Renegotiation and Dynamic Inconsistency: Contracting with Non-Exponential Discounting*

Doruk Cetemen[†] Felix Zhiyu Feng[‡] Can Urgun[§]

February 2021

Abstract

This paper studies a continuous-time, finite-horizon contracting problem with renegotiation and dynamic inconsistency arising from non-exponential discounting. The problem is formulated as a dynamic game played among the agent, the principal and their respective future "selves", each with their own discount function. We identify the principal optimal renegotiation-proof contract as a Markov Perfect Equilibrium (MPE) of the game, prove such a MPE exists, and characterize the optimal contract via an extended Hamilton-Jacobi-Bellman system. We solve the optimal contract in closed-form when the discount functions of the selves are related by time difference, a property that is satisfied by common forms of non-exponential discounting such as quasi-hyperbolic discounting and anticipatory utility. In particular, quasi-hyperbolic discounting leads to a U-shaped action path and anticipatory utility leads to a hum-shaped path, both are qualitatively different from the monotonic action path that would arise under exponential discounting.

Keywords: continuous-time contracting, dynamic inconsistency, renegotiation, extended HJB system, non-atomic games

JEL code: D82, D86, D91

*Previously titled: *Contracting with Non-Exponential Discounting: Moral Hazard and Dynamic Inconsistency*. We thank Paulo Barelli, Jess Benhabib, Alessandro Bonatti, Gonzalo Cisternas, Dino Gerardi, Hari Govindan, Johannes Hörner, Ryota Iijima, Rohit Lamba, Erik Madsen, Roberto Marfè, Stephen Morris, Pietro Ortoleva, Romans Pancs, Wolfgang Pesendorfer, Debraj Ray, Bruno Strulovici, Juuso Toikka, and Leeat Yariv for useful discussions, and seminar participants at Bilkent University, Collegio Carlo Alberto, London School of Economics, New York University, Princeton University, Queen Mary University of London, University of Bristol, University of Warwick, Yale University, conference participants at ESSET(Gerzensee) 2019, NASMES 2019, SAET 2019, and SITE 2019 conference on Dynamic Games, Contracts and Markets for helpful comments.

[†]Collegio Carlo Alberto. Email: doruk.cetemen@carloalberto.org

[‡]University of Washington. Email: ffeng@uw.edu

[§]Princeton University. Email: curgun@princeton.edu

1 Introduction

Dynamic contracting explores how two parties can use inter-temporal incentives to mitigate agency frictions. Most existing contracting models assume inter-temporal preferences take the special form of exponential discounting. This assumption has the advantage of technical conveniences such as dynamic consistency and stationarity. However, there is extensive evidence – anecdotal and empirical – that "time preferences can be non-exponential" (see Laibson, 1997 and the literature that followed). For example, politicians may be more concerned about policies with immediate impact and less about those that take affect after the next election. Alternatively, when a group of non-dictatorial and dynamically-consistent decision makers (such as a board of directors, teams, committee of experts) follow unanimous decisions, Jackson and Yariv (2015) show that they collectively behave as a non-exponential, present-biased discounter. Therefore, it is natural to ask: how would dynamic contracting be different under non-exponential discounting?

This paper aims to provide a general answer to that question. To fix ideas, consider the following scenario between a CEO and a board of directors. The board has quasi-hyperbolic discounting (i.e. the $\beta - \delta$ preference in discrete time) and is thus present-biased, while the CEO has the standard exponential discounting. As usual, at the onset of their relationship the two parties consent to a long-term contract that provides the CEO the incentives to exert effort that is beneficial to the firm. At any point in time the board can renegotiate; offer a new contract that replaces the old one if the CEO agrees. If the CEO disagrees, the old contract stays in place, and the same protocol applies to any future possible renegotiation.

While renegotiation is widely observed in practice, it is not always explicitly considered among existing contracting models, with one reason being that it is not always necessary to do so when exponential discounting is used.[1] Indeed, Fudenberg, Holmstrom, and Milgrom (1990) and Rey and Salanie (1990) show that under reasonable assumptions, a full commitment contract is also renegotiation-proof if both contracting parties have exponential discounting. However, in the example of contracting between the CEO and the board, it is imperative that an optimal contract explicitly takes renegotiation into account when non-exponential discounting is present. A contract that appears "optimal" to the board given its current preference may appear sub-optimal in the future, not because of any inefficient punishments or constraints, but simply due to dynamic inconsistency arising from the board's present bias. This issue is further complicated by the fact that any renegotiation must be carried out by a future self of the board, but the future board's optimization problem is different from that of the current board. Thus, the usual recursive tools of dynamic

[1] Notable exceptions include DeMarzo and Sannikov (2006), Strulovici (2020), etc.

contracting, such as the Bellman operator, are no longer directly applicable.

In this paper, we develop novel tools to study the renegotiation problem and show that the optimal renegotiation proof contract is qualitatively different than that when all parties are time consistent. We begin with a standard continuous-time dynamic moral hazard model in which a principal hires an agent to manage a project over a finite horizon. The project's outcome is noisy, and the agent controls its drift with private actions. Departing from standard models, we allow the principal and the agent to have different, generic, time-varying, and non-exponential discount functions. Both the principal and the agent are sophisticated: they make decisions knowing exactly how their future preferences will change due to non-exponential discounting. Consequently, the dynamic contracting relationship can be formulated as an inter- and intra-personal game, in which each player is a time-t "self" of the principal or the agent.[2]

Under this general framework, we provide an intuitive definition for a *renegotiation-proof, incentive compatible, optimal long-term contract*. We formally characterize such a contract and prove its existence under fairly general conditions by providing a link between dynamic contracting and static non-atomic games. Our notion of renegotiation-proofness in this framework is equivalent to the equilibrium of the inter- and intra-personal games. As illustrated in the board-CEO example earlier, we allow the principal to offer a new contract at each moment in time. If that offer is rejected, the existing contract stays in place. At any moment, the time-t principal's offer takes into account all possible renegotiation in the future. The offer is evaluated based on the time-t preferences of the principal and the agent, given the preferences and equilibrium strategies of the future selves. Because both principal and agent are sophisticated, a contract that induces an equilibrium path of the game is renegotiation-proof (i.e. never renegotiated in the equilibrium). Nevertheless, renegotiation still plays a vital role because the possibility of renegotiation determines the equilibrium dynamics of the contract.

Unlike the CEO-board example, in the general model we allow the agent to be time inconsistent as well. Interestingly, this turns out to bear much less impact than a time-inconsistent principal does. Despite the fact that the dynamically inconsistent agent is also playing an intra-personal game with his future selves, the agent's incentive compatibility condition is characterized by a local pay-performance sensitivity similar to that in time-consistent benchmarks (e.g. Sannikov (2008)). The principal only needs to provide incentives

[2] We allow each self to have a different discounting function without assuming any kind of additional structure. This generalization captures a number of economic theories widely applicable in academic research and in practice. Examples include *present bias* (e.g. Ainslie (1975), Thaler (1981)), *anticipatory utility* (e.g. Loewenstein (1987), Caplin and Leahy (2001), Brunnermeier, Papakonstantinou, and Parker (2016)), and inherent value-attachment to specific dates (election dates, birthdays, anniversaries, etc).

for the current agent and not his future selves. This provides a much-needed simplification, allowing us to take the promised equilibrium value to the agent as a state variable and focus on the effect of dynamic inconsistency on the principal's side.

With the agent's problem solved, we present our main theorem characterizing the incentive-compatible, renegotiation-proof, optimal long-term contract in as broad generality as possible. The theorem represents two innovations. First, we demonstrate how to characterize the optimal contract using recursive techniques when the usual tool of Bellman equation is no longer applicable. We argue that given the agent's local incentive compatibility and our notion of renegotiation-proofness, the optimal contract is equivalent to the equilibrium of a simpler auxiliary game in which each t-self of the principal can only influence the agent's action, consumption, and the evolution of continuation utility at time t. The recursive characterization yields an *extended Hamilton-Jacobi-Bellman (HJB) system*. In contrast to the usual HJB equation in the time-consistent benchmark, the extended HJB system contains additional terms. These additional terms capture the equilibrium incentives of the principal, who takes into account the impact of the contract policies on the payoff for all her future selves.

The second innovation of our main theorem is the proof that such an optimal contract exists. This is known to be a thorny problem because the well-posedness of the extended HJB system is not well understood even in the mathematical literature.[3] We prove that such a system must have a solution by exploiting a unique connection between partial differential equations and static non-atomic games of incomplete information. This novel connection enables us to utilize known existence results in non-atomic games to establish existence of a solution to our extended HJB system. We also show that any equilibrium of the principal's auxiliary game must also be a contract that solves the extended HJB system. That is, our characterization of the optimal contract is complete. The generality of our framework opens the door for a broad range of applications.

We demonstrate the applicability of our general framework by solving a more specific class of contracting problems. For the purpose of simplification, we assume the agent has CARA utility and exponential discounting, and the set of discount functions of the principal belongs to a *time-difference family*: i.e., how the principal's t-self discounts a payoff at a future date s ($s > t$) is the same as how her $(t + k)$-self discounts a payoff at date $s + k$ ($\forall k > 0$).[4] These assumptions allow us to solve the optimal contract in closed-form and tease

[3] Harris and Laibson (2012) and Bernheim, Ray, and Yeltekin (2015) have shown the existence of solutions for special cases only. Björk, Khapko, and Murgoci (2017) and Yong (2012) laid down the extended HJB systems for more general cases but leave the crucial question of existence open and provide only a partial verification theorem.

[4] Most commonly-studied examples of non-exponential discounting in economics, such as hyperbolic dis-

out the specific effect brought by non-exponential discounting (on the principal's side). The results also suggest a strategy to identify potential time-inconsistent preferences in practice, which is usually difficult to observe directly. After the closed-form characterization with time difference families we analyze two special cases more closely.

Our first special case pertains to *quasi-hyperbolic* discounting. We adopt the model of Harris and Laibson (2012) in which a "hyperbolic drop" to the discounting function (β in the β-δ preferences in discrete time) occurs at an exponential rate. The exponential rate of arrivals enable us to consider the "survival probability" of the current self. In addition to the usual death rate that comes from discounting, there is an additional exponential replacement rate caused by the hyperbolic drop.[5] If this survival probability is low, the current self cannot credibly promise sufficient future payments as dynamic incentives, because those payments will be renegotiated by her impatient future self. In this case, a larger hyperbolic drop in discounting (lower β) decreases both the incentive power of the contract and the equilibrium wage to the agent. If the current principal's survival probability is high, and the contract horizon is sufficiently long, the current principal can front-load wage and incentive power in the contract before her more impatient future self arrives. A larger hyperbolic drop in discounting increases the prominence of such front-loading, causing the wage and incentive power to increase at the start of the contract. The front-loading of wage and incentive power leads to a *U-shaped* action path, in contrast to a monotonic path which arises if the principal has exponential discounting. As the horizon extends, the U-shaped path flattens. If the horizon approaches infinity, the contract under quasi-hyperbolic discount function converges to one under an exponential discount function.

Our second special case explores *anticipatory utility*. We adopt the model from Loewenstein (1987) assuming that the principal receives utility from the anticipation of future utility in addition to her current consumption. Such a utility function implies a "double counting" of future returns and yields a compounded discount factor. We find that under anticipatory utility the incentive power and equilibrium wage of the optimal contract is also potentially non-monotonic in time. Similar to quasi-hyperbolic discounting, the length of the contracting horizon shapes the equilibrium dynamics: when the contracting horizon is long enough, anticipation of future payoff leads to back-loaded wage and incentive power, resulting in a *hump-shaped* action path in contrast to the monotonic path under exponential discounting. As the horizon extends, the hump flattens, and the contract again converges to one under an exponential discount function if the horizon approaches infinity.

counting and anticipatory utility, belong in this category.

[5]Discounting corresponds to a "termination" whereas the hyperbolic drop is a mere "devaluation" by the rate of β. Thus, quasi-hyperbolic discounting is not a flat increase in impatience but a combination of two effects.

Literature Review: This paper is broadly related to several strands of research. First, it belongs to the thriving literature of continuous-time dynamic contracting. Benchmark models with time-consistent preferences such as DeMarzo and Sannikov (2006), Biais, Mariotti, Plantin, and Rochet (2007), Sannikov (2008) demonstrate the analytical convenience of a continuous-time formulation, which allows the derivation of economic insights from problems that are otherwise challenging to solve in discrete time.[6] Our paper differs in that we explicitly model dynamic inconsistency from non-exponential discounting for both the principal and the agent. Moreover, dynamic inconsistency necessitates a formal discussion of contract renegotiation, which is often ruled out or assumed away in benchmark models that assume full commitment power for the principal.

Our definition and analysis of renegotiation are related to both contract renegotiation in dynamic moral hazard problems as well as equilibrium renegotiation in dynamic games. Among the studies of contract renegotiation, the closest to our work are Strulovici (2020) and Watson, Miller, and Olsen (2019). The former explores renegotiation-proof contracts with infinite horizon and persistent private information in continuous time. In Strulovici (2020) renegotiation-proofness captures the following natural cognitive ability on behalf of the principal: if there was ever another contract that offered a different continuation given the same information (promised payoff and reported cash flow), the principal should not differ in her behavior and act consistently. In our setting, we model renegotiation-proofness as an interpersonal MPE, hence such a consistency notion is implied by the Markov perfection. The latter shares similarity with our model in terms of the procedure of renegotiation considered: once a renegotiation is proposed, if that proposal is rejected, the previous contract remains in place. The setting of Watson, Miller, and Olsen (2019) is a more general principal-agent setup that incorporates an explicit renegotiation phase where both parties have different bargaining weights. In comparison, we make the simplifying assumption that the principal has full bargaining power. Thus, our model also resembles Fudenberg and Tirole (1990), who explore contract renegotiation in a one-time setting with the principal having full bargaining power. In addition, our renegotiation-proof contract is characterized as a Markov Perfect Equilibrium, which is closely related to the notion of internal consistency in Bernheim and Ray (1989) for repeated games and Ray (1994) for dynamic games.

[6]Power of continuous time tools has been demonstrated in other problems, for instance, Faingold and Sannikov (2011) analyze a continuous-time reputation model and are able to prove the uniqueness of the equilibrium. Georgiadis (2014) demonstrates the tractability of continuous-time tools in delivering sharp comparative statics in public good contribution games – a intractable problem in discrete-time. Cisternas (2018) expands the linear career concerns model of Holmström (1999) to a general non-linear setup in continuous-time. Recently, Kolb and Madsen (2019) show that a continuous-time setting enables a tractable analysis of both multiple actions and imperfect monitoring in information design problems.

Our study is also related to contracting problems with behavioral preferences.[7] In particular, we assume the contracting parties are sophisticated regarding their time-inconsistency. This resembles Galperti (2015) which focuses on the optimal provision of commitment devices with sophisticated agents. In contrast, a substantial number of studies on behavioral preferences assume the contracting parties are naive or partially naive. For example, Gottlieb (2008) studies the optimal design of non-exclusive contracts and identifies different implications of immediate-costs goods and immediate-rewards goods for dynamically inconsistent but naive consumers. Gottlieb and Zhang (2020) study repeated contracting between a risk neutral firm and dynamically inconsistent but partially naive consumers, and find that at-will terminations may improve welfare if the level of dynamic inconsistency is sufficiently high. These studies focus on adverse selection problems (reporting), while we focus on a moral hazard (hidden effort) problem. Moreover, DellaVigna and Malmendier (2004) analyze the optimal two-part tariff of a firm facing a partially naive consumer with present-biased preferences. Heidhues and Kőszegi (2010) study a similar setting with naive agents and show that simple restrictions on contract forms may significantly improve welfare.

In general, the idea that preferences may be dynamically inconsistent as the result of non-exponential discounting can be traced to early work of Strotz (1955) and Pollak (1968). Since then, many studies have explored dynamically inconsistent preferences in various settings: including consumption-saving problems (Laibson, 1997; Krusell and Smith, 2003; Bernheim, Ray, and Yeltekin, 2015; Ray, Vellodi, and Wang, 2017; Bond and Sigurdsson, 2018; Cao and Werning, 2018); investment and asset allocation (Caplin and Leahy, 2001; Grenadier and Wang, 2007; Brunnermeier, Papakonstantinou, and Parker, 2016); monetary policy (Kydland and Prescott, 1977); fiscal policy (Halac and Yared, 2014); procrastination (O'Donoghue and Rabin, 1999), public finance (Bisin, Lizzeri, and Yariv, 2015; Harstad, 2016), etc. These studies are typically limited to a single party being dynamically inconsistent. One exception is Chade, Prokopovych, and Smith (2008), who provide a recursive characterization of a repeated game in which all players share the same quasi-hyperbolic discount function. In comparison, we allow generic, non-exponential, and potentially different discounting for all players, including quasi-hyperbolic discounting as a special case.

2 General Framework

In this section we present the general framework which introduces generic, non-exponential discounting functions into an otherwise standard dynamic principal-agent model. We lay out the foundations, solve the agent's problem, and discuss the precise role of commitment

[7] See Kőszegi (2014) and Grubb (2015) for a survey of contract theory with behavioral preferences.

and renegotiation in the model.

2.1 Basic Environment

A principal (she) contracts with an agent (he) over a fixed-time horizon $T < \infty$. Time is continuous and indexed by t. The principal is risk-neutral, with unlimited liability, and has a time zero outside option $\underline{V} \geq 0$. The agent is (weakly) risk-averse and has an outside option that he loses after time zero valued at $\underline{W}$.[8] The monetary cash flow, observable to both the principal and the agent, is given by

$$dM_t = \hat{a}_t dt + \sigma dZ_t, \quad \text{where } M_0 = 0, \tag{1}$$

where $\hat{a}_t$ is the agent's private action (effort), chosen from a compact set $\mathbb{A} := [\underline{a}, \bar{a}]$, $\sigma > 0$ is given. Z_t is a standard Brownian motion and $Z_0 = 0$.

Between time t and t' the principal's and the agent's discount rates are given by *self-indexed* functions $R^t(t')$ and $r^t(t')$, respectively. Such a specification allows for discounting to vary with both the current self t and an arbitrary future date t', not just the time difference. In particular, at time t, the principal uses the mapping defined by $R^t : [0, T] \to [0, 1]$ as the discount function. One can interpret $R^t(s)$ as if there are infinitely many discounting functions, one corresponding to each point t in time. Similarly, at time t, the agent uses the mapping defined by $r^t : [0, T] \to [0, 1]$. We assume that both the agent and the principal are sophisticated regarding their time preferences and the time preferences are common knowledge.

We impose the following conditions on $R^t(\cdot)$ and $r^t(\cdot)$:

Assumption 1 *For all $t \geq 0$, the discount functions $R^t(\cdot), r^t(\cdot)$ satisfies*

1. $R^t(s), r^t(s) = 1$ *for all* $s \leq t$ *and* $\lim_{s\to\infty} R^t(s) = 0$, $\lim_{s\to\infty} r^t(s) = 0$.
2. $R^t(s), r^t(s) > 0$ *for all* $s > t$.
3. $\int_t^\infty R^t(s)ds < +\infty$ *and* $\int_t^\infty r^t(s)ds < +\infty$ *a.s.*
4. *Regarded as bi-functions of t and s, $R^t(s)$ and $r^t(s)$ are uniformly Lipshitz continuous and three times differentiable in both arguments.*

The first condition states that current payoff is not discounted, while payoff in the infinitely far future has a present value of 0. The second condition states that any return in a finite future has some positive value, albeit potentially very small. The third condition ensures that

[8] Assumptions like this is common in the literature, see, e.g., Strulovici (2020).

the discounted value of any bounded stream of consumption remains finite. The last condition is a technical one precluding drastic changes and guaranteeing sufficient differentiability in the discount functions across and over time.

It is important to underline the difference between the general family of discount functions adopted here and a restricted family where the discounting functions of different selves are transformations of each other based on the time difference, a time-difference family. Exponential discounting is a particular time-difference family. In particular, the familiar form $e^{-\gamma(s-t)}$ for some $\gamma > 0$ does not depend on s or t but only on $s - t$. The functions explored here can readily capture discount functions that are tied by time difference when the following additional restriction is imposed $r^t(s) = r^{t+k}(s+k) = r(s-t)$ and $R^t(s) = R^{t+k}(s+k) = R(s-t)$ for all t, s, k. In such a case, we use the term discount function to refer to $r(s-t)$ and $R(s-t)$.

Denote the probability space as $(\Omega, \mathcal{F}, \mathcal{P})$ and the associated filtration generated by the cash flow process as $\{\mathcal{F}_t\}_{t\in[0,T]}$. Contingent on the filtration, a contract $\mathcal{C}$ specifies a consumption process $\{c_t\}_{t\in[0,T]}$ with final payment c_T to the agent and a sequence of recommended actions $\{a_t\}_{t\in[0,T]}$, in order to remain with the same utility function, we make the use of the convention that at the end of contracting period the agent gets a utility from the final payment c_T as if taking action $\underline{a}$. All quantities are assumed to be square integrable and progressively measurable under the usual conditions.

The agent's utility depends on his action a and consumption c received from the contract. We assume his utility function satisfies the following:

Assumption 2 *The agent has a weakly risk averse utility function of instantaneous consumption and action, $u(c, a)$, such that*

1. *u is continuous in both arguments and twice differentiable with bounded derivatives.*

2. *u is convex and decreasing in a, concave and increasing in c.*

3. *$\lim_{c\to\infty} u(c, a) = \infty$ and $\lim_{c\to-\infty} u(c, a) = -\infty$ for all $a \in \mathbb{A}$, $u(\cdot, a)$ is invertible with a continuous inverse function for all $a \in \mathbb{A}$.*

The first two requirements are fairly standard. The final requirement is largely a technical one that facilitates the proof of the general existence theorem (Theorem 1). It is sufficient but certainly not necessary. For special cases in which more specific assumptions of other elements of the model are made, our results can be derived under more generalized utility functions of the agent. We do not allow the agent to save unless stated otherwise.

Given a contract $\mathcal{C}$ which is followed by the agent, the agent's time 0 expected payoff is

$$E^{\mathcal{C}}\left[\int_0^T r^0(s)u(c_s,a_s)ds + r^0(T)u(c_T,\underline{a}_T)\right].$$

The principal's time 0 payoff equals the expected cash flows net of the consumption process offered to the agent, when the agent follows the contract payoff of principal is:

$$E^{\mathcal{C}}\left[\int_0^T R^0(s)\left(dM_s - c_s ds\right) - R^0(T)c_T\right] = E^{\mathcal{C}}\left[\int_0^T R^0(s)\left(a_s - c_s\right)ds - R^0(T)c_T\right]. \tag{2}$$

Because both the agent and the principal can have time-inconsistent preferences, for each point in time $t > 0$, we must treat the principal and the agent as different players (selves), who we refer to as "agent's t-self" and "principal's t-self". To characterize the contract, we define the agent's *t-self continuation utility*, the agent's *incentive compatibility*, and the principal's *t-self continuation utility under an incentive compatible contract* as the following:

Definition 1 *Given any contract $\mathcal{C} = (a_t, c_t)_{t\in[0,T]}$, and agent's action stream $\hat{a}$, the agent t's continuation utility is defined as $\tilde{W}(t,t)$, where first t denotes time and second t denotes self.*

$$\tilde{W}(t,t) = E_t^{\mathcal{C}}\left[\int_t^T r^t(s)u(c_s,\hat{a}_s)ds + r^t(T)u(c_T,\underline{a}_T)\middle|\mathcal{F}_t\right], \tag{3}$$

where $\mathcal{F}_t$ is the natural filtration generated by Z_t. In a similar fashion, agent t's continuation utility from time $k > t$ onwards from t perspective is given by:

$$\tilde{W}(k,t) = E_t^{\mathcal{C}}\left[\int_k^T r^t(s)u(c_s,\hat{a}_s)ds + r^t(T)u(c_T,\underline{a}_T)\middle|\mathcal{F}_t\right].$$

A contract $\mathcal{C}$ is incentive compatible if at each instant t, the effort choice of agent's t-self is the suggested action by the principal, $\hat{a}_t = a_t$, and maximizes $\tilde{W}(t,t)$ given that all other selves of the agent follow the actions recommended by the contract $\mathcal{C}$.

Under an incentive compatible contract $\mathcal{C}$, the principal's t-self continuation utility is defined as $f(t,t)$, where first t denotes time and second t denotes self.

$$f(t,t) = E_t^{\mathcal{C}}\left[\int_t^T R^t(s)\left(a_s - c_s\right)ds - R^t(T)c_T\middle|\mathcal{F}_t\right]. \tag{4}$$

In a similar fashion principal t's continuation utility from time $k > t$ onwards from t per-

spective is given by:

$$f(k,t) = E_t^{\mathcal{C}} \left[\int_k^T R^t(s) \left(a_s - c_s\right) ds - R^t(T) c_T \middle| \mathcal{F}_t \right].$$

2.2 The Agent's Problem

We first solve the agent's problem. In general, the agent's consumption at any time could depend on the entire path of the outputs. As in Bismut (1978) and Williams (2015) we resolve the history dependence by taking the key state variable to be the density of the output process rather than the output process itself. In particular, let Υ be the space of continuous functions mapping $[0,T]$ into $\mathbb{R}$. Let Z_t^0 be a Brownian motion on Υ, which can be interpreted as the distribution of output resulting from an effort policy which makes output a martingale. The following lemma characterizes the continuation utilities and the incentive compatibility condition:

Lemma 1 *Under Assumptions 1 and 2, given any contract $\mathcal{C} = (a_t, c_t)_{t\in[0,T]}$ and any sequence of the agent's choices, there exists a flow of processes $\tilde{\psi}(s,t)$ and an equilibrium value process W such that*

$$\tilde{W}(k,t) = r^t(T) u(c_T, \underline{a}_T) - \int_k^T r^t(s) u(c_s, \hat{a}_s) ds + \int_k^T \tilde{\psi}(s,t) dZ_s^a \tag{5}$$

for each t-self of the agent. The equilibrium value process satisfies

$$W(t) = r^t(T) u(c_T, \underline{a}_T) - \int_t^T r^t(s) u(c_s, \hat{a}_s) ds + \int_t^T \tilde{\psi}(s,t) dZ_s^a. \tag{6}$$

Where the Brownian Motion (Z^a) is defined via Girsanov theorem in the following way:

$$Z_t^a = Z_t^0 - \int_0^t \frac{a_s}{\sigma} ds.$$

The contract is incentive compatible if and only if

$$\psi_t = \frac{\tilde{\psi}(t,t)}{\sigma} = u_a(c_t, a_t). \tag{IC}$$

Compared to benchmark models with time-consistent preferences (e.g. Sannikov (2008)), the implications of Lemma 1 are two-fold: first, unlike the benchmark, the time-inconsistent preference implies that the agent's continuation utility becomes a *flow* of processes, one for each t-self of the agent. $\tilde{W}(k,t)$ represents the value of continuation from period k onward for the t-self. Setting $k = t$ yields $W(t)$, the continuation utility of the t-self agent, taking into account not only the change in time but also the change in the agent's future preferences. Consequently, $W_t \in \mathbb{R}$ which is the realization of $W(t)$ at time t is the relevant state variable for the principal when designing the optimal contract. Second, despite the existence of future selves, the IC condition (IC) is a local constraint similar to the time-consistent benchmark. This is because the agent's action does not have a persistent effect, and each self wants to maximize their own utility.

Remark 1 *The resemblance of the IC condition does not imply that the derivation of the continuation utility is a trivial task. On the contrary, the evolution of the agent's continuation utility is very different from Sannikov (2008) or other time-consistent benchmarks. More specifically, for a fixed self t, $\tilde{W}(k,t)$ is a stochastic process. As the agent's selves evolve over time, $\tilde{W}(k,t)$ becomes a flow of stochastic processes. As a result, instead of a single backward stochastic differential equation that arises in a dynamically consistent environment (as in Cvitanic and Zhang, 2012), the representation in equation (5) is a backward stochastic Volterra integral equation. In the Appendix, we specify the details of the derivation of $\tilde{W}(k,t)$, and show that the resulting process $W(t)$ can also be represented by a diffusion process. This diffusion is constructed via the decoupling field of the integral equation (6), which highlights the connection between the discretized model and the continuous time model. Intuitively, assume the contracting horizon is divided into a discrete set of intervals, and each self of the agent is the agent at the beginning of an interval. Thus, there are only finitely many selves, each with non-negligible impact, and an associated continuation value for each self. As the size of the interval becomes smaller, pasting those continuation values together following the technique in Wang and Yong (2019) yields the decoupling field used to construct the diffusion process that represents $W(t)$.*

2.3 Renegotiation and Commitment

The introduction of time-inconsistent preferences necessitates some formal discussions of the principal's ability to commit to the contract. In particular, a future self of the principal with time-inconsistent preferences may find full commitment sub-optimal, as the resulting contract would prevent potentially beneficial alterations, once the principal's own preference has changed over time. To address this issue, we allow *contract renegotiation* in our model. The principal can offer a new long-term contract at any point in time. If the agent rejects the newly offered contract, the old contract stays in place, but can still be changed in the future by a different self (as in Watson, Miller, and Olsen, 2019). We further assume that the agent breaks indifference in favor of acceptance.

A renegotiation is called *feasible* if the agent is willing to accept the contract given his preferences at the time and the contract is incentive compatible to all future selves. Formally:

Definition 2 *An incentive compatible contract $\mathcal{C}'$ is a feasible renegotiation of an incentive compatible contract $\mathcal{C}$ at time t valued at W given information $\mathcal{F}_t$ if the following inequality holds*

$$E_t^{\mathcal{C}'}\left[\int_t^T r^t(s)u(c_s', a_s')ds + r^t(T)u(c_T', \underline{a}_T)\middle| \mathcal{F}_t\right] \\ \geq E_t^{\mathcal{C}}\left[\int_t^T r^t(s)u(c_s, a_s)ds + r^t(T)u(c_T, \underline{a}_T)\middle| \mathcal{F}_t\right] = W.$$

We will call $\mathcal{C}'$ a feasible renegotiation of $\mathcal{C}$ at time t that promises $W' \geq W \in \mathbb{R}$ if, in addition,

$$E_t^{\mathcal{C}'}\left[\int_t^T r^t(s)u(c_s', a_s')ds + r^t(T)u(c_T', \underline{a}_T)\middle| \mathcal{F}_t\right] = W'.$$

Based on this definition, the new set of actions and consumption of each feasible renegotiation contract must generate weakly higher continuation utility for the agent.[9] Notice that a renegotiation could increase not only the current selves payoff but also payoff of the all or some future selves. Since future selves payoff and the current selves payoff are valuations of the same stream of consumptions and actions albeit with different weights. However, our next proposition shows that, for each feasible and incentive compatible contract $\mathcal{C}'$, there exists a *principal-preferred alternative* contract $\mathcal{C}''$ that is feasible and incentive compatible, induces the same actions and prescribes the same consumption to the agent as $\mathcal{C}'$ does, but keeps the currently renegotiating agent's continuation utility unchanged by lowering only the final consumption c_T. Formally:

Proposition 1 *Let $\mathcal{C}$ be an incentive compatible contract. For any feasible renegotiation $\mathcal{C}'$ of $\mathcal{C}$ at time t given information $\mathcal{F}_t$, there exists a principal-preferred alternative $\mathcal{C}''$ such that $(c_t', a_t') = (c_t'', a_t'')$ for all $t < T$ and*

$$E_t^{\mathcal{C}}\left[\int_t^T r^t(s)u(c_s, a_s)ds + r^t(T)u(c_T, \underline{a}_T)\middle| \mathcal{F}_t\right] = E_t^{\mathcal{C}''}\left[\int_t^T r^t(s)u(c_s'', a_s'')ds + r^t(T)u(c_T'', \underline{a}_T)\middle| \mathcal{F}_t\right].$$

[9]Our notion of renegotiation defines an incentive compatible contract for the remainder of the contractual relationship. This is a less restrictive notion than the recursively defined set of feasible renegotiations (as in, e.g. Benoit and Krishna, 1993), which explicitly includes the strategies of the principal and the agent's future selves. However, imposing such equilibrium restrictions on the *set* of feasible renegotiations does not change the renegotiation-proof contract because of the focus on a principal intra-personal equilibrium. More discussion can be found in Section 3.1 after Definition 5.

Moreover, all selves of the principal weakly prefer $\{c''_t, a''_t\}_{t\in[0,T]}$ *over* $\{c'_t, a'_t\}_{t\in[0,T]}$.

Proposition 1 implies that there is no *Pareto improvement* by offering the agent higher continuation utility. Furthermore, because $\mathcal{C}''$, the principal's-preferred alternative, induces the same paths of actions and prescribes the same consumption as $\mathcal{C}'$ does except for a lower final consumption c''_T, $\mathcal{C}''$ is "preferred" not only by the current principal but all the principal's past and future selves.

Proposition 1 is a generalization of Theorem 1 in Fudenberg, Holmstrom, and Milgrom (1990) with a similar intuition: while a renegotiation can induce different paths of consumption and actions with different probabilities, the principal-preferred alternative keeps those paths and their probabilities unchanged while decreasing the final payments on each path so that the final *utility* decreases by a fixed amount. This does not change the incentives for any self of the agent because his ranking of the utilities remain the same. The principal can reduce the promised payments this way until it hits the level of the existing contract. This increases the payoff to all principals future selves, because the only payoff that is in the utility function of every principal's future selves is the final payoff.

In the case of a dynamically consistent preference, Proposition 1 is sufficient to imply that there is no renegotiation with a higher payoff to the agent that achieves a Pareto improvement. In the case of a dynamically inconsistent preference, we need to also consider the strategic incentives for renegotiation. That is, whether a renegotiation if accepted, can potentially alter the renegotiations among the principal and agent's future selves. However, our next result, which is not covered in Fudenberg, Holmstrom, and Milgrom (1990), demonstrates that renegotiation can not achieve any *strategic improvement* for the principal and the agent either: neither of them can utilize renegotiation to expand or restrict the set of implementable actions and consumptions available to their future selves. Formally:

Proposition 2 *Given any value of continuation utility* $W \in \mathbb{R}$*, let* $\xi_t(W)$ *represent the set of feasible distributions over* $\{a_s, c_s\}_{s\in[t,T)}$ *that can be generated by a feasible renegotiation which promises the same value* W *to the* t *self of the agent. For any* $t \in [0, T)$ *and any arbitrary pair of* W', W''*,* $(W' \neq W'')$ *the set* $\xi_t(W')$ *is equal to* $\xi_t(W'')$*,* $\xi_t(W') = \xi_t(W'')$.

Proposition 2 argues that regardless of how a feasible renegotiation changes the level of continuation utility of the time-t agent, the available set of actions and consumption paths that are incentive compatible in the future until T remains unchanged. Therefore, each principal's time-t self *cannot* use renegotiation to strategically influence the behavior and policies of her future selves. Similarly, each agent's time-t self *cannot* use the acceptance or rejection of a renegotiation to strategically influence the behavior and policies that the principal's future selves can offer.

The intuition for Proposition 2 resembles the argument in Proposition 1. Suppose a principal's current self is considering a feasible renegotiation that is not principal-preferred for the purpose of forestalling a particular renegotiation in the future. Similar to Proposition 1 what the principal's future self could do is simply increasing the final payment on every possible path such that the continuation payoff at the time of renegotiation reaches this higher promise. That way, the future self can implement the consumption and action path she desires regardless of the continuation utility carried over, and she would indeed do so. Thus, any renegotiation that is not principal-preferred reduces the payoff for all of the principal's future selves without limiting the set of feasible action paths for any of her future selves.

The absence of this strategic role of renegotiation implies that in our setup, without loss of generality, we can limit our attention to principal-preferred renegotiations. That is, we only need to consider renegotiations that do not change the agent's continuation utility, only how such continuation utility is delivered. Any other renegotiation that would result in a strictly higher continuation utility to the agent is never optimal for the principal and therefore is never offered in any principal optimal renegoatiation.

Remark 2 *Besides applying our framework, one can address the issue of renegotiation in alternative, albeit more restrictive settings. One alternative is to directly rule out renegotiation by imposing the assumption of full commitment. In this case, the contract is never renegotiated but is also not necessarily renegotiation-proof, because the principal's future selves may evaluate the same payoff streams differently than her time-0 self. Another alternative is to construct a dynamically-consistent preference from non-exponential discounting, in which case the contract is renegotiation-proof because what is optimal to the principal at time-zero remains optimal to her future selves. However, this setting requires specific weights of the discount functions of each selves on different dates. Formally, it requires* $R^t(s)/R^t(s+k) = R^{t'}(s)/R^{t'}(s+k)$ *and* $r^t(s)/r^t(s+k) = r^{t'}(s)/r^{t'}(s+k)$ *for all* $t, t', s,$ *and* k. *In contrast, our objective is to explore the optimal long-term contract in a general setting in which the need for renegotiation is a natural consequence of dynamic inconsistency.*

3 The Optimal Contract

In this section, we define and solve the optimal contract in several steps. First, we establish a formal definition of what the "optimal contract" refers to in our setting. Next, we provide a heuristic derivation of the extended HJB system that characterizes such a contract. Finally, we state our main theorem regarding the characterization and existence of the optimal contract, and compare it to the optimal contract in dynamically consistent benchmarks.

3.1 Intra-personal Game and the Optimal Equilibrium Contract

In this subsection, we establish a formal and economically appropriate notion of what "optimal contract" refers to in this paper. First and foremost, dynamic inconsistency and the possibility of contract renegotiation imply that the contracting problem can be formulated as an intra-personal game played by the different selves of the principal and the agent. Each time-t self is a player of a game in which the players take into into account how the contract and their actions will affect the decision of all future selves. In light of Lemma 1, we focus on the Markov Perfect Equilibrium (MPE) of the intra-personal game with W_t (promised payoff) being the state variable. That is, we look for an equilibrium in which the the principal's t-self's strategy is a mapping of W_t into a contract — a contingent path of recommended actions and consumption that promises W_t – and the agent's t-self's strategy is to reject or accept the contract and choose an action. Formally,

Definition 3 *Let $\mathfrak{S}_T$ denote the set of $\mathcal{F}_T$ measurable random variables, a Markov strategy of principal's t-self is a mapping $s_t^P : \mathbb{R} \to \xi_t(W_t) \times \mathfrak{S}_T$, with the convention that $s_0^P : \mathbb{R} \to \xi_0(\underline{W}) \times \mathfrak{S}_T$.*[10] *A Markov strategy of agent's t-self's is a mapping $s_t^A : \mathbb{R} \to \{0,1\} \times [a, \bar{a}]^2$, where $\{0,1\}$ represents whether the t-self agent rejects or accepts the contract offered.*

The MPE therefore follows the standard definition of mutual best responses:

Definition 4 *$\{s_t^P, s_t^A\}_{t\in[0,T]} = (s^P, s^A)$ is a Markov Perfect Equilibrium (MPE) if s_t^P is a best response to (s^P, s^A) for all W_t and s_t^A is a best response to (s^P, s^A) for all W_t.*

For any equilibrium tuple of strategies $\{s_t^P, s_t^A\}_{t\in[0,T]} = (s^P, s^A)$, there is a *realized* contract $\mathcal{C}_{(s^P,s^A)}$ resulting from the path of such equilibrium. For simplicity, unless stated otherwise, we use the resulting contract $\mathcal{C}, \mathcal{C}'$ to also represent their different corresponding strategies in the subsequent analysis. Notice that $\mathcal{C}_t$, a contract at time, may differ from s_t^P, because multiple strategies can induce the same equilibrium process. Nevertheless, the most natural and well-studied equilibrium is the one in which the principal at time-0 offers the equilibrium process directly as the contract, so that no further renegotiation is necessary. Such an equilibrium is equivalent to the principal offering an (incentive compatible) *renegotiation-proof* contract, defined as follows:

Definition 5 *A contract $\mathcal{C}$ is a renegotiation-proof contract if $\mathcal{C} = C_{(s^P,s^A)}$ for a MPE (s^P, s^A) and $s_t^P = \mathcal{C}$ for all t and $\mathcal{F}_t$.*

[10] Notice the set $\xi_t(W_t)$ captures the contract up to but not including T, hence we augment the final payoff which is useful for defining principal preferred renegotiations with the final payment. In other words a strategy maps a continuation payoff $W \in \mathbb{R}$ to a feasible continuation contract and a final payment.

Definition 5 is the notion of renegotiation-proofness used in our model. Given Propositions 1 and 2, it is sufficient to verify that a contract is renegotiation-proof as long as there is no deviation that the principal prefers. This implies that a renegotiation-proof contract corresponds to a MPE, because it satisfies the intra-personal equilibrium of the principal, and is incentive compatible for all selves of the agent.[11] Altogether, these observations imply that the intra-personal game among the different selves of the principal and the agent can be simplified to an auxiliary game. In this game, the time-t principal takes the equilibrium continuation process as given, and chooses the contract terms assuming she only affects the process at time-t subject to the incentive compatibility constraint identified in Lemma 1.

Next, we denote the continuation payoff of the time-t principal's from period t onward ($f(t,t)$) given a tuple of MPE strategies (s^P, s^A) under contract $\mathcal{C}$ with a current promised value W as $f_{\mathcal{C}}(t,W,t)$,

$$f_{\mathcal{C}}(t,W,t) = E^{\mathcal{C}}_{t,W}\left[\int_t^T R^t(s)\,(a_s - c_s)\,ds + B(t,W_T)\right] \tag{7}$$

where $B(t,W_T) \equiv -R^t(T)c_T$.

An *equilibrium contract* of this auxiliary game is given as follows:

Definition 6 *Consider an incentive compatible contract process $\mathcal{C}$. For any initial point and state (t,W) and a "small" increment of time Δ, define a "deviation" strategy $\mathcal{C}_\Delta$ as:*

$$\mathcal{C}_\Delta = \begin{cases} \hat{\mathcal{C}}(t',W') & \text{for } t \le t' < t+\Delta \\ \mathcal{C}(t',W') & \text{for } t+\Delta \le t' \le T \end{cases}$$

*where $\hat{\mathcal{C}} \neq \mathcal{C}$ is another incentive compatible contract, and W' is the state corresponding to the future date t'. $\mathcal{C}$ is called an **equilibrium contract** if*

$$\liminf_{\Delta\to 0} \frac{f_{\mathcal{C}}(t,W,t) - f_{\mathcal{C}_\Delta}(t,W,t)}{\Delta} \ge 0 \tag{8}$$

*for all $\hat{\mathcal{C}}$. For an equilibrium contract $\mathcal{C}$, the corresponding value function $f_{\mathcal{C}}(t,W,t)$, is called as an **equilibrium value function**.*

The equilibrium contract in Definition 6 is also renegotiation-proof according to Definition 5.[12] The local optimality corresponds to the best response of each self to the choices of the

[11]This is because our definition of feasible renegotiation (Definition 2) incorporates incentive compatibility. Therefore, any renegotiated contract will be actually carried out by the agent's future selves and yields the correct continuation utility for the agent's current self. This holds even without the Markov refinement we impose in Definition 5.

[12]This notion of equilibrium is an extension of the so-called "closed loop" controls introduced Ekeland

other selves and thus characterizes the payoffs each self receives in the intra-personal game. Extending the definition of the equilibrium contract here to include global deviations would introduce joint actions by multiple selves and yield a cooperative solution, as opposed to the non-cooperative equilibrium in Definition 5. Such global deviations require different and unorthodox notions of sophistication where some selves form and maximize the payoff of coalitions. In our setting, the optimal contract boils down to the contract that maximizes the principal's continuation payoff at each point of time among the set of renegotiation-proof, incentive compatible equilibrium contracts $\mathcal{C}$. Formally,

Definition 7 *Let $\mathfrak{C}$ denote the set of equilibrium contracts. A contract $\mathcal{C} \in \mathfrak{C}$ is* ***optimal*** *if it maximizes $f_{\mathcal{C}}(t, W, t)$ for all t, W.*

According to this definition, we can write the value of an optimal equilibrium contract to the principal as follows:

$$V(t, W) = \sup_{\mathcal{C} \in \mathfrak{C}} f_{\mathcal{C}}(t, W, t) = \sup_{\mathcal{C} \in \mathfrak{C}} E_{t,W}^{\mathcal{C}} \left[\int_t^T R^t(s) \left(a_s - c_s \right) ds + B(t, W_T) \right]. \tag{9}$$

In contrast to the problem of a single decision maker, the value function in (9) represents optimality as a notion of an equilibrium, not simply what a single self of the principal deems optimal. In other words, Definition 6 introduces the notion of equilibrium, and 7 defines the optimal contract as the Pareto-efficient equilibrium. Note that W_t here is an endogenous, equilibrium object, as opposed to an exogenous process (as in Björk, Khapko, and Murgoci (2017)).

In summary, this subsection establishes a formal and economically appropriate notion of what "optimal contract" refers to in this paper. To resolve the issue of dynamic inconsistency, we meticulously laid out several critical definitions and results to reach the conclusion that the most appropriate definition of the optimal contract is a *principal-optimal equilibrium contract* (as in Definitions 6 and 7), which is also renegotiation-proof (as in Definition 5). In the subsequent analysis, we simply refer to this contract as the *optimal contract* or the *equilibrium contract.*

3.2 Optimal Contract: Heuristic Derivation

This subsection provides a heuristic derivation of the optimal contract with a moderate amount of technicality. The main purposes are three-fold: first, to introduce some notations

and Lazrak (2010) by Björk et al. (2017). Studies such as Wei et al. (2017) and Li et al. (2017) have shown that along the trajectory such controls remain time-consistent and locally optimal, and hence constitute an equilibrium.

that are not required in standard, time-consistent dynamic programming but are necessary when time-inconsistency is involved. Second, to remind readers the standard HJB equation for a time-consistent benchmark under our notations. Third, to heuristically establish the extended HJB systems that are critical for characterizing the optimal contract under time-inconsistency, and compare them to the standard HJB equation. The formal characterization of the optimal contract is given in Theorem 1 in Subsection 3.3.

3.2.1 Notations

Let us first introduce a few notations frequently used in this and the subsequent sections. We have defined the equilibrium value function $f_{\mathcal{C}}(t, W, t)$ in (7) and the optimal equilibrium value function $V(t, W)$ in (9), both of which are equilibrium objects. In particular, a change in t moves both the self and the time in those objects forward. However, when a fixed self makes forward-looking decisions, it is important to keep track of how the continuation value evolves over time for that particular self. We thus distinguish a changing self from a changing time for a fixed self by treating the third entry in $f_{\mathcal{C}}(t, W, t)$ as a parameter, denoted by a superscript $f^t_{\mathcal{C}}(t, W)$, to highlight that the self is fixed. That is, for a fixed self t, the continuation payoff under a contract $\mathcal{C}$ from time t' onward is denoted as

$$f^t_{\mathcal{C}}(t', W) = E^{\mathcal{C}}_{t',W}\left[\int_{t'}^{T} R^t(s)\,(a_s - c_s)\,ds + B(t, W_T)\right]. \tag{10}$$

While $f^t_{\mathcal{C}}(t, W)$ looks very similar to $f_{\mathcal{C}}(t, W, t)$ (which coincides with $V(t, W)$ under the optimal contract), they represent different objects in our framework and respond differently to changes in the contract. Specifically, following any changes in the contract $\mathcal{C}$, $f^t_{\mathcal{C}}(t + dt, W_{t+dt})$ captures how the t-self's own valuation from time $t + dt$ onward changes; $f_{\mathcal{C}}(t + d_t, W_{t+dt}, t+dt)$ captures how the $t+dt$'th self's valuation in the current equilibrium changes; and $V(t+dt, W_{t+dt})$ captures how the principal's optimal equilibrium value from $t+dt$ onward changes.

Let $f_{\mathcal{C},t}(t, W, t)$ denote the total derivative with respect to t, Ito's lemma implies:

$$\begin{aligned}
f^t_{\mathcal{C}}(t + dt, W_{t+dt}) &= f^t_{\mathcal{C},t}(t, W) + \mu^{\mathcal{C}}_W f^t_{\mathcal{C},W}(t, W) + \frac{1}{2}(\sigma^{\mathcal{C}}_W)^2 f^t_{\mathcal{C},WW}(t, W),\\
f_{\mathcal{C}}(t + dt, W_{t+dt}, t + dt) &= f_{\mathcal{C},t}(t, W, t) + \mu^{\mathcal{C}}_W f_{\mathcal{C},W}(t, W, t) + \frac{1}{2}(\sigma^{\mathcal{C}}_W)^2 f_{\mathcal{C},WW}(t, W, t),\\
V(t + dt, W_{t+dt}) &= V_t(t, W) + \mu^{\mathcal{C}}_W V_W(t, W) + \frac{1}{2}(\sigma^{\mathcal{C}}_W)^2 V_{WW}(t, W),
\end{aligned}$$

where $\mu^{\mathcal{C}}_W$ corresponds how principal controls the drift of the W and similarly $\sigma^{\mathcal{C}}_W$ corresponds how principal controls the volatility of W. For ease of exposition, we adopt the notation $\mathcal{A}^C$, known in the literature as the *controlled infinitesimal generator*, to simplify the functions above. That is,

$$\mathcal{A}^{\mathcal{C}} f^t_{\mathcal{C}}(t, W) \equiv f^t_{\mathcal{C}}(t + dt, W_{t+dt}), \tag{11}$$

$$\mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}(t, W, t) \equiv f_{\mathcal{C}}(t + dt, W_{t+dt}, t + dt), \tag{12}$$

$$\mathcal{A}^{\mathcal{C}} V(t, W) \equiv V(t + dt, W_{t+dt}). \tag{13}$$

Finally, if the principal implements a different contract $\tilde{\mathcal{C}}$ at a single point of time t, we use $\mathcal{A}^{\tilde{\mathcal{C}}} f_{\mathcal{C}}(t, W, t)$ and $\mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\mathcal{C}}(W, t)$ to capture the effect of such deviation from the contract $\mathcal{C}$ on $f_{\mathcal{C}}(t, W, t)$ and $f^t_{\mathcal{C}}(W, t)$, respectively.

3.2.2 Reminder: A Time-consistent Benchmark

Consider a simple benchmark of an infinite horizon dynamic contracting problem with exponential discounting. Using our notations, we can state the principal's optimization problem as follows:

$$V(t, W) = \sup_{\mathcal{C} \in \mathfrak{C}} E^{\mathcal{C}}_{t,W} \left[\int_t^\infty e^{-rs}(a_s - c_s) ds \right].$$

where W_t is the agents continuation utility process with value W at time t (see for example Sannikov (2008)). The optimal contract can be summarized using the standard HJB equation:

$$rV(t, W) = \sup_{\tilde{\mathcal{C}} \in \mathfrak{C}} \left\{ (\tilde{a}_t - \tilde{c}_t) + \mu^{\tilde{\mathcal{C}}}_W V_W + \frac{1}{2} \left(\sigma^{\tilde{\mathcal{C}}}_W \right)^2 V_{WW} \right\}.$$

Importantly, the HJB equation relies on the fact that when *both* the process W_t and the discounting function e^{-rt} are time-homogeneous, i.e., $\partial V / \partial t = -rV(t, W)$. If time is also a state variable of the contracting problem, a more general HJB equation takes the following form:

$$0 = \sup_{\tilde{\mathcal{C}} \in \mathfrak{C}} \left\{ (\tilde{a}_t - \tilde{c}_t) + V_t + \mu^{\tilde{\mathcal{C}}}_W V_W + \frac{1}{2} \left(\sigma^{\tilde{\mathcal{C}}}_W \right)^2 V_{WW} \right\}.$$

In our notation, the HJB equation for a standard time-consistent dynamic contracting problem can be written as

$$0 = \sup_{\tilde{\mathcal{C}} \in \mathfrak{C}} \left\{ (\tilde{a}_t - \tilde{c}_t)\, dt + \mathcal{A}^{\tilde{\mathcal{C}}} V(t, W) \right\}. \tag{14}$$

3.2.3 Time-inconsistency: The Extended HJB System

We now turn to our setting, observe that by construction we must have

$$
\begin{aligned}
f_{\mathcal{C}}(t, W, t) &= (a_t - c_t)\, dt + E^{\mathcal{C}}_{t+dt, W_{t+dt}} \left[\int_{t+dt}^{T} R^t(s)\,(a_s - c_s)\, ds + R^t(T)\, B(t, W_T) \right] \\
&= (a_t - c_t) + f_{\mathcal{C}}(t + dt, W_{t+dt}, t).
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
f_{\mathcal{C}}(t + dt, W_{t+dt}, t + dt) &= f_{\mathcal{C}}(t + dt, W_{t+dt}, t + dt) + f_{\mathcal{C}}(t, W, t) - f_{\mathcal{C}}(t, W, t) \\
&= f_{\mathcal{C}}(t + dt, W_{t+dt}, t + dt) + (a_t - c_t) + f_{\mathcal{C}}(t + dt, W_{t+dt}, t) \\
&\quad - f_{\mathcal{C}}(t, W, t).
\end{aligned}
$$

Recall from equation (9) that $V(t, W) = \sup_{\mathcal{C}\in\mathfrak{C}} f_{\mathcal{C}}(t, W, t)$. Adding and subtracting $f(t + dt, W_{t+dt}, t + dt)$ yields

$$
\begin{aligned}
0 = \sup_{\tilde{\mathcal{C}}\in\mathfrak{C}} \{&V(t + dt, W_{t+dt}) - V(t, W) + (\tilde{a}_t - \tilde{c}_t)\, dt \\
&+ f_{\mathcal{C}}(t, W, t) - f_{\mathcal{C}}(t + dt, W_{t+dt}, t + dt) + f_{\mathcal{C}}(t + dt, W_{t+dt}, t) - f_{\mathcal{C}}(t, W, t)\}\,.
\end{aligned}
$$

Taking $dt \to 0$ yields the following:

$$
0 = \sup_{\tilde{\mathcal{C}}\in\mathfrak{C}} \left\{ \mathcal{A}^{\tilde{\mathcal{C}}} V(t, W) + (\tilde{a}_t - \tilde{c}_t) - \mathcal{A}^{\tilde{\mathcal{C}}} f_{\mathcal{C}}(t, W, t) + \mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\mathcal{C}}(t, W) \right\}. \tag{15}
$$

In particular, $f^t(t, W)$ is a martingale because $f^s(t + dt, W_{t+dt})$ captures how the principal's s-selves total payoff changes over time.[13] Thus, under the optimal contract $\mathcal{C}$, we have

$$
\mathcal{A}^{\mathcal{C}} f^s_{\mathcal{C}}(t, W) + R^s(t)(a^{\mathcal{C}}_t - c^{\mathcal{C}}_t) = 0, \tag{16}
$$

for all $s \leq t$. Equations (15) and (16) are the two critical components defining the *extended HJB system* used to formally characterize the solution to the optimal contract in the next subsection.

3.3 Optimal Contract: Formal Characterization and Existence

We are now in the position to present the main result of this paper:

[13] As a reminder, $f(t, W, t) \neq f^t(t, W)$, given the notations specified in Section 3.2.1.

Theorem 1 *Under Assumptions 1 and 2, there exists a principal-optimal renegotiation-proof equilibrium contract (as in Definition 5 to 7)* $\mathcal{C} = (a_t, c_t)_{t\in[0,T]}$. *The principal's equilibrium value function* $V(t, W)$ *under this contract is given by:*

$$V(t, W) = E^{\mathcal{C}}_{t,W}\left[\int_t^T R^t(k)\,(a_k - c_k)\,dk + B(t, W_T)\right]. \tag{17}$$

$V(t, W)$ *satisfies the following* ***extended HJB system****:*

$$\sup_{\tilde{\mathcal{C}}\in\mathfrak{C}} \left\{\mathcal{A}^{\tilde{\mathcal{C}}}V(t, W) + (\tilde{a}_t - \tilde{c}_t) + \mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\mathcal{C}}(t, W) - \mathcal{A}^{\tilde{\mathcal{C}}} f_{\mathcal{C}}(t, W, t)\right\} = 0, \tag{18}$$

subject to (6), the IC condition (IC), and boundary condition:

$$V(T, W) = B(T, W_T) \text{ for all } W. \tag{19}$$

For each fixed s, $f^s(t, W)$ *is defined as the solution of the following equation:*

$$\mathcal{A}^{\mathcal{C}} f^s_{\mathcal{C}}(t, W) + R^s(t)\,(a_t - c_t) = 0, \tag{20}$$

and $f^s_{\mathcal{C}}(t, W) = f_{\mathcal{C}}(t, W, s)$.[14]

Theorem 1 highlights the novel features of the optimal contract under time-inconsistent preferences. While the principal's value function in a time-consistent benchmark is captured by a single HJB equation involving only $(\tilde{a}_t - \tilde{c}_t)$ and $\mathcal{A}^{\tilde{\mathcal{C}}}V$ (a reminder and a heuristic derivation is provided in the previous subsection for interested readers, specifically equation 14), the principal's value function with time-inconsistent preferences involves extra terms in (18). One of the extra terms, $\mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\mathcal{C}}(t, W)$, is the solution to a system of backward stochastic equations (20). These extra terms stem from the different objects the value function $V(t, W)$ represent. In a time-consistent case, the value function tracks how the principal evaluates her discounted future payoff. Changes in time do not change the principal or how the discounting is made. In our setup, the value function tracks how the principal *at each time* evaluates her discounted future payoff *in an optimal equilibrium* of her intra-personal game. Changes in time change the principal's relevant self, and the value function is consequently linked to the equilibrium of the intra-personal game between the different selves.

[14] Under the optimal controls $\mathcal{A}^{\tilde{\mathcal{C}}}V_{\hat{\mathcal{C}}}(t, W) = \mathcal{A}^{\tilde{\mathcal{C}}} f_{\hat{\mathcal{C}}}(t, W, t) \neq \mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\hat{\mathcal{C}}}(t, W)$, which allows a cancellation in equation (18). The cancellation is useful for the proof of the existence of the solution, which we utilize in the Appendix, but is actually inconvenient for calculation of the value function because the extended HJB system still has both forward and backward components. For specific utility or discount functions, more convenient simplifications of the entire system are possible, as shown in Remark 3 and in Section 4.

When the principal's t-self evaluates the contract terms, her evaluation is based on her current preferences, which yields the term $f_{\mathcal{C}}^t(t, W)$. Meanwhile, when evaluating a small and potentially sub-optimal change in the contract terms, the principal's t-self must also consider the effect of such a change on her future selves, especially the immediately following self. Such consideration is reflected by the term $f_{\mathcal{C}}(t, W, t)$, which treats the selves to be changing as a variable. Finally, in the equilibrium path all selves of the principal correctly anticipate the behavior of her future selves and her evaluation must take those future behaviors into account. This backward induction logic of the intra-personal game implies that the backward equation system (20) behaves as a martingale and hence always equals to 0.

Besides the heuristic explanation of the extended HJB system, we must establish three additional results to complete the theorem: first, at least one solution to the system exists; second, if multiple solutions exist, they must all be captured by the system we propose; finally, a verification theorem needs to be provided showing that the solution to the extended HJB is the value function of the underlying contracting problem. The existence problem is known to be particularly challenging. As noted in Björk et al. (2017), "[t]he task of proving existence and/or uniqueness of solutions to the extended HJB system seems (...) to be technically extremely difficult". To address this challenge, we adopt a novel approach connecting the extended HJB system to a static game. We prove the game has an equilibrium and therefore the system has a solution.[15]

More specifically, the extended HJB system is the result of the intra-personal game among principal's different selves, we treat this PDE system as a static non-atomic game in which the utility function of each player incorporates the solutions of each of the backward systems (equations 20). This is a game of incomplete information with a particular structure: first, the random variable of interest is the entire realization of the Brownian Motion. Second, the game is populated by a mass of players all with different *characteristics* (their discount function (R^t)) and *information* ($\mathcal{F}_t$) but the same action set. Third, each player's utility is a function of his/her own action, characteristics, information, as well as an aggregation of other players' actions, and an equilibrium of the game is a distribution over the players' utility functions, strategies, and information.

In the equilibrium, the marginal distribution over the characteristics of the players is uniform over $[0, T]$. Thus, each t-self is one player whose action is the best response to the expected valuations identified by the backward equation given the information $\mathcal{F}_t$. This intricate and novel transformation from the intra-personal game to a non-atomic game with

[15] The solution may not be unique under the general framework, a common property of dynamic programming problems in which the uniqueness of the solution usually requires problem-specific boundary conditions. Nevertheless, the solution is (without the loss of generality) optimal following Definition 7.

incomplete information converts the original dynamic problem into a static one, and the dynamic optimization of the original game is reduced to the information set of each player in the static game.[16] Utilizing Balder (1991), this non-atomic static game with incomplete information has an equilibrium in distributional strategies. Thus, the existence of an equilibrium in this non-atomic game implies the existence of a solution to extended HJB system. We provide a verification argument for Theorem 1 : any contract $\mathcal{C}$ that solves the extended HJB system must be an equilibrium contract. We prove that the converse is also true: any equilibrium contract $\mathcal{C}$ must also solve the extended HJB system. Consequently, Theorem 1 captures all principal-optimal renegotiation-proof equilibrium contracts in our model.

The assumption of finite contracting horizon offers several critical analytical conveniences. It allows us to represent continuation utility processes a solution to BSDEs with known regularity properties, such as the smoothness.[17] To glean some intuition of the smoothness, consider a discrete-time finite horizon version of our model, in which the solution to the BSDEs can be approximately understood as the value functions derived via backward induction. In the last period, it is clear that for any continuation value principal will suggest a continuous action and consumption profile. The penultimate principal will therefore also suggest a continuous action and consumption given the behavior of the final principal. Continuing this argument inductively, we conclude each self of the principal must use continuous strategies.[18]

The proof of existence and the verification theorem are critical components of Theorem 1 but are technically involved. Readers who are only interested in applying our framework to solve problems with dynamically inconsistent decision makers can adopt the results in Theorem 1, knowing that a solution exists and is well behaved. In the next section, we apply Theorem 1 to two special cases in which we obtain closed-form solutions and conduct straightforward comparative statics.

Remark 3 *In case that the principal's family of discount functions is a time-difference family (so that there is a discount function $R(t-s)$), the extended-HJB system can reduced to a single forward equation:*

$$\sup_{\mathcal{C}}\{\mathcal{A}^{\mathcal{C}}V + (a_t - c_t) + E^{\mathcal{C}}_{t,W}\left[\int_t^T R'(s-t)(a_s - c_s)ds\right] - R'(T-t)E^{\mathcal{C}}_{t,W}\left[c_T\right]\} = 0\ . \quad (21)$$

[16] To the best of our knowledge, such transformation is new and allows us to establish the existence of a solution to similar systems of PDEs that is of separate interest. We believe this proof technique can be fruitfully used to prove the existence of a larger class of dynamic systems. Readers interested in applying this proof of existence strategy can find the details of this technique in the Appendix.

[17] As shown in the appendix the individual value functions $f^t(t, W)$ are all well-posed.

[18] See Strulovici and Szydlowski (2015) for the smoothness of the controls with time-consistent preferences.

This reduction simplifies calculations at the cost of intuition as it cancels out some terms in the self valuations ($f_{\mathcal{C}}^t(t, W)$) and the equilibrium valuations ($f_{\mathcal{C}}(t, W, t)$). Hence, the comparison with the usual HJB equation becomes more difficult, which we provide in the next subsection. The details of this simplification is included in the Appendix.

Remark 4 *Before closing this section, we wish to highlight that the following seemingly natural approach of proving Theorem 1 will run into technical hurdles: selecting an arbitrary incentive compatible contract $\mathcal{C}$, defining $f^t(t, W)$ via the backward equations (20), solving $V(t, W)$ by maximizing the HJB equation, finding $\mathcal{C}'$ (the contract that generates the solved $V(t, W)$), and starting this process over from $\mathcal{C}'$ in an attempt to find a fixed point using standard arguments (e.g. Schauder fixed-point theorem). However, such an approach requires establishing technical conditions on the mapping of $\mathcal{C}$ to $\mathcal{C}'$, which are endogenous objects. Typically, it is difficult to ascertain which combination of the model primitives will guarantee those conditions.*

4 Application: Time-difference Discount Functions

In this section we demonstrate the applications of our general framework with a particular class of discounting functions: the *time-difference discount family*, defined as the following:

Definition 8 *A family of discount functions $R^t(s)$ is called a* time-difference discount family *if $R^t(s) = R^{t+k}(s + k) = R(s - t)$ for all s, t, k. We call the function $R(s - t)$ as a time difference discount function.*

The unique feature of this type of discount functions is that the discount family relies on $s-t$ but not t or s individually. Thus, with a slight abuse of notation, we use $R(s - t)$ instead of $R^t(s)$. Although a special case in our general framework, the time-difference discount family includes some of the well-studied time-inconsistent preferences, such as quasi-hyperbolic discounting (following Harris and Laibson, 2012); and anticipatory utility (following Loewenstein, 1987), two examples which we solve in closed-form later in this section.

Our analysis so far has shown that the main impact of dynamic inconsistency is on the principal side. Thus, we further simplify the environment and focus on the case in which only the principal is dynamically inconsistent. Specifically, throughout this section we maintain the following:

Assumption 3 *The agent has the following exponential discounting function:*

$$r^t(s) = e^{-\gamma(s-t)}, \forall t, s \in [0, T]$$

and a constant-absolute-risk-aversion (CARA) utility function:

$$u(c,a) = -\frac{1}{\eta} e^{-\eta\left(c-\frac{1}{2}a^2\right)}.$$

where $\eta > 0$ measures his degree of risk-aversion. The agent also has access to a private savings account of which the balance grows at rate γ.

Assumption 3 result in two sources of benefits: first, the principal optimal long-term contract would be renegotiation-proof if the principal were also an exponential discounter (as in Holmstrom and Milgrom, 1987). This allows us to isolate the impact of renegotiation-proofness for the ensuing analysis regarding non-exponential discounting. Secondly, Assumption 3 implies a simplification of the agent's problem that is commonly used the contracting literature.[19] In our model, this simplification manifests as follows:

Lemma 2 *Under Assumption 3 the agent's continuation utility satisfies*

$$dW_t = \gamma(W_t - u(c_t, a_t))dt + \psi_t(dM_t - a_t dt) \ , \tag{22}$$

$$u(c_t, a_t) = \gamma W_t \ . \tag{23}$$

The agent's incentive compatibility condition becomes

$$\psi_t = a_t. \tag{24}$$

The evolution of $\ln(W)$ *is given by*

$$E\left[\ln(-W_t)\right] = \ln(-W_0) + \frac{1}{2}\int_0^t \eta^2\gamma^2\sigma^2\psi_s^2 ds. \tag{25}$$

Finally, there is no private savings on the equilibrium path.

The proof of Lemma 2 is analogous to the proof in He (2011), and hence omitted. Note that ψ_t measures the contract's incentive power: the variations in the agent's continuation utility following the realized cash flows. In the examples below we explore in details the implication of different discount functions for the incentive power of the optimal contract.

Substituting the simplified agent's continuation utility from Lemma 2 and assuming a time-difference discount function for the principal in Theorem 1 yields the optimal contract:

[19] See, e.g. He (2011), Williams (2015), Gryglewicz and Hartman-Glaser (2019), Marinovic and Varas (2018), Bloedel, Krishna, and Strulovici (2020), etc.

Proposition 3 *When the principal has a time-difference discount function $R(t-s)$, under Assumptions 3, ψ_t, a_t, and c_t in the optimal contract are given by*

$$\psi_t = a_t = \frac{1}{1 + K + \eta\gamma^2\sigma^2 \int_0^t [R(T-s) + R'(T-s)]ds}, \tag{26}$$

where K is a constant of integration and

$$c_t = \frac{1}{2}a_t^2 - \frac{\ln(\gamma\eta)}{\eta} - \frac{1}{\eta}\ln(-W_t). \tag{27}$$

Proposition 3 immediately implies the following result:

Corollary 1 *The equilibrium effort path is monotonic if*

- $R(s) + R'(s) < 0 \;\; \forall s \in [0, T]$ *(increasing path) .*
- $R(s) + R'(s) > 0 \;\; \forall s \in [0, T]$ *(decreasing path) .*

The equilibrium effort path is non-monotone if the sum $R(\cdot) + R'(\cdot)$ ever changes sign.

These are very useful features of the optimal contract. They demonstrate that the equilibrium path of the agent's actions is deterministic, determined only by the discount function R and it's first derivative R'. Therefore, in the case of exponential discounting the equilibrium effort path is always monotonic in time (increasing if $r < 1.$ and decreasing if $r > 1$). In the following subsections section, we show how these results change in two examples of non-exponential discounting.

Results in this section offer a potential empirical strategy to identify time-inconsistent preferences in practice, which is usually difficult to observe directly. Recall Corollary 1 that under the optimal contract, the equilibrium action path is monotonic in time if both the agent and the principal have exponential discounting. Therefore, if we observe non-monotonic and continuous suggested action paths based on actual contracts, we may deduce that principal has a non-exponential discount function.[20] Although this strategy requires a few assumptions

[20] Admittedly one caveat of this identification strategy is the assumption of CARA utility function, as stochastic and non-monotonic action paths are possible under other utility functions (see Sannikov (2008)). It is important to note that one can generate non-monotone action paths when the principal has a predetermined change in her discount rate at some point during the contracting period, e.g. if the principal has a discount rate ρ_1 before some point $t^* < T$ and $\rho_2 \neq \rho_1$ afterwards. However, this kind of discounting would lead to discontinuous action paths. Therefore, if the action path can be observed in high frequency, such as the prevalent use of continuous performance management (CPM) in recent years, it is still possible to reasonably deduce time inconsistency based on observed actions. Details of the solutions of this model are available upon request.

in general, it could potentially be applicable in practice with weaker conditions under more specific settings.[21]

4.1 Quasi-Hyperbolic Discounting

Our first example explores the case of quasi-hyperbolic discounting. Such discounting has been explored in many examples, such as O'Donoghue and Rabin (1999), Thaler and Benartzi (2004), Harris and Laibson (2012), Jackson and Yariv (2014), Jackson and Yariv (2015), Bisin, Lizzeri, and Yariv (2015), etc. We follow Harris and Laibson (2012) by assuming that the quasi-hyperbolic discount function is a convex combination of a short-term discount function and a long-term discount factor. Formally, we assume

Assumption 4 *The principal has the following time difference discount function:*

$$R^t(s) = R(s-t) = (1-\beta)e^{-(\rho+\lambda)(s-t)} + \beta e^{-\rho(s-t)} \tag{28}$$

with $\beta \in (0,1)$ and $\gamma > \rho + \lambda$.

The representation above is a deterministic characterization of a principal who values "near present" returns with a higher discount factor (discounted by $e^{-\rho(s-t)}$), and "far future" returns with a lower factor (discounted by $\beta e^{-\rho(s-t)}$). That is, the principal becomes less patient over time. $\beta < 1$ captures the size of the drop in discount factor in the far future. The switch between the "near present" and the "far future" occurs stochastically with arrival intensity λ, and the overall discount function $R^t(s)$ incorporates this expected drop.

Because $R(T-s)$ is an exponential function in s, substituting it back into (26) yields a closed-form solution, which is given in Appendix A.6. This closed-form solution implies the following analytical characterization of the properties of the optimal contract:

Proposition 4 *Under Assumption 4 the optimal renegotiation-proof contract derived in Proposition 3 has the following properties*

1. *If $\lambda + \rho \geq 1$, the incentive power ψ_t and wage c_t of the optimal contract decreases if β decrease, for all t and all W_t.*

2. *$\lambda + \rho < 1$ then there exists $T(\lambda, \rho) > 0$ such that*

   (a) If $T \leq T(\lambda, \rho)$, the incentive power ψ_t and wage c_t of the optimal contract decreases if β decrease, for all t and all W_t.

[21] See Heidhues and Strack (2019) which also feature a partial identification strategy for time inconsistent discounting based on an optimal stopping problem.

(b) If $T > T(\lambda, \rho)$ then there exists $t(\lambda, \rho) < T(\lambda, \rho)$, such that

i. The incentive power ψ_t and wage c_t of the optimal contract increases if β decrease, for all $t < t(\lambda, \rho)$ and all W_t,

ii. The incentive power $\psi_{t'}$ and wage $c_{t'}$ of the optimal contract decreases if β decrease, for all $t' \geq t(\lambda, \rho)$ and all $W_{t'}$.

The proposition demonstrates the effect of quasi-hyperbolic discounting and the associated present bias on the dynamic characteristics of the optimal equilibrium contract. In the case of exponential discounting, the discount factor ρ can be interpreted as the death rate or hazard rate of the principal. Under quasi-hyperbolic discounting, the principal also faces the additional risk λ of being replaced by a more impatient self. Thus, her overall survival probability is negatively correlated with the sum $\lambda + \rho$. The higher the sum is, the lower the survival probability of a time-t principal. Note that this survival probability distinctly different interpretation than just a higher discount factor, as it is a combination of termination (standard discount) and devaluation(hyperbolic drop).

A decrease in β has two effects: an *internalization effect*, whereby the current self of the principal internalizes the decrease in patience of the future self, and all her selves become more impatient; an *equilibrium effect*, whereby the principal front-loads contract incentives, because she correctly predicts that she cannot credibly promise sufficient dynamic incentives in equilibrium in the future, and wants to do so before the impatient self arrives. If the survival probability is low ($\lambda + \rho \geq 1$) or the horizon T is short, the internalization effect dominates. The wage and incentive power of the contract decrease unambiguously with the level of patience of the replacing self. However, if T is long enough and the current self has a sufficiently high probability of survival ($\lambda + \rho < 1$), the equilibrium effect results in the front-loading of contract wage and incentive power. The prominence of the front-loading increases as β decreases. That is, the more impatient the principal's future self is, the more the current self increases the wage and incentive power of the contract towards the beginning periods of the contracting relationship, consequently lowering the wage and incentive power toward the end of the contracting horizon.

In addition, we obtain the following observations regarding how close the optimal contract resembles a time-consistent benchmark:

Proposition 5 *Under Assumption 4 the optimal renegotiation-proof contract derived in Proposition 3 has the following properties:*

1. *If $\beta = 1$ (or $\beta = 0$), the optimal contract is identical to that for a dynamically-consistent principal with discount rate ρ (or $\rho + \lambda$), and the incentive power ψ_t is (weakly) monotonic in time t.*

2. *As the contracting horizon becomes infinitely long (i.e. $T \to \infty$), the optimal contract converges to that for a time consistent-principal with discount rate $\frac{\rho(\lambda+\rho)}{\rho+\beta\lambda}$.*[22]

3. *For every $\rho \in (0,1)$, the optimal incentive power ψ_t and the equilibrium actions a_t are non-monotone in t as long as λ is high enough.*

The first property is straightforward: if $\beta = 1$ ($\beta = 0$) the principal has the same discount rate for "near present" payoff and "far future" payoff, which is equivalent to having exponential discounting. Following Holmstrom and Milgrom (1987), the optimal incentive power is monotonic in time, regardless of the contracting horizon.[23]

The second property demonstrates that the effect of quasi-hyperbolic discounting becomes indistinguishable from exponential discounting if the time horizon becomes arbitrarily large. If the contracting horizon is infinitely long, the solution to the dynamically inconsistent problem agrees with that of a dynamically consistent one with *exactly the same instantaneous utility function* but a *different exponential discount factor*. Intuitively, as the horizon increases, the problems faced by each of the principal's time-t selves become more similar. When the horizon is infinitely long, the principal's problem becomes completely stationary. Under quasi-hyperbolic discounting, stationarity is sufficient to ensure that the resulting optimal contract converges to that of a principal with the exact same instantaneous utility function but a different discount rate.

The third property highlights the possibility of non-monotonic incentive power under finite horizon. Based on the first two properties, we know that the incentive power is (weakly) monotonic in time without quasi-hyperbolic discounting or without a finite horizon. Under quasi-hyperbolic discounting and finite horizon, however, a "near-future" principal with the short-term discount factor and a "far-future" principal with long-term discount factor may prefer different levels of incentive power. Suppose that the former prefers a high power contract and the latter prefers low power. Because the switch between the two types of principals occurs stochastically, the sophisticated "near-future" principal anticipates that her less patient "far-future" self will arrive at some point. Thus, she designs a path of incentive power that gradually decreases towards the level preferred by her "near-future" self. However, at some point, as she approaches the end of the contracting horizon (the deadline), the probability of the switch happening before the deadline decreases over time as the contract "runs out of time." Consequently, the principal resembles more closely her

[22] More precisely, a contract $\mathcal{C}$ converges to another contract $\mathcal{C}'$, if the processes ψ_t, a_t, c_t under $\mathcal{C}$ converge pointwise to the processes under $\mathcal{C}'$.

[23] Whether incentive power is monotonically increasing, decreasing, or constant over time is the result of model-specific boundary conditions. See He (2011), He, Wei, Yu, and Gao (2017), and Marinovic and Varas (2018) for some recent examples.

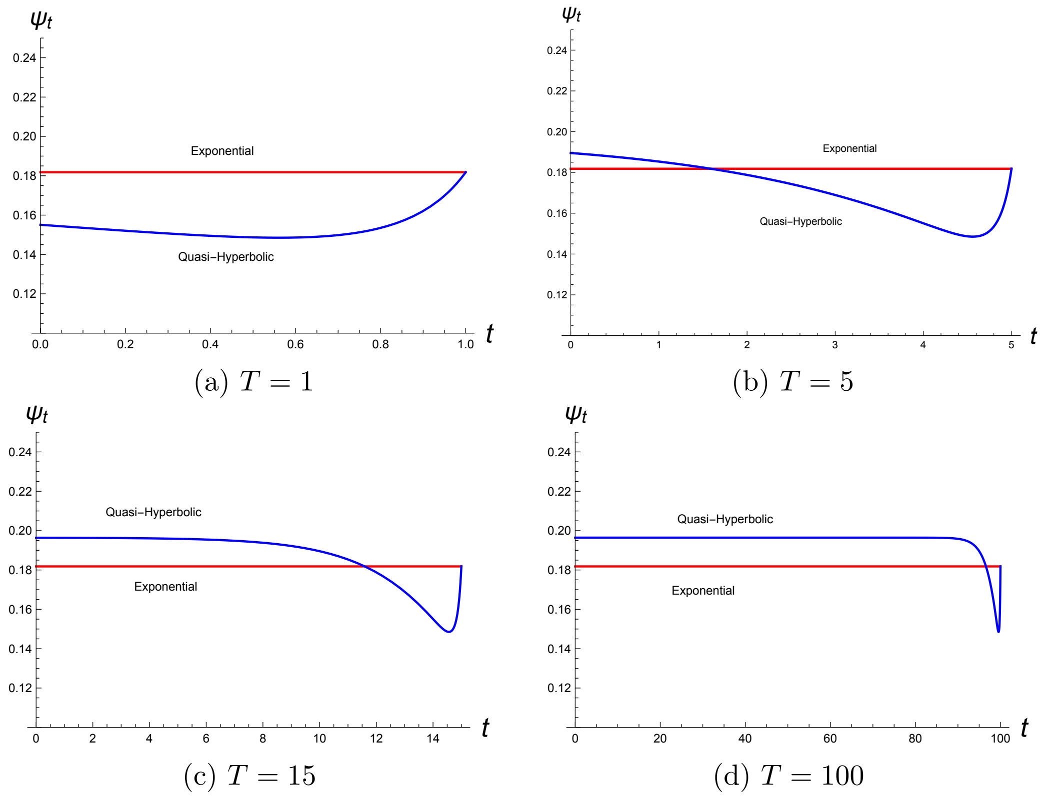

(a) $T = 1$ (b) $T = 5$

(c) $T = 15$ (d) $T = 100$

Figure 1: These plots illustrate the paths of the incentive power (also the paths of agent's actions given Proposition 3) under the optimal contract. Blue lines indicate the paths under quasi-hyperbolic discounting with $\eta = 2, \sigma = 1, \beta = 0.5, \lambda = 5, \rho = 0.5, \gamma = 1.5$. Red lines indicate the paths under exponential discounting with the same parameters except for $\rho = 1$. Each plot corresponds to a different time horizon (different T). If $\rho > 1$ ($\rho < 1$), the equilibrium path under exponential discounting would be monotonically increasing (decreasing).

"near future" patient self and reverts the incentive power of the contract until it converges to the time-consistent benchmark level at the deadline, causing a U-shaped path. We refer to this reverting of incentive power towards the end of the contracting horizon as the "deadline effect". The turning point is determined by λ, the arrival intensity of the drop in discount. The higher the λ, the sooner the action path changes course. If λ is sufficiently low, the drop in discount is very remote such that the "deadline effect" does not occur, and the path of optimal incentive power converges to the (weakly) monotonic case.

We illustrate the properties above and the comparison between quasi-hyperbolic discounting and the time-consistent benchmark in Figure 1.[24]

[24]When T is finite, discounting becomes irrelevant at $t = T$, and the optimal incentive power at that point (ψ_T) must be the same under both types of discounting. As T increases, the difference in the contracts induced by the two types of discounting is pushed further to the future. When $T \to \infty$, such difference is pushed infinitely far away, and the optimal contract under quasi-hyperbolic discounting coincides with that under exponential discounting but with a different discount factor (i.e. Property 2 in Proposition 5).

4.1.1 Anticipatory Utility

Our second example explores the case in which the principal has anticipatory utility. That is, in addition to immediate utility, the principal's well-being also depends on her expectations of future utility. Such expectations may arise from either "savoring" or "dreading" future consumption as in Loewenstein (1987), or from optimism, as in Brunnermeier, Papakonstantinou, and Parker (2016).[25]

We model anticipatory utility following Loewenstein (1987). In addition to the current payoff, the principal derives an additional $e^{-\zeta t}$ amount of utility from payoff she anticipates to receive at a future time t. Her discount of the future utility (both actual future payoffs and anticipation of more distant future payoffs) is ρ. Formally, we assume

Assumption 5 *The principal values consumption streams by:*

$$E\left[\int_0^T e^{-\rho t}\left(\Pi_t + a_t - c_t\right) d_t\right],$$

where Π_t *is given by*

$$\Pi_t = E\left[\int_t^T e^{-\zeta(s-t)}(a_t - c_s) ds | \mathcal{F}_t\right].$$

Here, Π_t captures the utility from the anticipated future payoff.By the Law of Iterated Expectations and a change in the order of integration, the principal's valuation of a stream of payoffs starting from any period t can be re-written as

$$E^{\mathcal{C}}\left[\int_t^T \left(\frac{e^{-\rho(s-t)} - e^{-\zeta(s-t)}}{\zeta - \rho} + e^{-\rho(s-t)}\right)(a_s - c_s) ds | \mathcal{F}_t\right].$$

As a result, the principal with the anticipatory utility can be understood as effectively having the following discounting function:

$$R^t(s) = R(s-t) = \frac{e^{-\rho(s-t)} - e^{-\zeta(s-t)}}{\zeta - \rho} + e^{-\rho(s-t)}.$$

The first term captures the discounted anticipation: letting $\zeta \to \infty$ yields the standard discounting as utility from anticipation disappears. We also make a technical assumption that $\zeta > \rho > 1$. The first inequality is necessary for transversality and the second is necessary to avoid corner solutions.[26]

[25] Other studies related to anticipatory utility include Caplin and Leahy (2001), Loewenstein, O'Donoghue, and Rabin (2003), Caplin and Leahy (2004), Takeuchi (2011), etc.

[26] Although not noted in Loewenstein (1987), an appropriate numerical relationship between ρ and ζ is

Similar to the case of quasi-hyperbolic discounting, because $R(T-s)$ is an exponential function in s, the integral in (26) has a closed-form solution, which is given in the Appendix A.7. Based on the closed-form solution, the implications of anticipatory utility can be summarized as follows:

Proposition 6 *Under Assumption 5 the optimal renegotiation-proof contract derived in proposition 3 has the following properties:*

1. *If $\zeta = \infty$, the optimal contract is identical to that for a dynamically-consistent principal with discount rate ρ and optimal incentive power ψ_t and the equilibrium actions a_t are (weakly) monotonic in time t.*

2. *As the contracting horizon becomes infinitely long (i.e. $T \to \infty$), the optimal contract converges to that for a time consistent-principal with discount rate $\frac{\rho\zeta}{1+\zeta}$.*

3. *For any finite ρ and ζ, there exists $\tilde{T}(\rho,\zeta)$ such that the optimal incentive power ψ_t and the equilibrium actions a_t are non-monotonic in t as long as $\tilde{T}(\rho,\zeta) < T < \infty$.*

The first property is straightforward: if $\zeta = \infty$ the principal receives no utility from anticipation. Her discounting is then exponential and the optimal contract follows that of Holmstrom and Milgrom (1987), with monotone incentive power regardless of the time horizon.

The second property demonstrates that the effect of anticipation disappears if the time horizon becomes arbitrarily large. The intuition is identical to why the effect of quasi-hyperbolic discounting disappears under an infinite horizon: as the horizon increases the problems faced by each of the principal's t-self become similar. When the horizon is infinitely long, the principal's problem becomes completely stationary. Thus, the solution converges to that of a dynamically consistent principal with a different but constant discount rate.

The third property highlights the possibility of non-monotonic incentive power under a finite horizon, hence a "deadline effect". Based on the previous two properties, the incentive power is (weakly) monotonic in time with exponential discounting or with an infinite horizon. With a finite horizon, the anticipation of future payoff leads the principal to back-load incentive power. However, the exact dynamics depend on the length of the contracting horizon: if the horizon is sufficiently long, back-loading causes incentives to gradually increase in time in the early phases. However, as the end of the contracting horizon approaches, the effect of anticipation diminishes, causing incentives to decline over time.

We illustrate the above properties in Figure 2.

economically crucial. Without further restrictions, utility of infinitely far future may still have a positive value today after discounting.

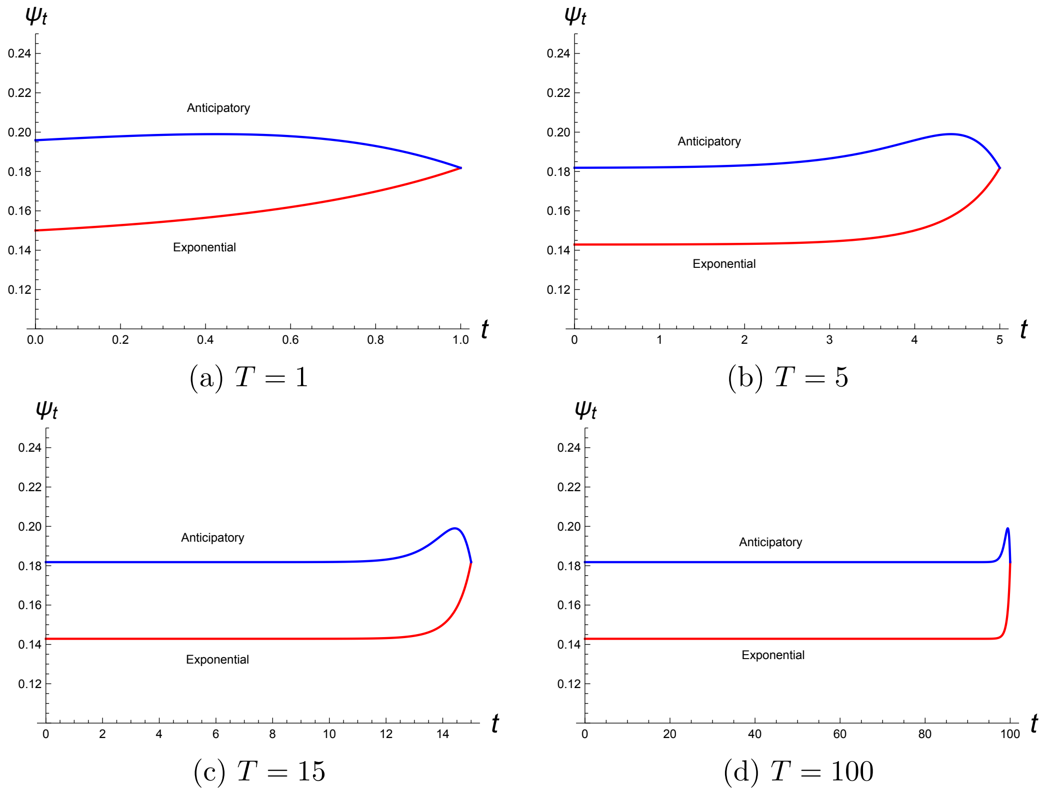

Figure 2: These plots illustrate the paths of the incentive power (also the paths of agent's actions given Proposition 3) under the optimal contract. Blue lines indicate the paths under anticipatory utility with $\eta = 2, \sigma = 1, \rho = 1.5, \zeta = 3, \gamma = 1.5$. Each plot corresponds to a different time horizon (different T).

5 Conclusion

Non-exponential discounting has been widely observed in practice and extensively studied in economic research. However, dynamic moral hazard and long-term contracting between sophisticated parties with non-exponential discounting has thus far eluded formal analysis due to several challenges. First, individuals with non-exponential discounting make dynamically inconsistent plans, and what currently appears optimal may become sub-optimal in the future. In such a scenario, what is the appropriate notion of optimality? Moreover, how does the contract provide incentive for a dynamically inconsistent agent, and how to ensure that a contract agreed by the agent and the principal today will actually be carried out or renegotiated in the future after their preferences have changed? Finally, does a long-term optimal contract always exist? If so, how is such a contract different from the optimal contract under exponential discounting?

In this paper, we answer these questions by establishing a broad framework of an *incentive compatible, renegotiation-proof, optimal long-term* contract under dynamically inconsistent preferences resulting from non-exponential discounting. We formulate the contracting problem as a dynamic, intra-personal game played among the agent, the principal and their respective future "selves". This formulation allows us to derive the optimal contract as the equilibrium of intra-personal game and prove its existence. We demonstrate the applicability of our general framework through two special examples of economically important behavioral patterns: present bias (quasi-hyperbolic discounting) and future bias (anticipatory utility). In both examples, our general framework yields closed-form solutions and generate testable implications, such as the precise impact of the "deadline effect" on the incentive power and compensation scheme of the optimal contract.

The agency friction of our general framework is adopted from the dynamic hidden effort problem in Sannikov (2008). A related strand of literature formulates the agency friction as a cash-flow diversion (CFD) problem. The analytical tools developed in this paper can be easily adapted to CFD problems and their various extensions to highlight the impact of non-exponential discounting on a number of applications, including but not limited to capital structure (DeMarzo and Sannikov, 2006), security design (Biais, Mariotti, Plantin, and Rochet, 2007), firm investment (DeMarzo, Fishman, He, and Wang, 2012), liquidity management (Bolton, Chen, and Wang, 2011), resource allocation (Feng and Westerfield, 2020), and risk management (Biais, Mariotti, Rochet, and Villeneuve, 2010).

Our framework can be expanded in different directions. In particular, the extended HJB system derived in this model and its variations can arise in many dynamically inconsistent optimal control problems for reasons other than non-exponential discounting. Our main theorem, which bridges dynamically inconsistent control problems with non-atomic games, may be modified to establish the solutions to these problems at both the theoretical and practical level. Examples of the practical applications include but are not limited to ambiguity, habit formation and mean-variance risk preferences. We leave these topics for future research.

A For Online Publication - Appendix

A.1 Proof of Lemma 1

Consider an equilibrium contract $\mathcal{C}$ offered by the principal. In general contract consumption and suggested action paths suggested by $\mathcal{C}$ could depend on the entire history of the outputs. To resolve such a history dependence we follow Williams (2015) and introduce a change of measure. Then, for any $a \in [\underline{a}, \bar{a}]$ suggested by $\mathcal{C}$ define the family of $\mathcal{F}_t$ predictable processes:

$$\Gamma_t(a) := exp\left(\int_0^t \frac{a_s}{\sigma} dZ_s^0 - \frac{1}{2}\int_0^t \left|\frac{a_s}{\sigma}\right|^2 ds\right),$$

with $\Gamma_0(a) = 1$. Also,

$$dM_t = \sigma dZ_t^0,$$

with M_0 is given. Given the linear output, the Novikov condition holds, and Γ_t is as martingale. By Girsanov theorem, we define the Brownian Motion (Z^a) in the following way:

$$Z_t^a = Z_t^0 - \int_0^t \frac{a_s}{\sigma} ds.$$

Then by suppressing a in Γ we can conduct the following change of measure:

$$E^a\left[\int_t^T r^t(s)u(c_s, a_s)ds + r^t(T)u(c_T, \underline{a}_T)\right] = E^0\left[\int_t^T \Gamma_s r^t(s)u(c_s, a_s)ds + \Gamma_T r^t(T)u(c_T, \underline{a}_T)\right].$$

Therefore, the optimal control problem of the agent can be formulated as follows:

$$J(t, \Gamma) = \sup_{\hat{a}_s,\, s\in[t,T]} E^0\left[\int_t^T \Gamma_s r^t(s)u(c_s, \hat{a}_s)ds + \Gamma_T r^t(T)u(c_T, \underline{a}_T)\right], \tag{29}$$

$$d\Gamma_s = \Gamma_s\left(\frac{\hat{a}_s}{\sigma}\right)dZ_s^0, \quad \Gamma_0 = 1. \tag{30}$$

Notice that the optimal control problem is still time inconsistent but with the change of measure we resolved the history dependence. Given the dynamic inconsistency of the problem, we search for a Markov Perfect $\hat{a}$. Under our regularity conditions (a originating from a compact space, σ_Γ being C^2 in a and Lipschitz) we can apply Theorem 1 of Yan and Yong (2019) to the time inconsistent control problem defined by 29 and 30.

Theorem 2 (Yan and Yong 2019) *Given a contract $\mathcal{C}(a, c)$ suppose that $(\hat{a}, \Gamma)$ is an agent intra-personal equilibrium pair (where each agent takes the suggested action by the principal) and suppose that for any given $t \in [0, T)$, the first order adjoint processes $W(\cdot, t), \tilde{\psi}(\cdot, t)$ and* *second order adjoint processes $P(\cdot,t),\Lambda(\cdot,t)$ are adapted solutions to the following BSDE's:*

$$\begin{aligned}
dW(s,t) &= -\left(\frac{\hat{a}_t}{\sigma}\tilde{\psi}(s,t)+r^t(s)u(c_s,\hat{a}_s)\right)ds+\tilde{\psi}(s,t)dZ^0_s \quad s\in[t,T],\\
W(T,t) &= r^t(T)u(c_T,\underline{a}_T),\\
dP(s,t) &= -\left(\left(\frac{\hat{a}_t}{\sigma}\right)^2 P(s,t)+2\left(\frac{\hat{a}_t}{\sigma}\right)\Lambda(s,t)\right)ds+\Lambda(s,t)dZ^0_s \quad s\in[t,T],\\
P(T,t) &= 0,
\end{aligned}$$

then almost surely for any deviation $\tilde{a}$ given the suggested path $a=\hat{a}$ we have the following global form of Pontyagrin maximum principle:

$$\begin{aligned}
0 \le &\lim_{\varepsilon\to 0}\frac{1}{\varepsilon}E_t\int_t^{t+\varepsilon}\left\langle\tilde{\psi}(s,t),\Gamma_s\left(\frac{\tilde{a}-\hat{a}}{\sigma}\right)\right\rangle ds+r^t(t)\Gamma_t\left(u(c_t,\tilde{a})-u(c_t,\hat{a})\right)\\
&+\left(\Gamma_t\frac{\tilde{a}-\hat{a}_t}{\sigma}\right)^2 P(t,t). \qquad (31)
\end{aligned}$$

If the agent had exponential discounting, there would be a single adjoint process. In our setup, we have a flow of adjoint processes. This flow of processes describes how different selves of the agent evaluate the equilibrium contract.

Given that the action set is convex, instead of (31), we can use a local version of the Hamiltonian (see Page 120 Case 2 of Yong and Zhou (1999) or Peng (1990)). Thus we can define the generalized Hamiltonian of time t for any time s as:

$$\mathcal{H}_t(s) = -r^t(s)\Gamma_s u(c_s,\hat{a}_s)ds+\tilde{\psi}(t,s)\Gamma_s\left(\frac{\hat{a}_s}{\sigma}\right). \qquad (32)$$

Hence at time t, we have

$$\mathcal{H}_t(t) = -r^t(t)\Gamma_t u(c_t,\hat{a}_t)ds+\tilde{\psi}(t,t)\Gamma_t\left(\frac{\hat{a}_t}{\sigma}\right). \qquad (33)$$

From (31), we know that the equilibrium actions maximize the generalized Hamiltonian of time t at time $s=t$. Thus we must have $\frac{\partial\mathcal{H}_t(t)}{\partial\hat{a}_t}=0$, recalling $r^t(t)=1$ which yields the incentive compatibility (IC) condition of the agent at time t:

$$u_a(c,\hat{a}) = \frac{\tilde{\psi}(t,t)}{\sigma} = \psi_t. \qquad (34)$$

Next, we derive the evolution of the agent's continuation utility. Observe that, first order adjoint process is equal to the equilibrium value function $J(t, \Gamma)$.

We write the equilibrium value function in the integral form as follows:

$$\tilde{W}(k,t) = r^t(T)u(c_T, \underline{a}_T) - \int_k^T \left(\frac{\hat{a}_s}{\sigma}\tilde{\psi}(s,t) + r^t(s)u(c_s, \hat{a}_s) \right) ds + \int_k^T \tilde{\psi}(s,t)dZ_s^0.$$

Taking $k \to t$ we have

$$\tilde{W}(t,t) = r^t(T)u(c_T, \underline{a}_T) - \int_t^T \left(\frac{\hat{a}_s}{\sigma}\tilde{\psi}(s,t) + r^t(s)u(c_s, \hat{a}_s) \right) ds + \int_t^T \tilde{\psi}(s,t)dZ_s^0.$$

Following Yan and Yong (2019), we can define the backward stochastic Volterra integral equation (BSVIE):

$$W(t) = r^t(T)u(c_T, \underline{a}_T) - \int_t^T \left(\frac{\hat{a}_s}{\sigma}\tilde{\psi}(s,t) + r^t(s)u(c_s, \hat{a}_s) \right) ds + \int_t^T \tilde{\psi}(t,s)dZ_s^0,$$

and under the regularity conditions we have $W(t) = \tilde{W}(t,t)$ and $W(t) = J(t, \hat{\Gamma})$ for all t. Using the fact that $Z_t^a = Z_t^0 - \int_0^t \frac{a_s}{\sigma} ds$, yields the representation provided in Lemma 1.

Since each agent t only acts once, $W(t)$ act as the principal's state variable. Therefore, it is natural to find a forward representation for $W(t)$. We will follow Wang and Yong (2019) and use the decoupling field of the forward-backward stochastic Volterra integral equation (FBSVIE) to pin down the evolution of $W(t)$. In particular under the Lipschitz continuity of both the forward and backward coefficients, we can represent the decoupling field θ as follows:

Theorem 3 (Wang and Yong 2019) *Under Assumption 2, the integral valued process $W(t)$ has a unique representation as a diffusion $\theta(t,t,\Gamma_t)$, where $\theta(t,s,\Gamma)$ is the unique, continuous solution to the following PDE:*

$$\theta_s(t,s,\Gamma) + \frac{1}{2}\left(\frac{a_s}{\sigma}\Gamma\right)^2 \theta_{\Gamma\Gamma}(t,s,\Gamma) + \Gamma\left(\frac{a_s}{\sigma}\right)^2 \theta_\Gamma(t,s,\Gamma) + r^t(s)u(c_s, a_s) = 0, \tag{35}$$

$$\theta(t,T,\Gamma) = r^t(T)u(c_T, \underline{a}_T)\Gamma. \tag{36}$$

This representation of a BSVIE as a diffusion is a direct application of Theorem 1 of Wang and Yong (2019), which requires two conditions noted as $H1$ and $H2$ in that paper. Here, $H2$ is directly satisfied by the boundedness of the derivative of the utility function. For $H1$ we observe that the difference in volatility of the Γ process is always linearly bounded since it is always less than $\bar{a}\Gamma$ thus satisfying the linear growth condition in $H1$. Therefore, the above PDE has a unique continuous solution and this solution is the decoupling field of the BSVIE; hence, $W(t)$ can be represented as a diffusion.

Notice that the decoupling field allows us to represent the integral equation defining the equilibrium value in terms of the forward component Γ_t. As also noted in Wei et al. (2017), the objective, $\int_t^T \Gamma_s r^t(s)u(a_s,c_s)ds$, does not rely on the backward terms; thus, for any flows of sensitivities $\psi(t,t)$ the PDE is linear with t as a parameter.[27] Furthermore, given that the forward part is a standard SDE, it allows us to represent the equilibrium value as a single process. By Ito's rule evolution of $W(t)$ can be written as:

$$dW(t) = \theta_1(t,t,\Gamma_t) + \theta_2(t,t,\Gamma_t) + \frac{1}{2}\left(\frac{a_s}{\sigma}\Gamma_s\right)^2 \theta_{\Gamma\Gamma}(t,t,\Gamma_t), \tag{37}$$

where $\theta_1(\theta_2)$ denotes the partial derivative with respect to the first (second) argument. If we follow the decoupling approach in the Sannikov (2008), we would have end up with the terms $\theta_1(t,t,\Gamma_t)$ and $\theta_{\Gamma\Gamma}$. The term $\theta_2(t,t,\Gamma_t)$ captures the effect of how different selves evaluates the equilibrium value.

A.2 Proof of Proposition 1

Let $t < T$ be arbitrary time and $\mathcal{F}_t$ be the information available at time t. Since $\mathcal{C}' = \{c'_t, a'_t\}_{t\in[0,T]}$ is a feasible renegotiation, it is incentive compatible and

$$\begin{aligned} &E_t^{\mathcal{C}'}\left[\int_t^T r^t(s)u(c'_s,a'_s)ds + r^t(T)u(c'_T,\underline{a}_T)\middle|\mathcal{F}_t\right] \geq \\ &E_t^{\mathcal{C}}\left[\int_t^T r^t(s)u(c_s,a_s)ds + r^t(T)u(c_T,\underline{a}_T)\middle|\mathcal{F}_t\right]. \end{aligned} \tag{38}$$

If the weak inequality holds with equality then the principal-preferred alternative is the feasible renegotiation itself and we are done. If the inequality is strict then let $k > 0$ be the difference in utility from agent t's perspective, that is

$$\begin{aligned} k = &E_t^{\mathcal{C}'}\left[\int_t^T r^t(s)u(c'_s,a'_s)ds + r^t(T)u(c'_T,\underline{a}_T)\middle|\mathcal{F}_t\right] \\ &- E_t^{\mathcal{C}}\left[\int_t^T r^t(s)u(c_s,a_s)ds + r^t(T)u(c_T,\underline{a}_T)\middle|\mathcal{F}_t\right]. \end{aligned} \tag{39}$$

Now define $\mathcal{C}''$ as follows, for any potential history up to time $(T,\mathcal{F}_T)$, we define $(c''_T|\mathcal{F}_T)$ such that $r^t(T)u(c''_T,\underline{a}_T|\mathcal{F}_T) = r^t(T)u(c'_T,\underline{a}_T|\mathcal{F}_T) - k$. Such c''_T exists and $c''_T < c'_T$ because u is continuous, increasing in c and unbounded below. For all $t < T$, we choose $\{c''_t, a''_t\}_{t\in[0,T)}$ such that $(c'_t,a'_t) = (c''_t,a''_t)$. Observe that compared to under $\mathcal{C}'$, the agent's selves *realized*

[27]See Wei et al. (2017) and Yong (2012) for more detailed discussions.

utility under $\mathcal{C}''$ is lower by a constant amount. This amount equals exactly k for the t-self. Given that $\mathcal{C}'$ is incentive compatible, each of the agent's selves will act the same way under $\mathcal{C}''$. Consequently, the new contract $\mathcal{C}''$ is also incentive compatible.

A.3 Proof of Proposition 2

Pick any $\{a_s^1, c_s^1\}_{s\in[t,T)} \in \xi_t(W_1)$. Given an arbitrary possible history path $\mathcal{F}_T$, let $(c_T^2|\mathcal{F}_T)$ be such that $r^t(T)u(c_T^2, \underline{a}_T|\mathcal{F}_T) = r^t(T)u(c_T^1, \underline{a}_T|\mathcal{F}_T) + W_1 - W_2$. Since u is continuous, increasing in c and unbounded, such $c_T^2 \neq c_T^1$ exists. For all $t < T$, let $\{c_t^2, a_t^2\}_{t\in[0,T)}$ be such that $(c_t^1, a_t^1) = (c_t^2, a_t^2)$. We observe that compared to the contract (c_t^1, a_t^1), (c_t^2, a_t^2) changes each of the agent's selves' *realized* utility for every path of history by a constant amount. This amount equals exactly $W_1 - W_2$ for the t-self. Given that (c_t^1, a_t^1) is incentive compatible, each of the agent's selves will act the same way under (c_t^2, a_t^2). Consequently, the new contract (c_t^2, a_t^2) is also incentive compatible. Furthermore, at time t, the contract (c_t^2, a_t^2) delivers W_2 and induces the path $\{a_s^1, c_s^1\}_{s\in[t,T)}$. Thus, $\{a_s^1, c_s^1\}_{s\in[t,T)} \in \xi_t(W_2)$.

An identical argument shows that any element of $\xi_t(W_2)$ is also an element of $\xi_t(W_1)$.

A.4 Proof of Theorem 1

We first prove the existence of a solution to the extended HJB system, and then provide a verification theorem.

A.4.1 The Existence of a Solution to the Backward System

Begin with an arbitrary incentive compatible contract $\hat{\mathcal{C}}$ and consider the second part of the HJB system, the backward equation:

$$\mathcal{A}^{\hat{\mathcal{C}}} f_{\hat{\mathcal{C}}}^s(t, W) + R^s(t)(\hat{a}_t - \hat{c}_t) = 0.$$

For each (t, W) the backward equation

$$\mathcal{A}^{\hat{\mathcal{C}}} f_{\hat{\mathcal{C}}}^t(t, W) + R^t(t)(\hat{a}_t - \hat{c}_t) = 0,$$

is a semi-linear parabolic partial differential equation (PDE). If the PDE had a solution $f_{\hat{\mathcal{C}}}^t(t, W)$, then for every s we could consider a path of W_s that reaches W by time t. This leads us to the notation $f_{\hat{\mathcal{C}}}^s(t, W)$, where we use the subscript $\hat{\mathcal{C}}$ to emphasize its dependence on the control $\hat{\mathcal{C}}$. Under this notation, an equivalent representation of the backward system is:

$$f^t_{\hat{\mathcal{C}}}(t,W) = f^t_{\hat{\mathcal{C}}}(T,W_T) - \int_t^T R^t(r)(\hat{a}_r - \hat{c}_r)dr - \int_t^T Y_r^{\hat{\mathcal{C}}} dZ_r,$$

for some adapted process $Y^{\hat{\mathcal{C}}}$. A solution to the backward system corresponds to a pair of processes $f^t_{\hat{\mathcal{C}}}(t,W)$ and $Y_t^{\hat{\mathcal{C}}}$. However, notice that for a given arbitrary incentive compatible contract $\hat{\mathcal{C}}$, consumption $\hat{c}_t$ is already pinned down by the IC condition. In particular, according to Assumption 2 for any given $\hat{a}_t$, $\hat{\psi}_t$, and W_t, the incentive compatible $\hat{c}_t$, denoted by $c_t^{IC}(a_t,\psi_t)$, is unique. Moreover given contract $\hat{\mathcal{C}}$, the backward system is accompanied by a forward system

$$dW(t) = \theta_1(t,t,\Gamma_t) + \theta_2(t,t,\Gamma_t) + \frac{1}{2}\left(\frac{a_s}{\sigma}\Gamma_s\right)^2 \theta_{\Gamma\Gamma}(t,t,\Gamma_t), \tag{40}$$

where θ denotes the decoupling field of the agent's BSVIE. Also notice that the solution to the backward system $f^t_{\hat{\mathcal{C}}}(t,W)$ and $Y_t^{\hat{\mathcal{C}}}$ does not appear in the forward system, hence this is a *decoupled* forward-backward stochastic differential equation (FBSDE).[28] Our goal is to ensure that the for any incentive compatible control $\hat{\mathcal{C}}$ the FBSDE system is well-posed and therefore has a unique and continuous solution.

Given that the system is decoupled, first we write down the generator of the backward system. Since we only consider Markovian controls that satisfy incentive compatibility, the generator only depends on the forward part:

$$R^t(s)(\hat{a}_s - \hat{c}_s).$$

Based on Cvitanic and Zhang (2012), Section 9.5, the FBSDE is well-posed if the both the forward and backward components have unique solutions and the solution to the forward component satisfies the Markov property. The uniqueness and Markov property of the forward component comes from Theorem 3; thus Lemma 1 admits an unique, Markovian solution. Regarding the backward component, we observe that it does not have the backward terms in the generator but only the forward ones, and the generator $R^t(s)(\hat{a}_s - \hat{c}_s)$ is trivially uniformly Lipschitz continuous in $f^t_{\hat{\mathcal{C}}}(t,W)$ and Y_t. Thus, by Theorem 9.3.5 in Cvitanic and Zhang (2012) the backward system has a unique solution.[29] By Theorem 3 the decoupling field is $C^{1,2}$, and since we assumed the discount function is C^3, the generator itself is C^3, thus by Corollary 2.9 of Pardoux and Peng (1992), $f^t_{\hat{\mathcal{C}}}(t,W)$ has continuous partial derivatives of

[28] For a textbook treatment of FBSDE, we refer the readers to Ma et al. (1999).

[29] The fact that backward terms do not appear in generator is not surprising when the backward system is a sort of "continuation utility", see Duffie and Epstein (1992), El Karoui et al. (1997).

order 1 and 2 in t and W respectively.[30]

A.4.2 Finding a Fixed Point to The Extended HJB System

The proceeding argument proves that for any incentive compatible control $\hat{\mathcal{C}}$ there exists a unique process $f^t_{\hat{\mathcal{C}}}(t, W)$ satisfying the backward system. Furthermore, for any time t under the optimal controls we have :

$$V_{\hat{\mathcal{C}}}(t, W) = f^t_{\hat{\mathcal{C}}}(t, W) = f_{\hat{\mathcal{C}}}(t, W, t), \tag{41}$$

$$\mathcal{A}^{\tilde{\mathcal{C}}} V_{\hat{\mathcal{C}}}(t, W) = \mathcal{A}^{\tilde{\mathcal{C}}} f_{\hat{\mathcal{C}}}(t, W, t) \neq \mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\hat{\mathcal{C}}}(t, W). \tag{42}$$

The second line stems from the fact that for $\mathcal{A}^{\tilde{\mathcal{C}}} V_{\hat{\mathcal{C}}}(t, W)$ and $\mathcal{A}^{\tilde{\mathcal{C}}} f_{\hat{\mathcal{C}}}(t, W, t)$, the infinitesimal generator $\mathcal{A}^{\tilde{\mathcal{C}}}$ changes both change both the self and time, whereas $\mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\hat{\mathcal{C}}}(t, W)$ is taken from a fixed self's perspective.[31] Therefore, under the suprema we can cancel out those terms; hence, self's t problem is reduced to the static optimization problem for each W,

$$\sup_{\tilde{\mathcal{C}}} \{(\tilde{a}_t - \tilde{c}_t) - \mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\hat{\mathcal{C}}}(t, W)\} = 0. \tag{43}$$

For any given control $\hat{\mathcal{C}}$, the solution to the equation above generates another control $\tilde{\mathcal{C}}$ using the value function generated from the backward system, $f^t_{\hat{\mathcal{C}}}(t, W)$.

Next, we define a probability space $(\Omega, \mathcal{P}, \mathcal{F})$ where $\Omega = Z_{[0,T]}$ and Z is the Brownian motion. We denote ω as an arbitary realized path of Z. Let $S = A \times C \times \mathbb{R}$ denote the space of strategies, $[0, T]$ the space of players, and $\mathcal{P}$ is common knowledge among all players. In particular, we treat the incomplete observations of the realized path by each each player as her differential information. The information of a player t is a sub $\sigma-$ algebra of $\mathcal{F}$, which is denoted by $\mathcal{F}_t$, and corresponds naturally to the filtration $\mathcal{F}$ given that players' different information is a result of the difference in time.

Let $\mathfrak{M}$ denote the space of all measures on S. Let $\mathfrak{S}$ denote the set of all measurable decision rules δ, $\delta : \Omega \to \mathfrak{M}$ (i.e. players are allowed to randomize). We equip $\mathfrak{S}$ with the weak topology, which is the weakest topology in which the functions identified below are continuous on $\mathfrak{S}$:

$$\delta \to \int_\Omega \phi(\omega) \left[\int_S c(s) \delta(\omega) ds \right] P(d\omega), \phi \in L^1(\Omega, \mathcal{P}, \mathcal{F}), c \in C_B(S).$$

[30]Since the decoupling field is $C^{1,2}$ by construction we do not need to make further assumptions about the drift and variance of the forward process.

[31]Note that under suboptimal but incentive compatible contracts the first equality does not hold either.

Here, $L^1(\Omega, \mathcal{P}, \mathcal{F})$ denotes the space of $\mathcal{P}$-integrable functions and $C_B(A)$ the space of bounded and continuous functions. We equip $(\Omega, \mathcal{P})$ with the usual L^1 norm.

With this notation, a pure strategy profile is a distribution over $\mathfrak{S}$, and we denote the set of all probability distributions over $\mathfrak{S}$ as $M(\mathfrak{S})$. For each t, we define $\mathfrak{S}_t \subset \mathfrak{S}$ as the set of all $\mathcal{F}_t$ measurable decision rules δ, $\delta : \Omega \to \mathfrak{M}$. Then, we define the set

$$D = \{(\mathcal{F}_t, \delta) : \mathcal{F}_t \in \{\mathcal{F}_t\}_{t\in[0,T]} \text{ and } \delta \in \mathfrak{S}_t\}.$$

Since $\{\mathcal{F}_t\}_{t\in[0,T]}$ is a filtration, $\mathcal{F}$ is the appropriate σ-algebra in which the conditional expectations are measurable. Due to lemma 2 of Balder (1991) D is $\mathcal{F} \times \mathcal{B}(\mathfrak{S})$ measurable and $\mathfrak{S}_t$ is a compact subset of $\mathfrak{S}$ for every $\mathcal{F}_t \in \{\mathcal{F}_t\}_{t\in[0,T]}$.[32] That is when using strategies from D, player t only uses information from $\mathcal{F}_t$.

Note that by definition any contract $\mathcal{C}$ defined as in the main text satisfies $\mathcal{C} \in M(\mathfrak{S})$ with the additional property that each (a_t, c_t) suggested by $\mathcal{C}$ such that $(a_t, c_t) \in \mathfrak{S}_t \ \forall t$. For any given t, ω and $\mathcal{F}_t$, using the backward system for any contract $\mathcal{C}$, following holds:

$$f^t_{\mathcal{C}}(t, W) = E^{\mathcal{C}}_{t,W}\left[\int_t^T R^t(s)(a_s - c_s)ds + B(t, W_T)\right].$$

The above equation only depends on $\mathcal{C}$ and ω and is calculated according to the information at $\mathcal{F}_t$. Meanwhile, for any player t, any $\delta \in \mathfrak{S}$ and $h \in M(\mathfrak{S})$, we can define the utility function as:

$$u_t(\delta, \mathcal{C}|\mathcal{F}_t) = \mathbb{E}\left[(a(\delta) - c(\delta)) - \mathcal{A}^{(a(\delta),c(\delta))} f^t_{\mathcal{C}}(t, W)|\mathcal{F}_t\right].$$

We observe that u_t is continuous in both its arguments and is measurable by definition. Furthermore, f is also continuous and differentiable since the backward system is well posed. Notice that with the randomness generated by the Z feed into the $f^t_{\mathcal{C}}(t, W)$, a player is identified by the characteristics $(t, \mathcal{F}_t)$, however, we do not restrict δ to be $\mathcal{F}_t$ measurable.

Analogous to Balder (1991) we can define the *utility function for the game* as a function $U : [0, T] \times D \times M(\mathfrak{S}) \to \mathbb{R}$. Therefore, a game is identified by $([0, T] \times \mathcal{F}, \mu, \{\mathfrak{S}_t\}_{t\in[0,T]}, U\}$, where $[0, T] \times \mathcal{F}$ denotes the set of characteristics, μ is a given distribution over characteristics, $\mathfrak{S}_t$ denotes the set of strategies available to a player with characteristics $(t, \mathcal{F}_t)$ and U is the utility function for the game.

Next, we can identify a *characteristic - strategy* (CS) distribution, which is a distribution over $[0, T] \times D$ specifying how the possible characteristic-strategy combinations are distributed in the game. If λ is a CS distribution then the respective marginals of the dis-

[32] Where $\mathcal{B}(\mathfrak{S})$ is the Borel sigma algebra defined on $\mathfrak{S}$.

tribution $\lambda_{|[0,T]\times\mathcal{F}}$ coincide with the distribution over characteristics and $\lambda_{|\mathfrak{S}}$ coincides with a distribution over the strategies.

Definition 9 *A CS distribution λ is an equilibrium if*

- *The marginal of $\lambda_{|[0,T]\times\mathcal{F}} = \mathcal{P} \times \lambda_U^T$,*
- $\lambda(\{((t,\mathcal{F}_t),\delta) \in [0,T] \times D : \delta \in \arg\max_{\mathfrak{S}_t} u_t(\delta, \lambda_{|\mathfrak{S}})\}) = 1,$

where λ_U^T denotes the uniform distribution over $[0,T]$.

We observe that in the game between the type t selves the distribution over $[0,T] \times \mathcal{F}$ is identified by the filtered probability space generated by $Z_{[0,T]}$. The marginal distribution over [0,T] λ is uniform since each principal is weighted equally. Now, we observe that $\mathfrak{S}$ is defined as above complete and separable, $\mathfrak{S}_t$ is compact as noted above, $U(t,\mathcal{F}_t,\cdot,\cdot)$ is continuous in both arguments and $U(\cdot,\cdot,\cdot,\nu)$ for $\nu \in \mathfrak{S}$ is measurable, hence by Theorem 1 in Balder (1991), there exists an equilibrium distribution. Hence there exists a contract $\mathcal{C}^*$ such that for every t, W the following is satisfied

$$(a^*, c^*) \in \arg\sup_{\tilde{\mathcal{C}}}\{\tilde{a}_t - \tilde{c}_t - \mathcal{A}^{\tilde{\mathcal{C}}} f^t_{\mathcal{C}^*}(t,W)\} = 0.$$

This leads to the conclusion that there exists a solution to the extended HJB system.

A.4.3 Verification Theorem

We divide the verification theorem into two parts. First, we show that if a contract $\mathcal{C}$ solves the extended HJB system, it must be an equilibrium of the game. Second, we show that if a contract $\mathcal{C}$ is an equilibrium of the game, then it must also solve the extended HJB system. Following proposition formalizes the verification result.

Proposition 7 *Assume a, c, V and $f_{\mathcal{C}}$ are C^2 with respect to W and C^1 with respect to t, then the following statements hold.*

1. *If a contract $\mathcal{C}$ solves the extended HJB equation, then it must be an equilibrium.*
2. *If a contract $\mathcal{C}$ is an equilibrium, then it also solves the extended HJB equation.*

Proof of Propoisition 7: The proof of the first statement is organized into two steps:
Step 1: By applying, Ito's formula to $f^s_{\mathcal{C}}(t,W)$ by 20 and using the boundary condition at time T, we obtain the representation

$$f^s_{\mathcal{C}}(t,W) = E^{\mathcal{C}}_{t,W}\left[\int_t^T R^s(r)(a_r - c_r)dr + B(t,W_T)\right].$$

From the extended HJB system

$$0 = \mathcal{A}^{\mathcal{C}} V(t,W) + \mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}^t(t,W) - \mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}(t,W,t) + a_t - c_t, \tag{44}$$

$$0 = a_t - c_t + \mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}^t(t,W), \tag{45}$$

which imply

$$\mathcal{A}^{\mathcal{C}} V(t,W) = \mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}(t,W,t),$$

for all t and W. Since V is smooth, by Ito's lemma:

$$EV(T,W_T) = V(t,W) + E\left[\int_t^T \mathcal{A}^{\mathcal{C}} V(s,W_s)ds\right],$$

which can be re-written as

$$EV(T,W_T) = V(t,W) + E\left[\int_t^T \mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}(s,W,s)ds\right].$$

By applying the same reasoning as above to $f^t(t,W)$ and using the boundary conditions for V and f time T, we conclude that

$$V(t,W) = f_{\mathcal{C}}(t,W,t).$$

Step 2: Next, we show that $\mathcal{C}$ is an equilibrium. Suppose the agent uses an arbitrary control law $\hat{\mathcal{C}}$ over period length $\Delta > 0$. Let $f_{\mathcal{C}_\Delta}(t,W,t)$ be the payoff of the principal under C_Δ; then in order to $\mathcal{C}$ be an equilibrium the following inequality must hold for any $\mathcal{C}_\Delta$

$$\liminf_{\Delta\to 0} \frac{f_{\mathcal{C}}(t,W,t) - f_{\mathcal{C}_\Delta}(t,W,t)}{\Delta} \geq 0. \tag{46}$$

We have

$$\frac{f_{\mathcal{C}}(t,W,t) - f_{\mathcal{C}_\Delta}(t,W,t)}{\Delta} = (a_t - c_t) - (a_t^\Delta - c_t^\Delta) - \mathcal{A}^{\mathcal{C}_\Delta} f_{\mathcal{C}}^t(t,W) + \mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}^t(t,W).$$

Since $\mathcal{C}$ solves the V equation $(a_t - c_t) + \mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}^t(t,W,t)$ is equal to zero. Since supremum is equal to zero $(a_t^\Delta - c_t^\Delta) + \mathcal{A}^{\mathcal{C}_\Delta} f_{\mathcal{C}}^t(t,W)$ must be negative which implies $f_{\mathcal{C}}(t,W,t) - f_{\mathcal{C}_\Delta}(t,W,t)$ is positive. Hence, contract $\mathcal{C}$ must be an equilibrium.

In this part, we prove the second statement of Proposition 7. Assume that $\hat{\mathcal{C}}$ is a continuous equilibrium contract and the corresponding value function V and f are sufficiently smooth ($V, f \in C^2$), then V solves the extended HJB equation. We begin with proving two auxiliary results. For a given contract $\mathcal{C}$, we define $V_{\mathcal{C}}(t,W)$, $f_{\mathcal{C}}^t(t,W)$, $f_{\mathcal{C}}(t,W,t)$. Using these notation we first establish a Feynman-Kac formulation of our problem as follows:

Lemma 3 *Consider a continuous contract $\mathcal{C}$ and moreover assume $f^t_{\mathcal{C}}$ is twice continuously differentiable and L^2 integrable. Then for every $s \in [0,T]$ $f^s_{\mathcal{C}}$ is a solution to the following PDE*

$$\mathcal{A}^{\mathcal{C}} f^s_{\mathcal{C}}(t,W) + R^s(t)\,(a_t - c_t) = 0 \quad (t,W) \in [0,T] \times \mathbb{R}.$$

Proof of Lemma 3: By construction, following is true

$$f^s_{\mathcal{C}}(t,W) = E^{\mathcal{C}}_{t,W}\left[\int_t^T R^s(k)\,(a_k - c_k)\,dk + B(t,W_T)\right], \tag{47}$$

and the boundary condition $f^s(T,W) = B(s,W)$ is satisfied. Fix Δ small enough then following is true

$$0 = E^{\mathcal{C}}_{t,W}\left[\int_t^{t+\Delta} R^s(r)\,(a_r - c_r)\,dr\right] + E^{\mathcal{C}}_{t,W}\left[f^s_{\mathcal{C}}(t+\Delta, W_{t+\Delta})\right] - f^s_{\mathcal{C}}(t,W).$$

Then, applying Ito's formula to the term $f^s_{\mathcal{C}}(t+\Delta, W_{t+\Delta})$ we reach

$$\begin{aligned} 0 = {} & E^{\mathcal{C}}_{t,W}\left[\int_t^{t+\Delta} R^s(r)\,(a_r - c_r)\,dr\right] + f^s_{\mathcal{C}}(t,W) + E^{\mathcal{C}}_{t,W}\left[\int_t^{t+\Delta} \mathcal{A}^{\mathcal{C}} f^s_C(r,W_r)dr\right] \\ & + E^{\mathcal{C}}_{t,W}\left[\int_t^{t+\Delta} \frac{\partial^2 f}{\partial^2 W}(\sigma^{\mathcal{C}}_W)^2 dZ_r\right] - f^s_{\mathcal{C}}(t,W). \end{aligned}$$

After simplification and using the fact that Ito integral is equal to zero on expectation, we obtain the following

$$0 = E^{\mathcal{C}}_{t,W}\left[\int_t^{t+\Delta} R^s(r)\,(a_r - c_r)\,dr\right] + E^{\mathcal{C}}_{t,W}\left[\int_t^{t+\Delta} \mathcal{A}^{\mathcal{C}} f^s_C(r,W_r)dr\right].$$

By dividing both side to Δ we reach

$$0 = E^{\mathcal{C}}_{t,W}\left[\frac{\int_t^{t+\Delta} R^s(r)\,(a_r - c_r)\,dr}{\Delta}\right] + E^{\mathcal{C}}_{t,W}\left[\frac{\int_t^{t+\Delta} \mathcal{A}^{\mathcal{C}} f^s_C(r,W_r)dr}{\Delta}\right].$$

Then taking limit $\Delta \to 0$, then we reach $R^s(t)(a_t - c_t) + \mathcal{A}^{\mathcal{C}} f^s_C(t,W) = 0$.

Based on the previous lemma we obtain the following result:

Lemma 4 *Consider an arbitrary contract $\mathcal{C}$ and an arbitrary deviation $\mathcal{C}_\Delta$, then following is true:*

$$\lim_{\Delta \to 0} \frac{f_{\mathcal{C}}\,(t,W,t) - f_{\mathcal{C}_\Delta}\,(t,W,t)}{\Delta} = -(a^\Delta_t - c^\Delta_t) - \mathcal{A}^{\mathcal{C}_\Delta} f^t_{\mathcal{C}}(t,W).$$

Proof of Lemma 4: Using Ito's rule and Lemma 3 we obtain the following equality,

$$\lim_{\Delta \to 0} \frac{f_{\mathcal{C}}(t, W, t) - f_{C_\Delta}(t, W, t)}{\Delta} = (a_t - c_t) - (a_t^\Delta - c_t^\Delta) - \mathcal{A}^{\mathcal{C}_\Delta} f_{\mathcal{C}}^t(t, W) + \mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}^t(t, W).$$

We observe that by Lemma 3 the following is true $(a_t - c_t) + \mathcal{A}^C f_{\mathcal{C}}^t(t, W) = 0$. Therefore, we end up with the term $-(a_t^\Delta - c_t^\Delta) - \mathcal{A}^{\mathcal{C}_\Delta} f_{\mathcal{C}}^t(t, W)$.

We now use the two previous results to complete the proof of Part 2. We need to show that if $\hat{\mathcal{C}}$ is an equilibrium for any other contract $\tilde{\mathcal{C}}$, we have the following

$$\mathcal{A}^{\tilde{\mathcal{C}}} V_{\hat{\mathcal{C}}}(t, W) + (\tilde{a}_t - \tilde{c}_t) - \mathcal{A}^{\tilde{\mathcal{C}}} f_{\hat{\mathcal{C}}}(t, W, t) + \mathcal{A}^{\tilde{\mathcal{C}}} f_{\hat{\mathcal{C}}}^t(t, W) \leq 0.$$

For all (t, W), $\hat{\mathcal{C}}$ being an equilibrium contract implies $\mathcal{A}^{\tilde{\mathcal{C}}} V_{\hat{\mathcal{C}}}(t, W) = \mathcal{A}^{\tilde{\mathcal{C}}} f_{\hat{\mathcal{C}}}(t, W, t)$. Therefore, following inequality is sufficient for showing that any other contract cannot solve the extended HJB system.

$$(\tilde{a}_t - \tilde{c}_t) + \mathcal{A}^{\tilde{\mathcal{C}}} f_{\hat{\mathcal{C}}}^t(t, W) \leq 0.$$

By defining $\tilde{\mathcal{C}}$ as deviating from $\hat{\mathcal{C}}$ over Δ amount of time. By Lemma 4, we have the following result:

$$\lim_{\Delta \to 0} \frac{f_{\hat{\mathcal{C}}}(t, W, t) - f_{\hat{\mathcal{C}}_\Delta}(t, W, t)}{\Delta} = -(\hat{a}_\Delta - \hat{c}_\Delta) - \mathcal{A}^{\hat{\mathcal{C}}_\Delta} f_{\hat{\mathcal{C}}}^t(t, W).$$

$\hat{\mathcal{C}}$ being an equilibrium implies

$$\lim_{\Delta \to 0} \frac{f_{\hat{\mathcal{C}}}(t, W, t) - f_{\hat{\mathcal{C}}_\Delta}(t, W, t)}{\Delta} \geq 0.$$

This allows us to conclude that, $(\hat{a}_\Delta - \hat{c}_\Delta) + \mathcal{A}^{\hat{C}_\Delta} f_{\hat{\mathcal{C}}}^t(t, W) \leq 0$. Hence any deviation from the equilibrium contract $\mathcal{C}$ can not satisfy (18).

In the last step we verify that $\mathcal{C}$ solves the extended HJB system. By construction $\mathcal{A}^{\hat{\mathcal{C}}} V_{\hat{\mathcal{C}}}(t, W) = \mathcal{A}^{\hat{\mathcal{C}}} f_{\hat{\mathcal{C}}}(t, W, t)$ and $(\hat{a}_t - \hat{c}_t) + \mathcal{A}^{\hat{\mathcal{C}}} f_{\hat{\mathcal{C}}}^t(t, W) = 0$. Therefore,

$$\mathcal{A}^{\hat{\mathcal{C}}} V_{\hat{\mathcal{C}}}(t, W) + (\hat{a}_t - \hat{c}_t) - \mathcal{A}\cdot\hat{\mathcal{C}} f_{\hat{\mathcal{C}}}(t, W, t) + \mathcal{A}^{\hat{\mathcal{C}}} f_{\hat{\mathcal{C}}}^t(t, W) = 0.$$

Thus, we complete the proof of Part 2 of Proposition 7.

A.5 Discounting Functions that Depend Only on the Time Difference

In the case of discounting that only depend on the time difference, that is $R(s - t)$ instead of $R^t(s)$ we can simplify the extended HJB system. This simplified system would be helpful for applications. In particular we have:

$$\begin{aligned}
\mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}(t,W,t) &= E_{t,W}^{\mathcal{C}}\left[\int_t^T \big(R(s-t)\mathcal{A}^{\mathcal{C}}(a_s-c_s) - R'(s-t)(a_s-c_s)\big)\, ds\right] \\
&- R(T-t)\mathcal{A}^{\mathcal{C}} E_{t,W}^{\mathcal{C}}\left[c_T\right] + R'(T-t) E_{t,W}^{\mathcal{C}}\left[c_T\right],
\end{aligned}$$

and

$$\mathcal{A}^{\mathcal{C}} f_{\mathcal{C}}^t(t,W) = E_{t,W}^{\mathcal{C}}\left[\int_t^T R(s-t)\mathcal{A}^{\mathcal{C}}(a_s-c_s)ds\right] + R(T-t)E_{t,W}^{\mathcal{C}}\left[\mathcal{A}^{\mathcal{C}} c_T\right].$$

Therefore, the extended HJB system can be reduced to

$$\sup_{\mathcal{C}}\{\mathcal{A}^{\mathcal{C}} V + (a_t - c_t) + E_{t,W}^{\mathcal{C}}\left[\int_t^T R'(s-t)(a_s-c_s)ds\right] - R'(T-t)E_{t,W}^{\mathcal{C}}\left[c_T\right]\} = 0. \quad (48)$$

Proof of Proposition 3 Under Assumptions 3, 4 the value function of the principal satisfies the following HJB system:

$$\begin{aligned}
\sup_{a_t} V_t + a_t - \left[\frac{1}{2}a_t^2 - \frac{\ln(\eta\gamma)}{\eta} - \frac{1}{\eta}\ln(-W)\right] + \frac{1}{2}\left(\gamma\eta W\psi\sigma\right)^2 V_{WW} \\
+ \int_t^T R'(s-t)\left(a_s - \left[\frac{1}{2}a_s^2 - \frac{\ln(\eta\gamma)}{\eta} - \frac{1}{\eta}\ln(-W_s)\right]\right) ds \\
+ \frac{\ln(-W_T)}{\eta} R'(T-t) = 0,
\end{aligned}$$

subject to (22) and incentive compatibility condition (24) and boundary condition

$$V(T,W) = B(T,W) \text{ for all } W. \quad (49)$$

We are going to guess the principal's value function has the following functional form:

$$V(t,W) = A_t \ln(-W) + B_t,$$

with the boundary condition $A_T = \frac{1}{\eta}$. Given the guess,

$$\dot{V}_t = \dot{A}_t \ln(-W) + \dot{B}_t, V_W = \frac{A_t}{W}, V_{WW} = -\frac{A_t}{W^2}.$$

Substituting these back into the extended HJB system yields:

$$1 - a_t - a_t\eta^2\gamma^2\sigma^2 A_t = 0 \Rightarrow a_t = \frac{1}{1+\eta^2\gamma^2\sigma^2 A_t}.$$

Collecting the terms with log:

$$\dot{A}_t = \frac{1}{\eta}\left(1 + \int_t^T R'(s-t)\,ds + R'(T-t)\right).$$

By integrating and using the fact that $R(0) = 1$, we reach:

$$\dot{A}_t = \frac{1}{\eta}\left(R(T-t) + R'(T-t)\right).$$

Using the boundary condition $A_T = \frac{1}{\eta}$,

$$A_t = \frac{1}{\eta}\int_0^t \left(R(T-s) + R'(T-s)\right) ds + K.$$

Where $K = \frac{1}{\eta}(2 - \int_0^T R(T-s)ds - R(T))$ plugging this back in to a_t yields:

$$a_t = \frac{1}{1 + K + \eta\gamma^2\sigma^2\left[\int_0^t \left(R(T-s) + R'(T-s)\right) ds\right]}.$$

A.6 Proofs for the Case of Quasi-Hyperbolic Discounting

Substituting R and R' with the quasi hyperbolic discounting function to the A.5 yields:

$$\dot{A}_t = \frac{1}{\eta}\left(\beta(1-\rho)e^{-\rho(T-t)} + (1-\beta)(1-(\rho+\lambda))e^{-(\rho+\lambda)(T-t)}\right),$$

with boundary condition $A_T = \frac{1}{\eta}$. Integrating the above equation yields:

$$A_t = \frac{1}{\eta}\left(\frac{\beta(1-\rho)}{\rho}e^{-\rho(T-t)} + (1-\beta)\,\frac{1-(\rho+\lambda)}{\rho+\lambda}e^{-(\rho+\lambda)(T-t)} + \frac{2\rho(\rho+\lambda) - (\beta\lambda+\rho)}{\rho(\rho+\lambda)}\right).$$

Thus equilibrium action is given by:

$$a_t = \frac{1}{1 + \eta^2\gamma^2\sigma^2\left(\frac{1}{\eta}\left(\frac{\beta(1-\rho)}{\rho}e^{-\rho(T-t)} + (1-\beta)\,\frac{1-(\rho+\lambda)}{\rho+\lambda}e^{-(\rho+\lambda)(T-t)} + \frac{2\rho(\rho+\lambda)-(\beta\lambda+\rho)}{\rho(\rho+\lambda)}\right)\right)}.$$

A.6.1 Proof of Proposition 4

First, notice that $\psi_t = a_t = \frac{1}{1+\eta^2\gamma^2\sigma^2 A_t}$. Hence $\frac{\partial \psi_t}{\partial \beta}$ has the opposite sign of $\frac{\partial A_t}{\partial \beta}$, and the latter can be solved explicitly:

$$\begin{aligned}
\frac{\partial A_t}{\partial \beta} &= \frac{1}{\eta}\left(\frac{(1-\rho)}{\rho}e^{-\rho(T-t)} - \frac{1-(\rho+\lambda)}{\rho+\lambda}e^{-(\rho+\lambda)(T-t)} + \frac{-\lambda}{\rho(\rho+\lambda)}\right), \\
&= \frac{(1-\rho)(\rho+\lambda)e^{-\rho(T-t)} - \rho(1-(\rho+\lambda))e^{-(\rho+\lambda)(T-t)} - \lambda}{\rho(\rho+\lambda)}, \\
&= \frac{(\rho-\rho^2-\rho\lambda+\lambda)e^{-\rho(T-t)} - (\rho-\rho^2-\rho\lambda)e^{-(\rho+\lambda)(T-t)} - \lambda}{\rho(\rho+\lambda)}, \\
&= \frac{(\rho-\rho^2-\rho\lambda)(e^{-\rho(T-t)} - e^{-(\rho+\lambda)(T-t)}) + \lambda e^{-\rho(T-t)} - \lambda}{\rho(\rho+\lambda)}, \\
&= \frac{-(\rho^2+\rho\lambda-\rho)(e^{-\rho(T-t)} - e^{-(\rho+\lambda)(T-t)}) + \lambda(1-e^{-\rho(T-t)})}{\rho(\rho+\lambda)}, \\
&= -\frac{(\rho^2+\rho\lambda-\rho)(e^{-\rho(T-t)} - e^{-(\rho+\lambda)(T-t)}) + \lambda(1-e^{-\rho(T-t)})}{\rho(\rho+\lambda)}, \\
&= -\frac{((\rho+\lambda)(\rho-1)+\lambda)(e^{-\rho(T-t)} - e^{-(\rho+\lambda)(T-t)}) + \lambda(1-e^{-\rho(T-t)})}{\rho(\rho+\lambda)}, \\
&= -\frac{((\rho+\lambda)(\rho-1))(e^{-\rho(T-t)} - e^{-(\rho+\lambda)(T-t)}) + \lambda(1-e^{-(\rho+\lambda)(T-t)})}{\rho(\rho+\lambda)}.
\end{aligned}$$

The sign of the numerator depends on the parameter values. We discuss their different cases below:

Case 1 $\rho + \lambda \geq 1$

There are two sub-cases in this scenario: i): $\rho \geq 1$, ii)$\rho < 1$. In subcase i) it is easy to see that $((\rho+\lambda)(\rho-1))(e^{-\rho(T-t)} - e^{-(\rho+\lambda)(T-t)}) + \lambda(1-e^{-(\rho+\lambda)(T-t)}) \geq 0$ since both the summands are positive, thus with the negative sign we have $\frac{\partial A_t}{\partial \beta} \leq 0$ for all t and W_t. In sub-case ii) rearranging the numerator yields the following summation, where the first term is positive and the second term is negative.

$$\frac{e^{-\rho(T-t)} - e^{-(\rho+\lambda)(T-t)}}{1-e^{-(\rho+\lambda)(T-t)}} + \frac{\lambda}{(\rho+\lambda)(1-\rho)}, \tag{50}$$

where the first summand is less than 1 and simple algebraic manipulations shows that the second summand is greater than 1 when $\rho+\lambda \geq 1$. Thus the numerator is positive, implying that $\frac{\partial A_t}{\partial \beta} > 0$, which ,in turn, implies $\frac{\partial A_t}{\partial \beta} \leq 0$ for all t and W_t.

Case 2 $\rho + \lambda < 1$

Re-arranging the numerator

$$\frac{e^{-\rho(T-t)} - e^{-(\rho+\lambda)(T-t)}}{1 - e^{-(\rho+\lambda)(T-t)}} + \frac{\lambda}{(\rho+\lambda)(1-\rho)}. \tag{51}$$

When $\rho + \lambda < 1$, both of the summands are always less than 1. For the first summand, note that for any T, one of the following must be true: if $\frac{e^{-\rho(T)} - e^{-(\rho+\lambda)(T)}}{1-e^{-(\rho+\lambda)(T)}} < \frac{\lambda}{(\rho+\lambda)(1-\rho)}$, then $\frac{\partial A_t}{\partial \beta}$ is always negative for all t and all W_t; if $\frac{e^{-\rho(T)} - e^{-(\rho+\lambda)(T)}}{1-e^{-(\rho+\lambda)(T)}} > \frac{\lambda}{(\rho+\lambda)(1-\rho)}$, then there exists some $t < T$ such that for all $t' \leq t$ and all $W_{t'}$, $\frac{\partial A_{t'}}{\partial \beta} \geq 0$ and for all $t'' > t$ and $W_{t''}$ $\frac{\partial A_{t''}}{\partial \beta} < 0$.

To prove the existence of $T(\lambda, \rho)$, we observe $\frac{e^{-\rho(T)} - e^{-(\rho+\lambda)(T)}}{1-e^{-(\rho+\lambda)(T)}} \to 0$ as $T \to \infty$, and $\to \lambda/(1+\lambda+\rho)$ (following L'hopital's Rule) as $\to 0$ with a derivative that does not change sign. Thus, for each ρ and λ, there exists a unique $T(\lambda, \rho)$ large enough such that the summands in (51) sum up to 0.

Similarly, given the equilibrium wage $c_t = \frac{1}{2}a_t^2 - \frac{\ln(\gamma\eta)}{\eta} - \frac{1}{\eta}\ln(-W_t)$. With a slight abuse of notation, letting $a_t(\beta)$ denoting the optimal a_t as a function of β we have

$$\frac{\partial c_t}{\partial \beta} = 2a_t(\beta)\frac{\partial a_t}{\partial \beta}.$$

We observe that $a_t(\beta) \geq 0$ for all β thus $dc/d\beta$ has the same sign as $\frac{a_t}{\beta}$ which is negatively correlated to A_t. Therefore, c_t is negatively related A_t.

A.6.2 Proof of Proposition 5

Part 1) By inspection, if $\beta = 1$, $\dot{A}_t = \frac{1}{\eta}(\rho + \lambda - 1)$. Similarly, if $\beta = 0$ $\dot{A}_t = \frac{1}{\eta}(\rho - 1)$ which is equivalent to time-consistent solutions.

Part 2) Fix any $t < T$,

$$\lim_{T\to\infty} a_t = \frac{1}{1 + \eta^2\gamma^2\sigma^2\frac{1}{\eta}\left(\frac{2\rho(\rho+\lambda)-(\beta\lambda+\rho)}{\rho(\rho+\lambda)}\right)}.$$

which corresponds to the solution in Holmstrom and Milgrom (1987) with a discount rate of $\frac{\rho(\lambda+\rho)}{\beta\lambda+\rho}$.

Parts 3) Recall that the closed-form solution for a_t is:

$$a_t = \frac{1}{1 + \eta^2\gamma^2\sigma^2\frac{1}{\eta}\left(\frac{\beta(1-\rho)}{\rho}e^{-\rho(T-t)} + (1-\beta)\frac{1-(\rho+\lambda)}{\rho+\lambda}e^{-(\rho+\lambda)(T-t)} + \frac{2\rho(\rho+\lambda)-(\beta\lambda+\rho)}{\rho(\rho+\lambda)}\right)}.$$

The dynamics of a_t are pinned down by the term $\beta\frac{(1-\rho)}{\rho}e^{-\rho(T-t)} + (1-\beta)\frac{1-(\rho+\lambda)}{\rho+\lambda}e^{-(\rho+\lambda)(T-t)}$.

It's time derivative equals to

$$\beta(1-\rho)e^{-\rho(T-t)} + (1-\beta)(1-(\rho+\lambda))e^{-(\rho+\lambda)(T-t)}. \tag{52}$$

Clearly, if the sign of (52) changes as time changes, it must happen at a unique point. In particular, for t small enough the first term dominates and (52) is always positive, but for t large enough, the second term dominates, and if $(1-(\rho-\lambda))(1-\beta)$ is negative, (52) can be negative.

A.7 Proofs for Anticipatory Utility Case

Under assumption 5 we can calculate how a payoff in period τ is valued in the stream $a_\tau - c_\tau$ as follows:

$$\int_t^\tau e^{-\rho(s-t)}e^{-\zeta(\tau-s)}E(a_\tau - c_\tau|\mathcal{F}_t)ds = \frac{e^{-\rho(\tau-t)} - e^{-\zeta(\tau-t)}}{\zeta-\rho}E(a_\tau - c_\tau|\mathcal{F}_t).$$

Therefore, anticipatory utility is equivalent to the following discounting:

$$R^t(s) = R(s-t) = \frac{e^{-\rho(\tau-t)} - e^{-\zeta(\tau-t)}}{\zeta-\rho}.$$

Substituting in R and R' from anticipatory utility to the equation A.5 yields

$$\dot{A}_t = \frac{1}{\eta}\left(\frac{(1-\rho)e^{-\rho(T-t)} - (1-\zeta)e^{-\zeta(T-t)}}{\zeta-\rho} + (1-\rho)e^{-\rho(T-t)}\right).$$

Using the boundary condition $A_T = 1/\eta$ and integrating with respect to t implies:

$$A_t = \frac{1}{\eta}\left(\frac{(\frac{1}{\rho}-1)e^{-\rho(T-t)} - (\frac{1}{\zeta}-1)e^{-\zeta(T-t)}}{\zeta-\rho} + \frac{(1-\rho)}{\rho}e^{-\rho(T-t)} + \frac{2\rho\zeta-1-\zeta}{\rho\zeta}\right).$$

Thus, equilibrium action a_t is thus given by:

$$a_t = \frac{1}{1+\eta\gamma^2\sigma^2\left(\frac{(\frac{1}{\rho}-1)e^{-\rho(T-t)}-(\frac{1}{\zeta}-1)e^{-\zeta(T-t)}}{\zeta-\rho} + \frac{(1-\rho)}{\rho}e^{-\rho(T-t)} + \frac{2\rho\zeta-1-\zeta}{\rho\zeta}\right)}.$$

A.7.1 Proof of Proposition 6

Part 1) Fix any $t < T$, then

$$\lim_{\zeta \to \infty} a_t = \frac{1}{1 + \eta\gamma^2\sigma^2 \frac{2\rho - 1 + (1-\rho)e^{-\rho(T-t)}}{\rho}},$$

which corresponds to the solution in Holmstrom and Milgrom (1987) with a discount rate equals ρ.

Part 2) As $T \to \infty$, solving for r yields:

$$\frac{2r-1}{r} = \frac{2\rho\zeta - 1 - \zeta}{\rho\zeta} \Rightarrow r = \frac{\rho\zeta}{1+\zeta}.$$

Part 3) Recall that the closed-form of a_t is given by

$$a_t = \frac{1}{1 + \eta\gamma^2\sigma^2 \left(\frac{(\frac{1}{\rho}-1)e^{-\rho(T-t)} - (\frac{1}{\zeta}-1)e^{-\zeta(T-t)}}{\zeta - \rho} + \frac{(1-\rho)}{\rho}e^{-\rho(T-t)} + \frac{2\rho\zeta - 1 - \zeta}{\rho\zeta} \right)}.$$

The dynamics of a_t are pinned down by the term

$$\frac{(\frac{1}{\rho} - 1)e^{-\rho(T-t)} - (\frac{1}{\zeta} - 1)e^{-\zeta(T-t)}}{\zeta - \rho} + \frac{(1-\rho)}{\rho}e^{-\rho(T-t)};$$

the time derivative equals:

$$\frac{(1-\rho)e^{-\rho(T-t)} - (1-\zeta)e^{-\zeta(T-t)}}{\zeta - \rho} + (1-\rho)e^{-\rho(T-t)}.$$

Observe that if the time derivative ever becomes 0 we must obtain

$$\begin{aligned}
(1-\rho)e^{-\rho(T-t)} - (1-\zeta)e^{-\zeta(T-t)} &= (\zeta - \rho)(\rho - 1)e^{-\rho(T-t)}, \\
(1 + \zeta - \rho)(1-\rho)e^{-\rho(T-t)} &= (1-\zeta)e^{-\zeta(T-t)}, \\
1 &= \frac{\zeta - 1}{\rho - 1}\frac{1}{1 + \zeta - \rho}e^{-(\zeta-\rho)(T-t)}.
\end{aligned}$$

Next, since we assume that $\zeta > \rho > 1$, we have $\frac{\zeta-1}{\rho-1}\frac{1}{1+\zeta-\rho} > 1$ and $1 \geq e^{-(\zeta-\rho)(T-t)} \geq e^{-(\zeta-\rho)T} > 0$. Observe that if T is small enough, the derivative is always negative. However, for any ζ, ρ if T is sufficiently large, there exists a $t^\star < T$ such that the derivative becomes negative at for all $t > t^\star$.

References

Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin 82*(4), 463.

Balder, E. J. (1991). On Cournot-Nash equilibrium distributions for games with differential information and discontinuous payoffs. *Economic Theory 1*(4), 339–354.

Benoit, J.-P. and V. Krishna (1993). Renegotiation in finitely repeated games. *Econometrica: Journal of the Econometric Society*, 303–323.

Bernheim, B. D. and D. Ray (1989). Collective dynamic consistency in repeated games. *Games and Economic Behavior 1*(4), 295–326.

Bernheim, B. D., D. Ray, and Ş. Yeltekin (2015). Poverty and self-control. *Econometrica 83*(5), 1877–1911.

Biais, B., T. Mariotti, G. Plantin, and J.-C. Rochet (2007). Dynamic security design: Convergence to continuous time and asset pricing implications. *Review of Economic Studies 74*, 345–390.

Biais, B., T. Mariotti, J.-C. Rochet, and S. Villeneuve (2010). Large risks, limited liability, and dynamic moral hazard. *Econometrica 78*, 73–118.

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

Bismut, J.-M. (1978). Duality methods in the control of densities. *SIAM Journal on Control and Optimization 16*(5), 771–777.

Björk, T., M. Khapko, and A. Murgoci (2017). On time-inconsistent stochastic control in continuous time. *Finance and Stochastics 21*(2), 331–360.

Bloedel, A., R. V. Krishna, and B. Strulovici (2020). Persistent private information revisited. Working paper.

Bolton, P., H. Chen, and N. Wang (2011). A Unified Theory of Tobin's q, Corporate Investment, Financing, and Risk Management. *Journal of Finance 66*, 1545–1578.

Bond, P. and G. Sigurdsson (2018). Commitment contracts. *Review of Economic Studies 85*(1), 194–222.

Brunnermeier, M. K., F. Papakonstantinou, and J. A. Parker (2016). Optimal time-inconsistent beliefs: Misplanning, procrastination, and commitment. *Management Science 63*(5), 1318–1340.

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

Caplin, A. and J. Leahy (2001). Psychological expected utility theory and anticipatory feelings. *Quarterly Journal of Economics 116*(1), 55–79.

Caplin, A. and J. Leahy (2004). The social discount rate. *Journal of Political Economy 112*(6), 1257–1268.

Chade, H., P. Prokopovych, and L. Smith (2008). Repeated games with present-biased preferences. *Journal of Economic Theory 139*(1), 157–175.

Cisternas, G. (2018). Two-sided learning and the ratchet principle. *Review of Economic Studies 85*(1), 307–351.

Cvitanic, J. and J. Zhang (2012). *Contract Theory in Continuous-Time Models.* Springer Science & Business Media.

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

DeMarzo, P., M. Fishman, Z. He, and N. Wang (2012). Dynamic Agency and the Q Theory of Investment. *Journal of Finance 67.*

DeMarzo, P. and Y. Sannikov (2006). Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance 61*, 2681–2724.

Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. *Econometrica*, 353–394.

Ekeland, I. and A. Lazrak (2010). The golden rule when preferences are time inconsistent. *Mathematics and Financial Economics 4*(1), 29–55.

El Karoui, N., S. Peng, and M. C. Quenez (1997). Backward stochastic differential equations in finance. *Mathematical Finance 7*(1), 1–71.

Faingold, E. and Y. Sannikov (2011). Reputation in continuous-time games. *Econometrica 79*(3), 773–876.

Feng, F. Z. and M. M. Westerfield (2020). Dynamic resource allocation with hidden volatility. *Journal of Financial Economics*. forthcoming.

Fudenberg, D., B. Holmstrom, and P. Milgrom (1990). Short-term contracts and long-term agency relationships. *Journal of Economic Theory 51*(1), 1–31.

Fudenberg, D. and J. Tirole (1990). Moral hazard and renegotiation in agency contracts. *Econometrica*, 1279–1319.

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

Georgiadis, G. (2014). Projects and team dynamics. *Review of Economic Studies 82*(1), 187–218.

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

Gottlieb, D. and X. Zhang (2020). Long-term contracting with time-inconsistent agents. *Econometrica*. Forthcoming.

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

Grubb, M. D. (2015). Overconfident consumers in the marketplace. *Journal of Economic Perspectives 29*(4), 9–36.

Gryglewicz, S. and B. Hartman-Glaser (2019). Investment timing and incentive costs. *Review of Financial Studies*. forthcoming.

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

Harris, C. and D. Laibson (2012). Instantaneous gratification. *Quarterly Journal of Economics 128*(1), 205–248.

Harstad, B. (2016). Technology and time inconsistency. Working paper.

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

He, Z., B. Wei, J. Yu, and F. Gao (2017). Optimal long-term contracting with learning. *Review of Financial Studies 30*(6), 2006–2065.

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 present-bias from the timing of choices. *Available at SSRN 3386017*.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *Review of Economic Studies 66*(1), 169–182.

Holmstrom, B. and P. Milgrom (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica: Journal of the Econometric Society*, 303–328.

Jackson, M. O. and L. Yariv (2014). Present bias and collective dynamic choice in the lab. *American Economic Review 104*(12), 4184–4204.

Jackson, M. O. and L. Yariv (2015). Collective dynamic choice: the necessity of time inconsistency. *American Economic Journal: Microeconomics 7*(4), 150–78.

Kolb, A. and E. Madsen (2019). Leaks, sabotage, and information design. Working paper.

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

Krusell, P. and A. A. Smith, Jr (2003). Consumption–savings decisions with quasi–geometric discounting. *Econometrica 71*(1), 365–375.

Kydland, F. E. and E. C. Prescott (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy 85*(3), 473–491.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics 112*(2), 443–478.

Li, X., Y.-H. Ni, and J.-F. Zhang (2017). On time-consistent solution to time-inconsistent linear-quadratic optimal control of discrete-time stochastic systems.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal 97*(387), 666–684.

Loewenstein, G., T. O'Donoghue, and M. Rabin (2003). Projection bias in predicting future utility. *Quarterly Journal of Economics 118*(4), 1209–1248.

Ma, J., J.-M. Morel, and J. Yong (1999). *Forward-backward stochastic differential equations and their applications.* Number 1702. Springer Science & Business Media.

Marinovic, I. and F. Varas (2018). CEO horizon, optimal pay duration, and the escalation of short-termism. *Journal of Finance*. forthcoming.

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

Pardoux, E. and S. Peng (1992). Backward stochastic differential equations and quasilinear parabolic partial differential equations. In *Stochastic partial differential equations and their applications*, pp. 200–217. Springer.

Peng, S. (1990). A general stochastic maximum principle for optimal control problems. *SIAM Journal on control and optimization 28*(4), 966–979.

Pollak, R. A. (1968). Consistent planning. *Review of Economic Studies 35*(2), 201–208.

Ray, D. (1994). Internally renegotiation-proof equilibrium sets: Limit behavior with low discounting. *Games and Economic Behavior 6*(1), 162–177.

Ray, D., N. Vellodi, and R. Wang (2017). Backward discounting. Technical report, Working Paper, New York University.[1546].

Rey, P. and B. Salanie (1990). Long-term, short-term and renegotiation: On the value of commitment in contracting. *Econometrica*, 597–619.

Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *Review of Economic Studies 75*, 957–984.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies 23*(3), 165–180.

Strulovici, B. (2020). Renegotiation-proof contracts with persistent private information. Working paper.

Strulovici, B. and M. Szydlowski (2015). On the smoothness of value functions and the existence of optimal strategies in diffusion models. *Journal of Economic Theory 159*, 1016–1055.

Takeuchi, K. (2011). Non-parametric test of time consistency: Present bias and future bias. *Games and Economic Behavior 71*(2), 456–478.

Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters 8*(3), 201–207.

Thaler, R. H. and S. Benartzi (2004). Save more tomorrow™: Using behavioral economics to increase employee saving. *Journal of Political Economy 112*(S1), S164–S187.

Wang, T. and J. Yong (2019). Backward stochastic volterra integral equations—representation of adapted solutions. *Stochastic Processes and their Applications*.

Watson, J., D. Miller, and T. E. Olsen (2019). Relational contracting, negotiation, and external enforcement.

Wei, Q., J. Yong, and Z. Yu (2017). Time-inconsistent recursive stochastic optimal control problems. *SIAM Journal on Control and Optimization 55*(6), 4156–4201.

Williams, N. (2015). A solvable continuous time dynamic principal–agent model. *Journal of Economic Theory 159*, 989–1015.

Yan, W. and J. Yong (2019). *Time-Inconsistent Optimal Control Problems and Related Issues*, pp. 533–569. Cham: Springer International Publishing.

Yong, J. (2012). Time-inconsistent optimal control problems and the equilibrium HJB equation. *arXiv preprint arXiv:1204.0568*.

Yong, J. and X. Y. Zhou (1999). *Stochastic controls: Hamiltonian systems and HJB equations*, Volume 43. Springer Science & Business Media.