

Sticky Incentives and Dynamic Agency

Optimal Contracting with Perks and Shirking

John Y. Zhu¹
January 10, 2011

Abstract

I explicitly derive the optimal dynamic incentive contract in a general continuous-time agency problem where inducing static first-best action is not always optimal. My framework generates two dynamic contracts new to the literature: (1) a “quiet-life” arrangement and (2) a suspension-based endogenously renegotiating contract. Both contractual forms induce a mixture of first-best and non-first-best action. These contracts capture common features in many real life arrangements such as “up-or-out”, partnership, tenure, hidden compensation and suspension clauses. In applications, I explore the effects of taxes, bargaining and renegotiation on optimal contracting. My technical work produces a new type of incentive scheme I call *sticky incentives* which underlies the optimal, infrequent-monitoring approach to inducing a mixture of first-best and non-first-best action. Furthermore, I show how differences in patience between the principal and agent factor into optimal contracting.

Keywords: Principal-agent problem, optimal contracts, hidden actions, agency, hidden compensation, perks, shirking, continuous-time, quiet life, suspension, renegotiation, infrequent monitoring

JEL Numbers: C61, C63, D82, D86, M51, M52, M55

¹U.C. Berkeley. The author would like to express his deepest gratitude to his advisor Professor Robert Anderson, for his guidance and the many hours he spent assisting in the revision of this paper. The author is indebted to Professors Adam Szeidl, Chris Shannon and Benjamin Hermalin for their many helpful suggestions and criticisms. The author would also like to thank Professor Bruno Biais, Professor Roberto Raimondo, Derek Horstmeier, Professor Kyna Fong and a number of anonymous referees for reading through and commenting on previous drafts of this paper and also all the participants in the Berkeley Theory Seminar. Lastly, the author appreciates the discussions he has had with Professor Sourav Chatterjee on some of the technical aspects of this work.

1 Introduction

The paper explicitly solves for the optimal contracts in a broad Brownian framework where inducing static first-best action is not always optimal.² This continuous-time setting models a dynamic principal-agent relationship where some asset (a firm, a project etc.) owned by a risk-neutral principal is contracted out to a risk-neutral agent to manage. The asset's variable cash flow is governed by a continuous stochastic process; the framework allows for this process to be either Brownian Motion or geometric Brownian Motion.³ The profitability of the asset is influenced by the hidden actions of the agent. At any moment in time, the agent can either choose the *first-best action* or a non-first-best action I call the *agency action*. Agency action provides a private benefit to the agent and can be naturally interpreted as perks consumption or shirking depending on the specific realizations of the fundamentals. The principal owns the cash flow. To properly motivate the agent, the principal writes a contract which stipulates a compensation plan contingent on observables along with a termination clause.

This framework and its related discrete-time counterparts have been considered by a number of previous papers including Biais et al (2005), DeMarzo and Sannikov (2006), and He (2009). All of these papers have two things in common: 1) they all solve for the best contract that is restricted to always inducing first-best action, providing insights into some high-powered incentive contracts used in practice, 2) they all find that this static first-best action contract is, in general, not always optimal. The second fact is not necessarily surprising. Empirical findings show perquisites play a large role in CEO contracts (Schwab and Thomas, 2006). Moreover, many contractual arrangements such as tenure, partnership, involuntary separation and suspension entail variable effort levels and do not fit neatly with the first-best action or even the stationary action viewpoint. In this paper, I demonstrate that the optimal contract of the Brownian framework reflects many of the aforementioned contractual features frequently observed in practice.

Theorem 3.1 states that the optimal contract takes on one of four forms depending on the fundamentals: 1) baseline form, 2) static form, 3) Quiet-Life form, 4) Renegotiating Baseline form. Baseline contracts only induce first-best action. In baseline contracts, the agent is rewarded with cash compensation, and his performance is constantly monitored by the principal. In static contracts the agent applies agency action forever with salary, and there is no monitoring. The main contributions of this paper involve the other two forms, both of which mix in periods of agency action with first-best action.

Quiet-Life contracts provide hidden compensation to the agent by inducing agency action. After a sustained period of good performance, a "Quiet-Life phase" is triggered. During this phase the principal *infrequently monitors* the agent, and the agent frequently consumes perks. The incentive scheme of the contract becomes somewhat unresponsive to asset perfor-

²Just to emphasize, in this setting an incentive-compatible contract that induces static first-best action is not the same as the first-best contract, which is usually not incentive-compatible.

³The paper and its results are presented in the Brownian Motion setting. In Section 7.E, I explain how the results translate into the geometric Brownian setting.

mance. After the Quiet-Life phase concludes, the principal returns to constantly monitoring the agent, and the agent applies first-best action. Then either sustained good performance triggers another round of the Quiet-Life phase or sustained poor performance brings about termination. Comparing Quiet-Life contracts with baseline contracts, I find that the hidden compensation packages of Quiet-Life contracts tend to be less lucrative than the cash compensation packages of baseline contracts. Termination also tends to be delayed in Quiet-Life contracts (Corollary 4.1).

The last optimal form is the Renegotiating Baseline form. Renegotiating Baseline contracts are baseline contracts where termination is replaced with suspension phases. There is a state variable that both keeps track of the agent's continuation payoff and serves as a dynamic rating of the agent's managerial performance. When poor performance pushes the agent's rating down to an endogenously determined low threshold, the contract triggers a suspension phase during which cash compensation is postponed, and the agent frequently exerts low effort. While the agent serves his suspension, the performance rating is "pegged" around the low threshold. Afterwards, the agent is forgiven for some of his prior poor performance, and the agent's rating is pushed upwards as the principal renegotiates some more slack for the agent. The underlying baseline contract of a Renegotiating Baseline contract is usually not renegotiation proof. In a baseline contract, it is important that the principal commits not to renegotiate in order to preserve the first-best action incentive scheme of the contract. A Renegotiating Baseline contract endogenously embeds some of the "renegotiable-ness" of the underlying baseline contract allowing both the principal and agent to internalize some of the value of renegotiation in an incentive-compatible way (Remark 5.1).

In this paper, I show the reason why the Quiet-Life and Renegotiating Baseline contracts are sometimes optimal has much to do with the relative patience of the principal. There is a contractual technique where the principal delays the enactment of a lucrative contract by first enacting a less lucrative period of agency action either to reward or punish the agent. This "*saving-the-best-for-last*" technique, which plays an important role in the construction of the Quiet-Life and Renegotiating Baseline contracts, is only useful if the principal is relatively patient (Remark 6.1). In fact, I show how in an equal patience setting, the inability of the principal to profitably utilize this technique depresses the value of the optimal contract. When the principal is more patient, saving-the-best-for-last implies a role for agency action in optimal contracting even when agency action is inefficient - just not too inefficient. Previous papers have already made note of this fact that the optimal contract may not always induce first-best action when agency action is not too inefficient. However, my paper is both the first to formalize the reason for this fact by analyzing the value of saving-the-best-for-last and the first to derive the optimal contract when agency action is not too inefficient.

With the optimal contract taking one of a number of forms, many natural comparative statics can be performed on the form of the optimal contract by shifting the fundamentals. I show how the optimal contract takes on the form that best highlights the contractual advantages of a particular realization of the agency action's value to the agent and cost to the principal (Section 7.A. The Domains of Optimality Theorem). For example, if agency action is very valuable to the agent and not too bad for the principal, the optimal contract is the

Quiet-Life contract which uses hidden compensation. I also show how the optimal contract changes depending on the relative bargaining power of the agent. In particular, I look at optimal contracting when 1) the agent can bargain for higher outside options and 2) when the agent can bargain for higher contract payoffs.⁴

In addition, I analyze how optimal contracting is affected by the presence of taxes. In this setting hidden compensation becomes more attractive due to the inefficiency of taxed cash compensation. Proposition 7.3 states how a tax hike induces a simple shift of emphasis from the two contract types that employ cash compensation (baseline and Renegotiating Baseline) to the two that don't (Quiet-Life and static). Specifically, taxes turn some situations where the baseline contract is optimal into situations where the Quiet-Life contract is optimal and also turn some situations where the Renegotiating Baseline contract is optimal to situations where the static contract is optimal.

Furthermore, I find that a subset of the Renegotiating Baseline contracts are renegotiation-proof. Contracts in this model, optimal or otherwise, are usually not renegotiation-proof. To design a renegotiation-proof contract inducing only first-best action, the principal would have to employ randomization and bound the value of the contract to be equal to the value of his outside option. However, renegotiation-proof Renegotiating Baseline contracts never employ randomization, nor must they be worth less than the principal's outside option. Furthermore, costly termination is never exercised in any Renegotiating Baseline contract, which is, again, in direct contrast to the renegotiation-proof static first-best action contracts (Remark 7.4).

While the four forms of the optimal contract are all structurally simple, a proper understanding of their value requires a number of technical advancements. The incentive scheme of the Quiet-Life and Renegotiating Baseline contracts follow what is known as Sticky Brownian Motion. This motion has not previously appeared in the contracting literature. And yet an understanding of Sticky Brownian Motion is a prerequisite to rigorously interpreting the results of this paper (Proposition 4.1, Lemma 4.1). For example, take a look back at the paragraph introducing the Quiet-Life contract. Much of the paragraph would not make formal sense without knowledge of the dynamics of Sticky Brownian Motion. The concepts of "infrequent" and "frequent", the term "somewhat unresponsive", and the comparative results between Quiet-Life and baseline contracts all have their formal roots in Sticky Brownian Motion.

The proof that the optimal contract only takes on the four listed forms is also technically nontrivial. Typically, the strategy in continuous-time agency problems is to deduce the optimal value function and extract from it the optimal contract. In this setting, the optimal value function is tied to two ODEs: the first-best action ODE deduced by DeMarzo and Sannikov (2006) and the agency action ODE which I explicitly solve in Lemma 6.2. The optimal value function pieced together from solutions to these two ODEs is, in general, not smooth.

⁴The emphasis is on agent bargaining power because the principal holds all the bargaining power in the base model.

So even applications of Ito's Lemma and optional sampling, which are straightforward in the previous literature, become complicated and require care. Moreover, the principal has great flexibility in switching between first-best and agency action. Virtually any action sequence can be made to be incentive-compatible. Whittling down this set can be challenging. For example, a major step in the optimality theorem is to simply show that there is at most one phase change point in the optimal value function - i.e. only one agent continuation payoff value in the optimal contract at which the principal will induce a switch of actions.

The techniques developed in this paper pave the way for more work to be done in dynamic agency models where agency cost is not prohibitive. The explicit optimal contracts derived in this paper represent a first step in the continuous-time literature to understanding the specific roles of perks and shirking in business. Overall, the theory of this paper provides some formal dynamic foundations for the agency action research that has already begun in earnest on the empirical side.

Related Literature

Grossman and Hart (1983) demonstrated that simultaneously determining the optimal action levels and the associated optimal incentive scheme is complex even in simple principal-agent models. One strain of the literature gets around this difficulty by deriving the optimal static first-best action contract or optimal stationary action contract, and providing some conditions under which such contracts are optimal. Such papers include Holmstrom and Milgrom (1987), Biais et al (2005), DeMarzo and Sannikov (2006), He (2009) and Edmans, Gabaix and Landier (2009). Another solution is to formulate some general principle about how to optimally induce arbitrary action sequences, but avoid the issue of finding the optimal action sequence (e.g. Edmans and Gabaix, 2009). Yet another way is to use the powerful machinery of stochastic calculus to produce theorems about the existence of optimal contracts in general settings at the expense of concreteness. The notable example of this strain of the literature is Sannikov (2008), which produces an existence theorem that implies a role for shirking without specifying the exact nature of that role. Lastly, one can restrict attention to static models. For example, two recent papers in the perks literature - Bennardo, Chiappori and Song (2010) and Kuhnen and Zwiebel (2009) - produce explicit contracts that involve perks consumption.

Thus there is a gap in the literature. On the one hand, papers like Sannikov (2008) (not to mention the myriad real-life contracts) tell us that agency action ought to play an important role in optimal contracting in general. But most of the explicit optimal dynamic contracts produced either induce only first-best action or some fixed stationary action. The need to close this gap is one of the primary motivations of my paper.

More broadly, this paper serves as a theoretical counterpart to a number of recent empirical papers investigating the role agency action in business. The debate over agency action started with the seminal papers Jensen and Meckling (1976) and Fama (1980). The agency cost versus ex-post settling up perspectives espoused by these two papers serve as the backdrop to the recent empirical work. Yermack (2006) shows that "the disclosed personal use of company aircraft by CEOs is associated with severe and significant underperformance

of their employers’ stocks.” However, the results do not unambiguously vindicate the agency cost perspective. For example, there does not seem to be a significant correlation between managers’ fractional stock ownership and personal aircraft use. It may be that the compound problem of “managerial shirking in the presence of lavish perks” is symptomatic of other human capital specific problems, and is not necessarily a blanket indictment on agency action. Bertrand and Mullainathan (2003) suggests that observable CEO preferences may indicate a desire for the “quiet life,” contrary to the active empire building theory that casts a pall on agency action. Rajan and Wulf (2006) finds that “the evidence for agency [cost] as an explanation of perks is, at best, mixed.” The paper argues that perks can be productivity improving and can also serve as a form of tax-free hidden compensation. Thus, perks can be justified as rational expenditures, and their findings point to more responsible practices of using agency action. My paper posits some of the ways a responsible principal can justify inducing agency action.

2 Model and Preliminaries

The paper is presented from the Brownian perspective. For a discussion about the geometric Brownian model, see section 7.E.

A.1 Setting

There is an asset belonging to a principal, for which he contracts an agent to manage. The asset produces a stochastic revenue stream. Over time, we assume that the cumulative revenue stream behaves as Brownian Motion with a drift process influenced by the hidden action applied by the contracted agent.

Formally, there is a stochastic process $Z = \{Z_t\}_{t \geq 0}$ defined on a probability space Ω with probability law P^μ . Under P^μ , Z is Brownian motion with drift μdt . At time t , Z_t is the cumulative revenue stream of the asset up to time t . The μdt drift corresponds to the default expected returns and can be interpreted as the intrinsic or maximum expected profitability of the asset.

A.2 Actions

The agent affects asset performance by selecting an action at each moment in time. Over time the agent’s action process $a = \{a_t\}_{t \geq 0}$ is a stochastic process taking values in a set $\{0, A\}$ with $A > 0$. $\{0\}$ is first-best action and $\{A\}$ is **agency action**. The action process a affects the underlying probability law: the default law P^μ changes to $P^{\mu-a}$, which is defined to be the law under which Z is Brownian motion with drift $(\mu - a_t)dt$.

The principal can choose a compensation scheme for the agent. Compensation is represented by a random nondecreasing process $I = \{I_t\}_{t \geq 0}$ started at zero that keeps track of the cumulative cash payments made to the agent up to time t . Termination is a stopping time τ .

A.3 Preferences

The principal is risk neutral, discounts at rate r , retains the cash flow of the asset, compensates the agent, and can exercise an outside option worth $L < \frac{\mu}{r}$ after the termination of the contract. His utility is

$$\mathbf{E}_{P^{\mu-a}} \left[\int_0^\tau e^{-rs} (dZ_s - dI_s) + e^{-r\tau} L \right]$$

The agent is risk neutral, discounts at rate γ , receives compensation from the principal, and can exercise an outside option worth $K \geq 0$ after the termination of the contract. The agent also receives an instantaneous utility flow $\phi a_t dt$ by applying action $a_t \in \{0, A\}$ at time t , where $\phi > 0$. His utility is

$$\mathbf{E}_{P^{\mu-a}} \left[\int_0^\tau e^{-\gamma s} (dI_s + \phi a_s ds) + e^{-\gamma\tau} K \right]$$

We assume for now that the principal is more patient: $r < \gamma$. Later on we will consider the implications of having an equally patient agent: $r = \gamma$. The assumption that the principal is at least as patient as the agent is an important one. Typically, one thinks of the agent as an actual individual like a CEO who is separated from a principal representing ownership. In certain cases ownership may also consist of a single individual. However, when there is separation between ownership and control, the more typical case is where ownership representation is in the form of an institution such as a board, shareholders, institutional investor etc. It is then reasonable to assume that such a permanent or semi-permanent entity would be relatively patient.

B. Incentive-Compatibility

Definition 2.1. A contract is a tuple (a, I, τ) consisting of an action process a , a compensation scheme I , and a termination clause τ .

Fix a contract (a, I, τ) . The **agent's continuation payoff** U_t is defined to be the agent's expected future utility given the history \mathcal{F}_t up to time t :

$$U_t = \mathbf{E}_{P^{\mu-a}} \left[\int_t^\tau e^{-\gamma(s-t)} (dI_s + \phi a_s ds) + e^{-\gamma(\tau-t)} K \middle| \mathcal{F}_t \right] \quad t \leq \tau$$

The evolution of U_t is the contract's incentives. The motion of U_t is characterized by the following stochastic differential equation:

Lemma 2.1. *There exists a stochastic process $\beta = \{\beta_t\}_{t \geq 0}$ such that*

$$dU_t = \gamma U_t dt - dI_t - \phi a_t dt + \beta_t (dZ_t - (\mu - a_t) dt)$$

Proof. Standard. □

The process β_t drops out of the martingale representation theorem, and represents how sensitive the contract's incentive scheme is to asset performance. High sensitivity will induce the agent to apply first-best action, and low sensitivity will mean the agent will apply agency action. To determine whether a contract is incentive-compatible requires comparing the contract's action process to the sensitivity process.

Lemma 2.2. (*Incentive-Compatibility Criterion*) A contract (a, I, τ) with sensitivity process β is incentive-compatible if and only if for all t , $U_t \geq K$ and

$$i) a_t = 0 \Rightarrow \beta_t \geq \phi \quad ii) a_t = A \Rightarrow \beta_t \leq \phi$$

Proof. Standard. □

The criterion tells us that in order to induce first-best action, the sensitivity factor needs to be at least ϕ . However, the greater the sensitivity, the more volatile the incentives, which entails a cost. Thus in optimality, whenever the principal wants to induce first-best action, he will always choose the lowest possible sensitivity: ϕ . Similarly, the best way to induce agency action is to select sensitivity 0. We can now pin down the two laws that will govern the incentives of the optimal contract:

Definition 2.2. When the optimal contract stipulates first-best action, the continuation payoff of the agent follows the **first-best action law**:

$$dU_t = \gamma U_t dt - dI_t + \phi(dZ_t - \mu dt)$$

which says to induce first-best action the continuation payoff of the agent needs to be sensitive to asset performance, and in expectation, compounds at the agent's discount rate less the cash dI_t delivered to the agent right now.

Similarly, when the optimal contract stipulates agency action, the continuation payoff of the agent follows the **agency action law**:

$$dU_t = \gamma U_t dt - dI_t - \phi A dt$$

which says the agent's continuation payoff is not sensitive to asset performance, and compounds at the agent's discount rate less the cash dI_t delivered to the agent right now and less the utility $\phi A dt$ the agent obtains from applying agency action A .

3 The Four Forms of the Optimal Contract

It turns out the optimal contract always follows a particular format which I describe in 3.A - The General Form of the Optimal Contract. This general form is a set of rules governing the motion of the optimal contract's agent continuation payoff U_t . In 3.B I state the optimality theorem which specifies the four realizations of the general form taken by the optimal contract. Two of the realizations have appeared in the previous literature - *baseline form* and *static form*, and I list a few relevant facts about them in 3.C.

A. The General Form of the Optimal Contract

The principal selects a good performance threshold U^{good} and a poor performance threshold U^{poor} subject to the condition $K \leq U^{poor} \leq U^{good}$. These thresholds will be the upper and lower bounds on the agent's continuation payoff U_t . Next, a value $U^{contract} \in [U^{poor}, U^{good}]$




Figure 1: Schematic diagram of the general form of the optimal contract.

is selected which is the total payoff of the contract to the agent. The continuation payoff of the agent is started at this value:

$$U_0 = U^{contract}$$

While the agent's continuation payoff U_t is in between U^{poor} and U^{good} it follows the first-best action law:

$$dU_t = \gamma U_t dt + \phi(dZ_t - \mu dt) \quad U_t \in (U^{poor}, U^{good})$$

When U_t reaches one of the thresholds we have the following possibilities:

At the good performance threshold the principal selects one of the following two options to reward the agent:

- 1) Provide the agent with cash compensation dI_t :
 - Cash compensation is chosen in such a way so that the Brownian U_t reflects downwards at U^{good} .
- 2) Induce agency action as a form of hidden compensation:
 - The law of U_t at the good performance threshold U^{good} switches to the agency action law $dU_t|_{U^{good}} = (\gamma U^{good} - \phi A)dt$.
 - This choice is available only if U^{good} is chosen to be less than or equal to $\frac{\phi}{\gamma}A$.⁵

At the poor performance threshold the principal selects one of the following two options to punish the agent:

- 1) Terminate the contract:
 - This choice is incentive-compatible only if $U^{poor} = K$.
- 2) Induce agency action as a form of suspension:
 - The law of U_t at the poor performance threshold U^{poor} switches to the agency action law $dU_t|_{U^{poor}} = (\gamma U^{poor} - \phi A)dt$.
 - This choice is available only if U^{poor} is chosen to be greater than or equal to $\frac{\phi}{\gamma}A$.

⁵If U^{good} was instead chosen to be greater $\frac{\phi}{\gamma}A$, then U_t would continue to go upwards at U^{good} , contradicting the upper bound assumption of U^{good} .

B. The Optimal Contract

Theorem 3.1. *In the optimal contract the principal selects two thresholds: a poor performance threshold U^{poor} and a good performance threshold U^{good} . The agent's continuation value is started between these two values and follows the first-best action law*

$$dU_t = \gamma U_t dt + \phi(dZ_t - \mu dt)$$

While following this law, U_t is sensitive to asset performance and serves as a **dynamic rating** of the agent's managerial performance. When good performance pushes the rating up to U^{good} the principal rewards the agent either through cash compensation or through hidden compensation by inducing agency action. When poor performance pushes the rating down to U^{poor} the principal either terminates the contract or punishes the agent by inducing agency action. The four different ways to choose between these options produce the four forms of the optimal contract which are summarized in the table below:

	U^{good}		
U^{poor}		cash compensation	agency action
termination		Baseline	Quiet-Life
agency action		Renegotiating Baseline	Static

Baseline contracts always induce first-best action. Static contracts always induce agency action. The Quiet-Life and Renegotiating Baseline contracts both induce agency action non-permanently in between periods of first-best action.

C. Some Remarks on the Baseline and Static Contracts

The baseline form and the conditions under which it is optimal have already been derived in the previous literature. It always induces first-best action and is identical to the credit limit contract of DeMarzo and Sannikov (2006). It is also the additive version of the no-shirking contract of He (2009) and the continuous-time version of the optimal contract of Biais et al (2007).

The following are a few relevant facts about the static contracts:

- 1) Static contracts induce agency action forever and may supplement the agent with a fixed salary sdt . Consequently, the good and poor performance thresholds coincide, and the agent's continuation payoff is permanently fixed at this value: $U^{poor,S} = U_t = U^{good,S} = \frac{\phi A + s}{\gamma}$.
- 2) The optimal static contract supplements the agent with a salary sdt just enough to prevent him from quitting: $s = \max\{0, \gamma K - \phi A\}$. The payoff to the agent is $\max\{\frac{\phi A}{\gamma}, K\}$ and the payoff to the principal is $\min\{\frac{\mu - A}{r}, \frac{\mu - A - (\gamma K - \phi A)}{r}\}$.

The rest of the paper is primarily concerned with the Quiet-Life form and the Renegotiating Baseline form. Both forms induce agency action *non-permanently*.

4 The Quiet-Life Contracts

This section analyzes the incentive scheme of Quiet-Life contracts, culminating in a rigorous characterization of Quiet-Life contracts in 4.B. Dynamics of the Quiet-Life Contracts. The concepts of *infrequent monitoring* (Proposition 4.1 and Definition 4.1) and *sticky incentives* (Lemma 4.1 and Proposition 4.2) are introduced. Infrequent monitoring and sticky incentives are the key properties of the Quiet-Life and Renegotiating Baseline contracts that differentiate these contracts from baseline contracts. These concepts imply certain appealing contractual characteristics (Corollary 4.1) which, combined with the discussion of patience in Section 6, help explain why Quiet-Life and Renegotiating Baseline contracts are sometimes better than baseline contracts.

A.1 Quiet-Life: Agency Action as Reward

Recall, a Quiet-Life contract induces agency action non-permanently when the agent's continuation payoff and performance rating U_t reaches the contract's good performance threshold. Termination is triggered when U_t drops down to the poor performance threshold. All Quiet-Life contracts satisfy $K = U^{poor} < U^{good} < \frac{\phi A}{\gamma}$. The first equality is due to the fact that termination means the agent exercises his outside option. The middle strict inequality is there because if the thresholds were equal the instructions of the contract would contradict. The last strict inequality comes from two observations: U^{good} cannot be greater than $\frac{\phi A}{\gamma}$ because such a threshold would constitute a promised continuation payoff greater than what agency action utility alone can deliver, nor can it be equal to $\frac{\phi A}{\gamma}$ since that would imply permanent agency action when U_t reaches U^{good} , contradicting the assumption that Quiet-Life contracts induce agency action non-permanently.

What $U_t \leq U^{good} < \frac{\phi A}{\gamma}$ implies is that when the agent finally reaches the good performance threshold, he receives an agency action flow, which if extended indefinitely, would represent a value greater than anything the contract actually promises. Thus agency action in Quiet-Life contracts serves to reward the agent, as a form of hidden compensation.

A.2 Quiet-Life: Infrequent Monitoring

What does a typical hidden compensation package look like? Let H denote the nondecreasing hidden compensation process (similar to the cash compensation process I), where H_t is the amount of agency action utility received by the agent up to time t . We already know for every moment t when $U_t = U^{good}$ the agent receives a fixed agency action utility flow $\phi A dt$. This implies:

$$dH_t = \begin{cases} 0dt & U_t < U^{good} \\ \phi A dt & U_t = U^{good} \end{cases}$$

Thus to characterize H it suffices to characterize the random set of hidden compensation times $\mathcal{T}^U(U^{good}) = \{t | U_t = U^{good}\}$.

Near the good performance threshold U^{good} , the continuation payoff and performance rating



Figure 2: A sample sequence of hidden compensation times.

U_t of the agent follows the first-best action law:

$$dU_t|_{U_t < U^{good}} = \gamma U_t + \phi(dZ_t - \mu dt)$$

which is sensitive to asset performance. This requires the principal to constantly monitor asset performance to properly adjust U_t . At the good performance threshold U^{good} , U_t follows the agency action law:

$$dU_t|_{U_t = U^{good}} = \gamma U^{good} - \phi A dt$$

and is no longer sensitive to asset performance. Consequently, the principal shuts down monitoring. Thus around the U^{good} threshold the principal mixes constant monitoring and no monitoring, producing a set of hidden compensation times with the following properties:

Proposition 4.1. *Any neighborhood of a hidden compensation time contains infinitely many other times of hidden compensation. Formally, the random set of hidden compensation times $\mathcal{T}^U(U^{good})$ almost surely satisfies the following three properties: 1) positive measure, 2) nowhere dense, 3) perfect.⁶ (See Figure 2)*

Proof. This is a simple consequence of Lemma 4.1 of the next subsection. □

The positive measure property implies that, in particular, $\mathcal{T}^U(U^{good})$ does not look like this:



Hidden compensation takes time. This is because a fixed utility flow over a set of times of measure 0 amounts to no utility at all. If a Quiet-Life contract's hidden compensation times were actually trivial then it would not be incentive-compatible.

The nowhere dense property implies that, in particular, $\mathcal{T}^U(U^{good})$ does not look like this:



By dispersing an interval of hidden compensation, the principal can in expectation fit more first-best action times before the end of a period of hidden compensation. The downside is that by mixing first-best action with agency action there is an added risk that termination may occur in the “gaps” between hidden compensation times. However, the risk is slight provided the dispersion is not too great. This gives us the third property, $\mathcal{T}^U(U^{good})$ is a perfect set.

⁶Compare with the random set of cash compensation times of a baseline contract which is also nowhere dense and perfect, but has zero measure. A set is *perfect* if it contains all of its limit points, and has no isolated points.

Definition 4.1. *The characterization of the hidden compensation times $\mathcal{T}^U(U^{good})$ implies a local (around U^{good}) monitoring structure where no monitoring times are inserted in between the monitoring times in a temporally nontrivial, nowhere dense and perfect way. Call this mixture **infrequent monitoring**. Call the periods of time in the Quiet-Life contract when the principal is infrequently monitoring, the **Quiet-Life phases**.*

A.3 Quiet-Life: Sticky Incentives

Given a Quiet-Life contract \mathcal{Q} with some good performance threshold U^{good} , we can design the companion baseline contract \mathcal{B} with the same threshold U^{good} . Let $U_t^{\mathcal{Q}}$ denote the Quiet-Life contract's agent continuation payoff process and define $U_t^{\mathcal{B}}$ similarly. Obviously these two contracts exhibit a large amount of structural similarity: they have the same performance thresholds,⁷ the agent continuation payoffs of the two contracts follow the same first-best action law on the open interval $(U^{poor} = K, U^{good})$, and both contracts terminate at $U^{poor} = K$. The only structural difference is at U^{good} where \mathcal{Q} induces agency action as a form of hidden compensation and \mathcal{B} delivers cash compensation.

Remark 4.1. *Comparing \mathcal{Q} and \mathcal{B} allows us to isolate and study the comparative advantages of using hidden-compensation-based versus cash-compensation-based incentives.*

We know from the previous literature that $U_t^{\mathcal{B}}$ is reflected Brownian motion.⁸ The technical term for $U_t^{\mathcal{Q}}$ is **Sticky Brownian motion**. The following is the crucial technical result on Sticky Brownian motion:

Lemma 4.1. *Sticky Brownian motion is reflected Brownian motion under a decelerated time change.*

Proof. See Harrison and Lemoine (1981). □

This immediately implies:

Proposition 4.2. *Hidden-compensation-based incentives are slower than those of cash compensation. We call this slower incentive scheme of the Quiet-Life contract **sticky incentives**. Formally, there is a decelerated time change which is a random nondecreasing process $S(t) \geq t$ such that*

$$U_{S(t)}^{\mathcal{Q}} =_d U_t^{\mathcal{B}}$$

Let $H_t^{\mathcal{Q}}$ denote the hidden compensation process of \mathcal{Q} , and $\tau^{\mathcal{Q}}$ denote the termination time of \mathcal{Q} . Similarly, let $I_t^{\mathcal{B}}$ denote the cash compensation process of \mathcal{B} , and $\tau^{\mathcal{B}}$ denote the termination time of \mathcal{B} . The following formalizes the value of slowing down incentives:

Corollary 4.1. *A sticky incentive scheme implies that hidden compensation is more modest than cash compensation:*

$$H_{S(t)}^{\mathcal{Q}} =_d I_t^{\mathcal{B}} \Rightarrow \mathbf{E}[H_t^{\mathcal{Q}}] < \mathbf{E}[I_t^{\mathcal{B}}] \quad \text{for all } t > 0$$

and delays termination:

$$\tau^{\mathcal{Q}} =_d S(\tau^{\mathcal{B}}) \geq \tau^{\mathcal{B}}$$

⁷Their poor performance thresholds are both K by assumption.

⁸For example, DeMarzo and Sannikov (2006).

B. Dynamics of the Quiet-Life Contracts

We can now give a precise description of the dynamics of a Quiet-Life contract. In a Quiet-Life contract the agent initially applies first-best action and his continuation payoff and performance rating U_t follows the first-best action law. Sustained good performance brings U_t up to the good performance threshold U^{good} . At this point the contract enters the *Quiet-Life phase* (Definition 4.1) where agency action is triggered as a form of *hidden compensation* (see subsection A.1), and the agent *frequently* (Proposition 4.1) consumes perks. During the Quiet-Life phase the principal *infrequently monitors* (Definition 4.1) the agent and as a result U_t *sticks around U^{good} for a while* (Proposition 4.2) following Sticky Brownian motion.

Eventually, poor performance brings U_t back down and the contract exits the Quiet-Life phase. The principal resumes constant monitoring of the agent and this dynamic remains until sustained good performance triggers another round of the Quiet-Life phase or sustained poor performance finally triggers termination.

The Quiet-Life arrangement has some comparative advantages over the cash-compensation arrangement of baseline contracts. In particular, termination tends to be delayed, and the agent receives less expected compensation (Corollary 4.1). When perks consumption is not too harmful to the principal, these advantages will imply that the optimally designed Quiet-Life contract outperforms the optimal baseline contract.

5 Renegotiating Baseline Contracts

Definition 5.1. *The **underlying baseline contract** of a Renegotiating Baseline contract \mathcal{R} is the baseline contract with the same good performance threshold as \mathcal{R} .*

This section shows how we can view a Renegotiating Baseline contract as the underlying baseline contract under repeated renegotiation. I explain how agency action can be induced as punishment in the form of suspension phases. The suspension phases allow the principal to credibly renegotiate without compromising the first-best incentive structure of the underlying baseline contract (Remark 5.1), and represent a potential advantage over baseline contracts which require commitment to not renegotiate. The concepts of sticky incentives and infrequent monitoring which were introduced in the previous section reappear. A characterization of the dynamics of Renegotiating Baseline contracts is found in 5.B.

A.1 Renegotiating Baseline: Agency Action as Punishment

In Renegotiating Baseline contracts, agency action is induced at the poor performance threshold U^{poor} which need not be equal to K . Mathematically, the dynamics of a Renegotiating Baseline contract's agent continuation payoff at U^{poor} is the mirror image of the dynamics of a Quiet-Life contract's agent continuation payoff at U^{good} . Therefore, the concepts of sticky incentives, infrequent monitoring, and the properties of the agency action times (e.g.

positive measure, nowhere dense, perfect) all translate over.

However, the *role* of agency action and the *value* of the associated infrequent monitoring and sticky incentives are different.

Unlike in Quiet-Life contracts, in Renegotiating Baseline contracts $\frac{\phi A}{\gamma} < U^{poor} \leq U_t$. This means that when the agent's continuation payoff and performance rating U_t drops down to the poor performance threshold U^{poor} , he receives an agency action flow, which if extended indefinitely, would represent a value strictly less than anything the contract actually promises. Thus agency action in Renegotiating Baseline contracts serves to punish the agent.

This is not to say the agent dislikes agency action. On the contrary, the application of agency action is simply the best the agent can do for himself in this arrested phase of the contract. The canonical example of this phenomenon is *suspension*. During a suspension, the agent's compensation is frozen, so he exerts low effort. Despite the agent's fondness for low effort, he would rather be working hard and receiving compensation than be stuck in this low state.

A.2 Renegotiating Baseline: Suspension and Renegotiation

The idea of contractual punishment is not new. A termination clause serves the same purpose. So why not just terminate like in a baseline contract?

In many baseline contracts (including the optimal one) when the agent's continuation payoff and performance rating U_t is near the poor performance threshold U^{poor} and termination is probabilistically imminent, the principal is better off giving the agent some more slack. The principal achieves this by simply shifting the performance rating U_t upwards, removing it from the vicinity of U^{poor} . By forgiving the agent for his poor performance, the principal is effectively renegotiating the baseline contract. Each time this is done the principal increases his own continuation payoff as well as that of the agent. However, the value of this renegotiation is predicated on the agent not expecting to be forgiven and applying first-best action throughout. Unfortunately, if the agent expects that the principal will renege on termination, then the incentives to apply first-best action will be destroyed. Thus such a renegotiation is not incentive-compatible, and it is imperative that the principal commits to not renegotiate.

However, the potential losses due to a premature end to the principal-agent relationship may be great. Thus it is important to find a way to both induce first-best action most of the time but still be able to back out of termination during periods of poor performance. The Renegotiating Baseline contract achieves this by picking a poor performance threshold and inducing agency action there as a suspension phase.

From our discussion of agency action phases in the Quiet-Life contracts we know two things will happen when the principal induces agency action at U^{poor} :

- 1) U_t will eventually leave the vicinity of U^{poor} after the end of the suspension phase.

2) The contract will spend a nontrivial amount of time at U^{poor} .

Remark 5.1. *That the Renegotiating Baseline contract naturally pushes the agent’s continuation payoff and performance rating U_t upwards (observation 1 above) after poor performance means that this contract is endogenously renegotiating the underlying baseline contract. That the renegotiation happens only after the suspension phase (observation 2 above) means that the first-best action incentives of the underlying baseline contract are not compromised by the renegotiation. The agent doesn’t get the extra slack of renegotiation for free. By having to first suffer through suspension, the agent is effectively “buying” the principal’s forgiveness through the postponement of the cash compensation promised by the underlying baseline contract.*

B. Dynamics of the Renegotiating Baseline Contracts

A Renegotiating Baseline contract begins as its underlying baseline contract, inducing first-best action and providing cash compensation whenever the agent’s continuation payoff and performance rating U_t hits the good performance threshold U^{good} .

However, when poor performance pushes U_t down to the poor performance threshold U^{poor} , a suspension phase is triggered. During suspension, the principal infrequently monitors the agent and the agent, lacking proper incentives to work, frequently exerts low effort. As a result U_t sticks to or is “pegged” around U^{poor} for a period of time, following the dynamics of Sticky Brownian motion.

Eventually, suspension ends, the agent is forgiven for some of his poor performance, and U_t is allowed to float again as the principal renegotiates some more slack for the agent. The contract returns to the first-best action incentives of the underlying baseline contract where the principal constantly monitors and good performance is rewarded with cash compensation. This dynamic remains the norm until sustained poor performance triggers suspension again.

6 Implementation

Given the optimality theorem, the most pertinent question is what form does the optimal contract take? Naturally, the answer will depend on the fundamentals, specifically, on agency action’s relative value to the agent and relative cost to the principal. This question will be answered in full in Section 7.A - The Domains of Optimality Theorem. In this section let us approach the general implementation problem by first exploring the following special case: what form does the optimal contract take when agency action is inefficient (see Definition 6.3)? The answer is Lemma 6.5. The results developed for this particular case provide the basic language and intuition used to tackle the general implementation problem in Section 7.A.

Specifically, this section shows that what role, if any, agency action plays in optimal contracting has to do with the usefulness of the contractual technique **saving-the-best-for-last** (Section 6.B). This technique arises from the following situation, in which the principal will often find himself: at some point in time the principal will have access to a lucrative contract (the “best” in saving-the-best-for-last) which he would like to exercise. However, the payoff of this contract to the agent will not match the agent’s promised continuation payoff at this particular moment in time. To achieve his goal of exercising the lucrative contract while still maintaining incentive-compatibility, the principal can *postpone* implementing that lucrative contract, and first induce agency action for a little while to get the agent continuation payoff right. This technique is what I call saving-the-best-for-last.

We now begin our approach to the implementation problem with a discussion, in 6.A, of some relevant value functions. These value functions, which will appear throughout the rest of the paper, are useful because they contain information about an important contracting process called the continuation payoff point process:

Definition 6.1. *Fix a contract with agent continuation payoff process U_t . We can define the corresponding principal continuation payoff process V_t . Together, (U_t, V_t) is the **continuation payoff point process** of the contract and (U_0, V_0) is the **payoff point** of the contract.*

A. Value Functions

Typically, the way optimal contracts are derived is by first deriving the *optimal value function* and then extracting from it the optimal contract. This process requires that we have two pieces of information:

1. The relevant differential conditions for the optimal value function.
2. The contractual interpretations of these relevant differential conditions.

There are two relevant differential conditions corresponding to the two types of actions:

- The **first-best action ODE**

$$ry = \mu + \gamma xy' + \frac{\phi^2}{2} y'' \tag{1}$$

- The **agency action ODE**

$$ry = \mu - A + (\gamma x - \phi A)y' \tag{2}$$

Lemma 6.1 then tells us how to interpret solutions to the first-best action ODE and Lemma 6.3 tells us how to interpret solutions to the agency action ODE. In 6.A.2, I also explicitly solve the agency action ODE (see Lemma 6.2).

A.1 Value Functions: A Review of the First-Best Action ODE

This subsection is distilled from the work of DeMarzo and Sannikov (2006).




Figure 3: The optimal baseline value function F^B and the extended optimal baseline value function $F^{ext,B}$.

Lemma 6.1. *Suppose there are two contracts, one with payoff point (U^1, V^1) and the other with payoff point (U^2, V^2) with $U^1 < U^2$. There is a unique solution f to first-best action ODE that connects these two points. Fix any point $U^{contract}$ between U^1 and U^2 . Then $(U^{contract}, f(U^{contract}))$ is the payoff point of the following contract:*

- Start agent's continuation payoff U_t at $U_0 = U^{contract}$.
- U_t follows the first-best action law until
 - 1) $U_t = U^1$ at which point the contract becomes the one with payoff point (U^1, V^1) .
 - 2) $U_t = U^2$ at which point the contract becomes the one with payoff point (U^2, V^2) .

This lemma implies a number of useful facts about static first-best action contracts and their value functions (also see Figure 3):

- 1) Assume $\mu > rL + \gamma K$. The *optimal baseline value function* F^B is a concave solution to the first-best action ODE on $[K, U^{good,B}]$ where F^B and $U^{good,B}$ are uniquely determined by a smooth pasting condition.⁹
- 2) The optimal static first-best action contract delivering payoff $x \in [K, U^{good,B}]$ to the agent exists. It is the baseline contract with good performance threshold $U^{good,B}$, the agent's continuation payoff is started at $U_0 = x$, and the payoff to the principal is $F^B(x)$. Call this contract the optimal baseline contract delivering payoff $x \in [K, U^{good,B}]$ to the agent.

⁹There exists a unique $U^{good,B}$ and a unique F^B such that $F^B(K) = L$, $F^{B \prime}(U^{good,B}) = -1$, and $F^{B \prime\prime}(U^{good,B}) = 0$.

- 3) One can extend F^B to values of $x > U^{good,B}$. More generally, the *extended optimal baseline value function* $F^{ext,B}$ is F^B with a straight line of slope -1 attached to the end.
- 4) The optimal static first-best action contract delivering payoff $x > U^{good,B}$ to the agent exists. It first delivers a lump sum $x - U^{good,B}$ to the agent. Then the contract becomes the optimal baseline contract delivering payoff $U^{good,B}$ to the agent. The payoff to the principal is $F^{ext,B}(x)$. Call this contract the optimal baseline contract delivering payoff $x > U^{good,B}$ to the agent.
- 5) The **optimal baseline contract** is the optimal baseline contract delivering payoff $\arg \max F^{ext,B} = \arg \max F^B$ to the agent. The payoff to the principal is $\max F^{ext,B} = \max F^B$.
- 6) Fix an optimal baseline contract delivering some payoff to the agent. At time t if the agent's continuation payoff is U_t then the principal's continuation payoff is $F^{ext,B}(U_t)$.
- 7) Cash compensation occurs when the principal's and agent's required expected cash flows exhaust expected returns:

$$\mu = rF^B(U^{good,B}) + \gamma U^{good,B}$$

- 8) Suppose $\mu \leq rL + \gamma K$. Then F^B is just the single point (K, L) and optimal baseline contract is simply to terminate right away. Also $F^{ext,B}$ is just the straight line of slope -1 starting at (K, L) .

If we are only interested in static first-best action contracts then we would be done. But since the optimal contract may induce agency action, we also need to analyze the ODE that governs the value function of agency action periods in contracts.

A.2 Value Functions: Solving the Agency Action ODE

In this subsection, I explicitly solve the agency action ODE (Lemma 6.2) and show how to contractually interpret it (Lemma 6.3).

Definition 6.2. The **agency action point** is defined to be $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$. This is the payoff point of the static contract with no salary - the agent receives a utility flow $\phi A dt$ forever and the principal receives an expected flow $(\mu - A) dt$ forever.

Lemma 6.2. The family of solutions to the agency action ODE is characterized as follows: Fix any point $(U, V) \neq (\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$. The unique solution to equation (2) going through (U, V) is the set of points (x, y) satisfying

$$\left(\frac{y - \frac{\mu - A}{r}}{V - \frac{\mu - A}{r}} \right)^{\frac{1}{r}} = \left(\frac{x - \frac{\phi A}{\gamma}}{U - \frac{\phi A}{\gamma}} \right)^{\frac{1}{\gamma}} \quad (3)$$

Proof. Clear. □

Thus solutions to the agency action ODE are just the power functions with power $\frac{r}{\gamma}$ and base point equal to the agency action point $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$. See Figure 4.




Figure 4: $r < \gamma$: Sample solutions to the agency action ODE. The middle point is the agency action point $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$. All solutions emanate from the agency action point. The horizontal and vertical lines are also solutions because they correspond to the limiting power functions with coefficients equal to 0 and ∞ respectively.

Lemma 6.3. *Suppose there is a contract with payoff point $(\tilde{U}, \tilde{V}) \neq (\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$. Let g be the unique solution to the agency action ODE going through (\tilde{U}, \tilde{V}) . Let (U_0, V_0) be any point lying on g in between (\tilde{U}, \tilde{V}) and $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$. Then (U_0, V_0) is the payoff point of an initial agency action period of length D followed by the contract with payoff point (\tilde{U}, \tilde{V}) where*

$$D = \frac{1}{\gamma} \log \left(\frac{\tilde{U} - \frac{\phi A}{\gamma}}{U_0 - \frac{\phi A}{\gamma}} \right) = \frac{1}{r} \log \left(\frac{\tilde{V} - \frac{\mu - A}{r}}{V_0 - \frac{\mu - A}{r}} \right) \geq 0$$

During the initial agency action period, the continuation payoff point process slides along g , deterministically heading toward (\tilde{U}, \tilde{V}) , which it reaches at the end of the agency action period.

Proof. This is a simple consequence of optional sampling and a little algebra. □

B. Saving the Best for Last

Now that we have some technical results about value functions, the next step is to apply them to understand the value of a contracting technique called saving-the-best-for-last. Remark 6.1 summarizes the potential usefulness of this technique. Determining the viability of this technique helps us solve the motivating problem of this section, which is the problem of determining the optimal contractual form when agency action is inefficient.

If the static contract with payoff point $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$ is not optimal, then there will be some other contract C with payoff point (\tilde{U}, \tilde{V}) where $\tilde{V} > \frac{\mu - A}{r}$.

Over the course of designing an optimal contract, the principal at some point will be faced with the problem of delivering some continuation payoff U_t to the agent which is between \tilde{U} and $\frac{\phi A}{\gamma}$. To solve this problem, the principal can write a contract mixing C and the static contract. One way to do this is to simply *randomize* over the two options. The other option is the **saving-the-best-for-last technique**:




Figure 5: Saving-the-best-for-last and Randomization - A Comparison

- First employ the static contract for some fixed duration of time D .
- Then employ the more lucrative contract C .

Randomization can generate any payoff point lying on the straight line between (\tilde{U}, \tilde{V}) and $(\frac{\phi^A}{\gamma}, \frac{\mu-A}{r})$. Lemma 6.3 tells us that the saving-the-best-for-last technique can generate any payoff point lying on the unique solution g to the agency action ODE between (\tilde{U}, \tilde{V}) and $(\frac{\phi^A}{\gamma}, \frac{\mu-A}{r})$.

How do these two options compare? The decision over which option is preferable is dictated by the relative patience of the principal, measured by the discount ratio $\frac{r}{\gamma}$:

Remark 6.1. *If the principal is more patient than the agent, saving-the-best-for-last is better than randomization because the principal does not mind waiting out the agency action period to get to the more desirable contract. On the other hand, if the principal were more impatient than the agent, randomization is better. With randomization there is a chance the principal can immediately enact the more desirable contract. Despite the risks involved (i.e. getting stuck with the static contract), it is more efficient for an impatient principal to gamble than wait out the predetermined agency action period required by saving-the-best-for-last.*

This is graphically confirmed by the relevant solution g to the agency action ODE representing the potential payoff points generated by saving-the-best-for-last (see Figure 5). The discount ratio $\frac{r}{\gamma}$ dictates the concavity of g . A patient principal implies $\frac{r}{\gamma} < 1$ and the concave g curves over the straight line representing the potential payoff points generated by randomization. Conversely, an impatient principal implies $\frac{r}{\gamma} > 1$ and the convex g curves under the straight line representing the potential payoff points generated by randomization.

C. Warmup to the General Implementation Problem

I now solve the motivating problem of this section: the implementation problem in the case when agency action is inefficient (Lemma 6.5).

Definition 6.3. Agency action is *inefficient* when the agency action point $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$ lies below the extended optimal baseline value function $F^{ext,B}$.

Recall there are two contracting techniques that incorporate agency action: randomization and saving-the-best-for-last. Because $F^{ext,B}$ is concave and because we have assumed that agency action is inefficient:

Employing agency action through randomization is not useful.

However, Remark 6.1 tells us that saving-the-best-for-last is better than randomization. This is not to say that saving-the-best-for-last is a surefire way to improve the optimal baseline contract. But if the principal is artful in how he employs this technique, it is *possible* that the optimal baseline contract can be beat. In particular, the principal needs to be mindful of the following:

- The “last” can’t be too far away - depriving oneself of the best for too long is not optimal.
- The “best” has to be good enough - saving-the-best-for-last is not worth using if the best is only marginally better than the alternative.
- The “best” can’t be too good - if it’s too good, waiting for the best is inefficient.

The principal can always control the duration of the agency action period in saving-the-best-for-last, so satisfying the first condition does not pose much of a challenge. That I have assumed agency action is inefficient in this section means that the second condition is not an issue either. The tricky part is the third condition. Now if agency action is “too” inefficient, then the “best” is too good to give up, even for a little while, and that’s when the optimal contract takes the baseline form. But when agency action is not too inefficient, the value of saving-the-best-for-last implies that the optimal baseline contract can be beat. This intuition is formalized in the following lemma:

Lemma 6.4. *The optimal contract takes the baseline form (i.e. is the optimal baseline contract) if and only if first-best action beats saving-the-best-for-last on the margin:*

$$\frac{d}{dx} F^{ext,B} \Big|_{x=U} \geq \frac{r}{\gamma} \frac{F^{ext,B}(U) - \frac{\mu - A}{r}}{U - \frac{\phi A}{\gamma}} \quad \text{for all } U \in [K, \infty) \quad (4)$$

Condition (4) is a modification of the one found in DeMarzo and Sannikov (2006) Proposition 8. Written in the form of (4), the condition can be understood to be a statement about the superiority first-best action over saving-the-best-for-last. The left hand side $\frac{d}{dx} F^{ext,B} \Big|_{x=U}$ simply denotes the marginal utility of delivering an extra unit of utility to the agent through first-best action. What about the right hand side? Recall, randomization produces payoff points lying on the straight line between some payoff point $(U, F^{ext,B}(U))$ and the agency action point $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$. Thus the marginal utility of delivering an extra unit of utility to the agent through randomization is $\frac{F^{ext,B}(U) - \frac{\mu - A}{r}}{U - \frac{\phi A}{\gamma}}$. Remark 6.1 tells us that saving-the-best-for-last introduces a value-added distortion by picking up the measure of relative patience of




Figure 6: The three regions under $F^{ext,B}$ determined by condition (4).

the principal: $\frac{r}{\gamma}$. The marginal utility of delivering an extra unit of utility to the agent by departing from first-best action and implementing saving-the-best-for-last is the right hand side of (4):

$$\frac{r}{\gamma} \frac{F^{ext,B}(U) - \frac{\mu-A}{r}}{U - \frac{\phi A}{\gamma}}$$

So condition (4) simply says that if the principal cannot profitably deviate on the margin from first-best action by implementing saving-the-best-for-last, then the optimal contract is the optimal baseline contract.

C.1 The Optimal Contract when Agency Action is Not Too Inefficient

Lemma 6.5. *Lemma 6.4 splits the region under $F^{ext,B}$ into three regions (see Figure 6). The bottom region is where condition (4) holds and Lemma 6.4 implies the optimal contract is the optimal baseline contract. If instead, the agency action point is in the right region then the optimal contract is the optimal Quiet-Life contract. Finally, if the agency action point is in the left region then the optimal contract is the optimal Renegotiating Baseline contract.*

I now give a heuristic proof that the optimal contract is the optimal Quiet-Life contract when the agency action point is in the right region. One can mirror this argument to give a heuristic proof that the optimal contract is the optimal Renegotiating Baseline contract when the agency action point is in the left region. The formal proof of the general implementation problem (The Domains of Optimality Theorem of 7.A) is given in the Appendix.

So suppose the agency action point is in the right region. That condition (4) is not satisfied implies the existence of a contract payoff point $(\tilde{U}, F^{ext,B}(\tilde{U}))$ such that if the principal properly employs saving-the-best-for-last with this point as the “best,” then he can outdo the optimal baseline contract. To understand this statement graphically, let g be the unique




Figure 7: On the left: saving-the-best-for-last with $(\tilde{U}, F^{ext,B}(\tilde{U}))$ as the “best” leads to an improvement over $F^{ext,B}$. On the right: use Markov intuition to get an almost Quiet-Life contract.

solution to the agency action ODE going through $(\tilde{U}, F^{ext,B}(\tilde{U}))$. Lemma 6.3 tells us that the achievable payoff points when the principal uses saving-the-best-for-last with $(\tilde{U}, F^{ext,B}(\tilde{U}))$ as the “best” are all the points on g between the agency action point and $(\tilde{U}, F^{ext,B}(\tilde{U}))$. Remark 6.1 tells us that g is concave. This upward curvature represents the extra value of saving-the-best-for-last over randomization which is due to the relative patience of the principal : $r < \gamma$. Graphically we see that the curvature of the relevant portion of g is great enough to “pierce” the extended optimal baseline value function and the part that lies above $F^{ext,B}$ are all the payoff points that represent improvements (see left graphic of Figure 7). For example, the point $(\bar{U}, g(\bar{U}))$ is achieved if the principal uses saving-the-best-for-last with $(\tilde{U}, F^{ext,B}(\tilde{U}))$ as the best and waits for a duration of $D = \frac{1}{\gamma} \log \left(\frac{\tilde{U} - \frac{\phi A}{\gamma}}{\bar{U} - \frac{\phi A}{\gamma}} \right) = \frac{1}{r} \log \left(\frac{\tilde{V} - \frac{\mu - A}{r}}{\bar{V} - \frac{\mu - A}{r}} \right)$ before enacting the lucrative contract. Thus saving-the-best-for-last delivers \bar{U} to the agent more efficiently than does the optimal baseline contract.

The next step is to apply a little Markov intuition. If initially inducing agency action for D -units of time is better than first-best action, then it should be done *every* time the agent’s continuation payoff U_t hits \bar{U} . The resultant contract is structurally almost a Quiet-Life contract - inducing agency action for a D -length duration every time the agent’s continuation payoff and performance rating U_t hits the “good” performance threshold \bar{U} . The value graph¹⁰ of this almost Quiet-Life contract is a union of a piece of a solution to the first-best action ODE and a piece of a solution to the agency action ODE (see right graphic of Figure 7). And now we see that if the principal starts the agent’s continuation payoff at the arg max of the value graph then he will have written a contract with a higher payoff than that of the optimal baseline contract. Notice this contract’s agent continuation payoff U_t travels through the interval (\tilde{U}, \bar{U}) in two ways over the course of the contract. During agency action periods, it travels down the interval, with the continuation payoff point process moving leftwards on the upper trajectory (which solves the agency action ODE). And during first-best action periods, it travels stochastically in the interval, with the continua-

¹⁰I use graph because technically the value “function” of this contract is not a function.




Figure 8: On the left: a Quiet-Life contract. On the right: the optimal Quiet-Life contract which is also the optimal contract.

tion payoff point process moving on the lower trajectory (which solves the first-best action ODE). The gap between these two paths represents an efficiency loss and an opportunity for improvement - the principal would always rather be on the top path. To eliminate the efficiency loss, the principal will shift \bar{U} downwards until $\bar{U} = \tilde{U}$. Now the gap is closed and the new contract induces agency action at the good performance threshold $\bar{U} = \tilde{U}$, and we have produced a true Quiet-Life contract. Notice that the solutions to the first-best action and agency action ODEs which were used to build the value graph of the almost Quiet-Life contract have now pasted together at a single point for the true Quiet-Life contract (see left graphic of Figure 8). Pasting means that the derivatives of the two solutions coincide. The value function of this true Quiet-Life contract is the first-best action solution going from (K, L) to the pasting point plus the pasting point itself.¹¹

Finally, optimality requires that the pasting point be smooth - the second derivatives must coincide. So the final step is to shift the good performance threshold \tilde{U} to the unique value $U^{good,Q}$ where the smooth pasting condition holds. The resultant value function is the optimal Quiet-Life value function F^Q (see right graphic of Figure 8). I can now extract from F^Q the optimal Quiet-Life contract which is also the optimal contract. The good performance threshold is $U^{good,Q}$, the agent's continuation payoff is started at $U_0 = U^{contract} = \arg \max F^Q$, and the payoff to the principal is $\max F^Q$.

D. Equal Patience Versus a More Patient Principal

In the previous subsection I solved the implementation problem when agency action is inefficient under the model assumption that the principal is more patient: $r < \gamma$. The goal of the present subsection is to give a formal statement (Lemma 6.6) of the importance of patience

¹¹I emphasize the pasting point itself so the reader is not tempted to think of this as purely a first-best action contract. The pasting point is part of the solution to the agency action ODE as well.

and saving-the-best-for-last by describing how the results of the previous subsection change when $r = \gamma$.

Throughout this subsection I will keep γ fixed and let $r \leq \gamma$ be variable.

Recall when $r < \gamma$, there is an extended optimal baseline value function $F^{ext,B,r}$ which governs, among other things, the optimal baseline contract B^r .¹²

When $r = \gamma$, there is no optimal baseline contract. However, there are arbitrarily close-to-optimal baseline contracts. The principal can simply import the optimal baseline contract B^r from a setting where $r < \gamma$. As $r \uparrow \gamma$, B^r becomes arbitrarily close-to-optimal in the $r = \gamma$ setting.¹³ The extended close-to-optimal baseline value function $F^{ext,B,\gamma}$ is defined to be

$$F^{ext,B,\gamma} = \lim_{r \uparrow \gamma} F^{ext,B,r}$$

$F^{ext,B,\gamma}$ governs the close-to-optimal baseline contracts in the $r = \gamma$ setting. Recall $F^{ext,B,r}$ is concave for all $r < \gamma$,¹⁴ and indeed, so is $F^{ext,B,\gamma}$.

Remark 6.2. $F^{ext,B,r}$ is concave for all $0 < r \leq \gamma$. Since $F^{ext,B,r}$ governs the optimal baseline contract, concavity means the optimal baseline contract is randomization-proof. Why is this so? Recall the two main effects of randomizing over contracts are:

- Increases variance of payoffs.
- Allows the principal to gamble on the possibility of immediately exercising the lucrative contract in situations where the principal could alternatively choose saving-the-best-for-last (see Remark 6.1).

Intuitively, since the principal and agent are both risk neutral and the principal is at least as patient as the agent, neither of these effects should be utility improving over an otherwise optimally designed baseline contract.

But this gives us the punchline for the $r = \gamma$ case. When $r = \gamma$, saving-the-best-for-last is utility-equivalent to randomization, and since $F^{ext,B,\gamma}$ is concave and therefore randomization-proof, we have:

Lemma 6.6. *When the principal and agent are equally patient, the best contracts are baseline contracts if and only if agency action is inefficient.*

For the complete domains of optimality result when $r = \gamma$ see Appendix Figure 15.

7 Comparative Statics and Applications

A. The General Implementation Problem - Domains of Optimality

The domain of the agency action point is $\{(X,Y) | X > 0 \text{ and } Y < \frac{\mu}{r}\}$. I now solve the

¹²An “r” superscript is added because r is allowed to vary in this subsection.

¹³The reason there is no B^γ is because as $r \uparrow \gamma$, the good performance threshold $U^{good,B,r} \uparrow \infty$.

¹⁴See fact 1 about baseline contracts in Section 6.A.1.




Figure 9: An example extended optimal static value function when agency action is efficient.

general implementation problem culminating in The Domains of Optimality Theorem. For simplicity, I assume $K = 0$.¹⁵

In Section 6, I solved the implementation problem when agency action is inefficient (see Lemma 6.5). Recall, I introduced a contractual technique called saving-the-best-for-last which is useful only if the “best” is both good enough and not too good. Lemma 6.4 formalized this intuition which is summarized in the following remark.

Remark 7.1. *When agency action is inefficient, the principal can always find a “best” that is good enough. The only problem is that this “best” might be too good. If the best is too good then the optimal contract is the optimal baseline contract. If the best isn’t too good then the optimal contract is either the optimal Quiet-Life or optimal Renegotiating Baseline contract.*

The second part of the implementation problem is when agency action is efficient (i.e. lies above $F^{ext,B}$). And we have a companion mirror-image intuition for the efficient agency action case:

Remark 7.2. *When agency action is efficient, the principal can always find a “best” that is not too good. The only problem is that this “best” might not be good enough. If the best is not good enough then the optimal contract is the optimal static contract. If the best isn’t too good then the optimal contract is either the optimal Quiet-Life or optimal Renegotiating Baseline contract.*

Just as Lemma 6.4 formalizes the intuition of Remark 7.1, so will I derive a lemma to formalize the intuition of Remark 7.2.

Definition 7.1. *For any point (X, Y) in the domain of the agency action point: $\{(X, Y) | X > 0 \text{ and } Y < \frac{\mu}{r}\}$, define $F_{(X,Y)}^{ext,B}$ to be the function that is the extended optimal baseline value function in the alternate universe where the outside option point (K, L) is equal to (X, Y) .*

Since $K = 0$, the optimal static contract provides no salary, and its payoff point is simply the agency action point $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$.¹⁶ Since the agency action point is assumed to be efficient

¹⁵Higher K ’s introduce boundary conditions that complicate the analysis. See Appendix for when $K > 0$.

¹⁶See fact 2 about static contracts in Section 3.C.




Figure 10: The domains of optimality that determine the form of the optimal contract. The pictured quadrant is the domain of the agency action point: $\{(X, Y) | X > 0 \text{ and } Y < \frac{\mu}{r}\}$.

The faint dotted curve near the bottom is the extended optimal baseline value function $F^{ext,B}$. There is a bold curve underneath $F^{ext,B}$. Any point (X, Y) below this bold curve satisfies Condition (4) of Lemma 6.4: $\frac{d}{dx} F^{ext,B}|_{x=U} \geq \frac{r}{\gamma} \frac{F^{ext,B}(U) - Y}{U - X}$ for all $U \in [K, \infty)$. If the agency action point lies in this region then it is too inefficient and the optimal contract takes the baseline form.

Above $F^{ext,B}$ there is a V-shaped bold curve. Any point (X, Y) above this V-shaped bold curve satisfies Condition (5) of Lemma 7.1: $\frac{d^-}{dx} F^{ext,(X,Y)}|_{x=X} \geq 0$ and $\frac{d^+}{dx} F^{ext,(X,Y)}|_{x=X} \leq 0$. If the agency action point lies in this region then the optimal contract takes the static form. Each branch of this V-shaped curve represents a boundary at which one of the two differential inequalities of Condition (5) holds with equality.

The remaining two regions are where saving-the-best-for-last can be utilized in a manner just like or mirroring the method of 6.C.1. These regions are the domains of the Quiet-Life and Renegotiating Baseline forms.

(i.e. lies above $F^{ext,B}$), we know the optimal contract, whatever form it may be, will employ agency action. The simplest method to incorporate agency action is to employ it in a permanent way: use first-best action incentives when the agent's continuation payoff is not equal to $\frac{\phi A}{\gamma}$, but when $U_t = \frac{\phi A}{\gamma}$ make a permanent switch to agency action. This method leads to an improvement because by assumption permanent agency action delivers payoff $\frac{\phi A}{\gamma}$ to the agent more efficiently than any first-best action contract:

$$\frac{\mu - A}{r} > F^{ext,B} \left(\frac{\phi A}{\gamma} \right)$$

The resultant value function is what I call the *extended optimal static value function*. Formally,

Definition 7.2. The *extended optimal static value function* $F^{ext,S}$ is the unique solution to the first-best action ODE going from (K, L) to the agency action point $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$ plus the agency action point itself plus $F_{(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})}^{ext,B}$ (see Figure 9).

More generally, for a point (X, Y) in the agency action domain, define $F^{ext,(X,Y)}$ to be the union of the unique solution to the first-best action ODE going from (K, L) to (X, Y) plus the point (X, Y) itself plus $F_{(X,Y)}^{ext,B}$.

I can now state the lemma that formalizes the intuition of Remark 7.2.

Lemma 7.1. When agency action is efficient, the optimal contract takes on the static form if and only if the principal can't find a good enough "best" with which to use saving-the-best-for-last either to the left or to the right of the agency action point:

$$\frac{d^-}{dx} F^{ext,S} \Big|_{x=\frac{\phi A}{\gamma}} \geq 0 \quad \text{and} \quad \frac{d^+}{dx} F^{ext,S} \Big|_{x=\frac{\phi A}{\gamma}} \leq 0 \quad (5)$$

Putting everything together, we have:

The Domains of Optimality Theorem. For any realization of the model parameters, the domain of the agency action point $\{(X, Y) | X > 0 \text{ and } Y < \frac{\mu}{r}\}$ can be split into four regions (see Figure 10). The boundaries of the four regions are determined by the differential conditions of Lemmas 6.4 and 7.1. The bottom region is the Baseline Domain and the top region is the Static Domain. The right region is the Quiet-Life Domain and the left region is the Renegotiating Baseline Domain. Whichever domain contains the agency action point, the optimal contract takes the corresponding form.

Proof. See Appendix. □

B. Optimal Contracting Under Taxes

Introduce a tax $\mathcal{T} \in [0, 1)$ on cash compensation, so that for every dollar paid by the principal the agent only receives a fraction $1 - \mathcal{T}$. To achieve a target agent continuation payoff, a tax forces the principal to inflate the portion of the cash flow paid to the agent.




Figure 11: The effect of a tax on the efficiency threshold and the extended optimal baseline value function.

Definition 7.3. The **taxed efficiency threshold** is the line $\mu = ry + \frac{\gamma x}{1-\mathcal{T}}$. The **taxed efficiency threshold** is a locus of efficient payoff points in the first-best action setting taking taxes into account (see fact 7 about baseline contracts in section 6.A.1). To achieve a continuation payoff point (U_t, V_t) lying on the taxed efficiency threshold, the principal with access to first-best returns μdt and subject to tax \mathcal{T} simply diverts a portion $\frac{\gamma U_t}{1-\mathcal{T}} dt$ of the flow to the agent.

If $\mathcal{T} = 0$ then we call the threshold simply the **efficiency threshold**.

The following is a straightforward generalization to the tax setting of Proposition 1, DeMarzo and Sannikov (2006) (proof omitted, see stronger version Lemma 9.1 in Appendix):

Lemma 7.2. The optimal static first-best action contract in the setting with tax \mathcal{T} exists and is a baseline contract: **the optimal taxed baseline contract**. Denote the corresponding optimal taxed baseline value function by $F_{\mathcal{T}}^B$, and the good performance threshold by $U_{\mathcal{T}}^{\text{good},B}$. The cash compensation point $(U_{\mathcal{T}}^{\text{good},B}, F_{\mathcal{T}}^B(U_{\mathcal{T}}^{\text{good},B}))$ of the optimal taxed baseline contract lies on the corresponding taxed efficiency threshold:

$$\mu = rF_{\mathcal{T}}^B(U_{\mathcal{T}}^{\text{good},B}) + \frac{\gamma U_{\mathcal{T}}^{\text{good},B}}{1-\mathcal{T}}$$

and is determined by a smooth-pasting condition. $F_{\mathcal{T}}^B$ is the unique solution to the first-best action ODE going from the outside option point (K, L) to $(U_{\mathcal{T}}^{\text{good},B}, F_{\mathcal{T}}^B(U_{\mathcal{T}}^{\text{good},B}))$ and it is concave.

The optimal taxed baseline contract has good performance threshold $U_{\mathcal{T}}^{good,B}$, the agent's continuation payoff is started at $U_0 = \arg \max F_{\mathcal{T}}^B$, and the payoff to the principal is $\max F_{\mathcal{T}}^B$.

If the principal exercised cash compensation either before or after $U_{\mathcal{T}}^{good,B}$, then the cash compensation point would lie below the taxed efficiency threshold and imply an efficiency loss. Graphically, we observe that a tax \mathcal{T} lowers the efficiency threshold, bringing down with it the optimal baseline contract's good performance threshold, principal's payoff, and agent's payoff. The optimal taxed baseline contract is also more susceptible to termination.

Not surprisingly, a tax \mathcal{T} alters the *domains of optimality*. Recall, the boundaries of the domains of optimality are defined by differential conditions on the extended optimal baseline and optimal static value functions (see conditions (4) and (5)). While the differential conditions remain unchanged, the extended value functions are affected by the tax,¹⁷ leading to altered boundaries (see Figure 11). A tax hike's effect on the domains of optimality can largely be summarized by two shifts from contracts emphasizing cash compensation (baseline and Renegotiating Baseline) to those that don't (Quiet-Life and static).

Remark 7.3. *With a tax hike, a subset of the agency action point values in the domain of the baseline form pre-tax now belong to the domain of the Quiet-Life form. This reflects the increased attractiveness of tax-free perks-based hidden compensation over taxed cash compensation.*

Mirroring this shift, a subset of the agency action point values in the domain of the Renegotiating Baseline form pre-tax now belong to the domain of the static form. As taxes increase, the underlying baseline contract that the Renegotiating Baseline contract is renegotiating becomes increasingly unprofitable, and not worth the trouble renegotiating. So the principal drops it and the Renegotiating Baseline contract degenerates into the static contract.

C. Optimal Contracting with Bargaining

When the principal's outside option L is very low (e.g. there is a threat of litigation by the agent for termination) the principal will employ a static contract or a Renegotiating Baseline contract to avoid termination. More generally, suppose the principal and agent can bargain for their outside options along some efficient **bargaining possibility frontier**, which is a concave decreasing function b . Assume that the principal can't bargain for more than the asset is worth ($b < \frac{L}{r}$) and the agent's outside option will be at least zero (domain of b is $[0, \infty)$).

Lemma 7.3. *Fix any bargaining possibility frontier and a tax \mathcal{T} . For all agents with sufficiently strong bargaining power, the optimal contract is either a low effort contract with salary (optimal static contract) or a high-effort contract with a suspension clause (optimal*

¹⁷Recall the extended optimal baseline value function $F^{ext,B}$ is defined to be F^B with the straight line of slope -1 attached to its end. In the tax \mathcal{T} setting, the extended optimal taxed baseline value function $F_{\mathcal{T}}^{ext,B}$ is similarly defined to be $F_{\mathcal{T}}^B$ with the straight line of slope $\frac{1}{\mathcal{T}-1}$ attached to its end. Just like $F^{ext,B}$ in the no tax setting, $F_{\mathcal{T}}^{ext,B}$ implies the structure of the optimal first-best action contract delivering payoff $x \in [K, \infty)$ to the agent in the tax setting.




Figure 12: (a) The optimal baseline value function and above it the optimal Renegotiating Baseline value function. The dotted lines are not part of the value function, but rather highlight the solutions to the first-best action and agency action ODEs used to construct the value function. (b) There are now three additional value functions. Of the three new value functions, the top one is not quite renegotiation-proof, the middle one is just barely renegotiation-proof, and the bottom one is “too” renegotiation-proof.

Renegotiating Baseline contract). The optimal choice is the low effort contract with salary if and only if agency action lies on or above the taxed efficiency threshold.

In particular, sufficiently high taxes and agent bargaining power means the optimal contract is the low effort contract with salary.

Proof. This is a consequence of the Domains of Optimality Theorem when $K > 0$ found in the Appendix. \square

Other bargaining models include if the agent is able to bargain for a fixed salary sdt throughout the duration of the contract. In this case the change in setting is isomorphic to a linear change-of-variables ($\mu \rightarrow \mu - s$, $\phi A \rightarrow \phi A + s$, etc.) and all the previous optimal contracting results apply.

Finally, the agent can bargain for a higher payoff than that of the optimal contract. The complete optimal contracting theorem for higher agent payoffs can be found in the Appendix 9.D.

D. Renegotiation-Proof Contracts

Recall baseline contracts are usually not renegotiation-proof. Indeed, the main reason the Renegotiating Baseline contract is sometimes optimal is because it allows the principal to renegotiate the underlying baseline contract to some extent. Not surprisingly, Renegotiating Baseline contracts are less renegotiable, and some are even renegotiation-proof.

A contract is renegotiation-proof if its value function is never upward sloping. Consider the setting in Figure 12a which depicts the value functions of the optimal baseline contract




Figure 13: The bottom value function is that of the optimal renegotiation-proof first-best action contract. The lightly dotted line attached to it represents the contract stipulation that termination be randomized at the poor performance threshold.

and the optimal Renegotiating Baseline contract (which is also the optimal contract). Both value functions have upward sloping portions so neither contract is renegotiation-proof. However, the upward slope of the Renegotiating Baseline value function is less steep than that of the baseline. Hence, the Renegotiating Baseline contract is closer to being renegotiation-proof. This is achieved because the principal has set the poor performance threshold of the Renegotiating Baseline contract to be higher than K - the poor performance threshold of the baseline contract. This alteration prevents the agent's continuation payoff from dropping too low, which precipitates the need to renegotiate.

In general, the higher the poor performance threshold, the more renegotiation that is embedded in the contract, and consequently, the closer the contract is to being renegotiation proof. At some point, the poor performance threshold is high enough that the value function no longer has an upward sloping portion. The value function simply starts with slope 0, then gradually decreases from there. The corresponding Renegotiating Baseline contract is then renegotiation-proof. (See Figure 12b)

Now the principal can continue to set even higher poor-performance thresholds, and the resultant Renegotiating Baseline contracts will also be renegotiation-proof. But "over-forgiving" the agent entails an efficiency loss, and these renegotiation-proof contracts are not as profitable. (See Figure 12b)

Recall from DeMarzo and Sannikov (2006), there also exist renegotiation-proof first-best action contracts which are basically modified baseline contracts. The poor performance threshold of baseline contracts is always K , and termination always occurs there. In the modified renegotiation-proof baseline contracts, the poor performance threshold is shifted upwards, and termination is randomized there.

As Figure 13 demonstrates, the renegotiation-proof contracts using agency action may dom-

inate those only inducing first-best action. Moreover:

Remark 7.4. *Renegotiation-proof contracts that use agency action do not require randomization or termination, and their value to the principal is not bounded by the value of the principal's outside option L . This is in direct contrast to renegotiation-proof contracts only inducing first-best action.*

E. The Geometric Brownian Setting

All the results of this paper can be translated over to the corresponding geometric Brownian setting. The model is as follows:

E.1 Setting

There is an asset belonging to a principal, for which he contracts an agent to manage. The asset produces a stochastic revenue stream. Over time, we assume that the cumulative revenue stream behaves as geometric Brownian Motion with a drift process influenced by the hidden action applied by the contracted agent.

Formally, there is a stochastic process $Z = \{Z_t\}_{t \geq 0}$ defined on a probability space Ω with probability law P^μ . Under P^μ , Z is Brownian motion with drift μdt . Upon Z_t is defined a geometric Brownian Motion:

$$dS_t = S_t dZ_t$$

At time t , S_t is the cumulative revenue stream of the asset up to time t . The μdt drift corresponds to the *scaled*, default expected returns and can be interpreted as the scaled, intrinsic or maximum expected profitability of the asset.

E.2 Actions

The agent affects asset performance by selecting an action at each moment in time. Over time the agent's action process $a = \{a_t\}_{t \geq 0}$ is a stochastic process taking values in a set $\{0, A\}$ with $A > 0$. $\{0\}$ is first-best action and $\{A\}$ is **agency action**. The action process a affects the underlying probability law: the default law P^μ changes to $P^{\mu-a}$, which is defined to be the law under which Z is Brownian motion with drift $(\mu - a_t)dt$.

The principal can choose a compensation scheme for the agent. Compensation is represented by a random nondecreasing process $I = \{I_t\}_{t \geq 0}$ started at zero that keeps track of the cumulative cash payments made to the agent up to time t . Liquidation is a stopping time τ .

E.3 Preferences

The principal is risk neutral, discounts at rate r , retains the cash flow of the asset, compensates the agent, and can retain a value LS_t after the liquidation of the asset. His utility is

$$\mathbf{E}_{P^{\mu-a}} \left[\int_0^\tau e^{-rt} (dS_t - dI_t) + e^{-r\tau} LS_\tau \right]$$

The agent is risk neutral, discounts at rate γ , receives compensation from the principal,

and retains a value KS_t after the liquidation of the contract. The agent also receives an instantaneous utility flow $\phi a_t S_t dt$ by applying action $a_t \in \{0, A\}$ at time t , where $\phi > 0$. His utility is

$$\mathbf{E}_{P^{\mu-a}} \left[\int_0^\tau e^{-\gamma s} (dI_t + \phi a_t S_t dt) + e^{-\gamma \tau} K S_t \right]$$

We assume that the principal is at least as patient: $r \leq \gamma$. We also require $\mu < r$, $L < \frac{1}{r-\mu}$, and $K \geq 0$.

E.4 Equivalence via Scaling

Despite the existence of *two* state variables now: the familiar agent continuation payoff U_t and the new geometric Brownian S_t , the key is to realize that there is only one *effective* state variable, which is the scaled agent continuation payoff $u_t = \frac{U_t}{S_t}$. This is the main point of He (2009). Once the model is scaled, all of the technical constructs from the Brownian model translate over: the first-best action and agency action laws fall out, as do the first-best action ODE and agency action ODE. As a result, the optimal contracts, the notions of sticky incentives and infrequent monitoring, the analysis of saving-the-best-for-last, the domains of optimality, and the rest of the comparative statics all translate over largely unchanged.

8 Conclusion

In this paper I have explicitly solved for the optimal contract in a general Brownian framework where agency action plays an integral role in optimal contracting. The framework underlies the models of DeMarzo and Sannikov (2006), Biais et al (2007) and He (2009). However, all of these papers focus on finding the optimal contract that is restricted to always inducing first-best action.

While static first-best action contracts have provided a great deal of insights into some contracting problems observed in real-life, many arrangements do not always employ static first-best action. Indeed, recent empirical work such as Yermack (2006) and Rajan and Wulf (2006) have pointed to possible uses of agency action in business.

In my dynamic model, I find that the optimal contract takes on one of four forms depending on fundamentals, including two that mix agency action phases in between periods of first-best action: the Quiet-Life form and the Renegotiating Baseline form. Quiet-Life contracts induce agency action as a form of reward and can be thought of as contracts that allow for efficient perks consumption. Renegotiating Baseline contracts are contracts that mostly induce first-best action but periodically trigger agency action phases as a form of punishment. These agency action phases can be thought of as suspension during which the agent applies low effort.

That the optimal contract may take one of these two forms helps demonstrate not only the value of agency action but also more broadly, the value of infrequent monitoring and slowing down incentives. Moreover, I show how taxes affect what form the optimal contract takes and how agency action can be utilized to produce optimal renegotiation-proof contracts.

Overall, the results of this paper help not only to bridge a gap in the dynamic contracting literature, but also provide a theoretical counterpart to the ongoing empirical research into the role of agency action.

9 Appendix

9.A Preliminary Results

It is best to skip directly to Section 9.B and to refer back to Section 9.A when needed.

The following lemma is an easy generalization of DeMarzo and Sannikov (2006) to the tax setting stated without proof.

Lemma 9.1. *Fix a setting with tax $\mathcal{T} \in (-\infty, 1)$. The following set of facts characterize optimal static first-best action contracting in the tax setting:*

- 1) *If the outside option point (K, L) is strictly below the taxed efficiency threshold:*

$$\mu > rL + \frac{\gamma K}{1 - \mathcal{T}}$$

then the optimal taxed baseline value function $F_{\mathcal{T}}^B$ is a concave solution to the first-best action ODE on $[K, U_{\mathcal{T}}^{\text{good}, B}]$ where $F_{\mathcal{T}}^B$ and $U_{\mathcal{T}}^{\text{good}, B}$ are uniquely determined by a smooth pasting condition: There exists a unique $U_{\mathcal{T}}^{\text{good}, B}$ and a unique $F_{\mathcal{T}}^B$ such that $F_{\mathcal{T}}^B(K) = L$, $F_{\mathcal{T}}^B'(U_{\mathcal{T}}^{\text{good}, B}) = \frac{1}{\mathcal{T}-1}$, and $F_{\mathcal{T}}^B''(U_{\mathcal{T}}^{\text{good}, B}) = 0$. $F_{\mathcal{T}}^B$ is strictly concave on $(K, U_{\mathcal{T}}^{\text{good}, B})$ and $F_{\mathcal{T}}^B'(x) \geq \frac{1}{\mathcal{T}-1}$ for all $x \in [K, U_{\mathcal{T}}^{\text{good}, B}]$.

- 2) *The optimal taxed static first-best action contract delivering payoff $x \in [K, U_{\mathcal{T}}^{\text{good}, B}]$ to the agent exists. It is the baseline contract with good performance threshold $U_{\mathcal{T}}^{\text{good}, B}$, the agent's continuation payoff is started at $U_0 = x$, and the payoff to the principal is $F_{\mathcal{T}}^B(x)$. Call this contract the optimal taxed baseline contract delivering payoff $x \in [K, U_{\mathcal{T}}^{\text{good}, B}]$ to the agent.*
- 3) *One can extend $F_{\mathcal{T}}^B$ to values of $x > U_{\mathcal{T}}^{\text{good}, B}$. More generally, the extended optimal taxed baseline value function $F_{\mathcal{T}}^{\text{ext}, B}$ is $F_{\mathcal{T}}^B$ with a straight line of slope $\frac{1}{\mathcal{T}-1}$ attached to the end. $F_{\mathcal{T}}^{\text{ext}, B} \in C^2[K, \infty)$.*
- 4) *The optimal taxed static first-best action contract delivering payoff $x > U_{\mathcal{T}}^{\text{good}, B}$ to the agent exists. It first delivers a lump sum $x - U_{\mathcal{T}}^{\text{good}, B}$ to the agent. Then the contract becomes the optimal taxed baseline contract delivering payoff $U_{\mathcal{T}}^{\text{good}, B}$ to the agent. The payoff to the principal is $F_{\mathcal{T}}^{\text{ext}, B}(x)$. Call this contract the optimal taxed baseline contract delivering payoff $x > U_{\mathcal{T}}^{\text{good}, B}$ to the agent.*

- 5) The **optimal taxed baseline contract** is the optimal taxed baseline contract delivering payoff $\arg \max F_{\mathcal{T}}^{\text{ext},B} = \arg \max F_{\mathcal{T}}^B$ to the agent. The payoff to the principal is $\max F_{\mathcal{T}}^{\text{ext},B} = \max F_{\mathcal{T}}^B$.
- 6) Fix an optimal taxed baseline contract delivering some payoff to the agent. At time t if the agent's continuation payoff is U_t then the principal's continuation payoff is $F_{\mathcal{T}}^{\text{ext},B}(U_t)$.
- 7) Cash compensation occurs when the principal's and agent's required expected cash flows exhaust expected returns taking taxes into account:

$$\mu = rF_{\mathcal{T}}^B(U_{\mathcal{T}}^{\text{good},B}) + \frac{\gamma U_{\mathcal{T}}^{\text{good},B}}{1 - \mathcal{T}}$$

- 8) If the outside option point lies on or above the taxed efficiency threshold then $F_{\mathcal{T}}^B$ is just the single point (K, L) and optimal taxed baseline contract is simply to terminate right away. Also $F_{\mathcal{T}}^{\text{ext},B}$ is just the straight line of slope $\frac{1}{\mathcal{T}-1}$ starting at (K, L) .
- 9) Let $\mathcal{T}_1 < \mathcal{T}_2$ be two taxes. For all $x > K$, $\frac{\mu}{r} > F_{\mathcal{T}_1}^{\text{ext},B}(x) > F_{\mathcal{T}_2}^{\text{ext},B}(x)$. Also $F_{\mathcal{T}_1}^{\text{ext},B}'(K) > F_{\mathcal{T}_2}^{\text{ext},B}'(K)$.

The Regularity Lemma. Let f_1 and f_2 be two distinct solutions to the first-best action ODE and $x^* \geq 0$. If

$$f_1(x^*) \geq f_2(x^*) \text{ and } f_1''(x^*) \leq f_2''(x^*)$$

then

$$f_1''(x) < f_2''(x) \text{ for all } x \in (x^*, \infty)$$

If

$$f_1(x^*) \geq f_2(x^*) \text{ and } f_1''(x^*) \geq f_2''(x^*)$$

then

$$f_1''(x) > f_2''(x) \text{ for all } x \in [0, x^*)$$

Proof. The straightforward, albeit tedious, proof of this lemma involves Euler's Method. Fix a set of initial conditions for the first-best action ODE: $(x^*, f(x^*), f'(x^*))$ with $x^* \geq 0$. Then

$$f''(x^*) = \frac{rf(x^*) - \mu - \gamma x^* f'(x^*)}{\phi^2/2}$$

and by Euler's Method, we have

$$\begin{aligned} f(x^* + \Delta x) &\approx f(x^*) + f'(x^*)\Delta x \\ f'(x^* + \Delta x) &\approx f'(x^*) + \frac{rf(x^*) - \mu - \gamma x^* f'(x^*)}{\phi^2/2}\Delta x \\ f''(x^* + \Delta x) &\approx \frac{rf(x^* + \Delta x) - \mu - \gamma(x^* + \Delta x)f'(x^* + \Delta x)}{\phi^2/2} \end{aligned}$$

$$\begin{aligned}
&= \frac{r(f(x^*) + f'(x^*)\Delta x) - \mu - \gamma(x^* + \Delta x)(f'(x^*) + \frac{rf(x^*) - \mu - \gamma x^* f'(x^*)}{\phi^2/2} \Delta x)}{\phi^2/2} \\
&= \frac{(1 - \gamma(x^* + \Delta x)\frac{\Delta x}{\phi^2/2})[rf(x^*) - \mu - \gamma x^* f'(x^*)] - (\gamma - r)\Delta x f'(x^*)}{\phi^2/2} \\
&= \left[1 - \gamma(x^* + \Delta x)\frac{\Delta x}{\phi^2/2} \right] f''(x^*) - \frac{(\gamma - r)\Delta x f'(x^*)}{\phi^2/2}
\end{aligned}$$

Now let f_1 and f_2 satisfy the hypothesis of the first half of the lemma at x^* and fix an arbitrary upper bound D with $x^* < D$. Let Δx be small enough so that $1 - \gamma(D + \Delta x)\frac{\Delta x}{\phi^2/2} > 0$. The assumptions imply $f_1'(x^*) > f_2'(x^*)$, and then it is easy to see that the Euler approximations of f_1 and f_2 satisfy the hypothesis of the first half of the lemma at $x^* + \Delta x$ as well. In fact, the second derivative of the Euler approximation of f_1 is now *strictly* less than that of the Euler approximation of f_2 at $x^* + \Delta x$. Then induction shows that the second derivative of the Euler approximation of f_1 is strictly less than that of the Euler approximation of f_2 at $x^* + n\Delta x$, so long as $x^* + n\Delta x \in (x^*, D]$. Letting $\Delta x \rightarrow 0$, we have

$$f_1''(x) < f_2''(x) \text{ for all } x \in (x^*, D]$$

Since D was arbitrary, the first half of the lemma holds.

Now suppose f_1 and f_2 satisfy the hypothesis of the second half of the lemma. If $f_1 > f_2$ on $[0, x^*)$ then the second half of the lemma must hold. Suppose not, then there is some $\tilde{x} \in [0, x^*)$ such that $f_1''(\tilde{x}) \leq f_2''(\tilde{x})$. But then the first half of the lemma implies that $f_1'(x^*) < f_2'(x^*)$. Contradiction.

So it suffices to prove $f_1 > f_2$ on $[0, x^*)$. The hypothesis of the second half of the lemma immediately implies that f_1 lies above f_2 in a left neighborhood of x^* . This means that if it is not true that $f_1 > f_2$ on $[0, x^*)$ then there must be some point \tilde{x} such that $f_1(\tilde{x}) = f_2(\tilde{x})$ and $f_1'(\tilde{x}) > f_2'(\tilde{x})$. But then this implies that $f_1''(\tilde{x}) < f_2''(\tilde{x})$ and once again the first half of the lemma implies a contradiction. \square

Corollary 9.1. *Let f_1 and f_2 be two distinct solutions to the first-best action ODE and $x^* \geq 0$. If*

$$f_1(x^*) \geq f_2(x^*) \text{ and } f_1'(x^*) \geq f_2'(x^*)$$

then

$$f_1'(x) > f_2'(x) \text{ for all } x \in (x^*, \infty)$$

If

$$f_1(x^*) \geq f_2(x^*) \text{ and } f_1'(x^*) \leq f_2'(x^*)$$

then

$$f_1'(x) < f_2'(x) \text{ for all } x \in [0, x^*)$$

Proof. If f_1 and f_2 satisfy the assumptions of the second half of the corollary, then $f_1''(x^*) < f_2''(x^*)$. Then the second half of the corollary follows from the second half of the Regularity

Lemma.

Now suppose $f_1(x^*) \geq f_2(x^*)$ and $f_1'(x^*) \geq f_2'(x^*)$ and there exists an $x \in (x^*, \infty)$ such that $f_1'(x) \leq f_2'(x)$. Without loss of generality, we may choose x so that $f_1(x) > f_2(x)$. But then the second half of the corollary implies a contradiction. \square

Corollary 9.2. *Let $\mathcal{T}_1 < \mathcal{T}_2$ be two taxes and $x \geq K$. Then*

$$F_{\mathcal{T}_1}^{ext,B}{}'(x) > F_{\mathcal{T}_2}^{ext,B}{}'(x)$$

Proof. $F_{\mathcal{T}_1}^{ext,B}{}'(x) = F_{\mathcal{T}_1}^{ext,B}{}'(x \wedge U_{\mathcal{T}_1}^{good,B}) > F_{\mathcal{T}_2}^{ext,B}{}'(x \wedge U_{\mathcal{T}_1}^{good,B}) \geq F_{\mathcal{T}_2}^{ext,B}{}'(x)$. If $x \wedge U_{\mathcal{T}_1}^{good,B}$ is in the domain of the linear branch of $F_{\mathcal{T}_2}^{ext,B}$ then the middle inequality comes from Lemma 9.1. If $x \wedge U_{\mathcal{T}_1}^{good,B}$ is in the domain of $F_{\mathcal{T}_2}^B$ then the middle inequality comes from Corollary 9.1. \square

Corollary 9.3. *The limit $F_{-\infty}^B$ of the extended optimal taxed baseline value functions:*

$$F_{-\infty}^B(x) = \lim_{\mathcal{T} \downarrow -\infty} F_{\mathcal{T}}^{ext,B}(x) \text{ for all } x \geq K$$

is a solution to the first-best action ODE. For all $x \geq K$, $F_{-\infty}^B{}'(x) > 0$ and

$$F_{-\infty}^B{}'(x) = \lim_{\mathcal{T} \downarrow -\infty} F_{\mathcal{T}}^{ext,B}{}'(x)$$

Proof. For any tax \mathcal{T} , define $s_{\mathcal{T}} \equiv F_{\mathcal{T}}^{ext,B}{}'(K)$. For any slope s , define f_s to be the unique solution to the first-best action ODE starting at (K, L) with initial slope s .

We will need to make use of the fact that $\frac{\mu}{r}$ is a constant solution to the first-best action ODE.

We begin by proving some results that show $f_{s_{\mathcal{T}}}$ and $F_{\mathcal{T}}^{ext,B}$ are close to each other.

Let us first show that $|f_{s_{\mathcal{T}}} - F_{\mathcal{T}}^{ext,B}| \leq \frac{1}{1-\mathcal{T}}$. Since $F_{\mathcal{T}}^B = f_{s_{\mathcal{T}}}|_{[K, U_{\mathcal{T}}^{good,B}]}$, $f_{s_{\mathcal{T}}}'(U_{\mathcal{T}}^{good,B}) = F_{\mathcal{T}}^{ext,B}{}'(x) = \frac{1}{\mathcal{T}-1}$ for all $x \geq U_{\mathcal{T}}^{good,B}$, so it suffices to show that $f_{s_{\mathcal{T}}}|_{[U_{\mathcal{T}}^{good,B}, \infty)}$ is decreasing, convex. $f_{s_{\mathcal{T}}}|_{[U_{\mathcal{T}}^{good,B}, \infty)}$ is certainly initially decreasing. If it is not always decreasing then there is some value $x^* > U_{\mathcal{T}}^{good,B}$ such that $f_{s_{\mathcal{T}}}(x^*) < \frac{\mu}{r}$ and $f_{s_{\mathcal{T}}}'(x^*) = 0$. Comparing $f_{s_{\mathcal{T}}}$ and $\frac{\mu}{r}$, Corollary 9.1 tells us that $f_{s_{\mathcal{T}}}$ is increasing on $[K, x^*)$. Contradiction. So $f_{s_{\mathcal{T}}}|_{[U_{\mathcal{T}}^{good,B}, \infty)}$ is always decreasing. Furthermore, we know that $f_{s_{\mathcal{T}}}''(U_{\mathcal{T}}^{good,B}) = 0$. Again comparing $f_{s_{\mathcal{T}}}$ and $\frac{\mu}{r}$, the Regularity Lemma tells us that $f_{s_{\mathcal{T}}}|_{[K, U_{\mathcal{T}}^{good,B}]}$ is convex.

From this result we can easily deduce the following two results. On any bounded interval $[K, D]$ we have $|f_{s_{\mathcal{T}}} - F_{\mathcal{T}}^{ext,B}| \leq \frac{D}{1-\mathcal{T}}$. Also, since $F_{\mathcal{T}}^B{}' \geq \frac{1}{\mathcal{T}-1}$ (Lemma 9.1), $F_{\mathcal{T}}^B{}'(U_{\mathcal{T}}^{good,B}) = \frac{1}{\mathcal{T}-1}$ (Lemma 9.1), and $f_{s_{\mathcal{T}}}$ is convex after $U_{\mathcal{T}}^{good,B}$, it must be that $f_{s_{\mathcal{T}}}'(x) \geq \frac{1}{\mathcal{T}-1}$ for all $x \geq K$.

Let us now show that $s_{-\infty} \equiv \lim_{\mathcal{T} \downarrow -\infty} s_{\mathcal{T}}$ is finite. Let $x > 0$ and let h_x be the solution to the first-best action ODE such that $h_x(x) = \frac{\mu}{r}$ and $h_x'(x) = 1$. By comparing h_x to $\frac{\mu}{r}$,

The Regularity Lemma tells that h_x is concave increasing on $[0, x]$. Now if $x = K + \frac{\mu}{r} - L$, then $h_x(K) < L$. By intermediate value theorem, there is some $x^* \in (K, K + \frac{\mu}{r} - L)$ such that $h_{x^*}(K) = L$. Thus we have found a solution to the first-best action ODE starting at (K, L) that crosses the line $y = \frac{\mu}{r}$. Clearly, $s_{\mathcal{T}} < h'_{x^*}(K)$ for all \mathcal{T} and we have shown $s_{-\infty} \leq h'_{x^*}(K) < \infty$.

Clearly, $f_{s_{-\infty}} = \lim_{\mathcal{T} \downarrow -\infty} f_{s_{\mathcal{T}}}$ and $f'_{s_{-\infty}} = \lim_{\mathcal{T} \downarrow -\infty} f'_{s_{\mathcal{T}}}$. Since we have already shown that $f'_{s_{\mathcal{T}}}(x) \geq \frac{1}{\mathcal{T}-1}$, it must be that $f'_{s_{-\infty}}(x) \geq 0$ for all $x \geq K$.

Let us now prove the stronger result claimed by the Corollary: $f'_{s_{-\infty}}(x) > 0$ for all $x \geq K$. Suppose not. Then there is some value x^* such that $f'_{s_{-\infty}}(x^*) = 0$. Now compare $f_{s_{-\infty}}$ to $\frac{\mu}{r}$. If $f_{s_{-\infty}}(x^*) \geq \frac{\mu}{r}$ then Corollary 9.1 tells us that $f_{s_{-\infty}}$ either is $\frac{\mu}{r}$ or always lies above $\frac{\mu}{r}$ on $[K, \infty)$. Both are contradictions. So $f_{s_{-\infty}}(x^*) < \frac{\mu}{r}$. But then $f''_{s_{-\infty}}(x^*) < 0$ which means that $f_{s_{-\infty}}$ is decreasing after x^* . Contradiction.

We can now prove the last part of the Corollary. Fix a bounded interval $[K, D]$. Then on this interval:

$$\lim_{\mathcal{T} \downarrow -\infty} |f_{s_{-\infty}} - F_{\mathcal{T}}^{ext, B}| \leq \lim_{\mathcal{T} \downarrow -\infty} |f_{s_{-\infty}} - f_{s_{\mathcal{T}}}| + \frac{D}{1 - \mathcal{T}} = 0$$

Since D is arbitrary, we have $f_{s_{-\infty}} = \lim_{\mathcal{T} \downarrow -\infty} F_{\mathcal{T}}^{ext, B}$. Moreover,

$$\lim_{\mathcal{T} \downarrow -\infty} |f'_{s_{-\infty}} - F_{\mathcal{T}}^{ext, B \prime}| \leq \lim_{\mathcal{T} \downarrow -\infty} |f'_{s_{-\infty}} - f'_{s_{\mathcal{T}}}| + \frac{1}{1 - \mathcal{T}} = 0$$

and so we have $f'_{s_{-\infty}} = \lim_{\mathcal{T} \downarrow -\infty} F_{\mathcal{T}}^{ext, B \prime}$. Thus $F_{-\infty}^B = f_{s_{-\infty}}$ and we are done. \square

Corollary 9.4. Recall $F_{(X, Y)}^{ext, B}$ is $F^{ext, B}$ in the alternate universe where the outside option point is (X, Y) (see Definition 7.1). Similarly define $F_{(X, Y)}^B$ and $U_{(X, Y)}^{good, B}$. Fix an X and let $Y_1 > Y_2$ be two numbers such that (X, Y_1) and (X, Y_2) are below the efficiency threshold. Let $x \in [X, U_{(X, Y_1)}^{good, B}]$. Then

$$F_{(X, Y_1)}^B \prime(x) < F_{(X, Y_2)}^B \prime(x)$$

Proof. We have $F_{(X, Y_1)}^B \prime(U_{(X, Y_1)}^{good, B}) < F_{(X, Y_2)}^B \prime(U_{(X, Y_1)}^{good, B})$. Corollary 9.1 implies the result. \square

Definition 9.1. The **first-best action inequality** is the first-best action ODE with equality replaced with a “ \geq ”:

$$ry \geq \mu + \gamma xy' + \frac{\phi^2}{2} y''$$

Similarly, the **agency action inequality** is:

$$ry \geq \mu - A + (\gamma x - \phi A) y'$$

The C^1 Maximum Principle. Let F^{ext} be some function on $[K, \infty)$ such that $F^{ext}(K) \geq L$, F^{ext} is concave, $F^{ext} \in C^1[K, \infty)$, $F^{ext \prime} \geq -1$, and $F^{ext \prime}$ is absolutely continuous. If

F^{ext} satisfies both the first-best action and agency action inequalities, then F^{ext} is an upper bound on the optimal value function.

More restrictive versions of the maximum principle have appeared in the previous literature. The main difference here is that the function F^{ext} is not assumed to be continuously *twice* differentiable, so it is not immediately apparent how to apply Ito’s lemma. However, there is a more general Ito’s Lemma for functions satisfying the hypotheses of this principle. See Theorem 22.5, Kallenberg (2001).

9.B Proof of Theorem 3.1 and The Domains of Optimality Theorem

The Domains of Optimality Theorem implies Theorem 3.1, so I will prove The Domains of Optimality Theorem. Also, Lemmas 6.4 and 7.1 pertaining to when the optimal contract takes on the baseline and static forms have already been proven in DeMarzo and Sannikov (2006). It remains to be shown that the two remaining unclaimed regions in Figure 10 in section 7.A are the domains of the Quiet-Life and Renegotiating Baseline contracts. For expositional simplicity, I assume $K = 0$ and the optimal baseline value function has an interior optimum.

The strategy of the proof is quite straightforward. When the agency action point is in the right domain, I begin by noting that if the principal was in an alternate universe with very negative taxes then the optimal taxed baseline contract would be the optimal contract. As I start shifting up the tax parameter, at some point the principal is going to be indifferent between sticking with the optimal taxed baseline contract and some other contract that uses agency action. The punchline is that the “other contract” is a Quiet-Life contract which is the optimal contract in the actual no tax universe. Similarly, when the agency action point is in the left domain I begin by noting that if the principal was in an alternate universe where his outside option was very high then the optimal baseline contract would be the optimal contract. As I start lowering the principal’s outside option parameter, at some point the principal is going to be indifferent between sticking with the optimal baseline contract and some other contract that uses agency action. The punchline is that the “other contract” is a Renegotiating Baseline contract which is the optimal contract in the actual universe.

Case 1: Quiet-Life Domain

Let the agency action point be in the right region.

For all sufficiently low (negative) taxes, $F_{\mathcal{T}}^{ext,B}$ satisfies the agency action inequality and is therefore an upper bound for the optimal value function. For example, Corollary 9.3 implies the existence of a unique (negative) tax $\bar{\mathcal{T}}$ such that:

$$F_{\bar{\mathcal{T}}}^{ext,B} \left(\frac{\phi A}{\gamma} \right) = 0$$

Since the agency action point is in the right region,

$$\frac{d^-}{dx} F^{\text{ext},S} \left(\frac{\phi A}{\gamma} \right) < 0$$

and since both $F_{\overline{\mathcal{T}}}^{\text{ext},B}$ and $F^{\text{ext},S}$ meet at the outside option point, so Corollary 9.1 implies that:

$$F_{\overline{\mathcal{T}}}^{\text{ext},B} \left(\frac{\phi A}{\gamma} \right) > \frac{\mu - A}{r}$$

This means that $F_{\overline{\mathcal{T}}}^{\text{ext},B}$ satisfies the agency action inequality.¹⁸ Now pick the least negative such tax \mathcal{T}^* . By continuity, there will be at least one point $(a, F_{\mathcal{T}^*}^{\text{ext},B}(a))$ that satisfies the agency action inequality with equality. This simply means:

$$g'_a(a) = F_{\mathcal{T}^*}^{\text{ext},B}{}'(a)$$

where g_a is the unique solution to the agency action ODE going through $(a, F_{\mathcal{T}^*}^{\text{ext},B}(a))$. Furthermore, a must be in the domain of $F_{\mathcal{T}^*}^B$. If not, then the pasting point $(a, F_{\mathcal{T}^*}^{\text{ext},B}(a))$ lies on the linear branch of $F_{\mathcal{T}^*}^{\text{ext},B}$ and $a < \frac{\phi A}{\gamma}$. But then for all \tilde{a} in a small left neighborhood of a , we have $g'_{\tilde{a}}(\tilde{a}) > F_{\mathcal{T}^*}^{\text{ext},B}{}'(\tilde{a})$ where $g_{\tilde{a}}$ is the unique solution to the agency action ODE going through $(\tilde{a}, F_{\mathcal{T}^*}^{\text{ext},B}(\tilde{a}))$. This then implies that $F_{\mathcal{T}^*}^{\text{ext},B}$ does not satisfy the agency action inequality at \tilde{a} . Contradiction.

By the least negativity property of \mathcal{T}^* , it must be that $\mathcal{T}^* > \overline{\mathcal{T}}$. Then Corollary 9.2 tells us that $F_{\mathcal{T}^*}^{\text{ext},B}$ must also be downward sloping to the right of $\frac{\phi A}{\gamma}$, which implies that $a < \frac{\phi A}{\gamma}$. Ito's Lemma and optional sampling imply $F^Q \equiv F_{\mathcal{T}^*}^{\text{ext},B}|_{[0,a]} = F_{\mathcal{T}^*}^B|_{[0,a]}$ is the value function of a Quiet-Life contract.

$F_{\mathcal{T}^*}^{\text{ext},B}$ satisfies the assumptions of the C^1 Maximum Principle, so it is an upper bound on the optimal value function. And since the F^Q portion is an actual value function and includes the maximum point of $F_{\mathcal{T}^*}^{\text{ext},B}$, F^Q must be the optimal value function. The optimal contract is then the Quiet-Life contract with good performance threshold $U^{\text{good},Q} = a$, the agent's continuation payoff is started at $U_0 = \arg \max F^Q$, and the payoff to the principal is $\max F^Q$.

Case 2: Renegotiating Baseline Domain

Let the agency action point be in the left domain.

Using arguments similar to before, for all sufficiently high alternate principal outside options l , $F_{(0,l)}^{\text{ext},B}$ (see Definition 7.1) satisfies the agency action inequality. Then the idea is the same as before: find the lowest alternate outside option l^* such that $F_{(0,l^*)}^{\text{ext},B}$ satisfies the agency action inequality. $l^* > L$ since by assumption $F^{\text{ext},B}$ does not satisfy the agency action inequality.

¹⁸Indeed, any concave function attaining its maximum at $\frac{\phi A}{\gamma}$ and with maximum greater than or equal to $\frac{\mu - A}{r}$ satisfies the agency action inequality.

Again like before, the minimality of l^* implies the existence of a pasting point $(a, F_{(0,l^*)}^{ext,B}(a))$. Mirroring the argument in Case 1, Corollary 9.4 implies $a > \frac{\phi A}{\gamma}$. Call the good performance threshold of $F_{(0,l^*)}^B$, $U^{good,R}$. Then define the function F^R to be $F_{(0,l^*)}^B|_{[a,U^{good,R}]}$. F^R is of course the value function of a Renegotiating Baseline contract.

The C^1 Maximum Principle implies that $F_{(0,l^*)}^{ext,B}$ is an upper bound of the optimal value function. The F^R portion contains the maximum point, so F^R is the optimal value function. The optimal contract is the Renegotiating Baseline contract with poor performance threshold $U^{poor,R} = a$, good performance threshold $U^{good,R}$, the agent's continuation payoff is started at $U_0 = \arg \max F^R$, and the payoff to the principal is $\max F^R$.

The Domains of Optimality Theorem when $K > 0$

Recall, the previous proof and section 7.A assumed that $K = 0$. How does the Domains of Optimality Theorem change when $K > 0$? There's not much change (see Figure 14):

When the agency action point is in the region $\{(X, Y) | X \geq K \text{ and } Y < \frac{\mu}{r}\}$ the boundaries are defined exactly like in the $K = 0$ case. Thus the only question is how to appropriately extend the boundaries to the region $O = \{(X, Y) | 0 < X < K \text{ and } Y < \frac{\mu}{r}\}$.

Recall in the $K = 0$ case there is a bottom curve that is a byproduct of Condition (4) Lemma 6.4. Since the value function $F^{ext,B}$ of the condition is independent of the agency action point, the bottom curve can be naturally extended using the same condition.

Now suppose the agency action point is in O , above the bottom curve. For any $l < \frac{\mu}{r}$, define g_l to be the unique agency action solution going through (K, l) and recall $F_{(K,l)}^{ext,B}$ (see Definition 7.1). The agency action point is above the efficiency threshold if and only if the optimal static payoff point $(K, \frac{\mu - A - (\gamma K - \phi A)}{r})$ is as well. If this is the case, then the C^1 Maximum Principle implies that $F_{(K, \frac{\mu - A - (\gamma K - \phi A)}{r})}^{ext,B}$ is an upper bound on the optimal value function and the optimal contract is the optimal static contract.

I now claim that if the agency action point is in O and between the bottom curve and the efficiency threshold, then the optimal contract is a Renegotiating Baseline contract. The method is almost identical to the $K = 0$ case. Because the agency action point is assumed to be below the efficiency threshold, $F_{(K, \frac{\mu - \gamma K}{r})}^{ext,B}$ satisfies the agency action inequality and is an upper bound on the optimal value function. Pick the minimal l^* such that $F_{(K,l^*)}^{ext,B}$ satisfies the agency action inequality. Then (K, l^*) is below the efficiency threshold which means that $F_{(K,l^*)}^{ext,B}$ (see Lemma 9.1) is nontrivial with a good performance threshold which I will call $U^{good,R}$. Furthermore, by the minimality of l^* there exists a pasting point $(a, F_{(K,l^*)}^{ext,B}(a))$. Then $F^R \equiv F_{(K,l^*)}^{ext,B}|_{[a,U^{good,R}]}$ is the optimal value function. The optimal contract is the Renegotiating Baseline contract with poor performance threshold $U^{poor,R} = a$, good performance threshold $U^{good,R}$, the agent's continuation payoff is started at $U_0 = \arg \max F^R$, and the




Figure 14: The Domains of Optimality Theorem when $K > 0$. The heavy bold lines are how the boundaries of the domains of optimality are extended to the region to the left of K .

payoff to the principal is $\max F^R$.

We can now extended the V-shaped curve: starting at the left most point of the old V-curve, go straight up until the efficiency threshold is hit, then travel up the efficiency threshold until reaching $(0, \frac{U}{r})$. See Figure 14.

The Domains of Optimality when $r = \gamma$

Recall that the good performance threshold U^{good} of a Quiet-Life contract is restricted to be $< \frac{\phi A}{\gamma}$. If however $U^{good} = \frac{\phi A}{\gamma}$ then $U_t = U^{good}$ is an absorbing event. I call such a contract a **tenure contract**. Similarly, the poor performance threshold U^{poor} of a Renegotiating Baseline contract is restricted to be $> \frac{\phi A}{\gamma}$. If however $U^{poor} = \frac{\phi A}{\gamma}$ then $U_t = U^{poor}$ is an absorbing event. I call such a contract an **inside option baseline contract**. The domains of optimality theorem when $r = \gamma$ is summarized in Figure 15.

Notice, in particular, if $K = 0$ then contracts that induce agency action non-permanently




Figure 15: The Domains of Optimality Theorem when $r = \gamma$.

(e.g. Quiet-Life and Renegotiating Baseline) are never optimal. This is a reflection of the intuition that saving-the-best-for-last is not useful. The reason why the Renegotiating Baseline contract is sometimes optimal when $K > 0$ is because the principal is required to keep the agent's continuation payoff $U_t \geq K$. This boundary condition prevents incentive-compatible contracts from using permanent agency action without salary.

The proof of this result is similar to the $r < \gamma$ case and is omitted.

9.C Smooth-Pasting

Definition 9.2. Two functions *paste* at a point (x, y) if both functions go through (x, y) and have the same derivative there.

Definition 9.3. A point (a, b) such that $a \neq \frac{\phi A}{\gamma}$ is called a **concave smooth pasting point** if the following condition is satisfied:

$$f''(a) = g''(a) < 0$$

where g is the unique solution to the agency action ODE and f is the unique solution to the first-best action ODE such that f and g paste at (a, b) .

The Concave Smooth Pasting Lemma. Consider the following **concave smooth-pasting function**

$$S(x) = \frac{\mu - A}{r} + \frac{\gamma}{r\phi} \frac{\phi}{2A} \frac{(x - \frac{\phi A}{\gamma})^2}{(1 - \frac{r}{\gamma}) - (x - \frac{\phi A}{\gamma})}$$

defined on the open interval $(-\infty, \frac{\phi}{2A}(1 - \frac{r}{\gamma}) + \frac{\phi A}{\gamma})$. A point (a, b) is a concave smooth pasting point or the agency action point if and only if $a \in (-\infty, \frac{\phi}{2A}(1 - \frac{r}{\gamma}) + \frac{\phi A}{\gamma})$ and $b = S(a)$.

S is a differentiable, convex function with unique interior minimum point equal to the agency action point $(\frac{\phi A}{\gamma}, \frac{\mu - A}{r})$.

Proof. Suppose a point (x, y) is a concave smooth pasting point. Let g, f be the associated functions. Then

$$g'(x) = \frac{r}{\gamma} \frac{y - \frac{\mu - A}{r}}{x - \frac{\phi A}{\gamma}}$$

$$g''(x) = \frac{r}{\gamma} \left(\frac{r}{\gamma} - 1 \right) \frac{y - \frac{\mu - A}{r}}{\left(x - \frac{\phi A}{\gamma} \right)^2}$$

$$f''(x) = \frac{ry - \mu - \gamma x f'(x)}{\phi^2/2} = \frac{ry - \mu - \gamma x \frac{r}{\gamma} \frac{y - \frac{\mu - A}{r}}{x - \frac{\phi A}{\gamma}}}{\phi^2/2}$$

Setting $f''(x) = g''(x) < 0$ and solving for y produces the function $S(x)$ and the associated domain. \square

Corollary 9.5. Let g_a denote the unique solution to the agency action ODE going through $(a, S(a))$ with $a < \frac{\phi A}{\gamma}$. Suppose g_a and some $F_{(X,Y)}^{ext,B}$ paste at some point $(x^*, g_a(x^*))$ with $x^* \in [a, \frac{\phi A}{\gamma})$. Then

$$F_{(X,Y)}^{ext,B} > g_a \text{ on } \left[a, \frac{\phi A}{\gamma} \right) - \{x^*\}$$

Proof. First assume $x^* \in (a, \frac{\phi A}{\gamma})$, continuity will take care of the $x^* = a$ case. Now note that in any pasting between some g_a and some $F_{(X,Y)}^{ext,B}$ satisfying the hypotheses of the the corollary, g_a is always strictly more concave at the pasting point. This is because S is convex, g_a is concave, and so $g(x^*) > S(x^*)$. This means the corollary holds locally.

Let $E : \{(x, y) \mid x < \frac{\phi A}{\gamma} - \epsilon\} \rightarrow \mathbb{R}^2$ be a (smooth) embedding that is both x -coordinate and orientation preserving (i.e. (x, y_1) is above $(x, y_2) \Rightarrow E((x, y_1))$ is above $E((x, y_2))$), and maps solutions to the agency action ODE into the horizontal lines. Assume ϵ is small enough so that $x^* \in (a, \frac{\phi A}{\gamma} - \epsilon)$

Now suppose that $F_{(X,Y)}^{ext,B}(\tilde{x}) = g_a(\tilde{x})$ for some $\tilde{x} \in [a, \frac{\phi A}{\gamma}) - \{x^*\}$. This implies that $E \circ$

$F_{(X,Y)}^{ext,B}(x^*) = E \circ F_{(X,Y)}^{ext,B}(\tilde{x})$. Furthermore, since the corollary holds locally, $(x^*, E \circ F_{(X,Y)}^{ext,B}(x^*))$ is a local minimum for $E \circ F_{(X,Y)}^{ext,B}$. That means there is a point \tilde{x}' in between x^* and \tilde{x} where $(\tilde{x}', E \circ F_{(X,Y)}^{ext,B}(\tilde{x}'))$ is a local maximum. The pullback of the horizontal line $y = E \circ F_{(X,Y)}^{ext,B}(\tilde{x}')$ is a solution \tilde{g} to the agency action ODE. The pullback of the local maximum condition says that $F_{(X,Y)}^{ext,B}$ and \tilde{g} paste at $(\tilde{x}', \tilde{g}(\tilde{x}'))$ but that \tilde{g} is *less* concave at the pasting point. But $g(\tilde{x}') > S(\tilde{x}')$. Contradiction. \square

Corollary 9.6. *Fix a point $(a, S(a))$ with $a \in [0, \frac{\phi A}{\gamma})$. Let g_a be the unique solution to the agency action ODE going through $(a, S(a))$. Suppose*

$$g'_a(a) \geq F_{(a,g_a(a))}^{ext,B}{}'(a)$$

Then there is unique point $(x^, g_a(x^*))$ such that $x^* \in [a, \frac{\phi A}{\gamma})$ and*

$$g'_a(x^*) = F_{(x^*,g_a(x^*))}^{ext,B}{}'(x^*)$$

Furthermore, $F_{(x^,g_a(x^*))}^{ext,B}$ satisfies the agency action inequality.*

Proof. Define the continuous function $\delta(x) = g'_a(x) - F_{(x,g_a(x))}^{ext,B}{}'(x)$. By assumption $\delta(a) \geq 0$ and clearly $\lim_{x \rightarrow \frac{\phi A}{\gamma}} \delta(x) = -\infty$. Existence of the pasting is implied.

Now suppose there were two distinct values x_1^*, x_2^* such that $x_i^* \in [a, \frac{\phi A}{\gamma})$ and

$$g'_a(x_i^*) = F_{(x_i^*,g_a(x_i^*))}^{ext,B}{}'(x_i^*) \quad i = 1, 2$$

By extending $F_{(x_i^*,g_a(x_i^*))}^{ext,B}$ leftwards through the first-best action ODE, one can find Y_1 and Y_2 such that

$$F_{(x_i^*,g_a(x_i^*))}^{ext,B} \subset F_{(0,Y_i)}^{ext,B}$$

Clearly, $F_{(0,Y_1)}^{ext,B}$ and $F_{(0,Y_2)}^{ext,B}$ are distinct, and therefore do not intersect. But both $F_{(0,Y_1)}^{ext,B}$ and $F_{(0,Y_2)}^{ext,B}$ paste with g_a and Corollary 9.5 implies that they both lie above g_a on $[a, \frac{\phi A}{\gamma})$. Contradiction.

Finally, recall the embedding function E of the previous corollary. We know $E \circ F_{(x^*,g_a(x^*))}^{ext,B}$ is locally increasing to the right of x^* . Indeed, Corollary 9.5 implies it cannot decrease on $[x^*, \frac{\phi A}{\gamma})$. The pullback of the nondecreasing condition says that $F_{(x^*,g_a(x^*))}^{ext,B}$ satisfies the agency action ODE on $[x^*, \frac{\phi A}{\gamma})$. That $F_{(x^*,g_a(x^*))}^{ext,B}$ satisfies the agency action inequality on $[\frac{\phi A}{\gamma}, \infty)$ comes from the more general fact that any decreasing concave function f defined on $[\frac{\phi A}{\gamma}, \infty)$ satisfies the agency action inequality if and only if $f(\frac{\phi A}{\gamma}) \geq \frac{\mu - A}{r}$ (which $F_{(x^*,g_a(x^*))}^{ext,B}$ satisfies by Corollary 9.5). \square

9.D Bargaining

Let F^{Opt} denote the optimal value function. Let $U^{contract} = \arg \max F^{Opt}$ and let U^{good} and U^{poor} be the good and poor performance thresholds of the optimal contract. The optimal value function gives us the optimal contracts delivering payoffs $x \in [U^{poor}, U^{good}]$ to the agent. I will now derive the extended optimal value function $F^{ext,Opt}$ which, in addition to governing optimal contracts delivering payoffs $x \in [U^{poor}, U^{good}]$ to the agent, also governs the optimal contracts delivering payoffs $x > U^{good}$ to the agent.

It is already known that when $F^{Opt} = F^B$ then $F^{ext,Opt} = F^{ext,B}$. Recall from 9.B F^R was defined as $F_{(K,l^*)}^{ext,B}|_{[a=U^{poor,R}, U^{good,R}]}$. Using similar reasoning one can easily show that if $F^{Opt} = F^R$ then $F^{ext,Opt} = F^{ext,R} \equiv F_{(K,l^*)}^{ext,B}|_{[U^{poor,R}, \infty)}$. Also if the optimal contract is the optimal static contract with payoff $(\frac{\phi A + s}{\gamma}, \frac{\mu - A - s}{r})$ where $s = \max\{0, \gamma K - \phi A\}$, then $F^{ext,Opt} = F_{(\frac{\phi A + s}{\gamma}, \frac{\mu - A - s}{r})}^{ext,B}$.

The only nontrivial case is when $F^{Opt} = F^Q$. Recall F^Q was defined as $F_{\mathcal{T}^*}^{ext,B}|_{[K, a=U^{good,Q}]}$. It was shown that

$$g'_a(a) = F_{\mathcal{T}^*}^{ext,B} '(a) < 0$$

where g_a is the unique solution to the agency action ODE going through $(a, F_{\mathcal{T}^*}^{ext,B}(a))$. In fact, $(a, F_{\mathcal{T}^*}^{ext,B}(a))$ is a concave smooth pasting point:

$$g''_a(a) = F_{\mathcal{T}^*}^{ext,B} ''(a)$$

This is because if

$$g''_a(a) > F_{\mathcal{T}^*}^{ext,B} ''(a) \text{ or } g''_a(a) < F_{\mathcal{T}^*}^{ext,B} ''(a)$$

then $F_{\mathcal{T}^*}^{ext,B}$ does not satisfy the agency action inequality in a neighborhood of a . For example, suppose $g''_a(a) < F_{\mathcal{T}^*}^{ext,B} ''(a)$, then pick a slightly higher solution \tilde{g} to the agency action ODE. Then \tilde{g} crosses $F_{\mathcal{T}^*}^{ext,B}$ twice in the neighborhood of a : once from the below and once from above. At the point where \tilde{g} crosses from above, the agency action inequality is not satisfied. This means that $g'_a(a) \geq F_{(a, g_a(a))}^{ext,B} '(a)$.

Now Corollary 9.6 implies the existence of a $F_{(x^*, g_a(x^*))}^{ext,B}$. $F^{ext,Q}$ is then $F_{\mathcal{T}^*}^{ext,B}|_{[K, a=U^{good,Q}]}$ plus $g_a|_{[a, x^*]}$ plus $F_{(x^*, g_a(x^*))}^{ext,B}$. This is because not only is every point on the function a payoff point of a contract but also because the function satisfies the assumptions of the C^1 Maximum Principle.

References

- [1] Bennardo, A., Chiappori, P.-A., and Song, J., (2010) "Perks as Second Best Optimal Compensations," working paper
- [2] Bertrand, M. and Mullainathan, S. (2003) "Enjoying the Quiet Life? Corporate Governance and Managerial Preferences," *Journal of Political Economy* vol. 111, pp. 1043 - 1075

- [3] Biais, B., Mariotti, T., Plantin, G., and Rochet, J.-C. (2007) “Dynamic Security Design: Convergence to Continuous-Time and Asset Pricing Implications,” *Review of Economic Studies* Vol. 74, pp. 345 - 390
- [4] DeMarzo, P. M. and Fishman, M. J. (2007) “Optimal Long-Term Financial Contracting,” *The Review of Financial Studies* Vol. 20, pp. 2079 - 2128
- [5] DeMarzo, P. M. and Sannikov, Y. (2006) Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model,” *Journal of Finance* Vol. 61, pp. 2681 - 2724
- [6] Durrett, S. *Probability: Theory and Examples*. Duxbury Press, Belmont, CA, 2008
- [7] Edmans, A. and Gabaix, X., (2009) “Tractability in Incentive Contracting,” working paper
- [8] Edmans, A., Gabaix, X., and Landier, A. (2009) “A Multiplicative Model of Optimal CEO Incentives in Market Equilibrium,” *Review of Financial Studies* Vol. 22, pp. 4881 - 4917
- [9] Fama, E., (1980) “Agency Problems and the Theory of the Firm,” *Journal of Political Economy* Vol. 88, pp. 288 - 307
- [10] Grossman, S. J. and Hart, O. D. (1983) “An Analysis of the Principal-Agent Problem,” *Econometrica* Vol. 51, pp. 7 - 45
- [11] Harrison, M. and Lemoine, A. (1981) “Sticky Brownian Motion as the Limit of Storage Processes,” *Journal of Applied Probability* Vol. 18, pp. 216 - 226
- [12] Holmstrom, B. and Milgrom, P., (1987) “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica* Vol. 55, pp. 303 - 328
- [13] Jensen, M. and Meckling, W., (1976) “Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure,” *Journal of Financial Economics* Vol. 3, 305 - 360
- [14] Kallenberg, O. *Foundations of Modern Probability*. Springer-Verlag, New York, 2001
- [15] Karatzas, I. and Shreve, S. *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York, 1991
- [16] Kuhnen, C. and Zwiebel, J., (2009) “Executive Pay, Hidden Compensation, and Managerial Entrenchment,” working paper
- [17] Rajan, R. and Wulf, J. (2006) “Are Perks Purely Managerial Excess?” *Journal of Financial Economics* Vol. 79, pp. 1 - 33
- [18] Sannikov, Y. (2008) “A Continuous-Time Version of the Principal-Agent Problem,” *Review of Economic Studies* Vol. 75, pp. 957 - 984

- [19] Schwab, S. and Thomas, R. "An Empirical Analysis of CEO Employment Contracts," *Washington and Lee Law Review* Vol. 63, pp. 231 - 270
- [20] Yermack, D. (2006) "Flights of Fancy: Corporate Jets, CEO perquisites, and Inferior Shareholder Returns," *Journal of Financial Economics* Vol. 80, pp. 211 - 242