

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 Köszegi, Shengwu Li, George Mailath, Pietro Ortoleva, Wolfgang Pendorfer, 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, Toulouse School of Economics, UC Berkeley, University of Colorado Boulder, Wharton, Warwick, 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 externalities 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

¹See, for example, Gruber and Kőszegi (2001); O’Donoghue and Rabin (2003); DellaVigna and Mal-mendier (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

²In 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

³There 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

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

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

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

where $\beta \in (0, 1)$, $\delta \in (0, 1]$ and $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ 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-*

⁴The 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.

⁵Continuity 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,

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

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

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

⁸See 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,

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

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:

$$\max_{c_1, c_2, c_3} \sqrt{c_1} + \beta (\sqrt{c_2} + \sqrt{c_3}) \quad (4)$$

subject to

$$c_1 + c_2 + c_3 = 1. \quad (5)$$

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

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

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

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

¹⁰Since 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:

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

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:

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

subject to

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

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

and

$$\begin{aligned} & u(c(s_\tau, (h^{\tau-1}, A))) + \beta E \left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, A, B, \dots, B))) \mid s_\tau \right] \\ & \geq u(c(s_\tau, (h^{\tau-1}, B))) + \beta E \left[\sum_{t>\tau} \delta^{t-\tau} u(c(s_t, (h^{\tau-1}, B, B, \dots, B))) \mid s_\tau \right]. \end{aligned} \quad (\text{IC})$$

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*:

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

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. *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:

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

subject to

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

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

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

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

$$u(c_3(A, A)) + \beta\delta u(c_4(A, A)) \geq u(c_3(A, B)) + \beta\delta u(c_4(A, B)), \quad (13)$$

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 $c_4(A, B)$:

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

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

$$c_2(A) \geq c_2(B), \quad (15)$$

and

$$c_3(A, A) \geq c_3(A, B). \quad (16)$$

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

$$c_2(B) = c_3(A, B) = 0. \quad (17)$$

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:

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

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, \dots, 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, \dots, B) to (A, B, \dots, 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, \dots, B) to (A, A, B, \dots, 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, \dots, 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, $\lim_{T \rightarrow +\infty} (W_T^C - W_T^I) = 0$.*

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

¹¹Along 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

¹²See 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%.

¹³In 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.

¹⁴To 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.

¹⁵When a non-lapsing constraint binds, there are also equilibria in which the consumer lapses and re-contracts 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:

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

subject to

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

and

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

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:

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

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

¹⁶To 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:

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

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

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

and the non-lapsing constraints

$$u(c_\tau(s_{\bar{\tau}}, (h_{\bar{\tau}}^{\bar{\tau}-1}, A))) + \beta E \left[\sum_{t > \bar{\tau}} \delta^{t-\bar{\tau}} u(c(s_t, (h_{\bar{\tau}}^{\bar{\tau}-1}, A, B, \dots, B))) \middle| s_{\bar{\tau}} \right] \geq V^I(s_{\bar{\tau}}), \quad (24)$$

$$u(c_\tau(s_{\bar{\tau}}, (h_{\bar{\tau}}^{\bar{\tau}-1}, B))) + \hat{\beta} E \left[\sum_{t > \bar{\tau}} \delta^{t-\bar{\tau}} u(c(s_t, (h_{\bar{\tau}}^{\bar{\tau}-1}, B, B, \dots, B))) \middle| s_{\bar{\tau}} \right] \geq \hat{V}^I(s_{\bar{\tau}}), \quad (25)$$

for all $s_{\bar{\tau}}$ following s_τ and all $h_{\bar{\tau}}^{\bar{\tau}}$ 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:

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

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):

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

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

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

and the perceived non-lapsing constraints,

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

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. *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, $\lim_{T \rightarrow +\infty} (W_T^C - W_T^I) = 0$.*

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

¹⁷In 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 $\underline{c} \equiv \{c(s_t) : s_t \in \mathbb{S}_t, t = 1, \dots, T\}$. 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

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

¹⁹Specifying 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:

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

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

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

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

²⁰If 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 unable to "exploit" consumers (as will be shown in Theorem 3).

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

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* pa-

rameters, 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:

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

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*:

$$\frac{D_{t+2}}{D_{t+1}} \geq \frac{D_{t+1}}{D_t}, \quad (29)$$

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

²¹As 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 $(\hat{D}_1, \hat{D}_2, \dots)$ 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 $t = 0, 1, \dots$ 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:*

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

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:

$$\begin{aligned} u(c_2(A)) + D_1u(c_3(A, B)) + D_2u(c_4(A, B)) &\geq u(c_2(B)) + D_1u(c_3(B, B)) + D_2u(c_4(B, B)), \\ u(c_3(A, A)) + D_1u(c_4(A, A)) &\geq u(c_3(A, B)) + D_1u(c_4(A, B)). \end{aligned}$$

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

²²See 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:

$$c_2(B) = c_3(B, B) = c_3(A, B) = 0. \quad (31)$$

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:

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

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

²³See 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 $W_T^H(c) \equiv u(c_1) + u(c_2) + \dots + u(c_T)$ denote the welfare from consumption stream $c = (c_1, \dots, c_T)$. 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

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

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

²⁴With 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 $t \in \{1, \dots, T\}$, the agent exerts effort e_t at a cost of $C(e_t)$. The cost of effort $C : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ 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.

²⁵Assumption 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, \dots, 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.

²⁶This 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.

²⁷The 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 $\lim_{T \rightarrow \infty} (W_T^I - V_T^A(\beta)) = \lim_{T \rightarrow \infty} (1 - \beta) E \delta^{T-1} u(c(s_T)) = 0$. Since the objective function is linear in β , it follows from the Envelope Theorem that $\frac{\partial V_T^A(\beta)}{\partial \beta} = E \delta^{T-1} u(c(s_T)) \geq \delta^{T-1} u(0)$. 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:

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

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: $c^0(s_T) = w(s_T)$. 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

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

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

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

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}_∞ exists.*

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

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

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:

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

subject to (IC), (PC), and

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

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:

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

subject to

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

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

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

subject to (IC), (PC), and

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

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:

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

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

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

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

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

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

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{c}' , for a time-consistent consumer such that the time-consistent consumer's utility is given by \underline{U}' :

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

The existence of \underline{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{c}' .

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

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

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

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

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

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

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

$$\frac{\partial \Pi_{T,\beta}}{\partial \beta} = E [\lambda \delta^{T-1} u'(c^A(s_T))] \geq \lambda \delta^{T-1} u'(0),$$

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

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

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

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

Note that $\Pi_{T,1} = \Pi_T^C(\underline{\mathbf{c}}')$ and $\Pi_{T,\beta} = \Pi_T^L(\underline{\mathbf{c}})$. 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,

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

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,

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

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{c}' and $\Pi_{T,\beta}$ is the problem for time inconsistent consumers with the outside option \underline{c} , it follows that $\lim_{T \nearrow \infty} |\Pi_T^C(\underline{c}') - \Pi_T^I(\underline{c})| = 0$.

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

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

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

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

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

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. \square

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

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

subject to the zero-profit condition, and

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

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

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

$$c_2(B) = c_3(B, B) = \dots = c_{T-1}(B, \dots, B) = 0. \quad (\text{A9})$$

Otherwise, we could lower $u(c_t(B, \dots, B))$ by $\frac{\epsilon}{D_{t-2}}$ and increase $u(c_T(B, \dots, B))$ 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: $c_{t+1}(\underbrace{A, \dots, A}_{t-1}, B) = \dots = c_{T-1}(\underbrace{A, \dots, A}_{t-1}, \underbrace{B, \dots, B}_{T-t-1}) = 0$.

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

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

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

subject to

$$\sum_{t=1}^T e_t(A, \dots, A) = E_T, \quad (\text{A10})$$

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

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

The (IC) at time $\tau = 2$ must be binding (otherwise, we could decrease $e_T(B, \dots, B)$, improving the objective). We claim that $e_t(B, \dots, B) = 0$ for $t \geq 3$. Suppose otherwise that $e_t(B, \dots, B) > 0$ for some t . Consider a perturbation that shifts effort in the baseline from period t to period 2: $C(e_2(B)) + \epsilon$ and $C(e_t(B, \dots, B)) - \frac{\epsilon}{D_{t-2}}$. 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 $t + 1$:

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

Substituting in the objective function, we obtain:

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

where the first and third equations come from $C(0) = 0$, and the second and fourth 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