FTG Working Paper Series

Contracting with a Present-Biased Agent: Sannikov meets Laibson

by

Alejandro Rivera

Working Paper No. 00128-00

Finance Theory Group

www.financetheory.com

Contracting with a Present-Biased Agent: Sannikov meets Laibson *

Alejandro Rivera†

UT Dallas

June 1, 2023

Abstract

This paper develops a methodology to solve dynamic principal-agent problems in which the agent features present-biased time preferences and naive beliefs. There are three insights. First, the problem has a recursive representation using the agent's *perceived* continuation value as a state variable (i.e., the remaining value the agent (wrongly) anticipates getting from the contract). Second, incentive compatibility corresponds to a volatility constraint on the agent's perceived continuation value. Finally, due to the agent's naïvete, a perceived action constraint needs to be satisfied. This constraint is accommodated by linking the agent's perceived effort policy and the volatility of his perceived continuation value. Novel economic insights regarding optimal time-varying incentives and the term-structure of compensation include a large signing bonus for the agent and more high powered and delayed incentives relative to the rational benchmark.

Keywords: Behavioral Finance, Recursive Contracts, Present-Bias.

JEL codes: D86, G40, C61.

*We are grateful to Subas Acharya, Luca Anderlini, Gilles Chemla, Joanne Juan Chen, Jaksa Cvitanic, Diogo Duarte, Felix Feng, David Jimenez-Gomez, Francois Gourio, Camilo Hernandez, Pablo Kurlat, Andrey Malenko, Nadya Malenko, Gustavo Manso, Peter Maxted (discussant), Simon Mayer, Benjamin Moll, Nam Nguyen, Christine Parlour, Facundo Piguillem, Dylan Possamai, Andrea Pratt, Alessandro Previtero, Jean-Charles Rochet, Lukas Schmid, Martin Szydlowski, Can Urgun, Stephane Villeneuve, Edward Van Wesep, Noah Williams, Xingtan Zhang and seminar participants at EIEF, Copenhagen Business School, Miami Herbert Business School, Florida International University, AFA 2023, BIRS 2023, and UT-Dallas for their insightful comments.

†Jindal School of Management, University of Texas at Dallas, alejorivera1@gmail.com

1 Introduction

Moral hazard is ubiquitous in most human economic relationships as pointed out all the way back by Adam Smith's "Theory of Moral Sentiments" (Smith (1822)). Economists rely on the principal-agent problem to model such situations, making it the cornerstone of a broad range of economic sub-fields (Bolton and Dewatripont (2004)). Moreover, the dynamic and repeated nature of these economic relationships requires modeling principal-agent settings with long horizons. However, the neoclassical literature has assumed, by and large, that economic participants are fully rational and discount the future exponentially. This assumption is problematic in light of the large body of field and experimental evidence documenting present-bias and the success of present-bias in endowing economic models with additional realism and explanatory power since Laibson (1997) introduced it to the profession. We fill this gap by asking the following questions: Can present-bias be incorporated into dynamic contracting models in a tractable manner? If so, what are the novel economic insights from such a model?

We argue that the answer to the first question is affirmative. In order to elaborate on the details it is worth recalling the progression in contract theory from its early static versions (e.g., Holmström (1979)) to modern dynamic contracting (e.g., Di Tella and Sannikov (2021); Cvitanić et al. (2018)). Contracting requires the solution of two optimization problems: an optimal contract design by the principal, and the optimal choice of an (unobservable) action by the agent. In the early 1980s the transition to dynamic contracts seemed daunting because the essential toolbox for dynamic optimization in economics, dynamic programming (Stokey and Lucas (1989)), seemed to have hit a natural limitation when addressing two intertwined dynamic optimization problems. This situation was sometimes referred to as "dynamic programming squared". The first major watershed in these developments was the realization that the principal's optimal contracting problem could be recursively characterized using the *agent's continuation value* from the contract as a state variable (Spear and Srivastava (1987)). In spite of this major insight, characterizing incentive compatible contracts remained challenging beyond assuming the first-order approach. Moreover, analytical progress remained elusive.

The second milestone occurred after dynamic contracting in continuous-time was mathematically formalized and martingale methods were used to solve the agent's problem (Sannikov (2008)).[1] This approach rendered the characterization of incentive compatibility (IC) as a simple constraint on the volatility of the agent's continuation value, thereby transforming the principal problem into a standard stochastic control problem. The enhanced analytical and numerical tractability from these developments set the stage for the growth in the dynamic contracting literature and its various applications across many fields in economics and finance.[2]

Meanwhile, integrating present-bias and naïvete about self-control into a principal-agent posed the unique challenge of incorporating in the contract design the action that the principal (actually) implements and the action the agent (wrongly) anticipates he will take in the future. Such inconsistency arises because the agent can display naïvete about his future present-bias and anticipates his discount rate to be $\hat{\beta}$; while in reality his discount rate is $\beta < \hat{\beta}$, leading him to take a different action than what he had anticipated to take. Heidhues and Kőszegi (2010) address this challenge by suitably expanding the usual contract space to include both a suggested action and a perceived action. Importantly, in order to ensure consistency between the agent's anticipated self control problems and his perceived actions, a new constraint referred to as the perceived choice constraint (PCC) needs to be added to the principal's optimization problem. That is, the principal optimization problem has to be solved subject to the (IC) and (PCC) constraints.

The main contribution of this paper is to adapt the above methodologies to solve dynamic contracting problems when the agent displays present-bias. The first insight is to realize that from the perspective of the principal, she no longer needs to keep track of the agent's (actual) continuation value, but instead she must keep track of his *perceived* continuation value, thereby making the latter the required state variable in the recursive formulation of the optimal contracting problem. Second, the (IC) constraint now links the agent's (actual) action to the volatility of the perceived continuation value, discounted by his actual discount rate β. Finally, the (PCC) corresponds to a constraint linking the perceived action to the volatility of the perceived continuation value discounted by his perceived discount rate $\hat{\beta}$. This methodology transforms the optimal contracting problem

[1] Contemporaneously, Williams (2004, 2011) relied instead on the stochastic maximum principle to solve the agent's problem, and Biais et al. (2007) tackled the problem as a limiting case of the discrete-time model.

[2] See surveys by Sannikov (2013); Biais et al. (2013) and a textbook treatment by Cvitanic and Zhang (2012).

into a new, but otherwise standard, stochastic control problem. Table 1 summarizes the above discussion and frames our contribution relative to the existing contract theory literature.

Setting	Two-period model	Continuous-time model
Exponential discounting	(IC)-constraint: Reward agent with higher consumption if "high" output is realized. Holmström (1979).	(IC)-constraint: Use sensitivity of agent's continuation value to output to incentivize effort. Sannikov (2008).
Present-biased agent	(PCC)-constraint: Rewards incentivize agent's perceived choice under his (wrongly) anticipated future present-bias $\hat{\beta}$. Heidhues and Kőszegi (2010).	(PCC)-constraint: Use sensitivity of agent's *perceived* continuation value to incentivize agent's perceived choice using $\hat{\beta}$ as discount factor. This paper.

Table 1: Contract theory with present-bias and in continuous-time.

In order to make analytical progress in answering the second question we specialize our principal-agent problem to the risk-neutral with limited liability setting of DeMarzo and Sannikov (2006). Our methodology, however, can be easily adapted to other applications and settings without major challenges. There are four main economic insights. First, the term-structure of the contract changes: i) it features a large initial payment to the agent in order to attract him into the venture, but ii) upon this initial payment the remaining compensation is further back-loaded relative to that of an exponential agent. Therefore, our model jointly rationalizes the use of both: front-loaded signing bonuses (Van Wesep, 2010; Parsons and Van Wesep, 2013; Xu and Yang, 2016) and back-loaded incentive schemes (Edmans et al., 2017) in managerial compensation packages. Second, because the principal relies on (future) rewards conditional on good performance to incentivize effort, present-bias increases incentive costs, decreases implemented effort, and reduces principal value. As a result, present-bias may preclude the formation of positive NPV ventures.

Third, principal value is increasing in the agent's naïvete about his present-bias. Such is the case because the principal exploits the agent's naïvete by offering him a contract with high powered incentives that induces a large wedge between the perceived action and the implemented action. Moreover, the principal optimally delays further the compensation of a naive agent in order to lengthen the life of the contract and exploit his naïvete for longer. In other words, the principal offers the agent a contract with generous future rewards conditional on good performance. The agent anticipates working hard in the future and reaping the benefits of this high powered incentives. In reality, his actual effort is lower than anticipated and a big share of the expected rewards do not materialize. Finally, the extent to which the principal exploits the agent, measure by the distance between the perceived continuation value and the actual continuation value, is time-varying. Exploitation is largest early in the life of the contract, but decreases with time as the agent (on average) reports good performance and (finally) starts receiving his promised back-loaded rewards.

1.1 Literature Review

Our paper is closely related to three strands of the literature. First, we contribute to the literature on dynamic contracting with non-standard preferences. Adrian and Westerfield (2008) developed the first dynamic contracting model with heterogeneous beliefs. They show that disagreement can lead to gains from trade by transferring payments in states that the agent deems relatively more likely to occur. Miao and Rivera (2016) characterize the optimal contract when the principal is ambiguity-averse and show that ambiguity generates endogenous belief heterogeneity inducing a trade-off between incentives and ambiguity sharing in such a way that the incentive constraint does not always bind.[3] Malenko (2019) characterizes the optimal contract when the manager has empire-building preferences. He shows that empire building preferences lead to a budgeting mechanism with threshold separation of financing. Finally, Cetemen et al. (2021) studies dynamic contracting with non-exponential discounting under sophistication. Our paper is the first to consider situations in which the agent's actions are different than what he had anticipated to take. We show that the relevant state variable for the recursive formulation of this problem is the *perceived* continuation value of the agent, and that the perceived

[3] Relatedly, Szydlowski and Yoon (2021) show that the optimal contract when the principal is ambiguity averse about the agent's cost of effort generates a seemingly excessive pay-performance sensitivity, thereby rationalizing the use of performance-sensitive debt. Dicks and Fulghieri (2021) study a model in which both principal and agent display ambiguity aversion and apply this model to optimal compensation of division managers within an organization.

choice constraint can be formulated as a constraint on the volatility of this state variable.

Second, we also contribute to the literature on behavioral contracting with present-bias. DellaVigna and Malmendier (2004) study optimal contract design by a profit-maximizing firm facing present-biased customers. The theoretically predicted contracts match the empirical contract design across the credit card, health club, gambling, and life insurance industries. Heidhues and Kőszegi (2010) study contract choice and repayment behavior in competitive credit markets when customers display present-bias. They show that, consistent with subprime mortgages and most credit cards, the baseline repayment terms are inexpensive, but they are also inefficiently front-loaded and delays require paying large penalties. On the technical side, they formally show how to enlarge the contract space to incorporate perceived choices and actual choices when agents display naïvete about their self-control. Most recently, Gottlieb and Zhang (2021) study contracts between risk-neutral firms and present-biased consumers. Their analysis shows that the welfare loss from present-bias vanishes as the contracting horizon grows. Our paper enriches this literature by unlocking the tractability of continuous-time techniques to study fully dynamic moral hazard settings and provide an analytical characterization of the impact of present-bias on the optimal contract.

Finally, our paper also contributes to the recent literature harnessing the tractability of continuous-time to model present-bias following the Instantaneous Gratification (IG) Model pioneered by Harris and Laibson (2013). This literature explores the impact of present-bias in various economic settings: Grenadier and Wang (2007) characterize the impact of present-bias on the optimal exercise of real options within a corporate setting and show that present-bias can sub-optimally hasten investment. Laibson et al. (2021) incorporate present-bias into a heterogeneous agent macroeconomic model and study fiscal and monetary policy within this setting. They show that present-bias increases the impact of fiscal policy and amplifies the effect of monetary policy.[4] Our paper complements this literature by developing a tractable framework to study the impact of

[4]Other recent contributions to this literature include Acharya et al. (2020) which characterizes the value of sophistication when investors are risk-averse and present-biased. Maxted (2020) generalizes the IG model to provide analytical characterizations in a very general setting with liquid and illiquid assets. Shigeta (2020) develops a portfolio choice problem in which the agent has present-bias and stochastic differential utility. Tian (2016) characterizes a firm's optimal capital structure under present-bias. Finally, Hernández and Possamaï (2020) provide a general theory to model time-inconsistent sophisticated agents and establish a direct and rigorous proof of an extended dynamic programming principle similar to that of the classical theory.

present-bias on dynamic principal-agent settings.

2 Model

In this section we develop a continuous-time principal-agent model in which the agent features present-bias preferences. In section 2.1 we lay out the risk-preferences and technology of our principal-agent model following the corporate finance application of DeMarzo and Sannikov (2006). In section 2.2 we model the present-bias preferences of the agent in our continuous-time setting following the instantaneous gratification model of Harris and Laibson (2013). To reiterate, we use the DeMarzo and Sannikov (2006) formulation of the continuous-time principal-agent to facilitate conveying the gist of our methodological contribution. However, our methodology can be adapted to any continuous-time contracting problem with "hidden action" as long as there are no "hidden states".[5]

2.1 Output and Technology

A risk-neutral principal (she) needs to hire an agent (he) to manage a project. The principal is risk-neutral and discounts the future at rate r. The agent is also risk-neutral, has limited liability, and displays $\beta - \gamma$ present-bias preferences (discussed below) à la Laibson (1997). Fixing an effort policy $a \in A$, the project's cumulative output Y_t follows a controlled arithmetic Brownian motion process:

$$dY_t = a_t \mu dt + \sigma dZ_t^a, \tag{1}$$

where Z_t^a is a standard Brownian motion under the probability measure $\mathbb{P}^a$. Everyone observes the realization of output Y_t, but effort is the agent's private information. Exerting effort is costly to the agent and is given by $g(a) = \theta a^2/2$. Thus, moral hazard arises and incentives need to be provided.

Because the agent features naïvete about his present-bias, there will be a discrepancy between his actual effort and his perceived effort policies. As a result, a contract

[5]We conjecture that in the special case of CARA preferences and unlimited liability (e.g., He (2011); Hackbarth et al. (2021)), our methodology could still be used to accommodate a present-biased agent with "hidden savings" as long as we impose the Hyperbolic Euler Equation (see Harris and Laibson (2001)) as a constraint in lieu of the Euler Equation to ensure a no-private savings condition.

$\Gamma = (C, \tau, a, \hat{a})$ specifies consumption C, a termination clause τ, an implemented effort policy a, and a perceived effort policy $\hat{a}$, all as functions of the observed history of output.[6] Upon termination the principal obtains L and the agent R, which we normalize to 0 for simplicity. All quantities are assumed to be square integrable and progressively measurable under the usual conditions.

2.2 Present-Bias

We now describe the agent's discount function that models his present-bias preferences in continuous-time. All periods, present and future, are discounted exponentially with discount factor $0 < \gamma < 1$. However, future periods are additionally discounted with uniform weight $0 < \beta \leq 1$. As a result, the present period receives full weight, while future periods are given weight $\beta e^{-\gamma t}$.

We model the agent as a sequence of selves. Call the self born at time $s_0 = 0$ "self 0". The lifetime of self 0 is split into the present, which lasts from s_0 to $s_0 + \tau_0$, and the future, which lasts from $s_0 + \tau_0$ to ∞. The present can be thought of as the interval during which control is exercised by self 0, while the future is the interval during which control is exercised by future selves. The length of the time interval τ_0 is stochastic and exponentially distributed with hazard rate $\lambda \in [0, \infty)$.

Once the present of self 0 ends at $s_0 + \tau_0$, self 1 is born and takes control. The preferences of self 1 are identical to those of self 0, and his present lasts from $s_1 = s_0 + \tau_0$ to $s_1 + \tau_1$. Proceeding in this manner we obtain an infinite sequence of selves $\{0, 1, 2, ...\}$ born respectively at dates $\{s_0, s_1, s_2, ...\}$. Each self applies a discount factor $D_n(t)$ to the utility flow at time $s_n + t$, where

$$D_n(t) = \left\{ \begin{array}{c} e^{-\gamma t} \textit{ if } t \in [0, \tau_n) \\ \beta e^{-\gamma t} \textit{ if } t \in [\tau_n, \infty) \end{array} \right\}.$$

For tractability reasons we will focus in the limiting case when $\lambda \to \infty$, known as the Instantaneous Gratification (IG) Model developed by Harris and Laibson (2013). In this

[6]Expanding the standard contract space to include the perceived effort policy $\hat{a}$ is inspired by the two-period principal-agent formulation detailed in Section 2.2 of Koszegi (2014) applied to credit markets, in which $(\hat{r}_1, \hat{r}_2)$ correspond to the repayment plan the agent believes he will chose (i.e., to $\hat{a}$ in our setting) and (r_1, r_2) to the repayment plan he actually chooses (i.e., to a in our setting).

case the discount function exhibits a discrete discontinuity at $t = 0_+$, so that

$$D(t) = \left\{ \begin{array}{c} 1 \text{ } \textit{if} \, t = 0 \\ \beta e^{-\gamma t} \text{ } \textit{if} \, t \in (0, \infty) \end{array} \right\}.$$

The assumption that each self lives only for an extremely short instant is made for mathematical tractability. However, Laibson and Maxted (2020) show that IG preferences closely approximate discrete-time models with period lengths that are psychologically plausible (e.g., one week or less).

Finally, a key consideration in our analysis is whether the agent correctly anticipates the present-bias he will have in the future. Following O'Donoghue and Rabin (2001), we suppose that the agent believes his future selves β will actually be $\hat{\beta} \geq \beta$. That is, the agent can display naïvete about his future present-bias, and incorrectly underestimate his future self-control problems. Our model, nests the sophisticated case ($\beta = \hat{\beta}$), whereby the investor understands his future selves will behave differently than he would like to, and the standard exponential discounting case ($\beta = \hat{\beta} = 1$) used in most economic models, whereby the agent is not present-biased.

3 The Agent's Problem

Consider now the problem faced by an agent in light of a contract $\Gamma = (C, \tau, a, \hat{a})$. Suppose the agent anticipates his future selves will take action $\hat{a}$. Then, he can proceed to compute his (perceived) continuation utility $\hat{V}$ from this contract under exponential discounting:

$$\hat{V}_t = E_t^{\hat{a}} \left[\int_t^\tau e^{-\gamma(s-t)} (dC_s - g(\hat{a}_s)) ds \right]. \tag{2}$$

We note the expected value is computed under the probability measure $\mathbb{P}^{\hat{a}}$, since the agent (incorrectly) anticipates his future selves to exert effort policy $\hat{a}$.

Using similar arguments as in Sannikov (2008) we can apply the Martingale Representation Theorem to show that there exists stochastic process ϕ_t, adapted to the filtration generated by output Y, such that the evolution of $\hat{V}$ is given by

$$d\hat{V}_t = \gamma \hat{V}_t dt - (dC_t - g(\hat{a}_t) dt) + \phi_t \left(dY_t - \hat{a}_t \mu dt \right). \tag{3}$$

The first term captures the appreciation in the perceived continuation value of the agent due to long-term exponential discounting. The second term captures the change resulting from the per-period utility anticipated by the agent from his consumption net of effort costs: $dC_t - g(\hat{a}_t)dt$. Finally, the last term captures the sensitivity with respect to the realization of output. That is, $\phi_t = d\hat{V}_t/dY_t$ is a measure of the contract's incentives.

We proceed to define incentive compatibility in our context.

Definition 1. *We say that a contract $\Gamma = (C, \tau, a, \hat{a})$ is incentive compatible (IC) if it is optimal for the agent's current self at time t to exert effort a_t when it anticipates his future selves to exert effort $\hat{a}_s$, for all $s > t$.*

Lemma 1. *A contract $\Gamma = (C, \tau, a, \hat{a})$ is incentive compatible (IC) if and only if:*

$$g'(a_t) = \beta\phi_t\mu \iff a_t = \frac{\beta\mu\phi_t}{\theta} \tag{IC}$$

for all t, where ϕ comes from the dynamics of $\hat{V}$ given in equation (3).

Equation (IC) characterizes incentive compatibility as a link between the marginal cost of effort and the volatility of the agent's perceived continuation value, which is proportional to the marginal benefit of effort. That is, for a contract to be (IC) the implemented effort a has to be optimal for any current t-self when he is trading off the cost of increasing his effort (born in the present) versus the anticipated benefits of effort resulting from a more generous contract when high output is realized. The benefits are enjoyed in the future, and therefore discounted with parameter β reflecting the current t-self present-bias.

Lemma 1 generalizes the characterization in Sannikov (2008) for naive present-biased agents by noting that, for a rational agent, his perceived continuation value is his actual continuation value. In particular, note equation (4) in Sannikov (2008) is essentially the same as equation (IC) above, except that the volatility of the continuation value is replaced by the volatility of the perceived continuation value.

We now proceed to address the fact that the agent's expectation about the effort policy chosen by his future selves has to be consistent with his beliefs about his future self-control problems. In a two-period setting, Heidhues and Kőszegi (2010) introduce the perceived choice constraint as an additional constraint on the principal's optimization problem, which we adapt below to our fully dynamic environment:

Definition 2. *We say that a contract $\Gamma = (C, \tau, a, \hat{a})$ satisfies the perceived choice constraint (PCC) if the 0-self agent thinks it will be optimal for all his future selves to choose effort policy $\hat{a}$ for all $t > 0$.*

Lemma 2. *A contract $\Gamma = (C, \tau, a, \hat{a})$ satisfies the perceived choice constraint (PPC) if and only if:*

$$g'(\hat{a}_t) = \hat{\beta}\phi_t\mu \iff \hat{a}_t = \frac{\hat{\beta}\mu\phi_t}{\theta} \tag{PCC}$$

for all t, where ϕ comes from the dynamics of $\hat{V}$ given in equation (3).

Equation (PCC) is new in the literature. It characterizes the perceived action constraint by equating the marginal cost of effort with its marginal benefit, which is proportional to the volatility of the agent's perceived continuation value and the agent's (incorrect) beliefs about his future self-control problems $\hat{\beta}$.

To gain intuition, consider the 0-self anticipation of the optimization problem that his future t-self will be faced with. The 0-self anticipates his t-self trading off the cost of increasing his effort (born at time t, i.e., in the present, and therefore receiving full weight) versus the anticipated benefits of exerting effort resulting from a more generous contract when high output is realized. The benefits are enjoyed in the future (i.e., after time t) and which the 0-self imagines his t-self will discount with parameter $\hat{\beta}$, reflecting the 0-self underestimation of his t-self present-bias. We note that equation (PCC) represents a continuous-time analogue of the (PCC) constrained stated in equation (1) of Heidhues and Kőszegi (2010).

Importantly, the IC constraint and the PCC coincide with each other when the agent is sophisticated, namely $\beta = \hat{\beta}$. This redundancy of the PCC is expected because the sophisticated agent has correct expectations about his future selves. Therefore, what he perceives as his future choice will be his actual choice. In other words, the implemented effort policy and the perceived effort policy must coincide (i.e., $a = \hat{a}$), which yields the perceived effort policy redundant as part of the contract space. Consequently, when dealing with a sophisticated agent it is sufficient to specify the suggested effort policy and the IC constraint.

Finally, the agent's participation constraint (PC) states that the perceived payoff from the contract at $t = 0$ must be larger than an exogenous initial outside option denoted $\underline{\hat{V}}$:

$$\beta E^{\hat{a}}\left[\int_{0_+}^{\tau} e^{-\gamma s}(dC_s - g(\hat{a}_s))ds\right] + dC_0 = \beta\hat{V}_{0_+} + dC_0 \geq \underline{\hat{V}}. \tag{PC}$$

The characterization of the incentive-compatibility constraint via equation (IC), the perceived choice constraint via equation (PCC), and the participation constraint via equation (PC) will be very useful when solving the principal's problem. Anticipating, these characterizations allow us to write the principal's problem recursively and thereby transform it into a standard Markovian stochastic control problem with the agent's perceived continuation value $\hat{V}$ as a state variable.

4 The Principal's Problem

The principal's problem consists of choosing a contract $\Gamma = (C, \tau, a, \hat{a})$ that maximizes the net present value of the cashflows generated by the project net of the agent's compensation, subject to the agent's IC, PCC, and PC constraints.[7] Formally, the principal solves

$$\max_{\Gamma} \mathbb{E}^a\left[\int_0^{\tau} e^{-rt}\left(dY_t - dC_t\right) + e^{-r\tau}L\right] \tag{4}$$

subject to (IC), (PCC), and (PC).

Before proceeding to solve the principal's problem, we restate the dynamics of $\hat{V}$ from the perspective of the probability measure $\mathbb{P}^a$ used by the principal when forming her expectations (i.e., the probability induced by the implemented effort level a):

$$d\hat{V}_t = \gamma\hat{V}_t dt - (dC_t - g(\hat{a}_t)dt) + \phi_t\mu(a_t - \hat{a}_t)dt + \phi_t\sigma dZ_t^a, \tag{5}$$

where dZ_t^a is Brownian motion process under $\mathbb{P}^a$. It is worth it at this stage to compare the dynamics of $\hat{V}$ as expected by the agent (equation (3)) and as expected by the principal (equation (5)). Because the principal is aware that the agent will choose effort a instead of $\hat{a}$, the perceived continuation value features an additional term $\phi_t\mu(a_t - \hat{a}_t)dt$ relative to the dynamics perceived by the agent. Importantly, the principal knows the agent's

[7] Implicit in the formulation below, and in line with the dynamic contracting literature, once the principal offers a contract to the agent, the principal commits to delivering to the agent the consumption stated in the contract as a function of the history of observed output. That is the principal can commit to the contract.

perceived continuation value will not grow as fast as the agent anticipates it to, since his effort level will be lower than what he anticipates (i.e., since $a_t \leq \hat{a}_t$). This observation plays a critical role in section 5 where we characterize the principal's optimal contract and the extent to which she can exploit the agent's present-bias to increase her profit.

4.1 Recursive Formulation for $t > 0$

In this section we study the problem faced by a principal for any $t > 0$. Suppose that the agent's perceived continuation value at the time is $\hat{V}_t$ (i.e., the value that a previous self anticipates to get from the continuation contract when discounting the payoffs exponentially). Denote the value function for the principal as $F(\hat{V})$. Because both principal and agent are risk neutral, by a similar intuition as in DeMarzo and Sannikov (2006), we expect the optimal contract to feature $dC_t > 0$ only after sufficiently good performance. In particular, the perceived continuation value evolves according to (3) (with $dC_t = 0$ whenever $\hat{V}_t \in [0, \bar{V})$) and is reflected down by $dC_t > 0$ whenever $\hat{V}_t = \bar{V}$. The contract is terminated the first time the perceived continuation value goes down to zero (i.e., $\tau = \min\{t : \hat{V}_t = 0\}$).

Hence, we conjecture that the value function satisfies the following ODE for $\hat{V} \in [0, \bar{V}]$:

$$rF(\hat{V}) = \max_{\phi} \left\{a(\phi)\mu + F'(\hat{V})(\gamma\hat{V} + g(\hat{a}(\phi)) + \phi\mu(a(\phi) - \hat{a}(\phi))) + \frac{1}{2}F''(\hat{V})\phi^2\sigma^2\right\} \tag{6}$$

$$F(0) = L, \qquad F'(\bar{V}) = -1, \qquad F''(\bar{V}) = 0, \tag{7}$$

where $a(\phi) = \frac{\beta\mu\phi}{\theta}$ and $\hat{a}(\phi) = \frac{\hat{\beta}\mu\phi}{\theta}$ as given by the IC and PCC constraints, respectively. The first boundary condition in (7) states that the principal gets L upon termination. The second boundary condition states that the agent is paid whenever the marginal benefit to the principal of promising a payment equals his marginal cost of 1. Finally, the last boundary condition corresponds to a smooth-pasting condition which pins down the optimal payout boundary $\bar{V}$.

Upon substituting the first order condition for the optimal ϕ:

$$\phi(\hat{V}) = -\frac{\beta\mu^2}{\mu^2(2\beta - 2\hat{\beta} + \hat{\beta}^2)F'(\hat{V}) + \theta\sigma^2 F''(\hat{V})}, \tag{8}$$

the system (6)-(7) can be solved numerically using any standard ODE solver. Moreover, many of the properties of the value function, as well as that of the optimal contract are analytically tractable via probabilistic methods. We turn to such analysis of the optimal contract in the next section. The proposition below formalizes the analysis for this subsection.

Proposition 1. *Consider the contracting problem faced by the principal at $t = 0_+$. Suppose that there exists a unique twice continuously differentiable solution $F(\hat{V})$ to the ODE (6) on $[0, \bar{V}]$ with boundary conditions (7). Then, i) $F(\hat{V})$ is the value function for the principal when she has to deliver a perceived continuation value $\hat{V}$ to the agent, ii) optimal incentives $\phi(\hat{V})$ are given by (8), iii) the contract delivers a perceived value $\hat{V} \in [0, \bar{V}]$ to the agent whose perceived continuation value follows the dynamics (3) for $t \in [0, \tau]$, where the optimal payments dC reflect the process $\hat{V}_t$ at $\bar{V}$, and iv) termination occurs at time $\tau = \inf\{t \geq 0 : \hat{V}_t = 0\}$.*

4.2 Value Function at $t = 0$

We now solve for the optimal initial payment to the agent dC_0. The reason for characterizing this payment separately from the rest of this contract is because the fact that the agent displays present-bias modelled following the IG model implies that the participation constraint disproportionaly values the instantaneous utility of the 0-self, as seen in equation (PC). Formally, the optimal initial payment solves

$$\max_{dC_0} F(\hat{V}_{0^+}) - dC_0, \tag{9}$$

subject to the participation constraint (PC). Because, at the optimum, the participation constraint binds, we can substitute it into (9). Taking first order conditions implies that the optimal initial payment is given by

$$dC_0 = \begin{cases} 0, & \text{if } 0 \leq \underline{\hat{V}} < \tilde{V}, \\ \underline{\hat{V}} - \tilde{V}, & \text{if } \underline{\hat{V}} \geq \tilde{V}, \end{cases} \tag{10}$$

where $\tilde{V}$ solves $F'(\tilde{V}) = -\beta$. The intuition for the initial payment is straightforward. Because the agent discounts the future at rate β, in order to satisfy the PC constraint, the principal finds it optimal to pay him "early" in order to entice him to participate in the contract. The following proposition characterizes the value function for the principal at $t = 0$.

Proposition 2. *The value for the principal's optimization problem* (4) *is equal to* $F\left(\frac{\hat{V}-dC_0}{\beta}\right)$, *where* dC_0 *is given by* (10).

5 Economic Insights

In this section we study the novel economic insights delivered by enriching, an otherwise standard continuous-time principal-agent model, with the possibility that the agent features present-bias preferences. In section 5.1 we show the impact of present-bias on the level of incentives and effort implemented by the optimal contract. In section 5.2 we characterize the effect of present-bias on the optimal term-structure of compensation (i.e., optimal amount of back-loading and front-loading of incentives). Finally, in section 5.3 we characterize the additional profit that a principal earns by exploiting the agent's present-bias, and explore the situations in the state space for which such exploitation is most prominent.

5.1 Value Function, Incentives, and Effort

In this section we explore the impact of the present-bias parameters on the principal's value function, the contract's incentives, implemented effort, and perceived effort. Proposition 3 below constitutes a natural starting point for this analysis by providing probabilistic representations for the comparative statics of the value with respect to β and $\hat{\beta}$.

Proposition 3. *The comparatives statics of the value function with respect to* β *and* $\hat{\beta}$ *are given by:*

$$\frac{\partial F(\hat{V};\hat{\beta})}{\partial \hat{\beta}} = \mathbb{E}^a\left[\int_t^\tau e^{-r(s-t)}F'(\hat{V})(\hat{\beta}-1)\frac{\mu^2\phi(\hat{V})^2}{\theta}|\hat{V}_t = \hat{V}\right] \tag{11}$$

$$\frac{\partial F(\hat{V};\beta)}{\partial \beta} = \mathbb{E}^a\left[\int_t^\tau e^{-r(s-t)}[1+F'(\hat{V})\phi(\hat{V})]\frac{\mu^2\phi(\hat{V})}{\theta}|\hat{V}_t = \hat{V}\right] \tag{12}$$

These expressions are intuitive. Start with the comparative static with respect to $\hat{\beta}$. An increase in the agent's perceived present-bias parameter means that the agent anticipates having less present-bias in the future than he actually will. That means that he will, on average, under-perform his expectations by a larger margin since he will exert less effort than anticipated. As a result, his perceived continuation value will tend to decrease. The payoff for the principal from such reduction in the agent's perceived continuation value is proportional to $F'(\hat{V})$ multiplied by the reduction in the drift induced by a larger $\hat{\beta}$ at each point in time (i.e., we integrate this effect over time and take expectations).

Panel A in Figure 1 shows comparative statics for the value function with respect to three values of $\hat{\beta}$ and shows that in general the value function is increasing in $\hat{\beta}$. Mathematically, this comes from the fact that most of the time $F'(\hat{V}) < 0$ (i.e., the principal value function is decreasing in her promises to the agent). Intuitively, this comparative static says that having a more naive agent is valuable to the principal, because she can exploit this naïvete and increase her value. We refer to this mechanism as the *exploitation effect*.

How does the principal exploit the agent's naivite? Because the principal knows that the agent is underestimating his future present-bias, she offers him a contract with more high powered incentives (panel B). The agent, values these incentives, because he thinks he will exert a lot of effort (panel D), and in turn earn large rewards in the future. His actual effort, however, will be much lower than anticipated (panel C). Thus, the principal will end up paying him a lot less. By saving on these future compensation, the principal earns a higher profit when dealing with a more naive agent. Section 5.3 provides additional insights on the magnitude of the exploitation effect.

Similarly, the expression for the comparative static with respect to the present-bias parameter β is also intuitive. There are two effects. First, as β increases, the agent discounts the future less, and exerts more effort for a given level of incentives. Second, higher β implies the agent is relatively less naive since $\hat{\beta} - \beta$ decreases. As discussed above, lower naïvete decreases the value to the principal in proportion to $F'(\hat{V})$.

Panel A in Figure 2 shows comparative statics for the value function with respect to three values of β. It shows that for low values of $\hat{V}$ the first effect dominates and the value function is increasing in β. By contrast, because $F'(\hat{V}) < 0$ for high values of $\hat{V}$, the second effect dominates, rendering the value function decreasing in β. Moreover, panel B shows that as β increases it is easier to incentivize the agent, and thus incentives are decreasing in β. By a similar mechanism as before, panel C shows that effort is increasing in β for low

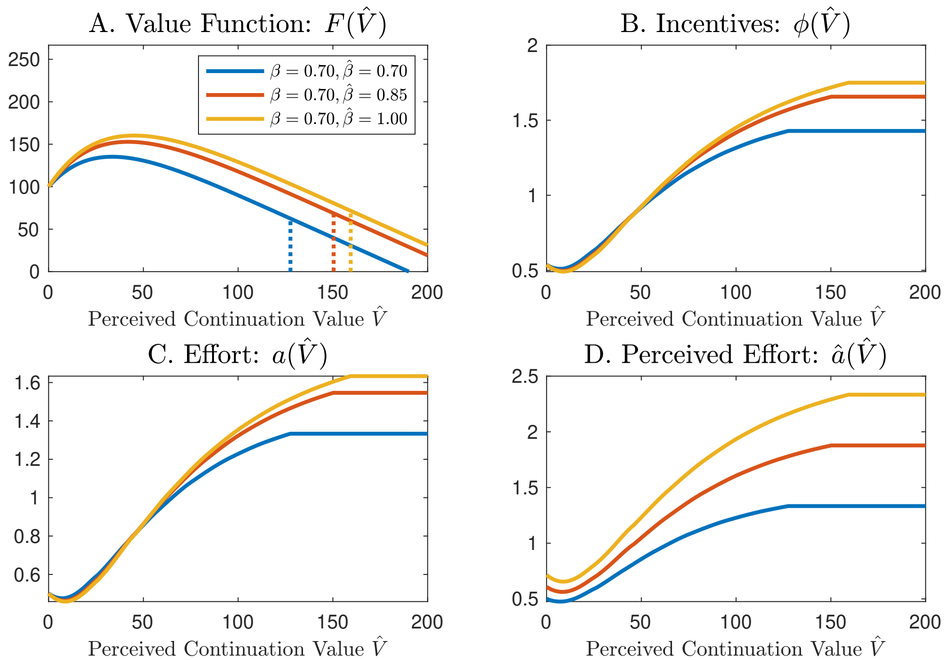

Figure 1: **Comparative statics with respect to $\hat{\beta}$.** Other parameter values are $r = 0.05, \gamma = 0.08, \mu = 20, \sigma = 25, \theta = 15, L = 100$.

values of $\hat{V}$ (because it is easier to incentivize the agent), but decreasing for high values of $\hat{V}$ (because the agent is less naive, which reduces the exploitation effect). Finally, panel D shows that because increasing β makes the agent less naive, his perceived effort will decrease.

Finally, we discuss what happens when increasing β and $\hat{\beta}$ at the same time. This can be thought of as reducing the agent's present bias while keeping his naïvete unchanged. Adding both effects yields the following expression:

$$\frac{\partial F(\hat{V};\hat{\beta})}{\partial \hat{\beta}} + \frac{\partial F(\hat{V};\beta)}{\partial \beta} = \mathbb{E}^a \left[\int_t^\tau e^{-r(s-t)} [1 + \hat{\beta} F'(\hat{V}) \phi(\hat{V})] \frac{\mu^2 \phi(\hat{V})}{\theta} | \hat{V}_t = \hat{V} \right]. \tag{13}$$

Again, there are two effects. The first effect, as before, simply captures the higher output generated by the greater effort exerted by a less present-biased agent. The second effect captures the rise in the drift of the agent's perceived continuation value needed to compensate him for his higher effort. The impact of the latter effect is positive for the principal when $F'(\hat{V}) > 0$ ($\hat{V}$ is low), but negative when $F'(\hat{V}) < 0$ ($\hat{V}$ is high). Panel A in Figure 3 provides a numerical illustration in which the first effect dominates, and the value function increases as present-bias decreases.

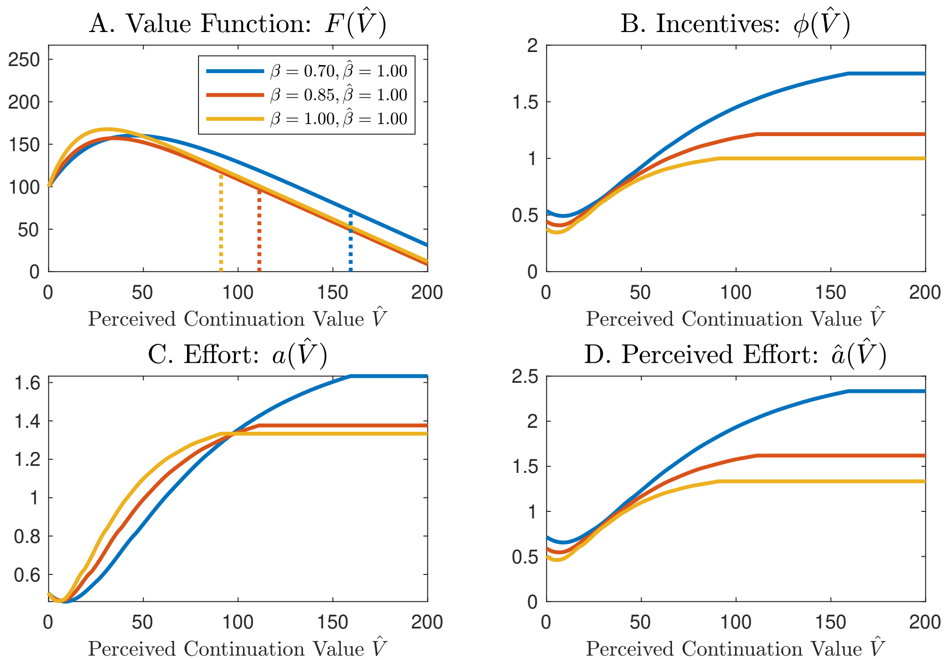

Figure 2: **Comparative statics with respect to** β**.** Other parameter values are $r = 0.05, \gamma = 0.08, \mu = 20, \sigma = 25, \theta = 15, L = 100$.

Furthermore, panel B shows that a less present-biased agent is easier to incentivize, and thus requires less powerful incentives. However, because the principal finds it cheaper to incentivize the agent, she implements a higher level of effort (which equals perceived effort in this case since $\hat{\beta} = \beta$), as depicted in panel C (resp. panel D).

5.2 Payout Boundary and Initial Bonus

We now turn our attention to explore the impact of the present-bias parameters on the payout boundary and the initial bonus offered by the principal to recruit the agent. Proposition 4 below provides probabilistic representations for the comparative statics of the payout boundary with respect to β and $\hat{\beta}$.

Proposition 4. *The comparatives statics of the payout boundary with respect to* β *and* $\hat{\beta}$

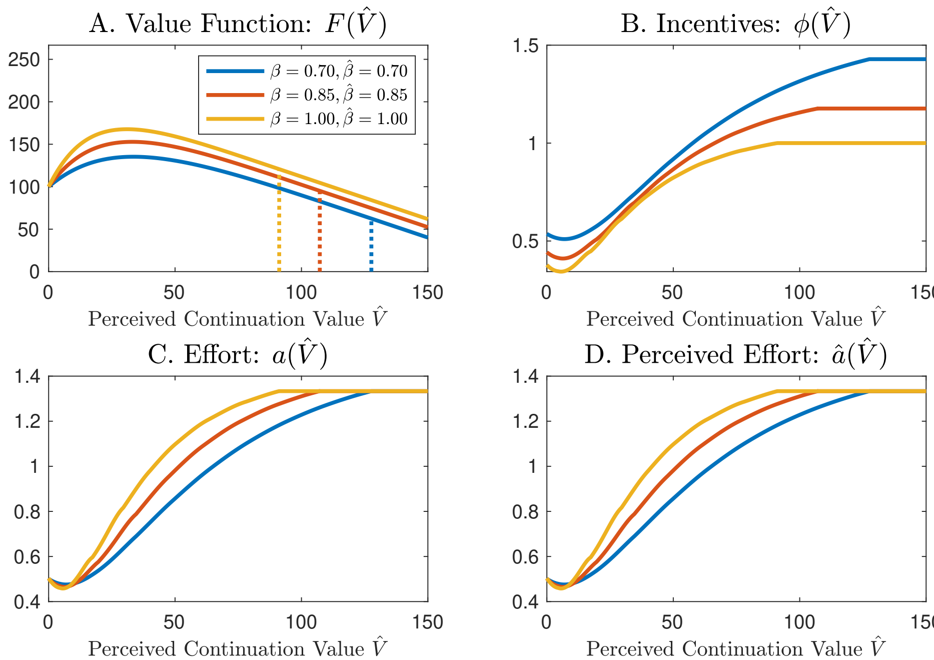

Figure 3: **Comparative statics with respect to β and $\hat{\beta}$ simultaneously.** Other parameter values are $r = 0.05, \gamma = 0.08, \mu = 20, \sigma = 25, \theta = 15, L = 100$.

are given by:

$$\frac{\partial \bar{V}}{\partial \hat{\beta}} = \frac{r}{\gamma - r}\left[\underbrace{\frac{(1-\hat{\beta})\beta^2\mu^2}{r\theta(2\beta - 2\hat{\beta} + \hat{\beta}^2)^2}}_{(+)} - \underbrace{\mathbb{E}^a\left[\int_t^\tau e^{-r(s-t)} F'(\hat{V})(\hat{\beta}-1)\frac{\mu^2\phi(\hat{V})^2}{\theta}|\hat{V}_t = \bar{V}\right]}_{(+/-)}\right], \tag{14}$$

$$\frac{\partial \bar{V}}{\partial \beta} = \frac{r}{\gamma - r}\left[\underbrace{\frac{(\beta - 2\hat{\beta} + \hat{\beta}^2)\beta\mu^2}{r\theta(2\beta - 2\hat{\beta} + \hat{\beta}^2)^2}}_{(-)} - \underbrace{\mathbb{E}^a\left[\int_t^\tau e^{-r(s-t)}[1 + F'(\hat{V})\phi(\hat{V})]\frac{\mu^2\phi(\hat{V})}{\theta}|\hat{V}_t = \bar{V}\right]}_{(+/-)}\right]. \tag{15}$$

Consider first the effect of $\hat{\beta}$ on the payout boundary. The first term constitutes a direct effect that captures the increased value for a principal from dealing with a more naive agent. In order to exploit this naïvete the principal finds it optimal to postpone the agent's compensation. The second term, constitutes an indirect effect capturing the effect of higher naïvete on the drift of the agent's continuation value, which in turn affects the relative benefit of postponing the agent's compensation. The sign of this effect is ambiguous and depends on the sign of $F'(\hat{V})$ throughout the life of the contract. Panel A in Figure 4 provides an illustration in which the exploitation effect renders it optimal for the principal to postpone the compensation of the agent as his naïvete increases. Intuitively, the longer the principal can interact with a naive agent, the more he can exploit his naïvete. Thus, postponing his compensation reduces the probability of termination, and lengthens his profitable interaction with the agent. By the same token, the principal will reduce the size of the initial bonus of a more naive agent, in order to reduce the probability of terminating his relationship with this more naive agent (i.e., the initial bonus dC_0 is decreasing in $\hat{\beta}$, as shown in panel D).

Next, we study the effect of β on the payout boundary (while keeping $\hat{\beta}$ constant). There are two terms in this expression. The first term is negative, capturing the reduction in value for the principal from the fact that the agent's beliefs about his future present-bias are less distorted (i.e., from β moving closer to $\hat{\beta}$). The second term, as before, is ambiguous.

Intuitively, there are two effects. First, as present-bias declines, the agent is willing to exert more effort in exchange for promises of future rewards, an incentive effect. Hence, the volatility of the agent's continuation value is lower and termination is less likely. Because termination is less of a concern, the principal finds it optimal to pay earlier and make a larger initial bonus. Additionally, the second effect is what we have called the exploitation effect. As his perceived present-bias moves closer to his actual present-bias, there is less room for the principal to exploit the agent. This means the principal is less concerned about termination. As a result, earlier compensation (panel B) and a higher initial bonus are optimal (panel E).

Finally, panels C and F shut down the exploitation effect by increasing β and $\hat{\beta}$ at the same time. The incentive effect reduces the payout boundary and increases the initial bonus. However, their magnitudes are smaller relative to those of panels B and E, since the exploitation motive is unchanged in this exercise.

5.3 Exploitation Effect

In this section we compute the severity of the exploitation effect induced by the optimal contract. Recall that, because the agent is naive, his actual effort policy will be different from what he anticipates. Moreover, the expected path of output, and therefore his compensation will also be different. Hence, to understand the magnitude of the exploitation effect, we need to compute the (actual) continuation value V_t:

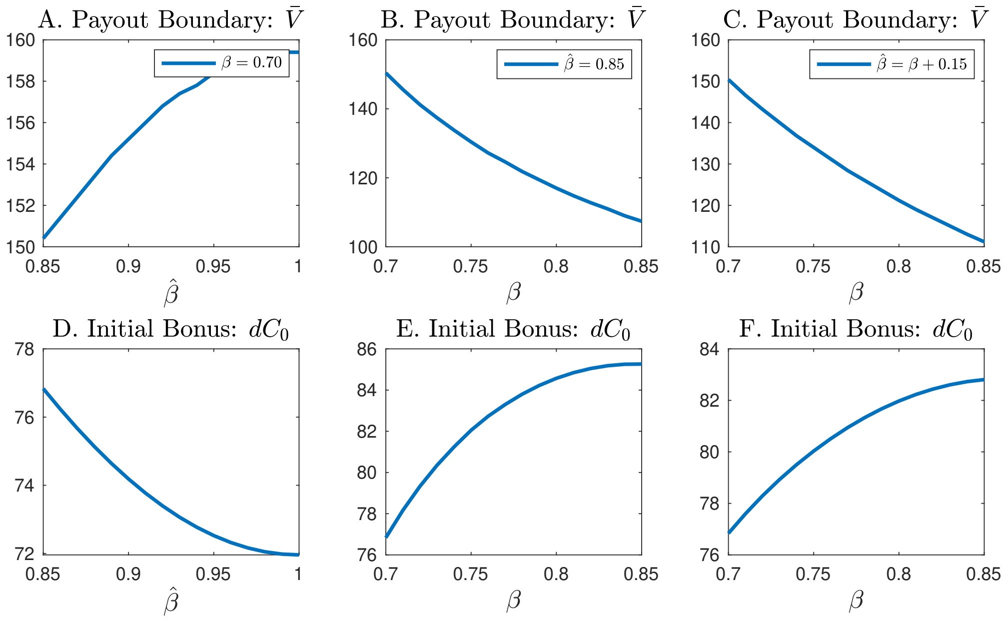

Figure 4: **Comparative statics for the payout boundary and initial bonus.** Other parameter values are $r = 0.05, \gamma = 0.08, \mu = 20, \sigma = 25, \theta = 15, L = 100, \underline{\hat{V}} = 150$.

$$V_t = E_t^a \left[\int_t^{\tau} e^{-\gamma(s-t)} (dC_s - g(a_s)) ds \right], \tag{16}$$

where we note that, in contrast to expression (2), the expected value is computed under the probability measure $\mathbb{P}^a$ (i.e., the probability measure induced by the actual effort policy chosen by the agent).

In general, we expect the agent's continuation value to be lower than his perceived continuation value. After all, the principal is aware of the agent's naïvete, and it is in her interest to design a contract that exploits this naïvete: $V_t < \hat{V}_t$. Panel A in Figure 5 compares V_t and $\hat{V}_t$ for three different values of $\hat{\beta}$. The blue line corresponds to the 45 degree line, because when the agent is sophisticated (i.e., when $\hat{\beta} = \beta$), his perceived continuation value equals his (actual) continuation value. For the red and the yellow lines, we see that his perceived continuation value is lower than his (actual) continuation value due to the exploitation effect. Moreover, and in line with intuition, as the agent becomes more naive (i.e., as $\hat{\beta}$ grows), the wedge between V_t and $\hat{V}_t$ grows larger.

Panel B depicts the ratio of the continuation value to the perceived continuation value. This ratio is equal to one in the case of sophistication (blue line), but it is below one when the agent is naive. Moreover, this ratio is increasing in $\hat{V}$. The intuition for this

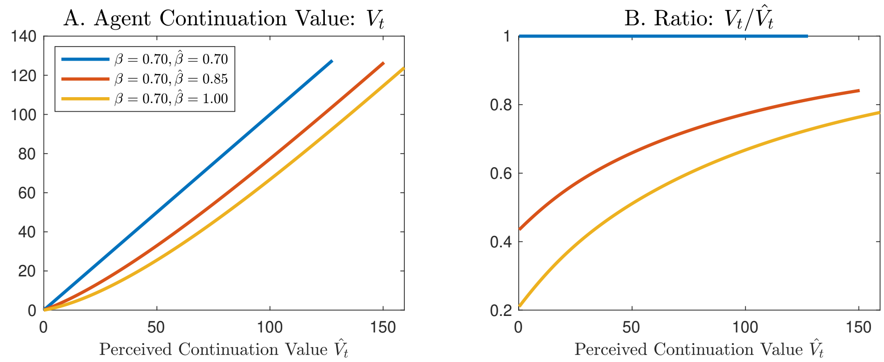

Figure 5: **Continuation value versus perceived continuation value.** Other parameter values are $r = 0.05, \gamma = 0.08, \mu = 20, \sigma = 25, \theta = 15, L = 100, \underline{\hat{V}} = 150$.

pattern stems from the fact that for low values of $\hat{V}$ the agent's compensation is far in the future. Because naïvete manifests in the form of optimism about future performance, he overestimates his compensation by a larger margin when he needs a long sequence of positive cash flows to get paid, than when his compensation is imminent, which is the case when $\hat{V}$ is high.

6 Conclusion

We analyze a dynamic principal-agent problem in which the agent displays present-bias. We develop a methodology that makes this problem numerically and analytically tractable upon converting it into a standard stochastic control problem. Our approach mimics the dynamic contracting literature by uncovering the recursive structure of the problem upon realizing that the right state variable is the perceived continuation value of the agent, and that both, the (IC) and (PCC) constraints correspond to links between the agent's action and perceived action with the volatility of the perceived continuation value.

Our results suggest several fruitful avenues for future research. First, given the prevalence of present-bias borrowers in the mortgage market, can our model provide additional insights to the design of optimal mortgages in the spirit of Piskorski and Tchistyi (2010)? Second, what are the normative implications for behavioral taxation (Farhi and Gabaix (2020)) when households display present-bias? Lastly, can we extend the analysis of Heidhues and Strack (2021) to jointly identify present-bias in a setting in which effort, timing, and incentives are all endogenously determined? These and other questions are the subject of ongoing research.

7 Appendix

Proof of Lemma 1. The agent's current self chooses effort in order to maximize the sum of i) his instantaneous utility plus ii) his perceived continuation utility discounted by the present-bias parameter β:

$$\max_{\tilde{a}}(dC_t - g(\tilde{a})dt) + \beta E_t^{\tilde{a}}[d\hat{V}_t] \tag{17}$$

$$= \max_{\tilde{a}}(dC_t - g(\tilde{a})dt) + \beta[\gamma \hat{V}_t dt - (dC_t - g(\hat{a}_t)dt) + \phi_t \mu\,(\tilde{a}dt - \hat{a}_t dt) + E_t^{\tilde{a}}[\sigma dZ_t^{\tilde{a}}]] \tag{18}$$

$$= \max_{\tilde{a}}\{\beta\phi_t\mu\,(\tilde{a}dt - \hat{a}_t dt) - g(\tilde{a})dt\} + \beta[\gamma\hat{V}_t dt - g(\hat{a}_t)dt], \tag{19}$$

which, by convexity of the effort cost function $g''(a) \geq 0$, implies that the contract $\Gamma = (C, \tau, a, \hat{a})$ is incentive compatible if and only if for all t:

$$g'(a_t) = \beta\phi_t\mu \iff a_t = \frac{\beta\mu\phi_t}{\theta}. \tag{20}$$

□

Proof of Lemma 2. The agent's 0-self anticipates the optimization problem that his future t-self will solve as choosing effort in order to maximize the sum of i) t-self's instantaneous utility plus ii) t-self's perceived continuation utility discounted by 0-self's (underestimated) present-bias parameter for his t-self $\hat{\beta}$:

$$\max_{\tilde{a}}(dC_t - g(\tilde{a})dt) + \hat{\beta} E_t^{\tilde{a}}[d\hat{V}_t] \tag{21}$$

$$= \max_{\tilde{a}}(dC_t - g(\tilde{a})dt) + \hat{\beta}[\gamma \hat{V}_t dt - (dC_t - g(\hat{a}_t)dt) + \phi_t \mu\,(\tilde{a}dt - \hat{a}_t dt) + E_t^{\tilde{a}}[\sigma dZ_t^{\tilde{a}}]] \tag{22}$$

$$= \max_{\tilde{a}}\{\hat{\beta}\phi_t\mu\,(\tilde{a}dt - \hat{a}_t dt) - g(\tilde{a})dt\} + \hat{\beta}[\gamma\hat{V}_t dt - g(\hat{a}_t)dt], \tag{23}$$

which, by convexity of the effort cost function $g''(a) \geq 0$, implies that the contract $\Gamma = (C, \tau, a, \hat{a})$ satisfies the PCC constraint if and only if for all t:

$$g'(\hat{a}_t) = \hat{\beta}\phi_t\mu \iff \hat{a}_t = \frac{\hat{\beta}\mu\phi_t}{\theta}. \tag{24}$$

□

Proof of Proposition 1. Fix an arbitrary contract $\Gamma' = (C', \tau', a', \hat{a}')$ that satisfies the (IC) and (PCC) constraints. Applying lemmas 1 and 2 it must be the case that $a_t^{'} = \frac{\beta\mu\phi_t^{'}}{\theta}$ and $\hat{a}_t^{'} = \frac{\hat{\beta}\mu\phi_t^{'}}{\theta}$ for all t. Next, consider the following gain process:

$$G_t^{\Gamma^{'}} = \int_0^{t\wedge\tau^{'}} e^{-rs}\left(a_s^{'}\mu - dC_s^{'}\right) + e^{-r(t\wedge\tau^{'})}F(\hat{V}_{t\wedge\tau^{'}}).$$

Applying Ito's lemma to $dG_t^{\Gamma^{'}}$ combined with equation (6) imply that $G_t^{\Gamma^{'}}$ is a $\mathbb{P}^{a^{'}}$ supermartingale. Integrating from 0 to T and taking expectations gives:

$$G_0^{\Gamma^{'}} \geq \mathbb{E}^{a^{'}}\left[\int_0^{T\wedge\tau^{'}} e^{-rt}\left(dY_t - dC_t\right) + e^{-r(T\wedge\tau^{'}))}L + e^{-r(T\wedge\tau^{'})}F(\hat{V}_{T\wedge\tau^{'}})\right]. \tag{25}$$

Letting $T \to \infty$ and noticing that because $\Gamma^{'}$ satisfies the (PC) constraint then the transversality condition $\lim_{T\to\infty}\mathbb{E}^{a^{'}}\left[e^{-r(T\wedge\tau^{'})}F(\hat{V}_{T\wedge\tau^{'}})\right] = 0$ holds, we obtain that:

$$F(\hat{V}_0) = G_0^{\Gamma^{'}} \geq \mathbb{E}^{a^{'}}\left[\int_0^{T\wedge\tau^{'}} e^{-rt}\left(dY_t - dC_t\right) + e^{-r(T\wedge\tau^{'}))}L\right], \tag{26}$$

where the first equality follows from the fact that the conjectured optimal contract $\Gamma = (C, \tau, a, \hat{a})$ satisfies equation (6) with equality, making G_t^{Γ} a $\mathbb{P}^a$ martingale. Noting that (26) holds for any arbitrary contract $\Gamma^{'}$ that satisfies the (PC), (PCC), and (IC) constraints completes the proof. □

Proof of Proposition 2. Since the principal value function for 0_+ has already been shown to be $F(\hat{V}_{0_+})$, the problem for the principal at time $t = 0$ corresponds to optimally choosing dC_0 and "stitching" the remainder of the contract as described by proposition 1. Because the PC constraint binds, the principal problem at $t = 0$ becomes:

$$\max_{dC_0} F\left(\frac{\hat{\underline{V}} - dC_0}{\beta}\right) - dC_0. \tag{27}$$

Finally, a standard proof by contradiction shows that $F(\hat{V})$ is concave, which implies that the maximizer in (27) is uniquely given by expression (10).

□

Proof of Proposition 3. We begin by rewriting ODE (6) and boundary conditions (7) by making explicit the dependence of the value function $F(\hat{V}; \hat{\beta})$ on the present-bias parameter $\hat{\beta}$, denoting the optimal control by $\phi(\hat{V})$, and substituting for a and $\hat{a}$ from (IC) and (PCC) respectively gives:

$$
\begin{aligned}
rF(\hat{V};\hat{\beta}) =& \frac{\beta\mu\phi(\hat{V})}{\theta} + F'(\hat{V};\hat{\beta})(\gamma\hat{V} + \frac{\hat{\beta}^2\mu^2\phi(\hat{V})^2}{2\theta} + \frac{\mu^2\phi(\hat{V})^2(\beta-\hat{\beta})}{\theta}) + \frac{1}{2}F''(\hat{V};\hat{\beta})\phi(\hat{V})^2\sigma^2. \\
F(0;\hat{\beta}) =& L, \qquad F'(\bar{V}(\hat{\beta});\hat{\beta}) = -1, \qquad F''(\bar{V}(\hat{\beta});\hat{\beta}) = 0.
\end{aligned}
$$

Differentiating with respect to $\hat{\beta}$ yields:

$$
\begin{aligned}
r\frac{\partial F(\hat{V};\hat{\beta})}{\partial\hat{\beta}} =& \frac{\partial F'(\hat{V};\hat{\beta})}{\partial\hat{\beta}}(\gamma\hat{V} + \frac{\hat{\beta}^2\mu^2\phi(\hat{V})^2}{2\theta} + \frac{\mu^2\phi(\hat{V})^2(\beta-\hat{\beta})}{\theta}) + \frac{1}{2}\frac{\partial F''(\hat{V};\hat{\beta})}{\partial\hat{\beta}}\phi(\hat{V})^2\sigma^2 \\
& + F'(\hat{V};\hat{\beta})\frac{\mu^2\phi(\hat{V})^2}{\theta}(\hat{\beta}-1). \qquad (28) \\
\frac{\partial F(0;\hat{\beta})}{\partial\hat{\beta}} =& 0, \qquad \frac{\partial F'(\bar{V};\hat{\beta})}{\partial\hat{\beta}} = 0
\end{aligned}
$$

Applying a standard Feynman-Kac type of argument shows that the probabilistic representation of $\frac{\partial F(\hat{V};\hat{\beta})}{\partial\hat{\beta}}$ which satisfies the system (28) is given by:

$$
\frac{\partial F(\hat{V};\hat{\beta})}{\partial\hat{\beta}} = \mathbb{E}^a\left[\int_t^\tau e^{-r(s-t)}F'(\hat{V})(\hat{\beta}-1)\frac{\mu^2\phi(\hat{V})^2}{\theta}|\hat{V}_t = \hat{V}\right],
$$

as desired. The second part of the proposition follows from an identical argument. □

Proof of Proposition 4. First, we substitute the boundary conditions (7) and the expression for $\phi(\hat{V})$ from (8) into (6). Next, we set $\hat{V} = \bar{V}$ and differentiate with respect to $\hat{\beta}$. Finally, we solve for $\frac{\partial\bar{V}}{\partial\hat{\beta}}$, which gives:

$$
\frac{\partial\bar{V}}{\partial\hat{\beta}} = \frac{r}{\gamma-r}\left[\frac{(1-\hat{\beta})\beta^2\mu^2}{r\theta(2\beta-2\hat{\beta}+\hat{\beta}^2)^2} - \frac{\partial F(\hat{V};\hat{\beta})}{\partial\hat{\beta}}\right]. \qquad (29)
$$

Substituting for $\frac{\partial F(\hat{V};\hat{\beta})}{\partial\hat{\beta}}$ from (11) yields (14). A similar argument shows (15). □

References

Acharya, Subas, David Jimenez Gomez, Dmitrii Rachinskii, and Alejandro Rivera, "Present-Bias and the Value of Sophistication," *Available at SSRN 3722608*, 2020. 4

Adrian, Tobias and Mark M Westerfield, "Disagreement and learning in a dynamic contracting model," *The Review of Financial Studies*, 2008, *22* (10), 3873–3906. 1.1

Biais, Bruno, Thomas Mariotti, and Jean-Charles Rochet, "Dynamic financial contracting," *Advances in economics and econometrics*, 2013, *1*, 125–171. 2

_ , _ , Guillaume Plantin, and Jean-Charles Rochet, "Dynamic security design: Convergence to continuous time and asset pricing implications," *The Review of Economic Studies*, 2007, *74* (2), 345–390. 1

Bolton, Patrick and Mathias Dewatripont, *Contract theory*, MIT press, 2004. 1

Cetemen, Doruk, Felix Zhiyu Feng, and Can Urgun, "Renegotiation and dynamic inconsistency: contracting with non-exponential discounting," *Available at SSRN 3442367*, 2021. 1.1

Cvitanic, Jakša and Jianfeng Zhang, *Contract theory in continuous-time models*, Springer Science & Business Media, 2012. 2

Cvitanić, Jakša, Dylan Possamaï, and Nizar Touzi, "Dynamic programming approach to principal–agent problems," *Finance and Stochastics*, 2018, *22* (1), 1–37. 1

DellaVigna, Stefano and Ulrike Malmendier, "Contract design and self-control: Theory and evidence," *The Quarterly Journal of Economics*, 2004, *119* (2), 353–402. 1.1

DeMarzo, Peter M. and Yuliy Sannikov, "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," *The Journal of Finance*, 2006, *61* (6), 2681–2724. 1, 2, 4.1

Dicks, David L and Paolo Fulghieri, "Uncertainty, Contracting, and Beliefs in Organizations," *Kenan Institute of Private Enterprise Research Paper*, 2021, (19-1). 3

Edmans, A, X. Gabaix, and D Jenter, "Executive Compensation: A Survey of Theory and Evidence," *383-539*, 2017, (4). 1

Farhi, Emmanuel and Xavier Gabaix, "Optimal taxation with behavioral agents," *American Economic Review*, 2020, *110* (1), 298–336. 6

Gottlieb, Daniel and Xingtan Zhang, "Long-Term Contracting With Time-Inconsistent Agents," *Econometrica*, 2021, *89* (2), 793–824. 1.1

Grenadier, Steven R. and Neng Wang, "Investment under uncertainty and time-inconsistent preferences," *Journal of Financial Economics*, 2007, *84* (1), 2–39. 1.1

Hackbarth, Dirk, Alejandro Rivera, and Tak-Yuen Wong, "Optimal Short-Termism," *Management Science*, 2021, *74* (4), 1707–1751. 5

Harris, Christopher and David Laibson, "Dynamic choices of hyperbolic consumers," *Econometrica*, 2001, *69* (4), 935–957. 5

_ **and** _ , "Instantaneous Gratification," *The Quarterly Journal of Economics*, 2013, *128* (1), 205–248. 1.1, 2, 2.2

He, Zhiguo, "A model of dynamic compensation and capital structure," *Journal of Financial Economics*, 2011, *100* (2), 351–366. 5

Heidhues, Paul and Botond Kőszegi, "Exploiting naivete about self-control in the credit market," *American Economic Review*, 2010, *100* (5), 2279–2303. 1, 1.1, 3, 3

_ **and Philipp Strack**, "Identifying present bias from the timing of choices," *American Economic Review*, 2021, *111* (8), 2594–2622. 6

Hernández, Camilo and Dylan Possamaï, "Me, myself and I: a general theory of non-Markovian time-inconsistent stochastic control for sophisticated agents," *arXiv preprint arXiv:2002.12572*, 2020. 4

Holmström, Bengt, "Moral hazard and observability," *The Bell journal of economics*, 1979, pp. 74–91. 1

Koszegi, Botond, "Behavioral contract theory," *Journal of Economic Literature*, 2014, *52* (4), 1075–1118. 6

Laibson, David, "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443–478. 1, 2.1

___ **and Peter Maxted**, “The $\beta - \delta - \Delta$ sweet spot,” Technical Report, Tech. rep., Mimeo 2020. 2.2

___ **, ___ , and Benjamin Moll**, “Present bias amplifies the household balance-sheet channels of macroeconomic policy,” Technical Report, National Bureau of Economic Research 2021. 1.1

Malenko, Andrey, “Optimal dynamic capital budgeting,” *The Review of Economic Studies*, 2019, *86* (4), 1747–1778. 1.1

Maxted, Peter, “Present bias in consumption-saving models: A tractable continuous-time approach,” Technical Report, Tech. rep., Mimeo 2020. 4

Miao, Jianjun and Alejandro Rivera, “Robust contracts in continuous time,” *Econometrica*, 2016, *84* (4), 1405–1440. 1.1

O'Donoghue, Ted and Matthew Rabin, “Choice and Procrastination,” *The Quarterly Journal of Economics*, feb 2001, *116* (1), 121–160. 2.2

Parsons, Christopher A and Edward D Van Wesep, “The timing of pay,” *Journal of Financial Economics*, 2013, *109* (2), 373–397. 1

Piskorski, Tomasz and Alexei Tchistyi, “Optimal mortgage design,” *The Review of Financial Studies*, 2010, *23* (8), 3098–3140. 6

Sannikov, Yuliy, “A Continuous- Time Version of the Principal: Agent Problem,” *The Review of Economic Studies*, 2008, *75* (3), 957–984. 1, 3, 3

___ , “Contracts: The Theory of Dynamic PrincipalâAgent Relationships and the,” in “Advances in Economics and Econometrics: Volume 1, Economic Theory: Tenth World Congress,” Vol. 49 Cambridge University Press 2013, p. 89. 2

Shigeta, Yuki, “Quasi-Hyperbolic Discounting under Recursive Utility and Consumption-Investment Decisions,” *Available at SSRN 3672999*, 2020. 4

Smith, Adam, *The theory of moral sentiments*, Vol. 1, J. Richardson, 1822. 1

Spear, Stephen E and Sanjay Srivastava, “On repeated moral hazard with discounting,” *The Review of Economic Studies*, 1987, *54* (4), 599–617. 1

Stokey, Nancy and Robert Lucas, *Recursive methods in economic dynamics*, Harvard University Press, 1989. 1

Szydlowski, Martin and Ji Hee Yoon, "Ambiguity in dynamic contracts," *Journal of Economic Theory*, 2021, p. 105229. 3

Tella, Sebastian Di and Yuliy Sannikov, "Optimal asset management contracts with hidden savings," *Econometrica*, 2021, *89* (3), 1099–1139. 1

Tian, Yuan, "Optimal capital structure and investment decisions under time-inconsistent preferences," *Journal of Economic Dynamics and Control*, 2016, *65*, 83–104. 4

Wesep, Edward Dickersin Van, "Pay (be) for (e) performance: the signing bonus as an incentive device," *The Review of Financial Studies*, 2010, *23* (10), 3812–3848. 1

Williams, Noah, "On Dynamic Principal-Agent Problems in Continuous Time," Technical Report, UCLA Department of Economics 2004. 1

_ , "Persistent private information," *Econometrica*, 2011, *79* (4), 1233–1275. 1

Xu, Jin and Jun Yang, "Golden hellos: Signing bonuses for new top executives," *Journal of Financial Economics*, 2016, *122* (1), 175–195. 1