA33) from (IC), (sA34) from the actual non-lapsing constraints (NL), (sA35) follows from a revealed preference argument since $\hat{c}$ is also feasible in program $V(s_\tau)$, (sA36) follows from simple algebra, (sA37) uses (sA31), and (sA38) follows from $\hat{c}(s_\tau, h_\tau^\tau) \geq 0$.

By the previous argument, the perceived choice constraints and the perceived non-lapsing constraints can be ignored, so the program reduces to:

$$
\max E \sum_{t=1}^{T-1} \delta^{t-1} u(c(s_t, (A, \cdots, A))) + \beta \delta^{T-1} u(c(s_T, (A, \cdots, A))),
$$

subject to the zero-profit condition and the non-lapsing constraints. Since the objective is the same as the utility of a dynamically consistent consumer, we can replace the non-lapsing constraints by front-loading constraints. So $c^{1E} = c^{1A}$. □

Corollary 1 with One-Sided Commitment

This appendix generalizes Corollary 1 for settings with one-sided commitment, as mentioned in footnote 17:

Corollary 2. *Consider the model with one-sided commitment. There exists a perception-perfect equilibrium that does not depend on the consumer's naiveté* $\hat{\beta} \in (\beta, 1]$*. Moreover, any perception-perfect equilibrium has the same consumption path, which is continuous in* $\beta \in (0, 1]$*.*

Proof. By Lemma 4, we can focus on the auxiliary program with one-sided commitment. Let $x(s_t) \equiv u(c(s_t))$ denote the agent's utility from the consumption he gets in state s_t, and consider the dual program:

$$\max_{\{x(s_t)\}} \sum_{t=1}^{T} \sum_{s_t \in \mathbb{S}_t} p(s_t|s_1) \frac{w(s_t) - u^{-1}\left(x(s_t)\right)}{R^{t-1}}, \tag{sC1}$$

subject to

$$\sum_{t=1}^{T-1} \sum_{s_t \in \mathbb{S}_t} \delta^{t-1} p(s_t|s_1) x(s_t) + \beta \sum_{s_T \in \mathbb{S}_T} \delta^{T-1} p(s_T|s_1) x(s_T) \geq \underline{u}, \tag{sC2}$$

and

$$\sum_{t \geq \tilde{\tau}} \sum_{s_t \in \mathbb{S}_t} \delta^{t-\tilde{\tau}} p(s_t|s_{\tilde{\tau}}) x(s_t) + \beta \sum_{s_T \in \mathbb{S}_T} \delta^{T-\tilde{\tau}} p(s_T|s_{\tilde{\tau}}) x(s_T) \geq V_T^A(s_{\tilde{\tau}}) \;\; \forall s_{\tilde{\tau}} \in \mathbb{S}_{\tilde{\tau}}(s_\tau), \forall \tau, \tag{sC3}$$

This program corresponds to the maximization of a strictly concave function over a convex set, so, by the Theorem of the Maximum, the solution is unique and continuous in $\beta \in$

$(0, 1]$. Moreover, since the program does not involve $\hat{\beta}$, the equilibrium consumption path is not a function of the consumer's naiveté.

Once we pin down the unique consumption path, the baseline options are either zero or determined by the binding IC constraints and non-lapsing constraints (See the proof of Lemma 4). In particular, these constraints do not depend on the consumer's naiveté. So the equilibrium consumption vector is not a function of the consumer's naiveté. □

Removing Commitment Power

This appendix presents the formal analysis of the welfare effect of removing commitment power, as described in Subsection 3.1. We show that, for a fixed contract length, removing commitment power can make the consumer better off. To formalize the argument given in the text, let $\mathcal{V}_T^S$ denote the agent's welfare from smoothing consumption perfectly in the first $T-1$ periods and consuming zero in the last period:

$$\mathcal{V}_T^S \equiv \max_{\{c(s_t)\}} \sum_{t=1}^{T-1} E\left[\delta^{t-1} u\left(c(s_t)\right)\right] + \delta^{T-1} u(0),$$

subject to

$$\sum_{t=1}^{T-1} E\left[\frac{c(s_t)}{R^{t-1}}\right] \leq \sum_{t=1}^{T} E\left[\frac{w(s_t)}{R^{t-1}}\right].$$

Let $\mathcal{V}_T^{NS}$ denote the agent's welfare from consuming the endowment in each state:

$$\mathcal{V}_T^{NS} \equiv \sum_{t=1}^{T} E\left[\delta^{t-1} u\left(w(s_t)\right)\right].$$

Proposition 10. *Suppose agents are time inconsistent and $\mathcal{V}_T^{NS} > \mathcal{V}_T^S$. There exists $\bar{\beta} > 0$ such that if $\beta < \bar{\beta}$, the welfare with one-sided commitment is greater than the welfare with two-sided commitment.*

Proof. First, note that the welfare with two-sided commitment approaches to $\mathcal{V}_S$ as β approaches to zero. It suffices to show that the welfare with one-sided commitment is bounded below by $\mathcal{V}_{NS}$. In the remainder of the proof, we will therefore focus on the equilibrium with one-sided commitment.

We claim that for β close to zero, the equilibrium consumption equals the endowment in all last-period states: $c(s_T) = w(s_T), \forall s_T \in \mathbb{S}_T(s_1)$. To see this, consider a perturbation that shifts consumption from a state in the last period to the preceding state, that is, it increases $c(s_{T-1})$ by $\epsilon > 0$ and reduces $c(s_T)$ by $\frac{\epsilon R}{p(s_T|s_{T-1})}$ for some $s_T \in \mathbb{S}_T$ with $p(s_T|s_{T-1}) > 0$. Let W_{s_T} denote the future value of all income up to state s_T. The amount W_{s_T} is how much the agent would be able to consume at state s_T if he saves all his income from all periods for the last one. It therefore gives an upper bound on how much the agent can consume in the last period. Since there are finitely many states and $W_{s_T} < \infty$ for all s_T, we can take the uniform bound $W \equiv \max_{s_T} W_{s_T}$. This perturbation affects the LHS of the non-lapsing constraint at state s_t by

$$p(s_{T-1}|s_t)\left[u'(c(s_{T-1})) - \beta R\delta u'(c(s_T))\right]\delta^{T-1-t}\epsilon$$
$$> p(s_{T-1}|s_t)\left[u'(0) - \beta R\delta u'(W_{s_T})\right]\delta^{T-1-t}\epsilon,$$

which is positive whenever

$$\frac{u'(0)}{R\delta u'(W)} > \beta. \tag{sC1}$$

The perturbation has exactly the same effect on the objective function (scaled down by δ^t and multiplied by the probability of reaching state s_{T-1}). Thus, as long as β satisfies (sC1), the equilibrium will have the smallest consumption possible in the last period, which is determined by the non-lapsing constraint.

Substituting $c(s_T) = w(s_T)$ in the auxiliary program, it becomes analogous to the program of a time-consistent agent except that the contracting problem ends at period $T-1$ instead of period T:

$$\max_{\{c(s_t)\}} \sum_{t=1}^{T-1} \sum_{s_t \in \mathbb{S}_t(s_1)} \delta^{t-1} p(s_t|s_1) u\left(c(s_t)\right),$$

subject to

$$\sum_{t=1}^{T-1} \sum_{s_t \in \mathbb{S}_t(s_1)} p(s_t|s_1) \frac{w(s_t) - c(s_t)}{R^{t-1}} = 0,$$

and

$$\sum_{t=\tilde{\tau}}^{T-1} \sum_{s_t \in \mathbb{S}_t(s_{\tilde{\tau}})} \delta^{t-\tilde{\tau}} p(s_t|s_{\tilde{\tau}}) u\left(c(s_t)\right) \geq (V')^C(s_{\tilde{\tau}}) \text{ for all } s_{\tilde{\tau}},$$

for all $\tilde{\tau} = 2, ..., T$, where $(V')^C(s_{\tilde{\tau}})$ denotes the outside option for the time-consistent agent in this $(T-1)$-period economy.

It is straightforward to verify that $(V')^C(s_1)$ is bounded below by the utility from consuming the endowment in all states. If the endowment already satisfies the non-lapsing constraints, then the result follows from revealed preference because the endowment also satisfies zero profits. If the endowment does not satisfy the non-lapsing constraints, any renegotiation of the endowment satisfies the zero-profits condition and gives the time-consistent agent a strictly higher utility conditional on that state. So, replacing the endowment by the solution of the continuation program in all states where the non-lapsing constraints are violated leads to a profile of consumption that satisfies the constraints and gives a utility greater than the utility of consuming the endowment in each period. It thus follows by revealed preference that the solution of the program also gives a higher utility than consuming the endowment in all states.

Since the solution of a naive agent coincides with the solution of this auxiliary program, their welfare is also bounded below by the welfare from consuming their endowment in all periods $\mathcal{V}_{NS}$ when (sC1) holds. Therefore, by continuity, if $\mathcal{V}_{NS} > \mathcal{V}_S$, there exists $\bar{\beta}_N$ such that if $\beta < \bar{\beta}_N$, the welfare with one-sided commitment dominates the welfare with two-sided commitment. □

Notice that for generic endowment paths, the condition that $\mathcal{V}_T^{NS} > \mathcal{V}_T^S$ fails when T is large enough. So, as the contracting length grows, it becomes increasingly hard to satisfy the conditions for the time-inconsistent consumer to obtain higher welfare without commitment, as described in the text.

Proof of Claim in Section 3.5

In this appendix, we establish that a naive agent saves more than a sophisticate in the first period. Given a vector $x = (x_1, x_2, \cdots, x_T)$ with $x_1 = 1$ consider the program:

$$\max_{(c_1,\cdots,c_T)} \sum_{t=1}^{T} x_t u(c_t), \tag{sD1}$$

subject to

$$\sum_{t=1}^{T} \frac{c_t}{R^{t-1}} = \sum_{t=1}^{T} \frac{w}{R^{t-1}}. \tag{sD2}$$

The first-order conditions of this program are:

$$R^{t-1} x_t u'(c_t) \leq \lambda, \ \ \forall t, \tag{sD3}$$

where λ is the Lagrangian multiplier on the zero-profits condition (sD2). The inequality becomes equality whenever $c_t > 0$.

The consumption path of a naive agent solves this program for

$$x^N = \left(1, \frac{D_{T-1}}{D_{T-2}}, \frac{D_{T-1}}{D_{T-3}}, \cdots, \frac{D_{T-1}}{D_1}, D_{T-1}\right),$$

whereas the consumption path of a sophisticated agent solves this program for vector

$$x^S = (1, D_1, D_2, \cdots, D_{T-1}).$$

Let λ^N and λ^S denote the Lagrangian multiplier in the case of a naive agent and a sophisticated agent, respectively. Recall from equation (29), $x_t^N \geq x_t^S$ for all $t = 2, \cdots, T$.

We argue by contradiction and suppose the naive agent consumes strictly more than the sophisticate in the first period, i.e., $c_1^N > c_1^S$. We claim that $c_t^N \geq c_t^S, \forall t$. Then the claim together with $c_1^N > c_1^S$ would violate the zero-profits condition.

To prove the claim, first note that the claim trivially holds if $c_t^S = 0$. Now suppose $c_t^S > 0$, then $R^{t-1} x_t^S u'(c_t^S) = \lambda^S \geq u'(c_1^S)$. Since $c_1^N > c_1^S \geq 0$, the FOC at c_1^N is an equality: $u'(c_1^N) = \lambda^N$.

Note that for any $t = 2, \cdots, T$,

$$R^{t-1}x_t^S u'(c_t^S) = \lambda^S \geq u'(c_1^S) \geq u'(c_1^N) = \lambda^N \geq R^{t-1}x_t^N u'(c_t^N) \geq R^{t-1}x_t^S u'(c_t^N),$$

where the last inequality comes from $x_t^N \geq x_t^S$. It follows that $u'(c_t^S) \geq u'(c_t^N)$, i.e., $c_t^N \geq c_t^S, \forall t = 2, \cdots, T$. Together with $c_1^N > c_1^S$, it contradicts to the zero-profits condition. As a result, the naive agent must consume weakly less than the sophisticate in the first period (i.e., the naive agent saves weakly more than the sophisticate in the first period).

Moreover, the naive agent must consume *strictly* less than the sophisticate in the first period if $\lim_{c\searrow 0} u'(c) = +\infty$. In this case, we have an interior solution, and (sD3) becomes equality because consumption is always strictly positive. To see that $c_1^N < c_1^S$, we need to show that there is a contradiction when $c_1^N = c_1^S$. We recall that $x_t^N > x_t^S$ for all $t = 2, \cdots, T-1$. Now for $t = 2, \cdots T-1$,

$$R^{t-1}x_t^S u'(c_t^S) = u'(c_1^S) = u'(c_1^N) = R^{t-1}x_t^N u'(c_t^N) > R^{t-1}x_t^S u'(c_t^N),$$

which implies that $c_t^N > c_t^S$, for $2 \leq t \leq T-1$. We still have $c_T^N \geq c_T^S$. Together, we have a contradiction to the zero-profits condition. So if $\lim_{c\searrow 0} u'(c) = +\infty$, the naive agent must consume *strictly* less than the sophisticated agent in the first period.