Building Trust in Relational Contracting

Melis Kartal*

Department of Economics, New York University
19 West 4th Street, New York, NY 10012
E-mail: mk2672@nyu.edu

Abstract

I study a relational contracting model, in which the agent's discount factor is fixed and known, whereas the discount factor of the principal is her private information. This represents the situation where the agent does not fully know the motives and the preferences of the principal in the relationship. The principal's discount factor can be one of two types: high or low. The high discount factor is associated with more cooperative behavior whereas a low discount factor typically results in short-sighted and opportunistic behavior. Therefore, the discount factor is a proxy for measuring the trustworthiness and the commitment of the principal. I focus on separating equilibria as well as pooling equilibria. I show that a separating equilibrium with signaling only in the initial period is impossible with some parameter values. Nevertheless, there always exists a separating equilibrium, and it sometimes requires dynamic costly signaling for an extended period of time (i.e., costly signaling may have to continue even after types separate in the initial period). I characterize the optimal separating contract and show that the bonus payment for high performance, the agent effort and the surplus in the relationship all increase gradually in the optimal contract of the high type. Hence, the separating contract is characterized by gradualism, which is a phenomenon that is observed in many real-world economic relationships. Finally I show that the optimal separating contract generates higher surplus than the optimal pooling contract regardless of the prior, which is different than the outcome in the standard signaling models. *Journal of Economic Literature* Classification Numbers: C73, D82, L14

*I am truly indebted to Alessandro Lizzeri for his guidance and encouragement. I would like to thank Guillaume Frechette, John Leahy, Debraj Ray, Andrew Schotter, Ennio Stacchetti, Matan Tsur, Kemal Yildiz and especially Ariel Rubinstein for their suggestions and comments.

1 Introduction

This paper studies a relational contracting model with hidden information. In particular, it is assumed that the discount factor of the principal is her private information. This represents the situation where the agent does not fully know the motives and the preferences of the principal at the beginning, which is a problem in relational contracting. I analyze how trust can be built in such an environment.

In the standard principal-agent model, the contract between the two parties is a legally-enforceable agreement that ties compensation to verifiable performance measures. Yet, performance measures may conceivably be subjective and nonverifiable in numerous instances. For instance, the agent's actions may be observable to the principal but not to the court. Similarly, arrangements for recording output may be informal, which makes output difficult to verify in the court. Even if the principal has only nonverifiable information regarding the agent's performance, a cooperative relationship can still be sustained in a repeated setting, as is well known. This is in line with what we observe: Economic interactions involve significant amount of trust and many *informal* contracts are honored voluntarily. Even though there is no legal sanction against the breach of an informal contract, a mutually beneficial trade can still be enforced if the relationship is valuable for both parties.

The literature on relational contracts has usually relied on the following critical assumption: Each player has complete information about the other's preferences. In reality, trading parties may well be uncertain about each other's motives at the beginning of the relationship. In particular, the agent may not know how "trustworthy" or "committed" the principal is. This is a problem in relational contracting since the enforceability of an informal contract depends precisely on this information. A dishonest principal would like to be thought of as trustworthy so that she can exploit the agent. Therefore, a trustworthy principal needs to build up a reputation. Otherwise the uncertainty that the uninformed agent has to bear reduces the surplus from trade and the relationship may even break down. There is indeed asymmetric information about the trustworthiness of firms in the business world. This is why Forbes has been announcing "America's most trustworthy companies" every year. There is also a stark difference among the firms regarding the value they give to their workforce.[1] This has become especially prominent in the aftermath of crisis periods such as the recent financial crisis and the September 11 attacks. Southwest Airlines famously declined to quit its no-layoff policy and kept the employee morale high even in difficult times whereas numerous corporations from different sectors have chosen to renege on their no-layoff promises in order to cut costs.

It is well known from the theory of repeated games that a high discount factor is associated with more cooperative behavior, whereas a low discount factor typically results in myopia and opportunistic behavior. In particular, the higher the principal's discount factor the easier it is to keep the promises in relational contracts because it makes the relationship more valuable. Therefore the discount factor is a very reasonable proxy for trustworthiness and commitment. I assume hat the principal's discount factor is her private information whereas the agent's discount factor is fixed and known.[2] The principal can be one of two types: "high" type, meaning that the principal is patient, and "low" type. Using the discount factor as a proxy, I address the issues arising from the uncertainty of the agent regarding the true motives of the principal. The principal learns her type at the beginning of the game. At the beginning of each period, the principal makes a compensation offer to the agent, which the agent either accepts or rejects. The offer consists of an enforceable fixed wage w, and performance-contingent bonus transfer b, which is not enforceable. If the offer is accepted, the agent chooses an effort level, which cannot be observed by the principal. Exerting effort is costly and it generates a stochastic output. Output is observed by both the principal and the agent, but it cannot be verified by a third party. At the end of each period, the party responsible for making the bonus transfer decides whether or not to honor it. If the agent rejects the principal's offer in a period, both parties receive their outside options in that period.

First, I present the symmetric information benchmark. It is different from the previous literature to some extent, because the principal and the agent differ in their discount factors (for at least one principal type). In Levin (2003a), the surplus generated by the optimal contract can be divided in many different ways since both parties are risk-neutral

[1]CNNMoney has recently reported that more than 7,000 lawsuits were filed against employers in 2011 in the U.S. alleging wage and hour violations (May 29, 2012).

[2]Halac (2012) is, to my knowledge, the only other paper that analyzes relational contracting with persistent private information about the principal. I explain the distinction between my work and Halac in subsequent sections.

and have the same discount factor. On the other hand, the sharing of surplus is pinned down uniquely in the optimal contract if the discount factor of the agent is different from that of the principal. So, the indeterminacy of the optimal contract in Levin stems from the knife-edge case where the two parties have identical discount factors.

Introducing hidden information to the benchmark setting further exacerbates the problems associated with the nonverifiability of information. Since a bonus transfer is not legally enforceable the commitment of the principal depends on the value of the future trade. The high type values future trade more, and the agent can trust her more to deliver on her promises. Therefore the low type would be better off being thought of as the high type, and a credible signaling mechanism is required for separation to take place. I characterize the perfect public Bayesian Nash equilibrium of the game. I focus on both pooling equilibria and separating equilibria, in which the two principal types separate through their contract choices. I characterize the separating equilibrium with a property that allows for dynamic costly signaling even after types fully separate.[3] I first show that it is not generally possible to obtain immediate separation, where only initial period choices are distorted. Next, I show that there *always* exists a separating equilibrium: Such an equilibrium sometimes requires dynamic costly signaling for an extended period of time.

I analyze the properties of the optimal separating contract. In equilibrium, the low type principal immediately offers her optimal symmetric information contract and the game continues as in the symmetric information setting. On the other hand, the high type principal engages in dynamic costly signaling that evolves until reaching the symmetric information benchmark in finite time. Along the high type's equilibrium path, the effort exerted by the agent and the surplus in the relationship increase gradually.[4] This is very intuitive because time has different effects on the prospective payoffs of the two types. Being more patient, the high type must forgo earlier profits for the benefits of later and higher cooperation, thereby making imitation less tempting for the low type. This signaling dynamic provides a theoretical foundation for the phenomenon of "grad-

[3]In the next section, I explain this property and discuss the literature that has applied it.

[4]One caveat is that I could only show that this is optimal among all time-dependent contracts, and there may exist superior contracts that condition in a richer way on prior history. On the other hand, I could show that this is the (unrestricted) optimal solution if agent effort is observable. I will discuss these in more detail below.

ualism", which is observed in many real-world economic relationships (see, for example, the discussion about informal credit markets in Ghosh and Ray (1996)). A number of studies have obtained similar results regarding "gradual building of trust" under different assumptions and in different settings. However, the equilibrium mechanism that gives rise to the gradual building of trust in my paper is novel.[5]

One interesting testable implication of the model is as follows: In the equilibrium contract for the high type, the pay-for-performance sensitivity and the surplus increase jointly as the high type establishes her reputation and the relationship matures. This has some empirical support as it is consistent with the findings of a number of studies investigating the relationship between managerial pay-for-performance sensitivity and firm performance.[6] These papers argue that high performance results from high pay-for-performance sensitivity but as my model shows, firm reputation is an essential (but an overseen) component of the argument. With asymmetric information, high pay-for-performance sensitivity is not simply a choice but also an outcome, itself. Good reputation (which is inherently linked to good prospects for a firm) makes high pay-for-performance sensitivity possible, which in turn increases profits (and justifies good prospects).

Finally, I characterize the optimal pooling contract. The optimal pooling contract is the low type's optimal symmetric information contract. In my model, any separating equilibrium generates higher surplus than the optimal pooling equilibrium regardless of the prior beliefs. This is different from the results of the standard signaling models. For instance, in the classical job-market signaling model due to Spence (1974), whether pooling is better or worse than separation depends on the prior beliefs.

The paper is organized in the following way. The next section discusses the related literature. Section 3 introduces the model and presents the analysis for the symmetric information model. Section 4 shows the existence of a separating contract and shows

[5] On the one hand, all the learning takes place in the initial period. On the other hand, the costly signaling is spread across periods in a monotone fashion.

[6] Abowd (1990) finds a positive relation between managerial pay for performance sensitivity and firm performance in his 1981–86 sample. Mehran (1995) finds that firm value increases in the fraction of non-cash (i.e., incentive) compensation in his 1979-80 sample. Mishra, McConaughy and Gobeli (2000) find a positive but diminishing relationship between pay performance sensitivity and firm performance in their 1974–1988 sample.

that it requires repeated costly signaling with some parameter values. Also, the optimal contract is characterized within a large class. Final section concludes. Proofs that are omitted in the main text are relegated to the Appendix.

2 Related Literature

This paper combines dynamic signaling with relational contracting. Below, I discuss the two literatures separately.

2.1 Dynamic Signaling

There is extensive work on dynamic signalling, and my work is closest to a subset of the literature in which out-of-equilibrium degenerate beliefs are allowed to change. This property allows costly signaling even after beliefs have become degenerate. It may be argued that once the agent attaches probability 1 to the principal being a certain type the subsequent game should essentially be one of symmetric information. Yet, this requires a restriction on beliefs as it is not implied by the definition of perfect Bayesian equilibrium. Moreover, such a restriction rules out intuitive equilibria in which the agent *tentatively* assigns a probability of 1 to the principal being a certain type and expects the principal to behave in a certain way so that his belief is sustained during the course of the game.[7] If the agent's expectations are based on the equilibrium strategy of the principal, then they are indeed fulfilled on the equilibrium path. However, off the equilibrium path beliefs of the agent may switch away from a degenerate distribution if costly signaling stops too early, as this indicates deception.

The concept of dynamic costly signaling after beliefs have become degenerate has been applied in numerous models such as Admati and Perry (1987), Noldeke and Van

[7]The continuous time job market signaling example presented in Noldeke and Van Damme (1990) illustrates very clearly why it is very plausible to have costly signaling even after beliefs have become degenerate. In the basic job market signaling model by Spence, the game is static and the more productive type separates from the less productive type simply by investing in education. Even though there is a multiplicity of PBE, only the Pareto-best separating equilibrium (the so-called Riley outcome) survives the Intuitive Criterion, as shown by Cho and Kreps (1987). Could the Riley outcome survive in a dynamic continuous setting? If it is to survive, that surely requires costly signaling after beliefs have become degenerate.

Damme (1990), Cramton (1992), and Kaya (2009). Admati and Perry and Cramton analyze a bargaining game whereas Noldeke and Van Damme extend Spence's job market signaling paper to a dynamic context. In these papers (except Kaya), the range of the signaling variable is too small, and hence, distortion of behavior in a single period is not sufficient to achieve separation. Kaya (2009) extends the concept of dynamic signaling to a class of repeated games where separation is achievable even if costly signaling lasts only one period. The multiplicity of separating equilibria results in a problem of equilibrium selection. Kaya characterizes the "least cost separating equilibrium" and shows that depending on the payoff functions of different types, the least cost separating equilibrium sometimes involves costly signaling for multiple periods. She argues that this is the only separating equilibrium that would survive the dynamic version of Intuitive Criterion. My model differs from the aforementioned papers in three aspects. First of all, the signal space is essentially endogenous (i.e., it depends on the surplus generated by the relationship). Moreover, it is unknown whether separation is achievable if behavior is distorted in only one period, unlike the case in Kaya. Finally, these papers do not display a dynamic buildup of trust and cooperation.

My results are also connected to (but qualitatively distinct from) signaling models such as Ghosh and Ray (1996), and Watson (1999, 2002). In these papers, the two matched players choose the level of cooperation in every period and decide individually whether or not to behave opportunistically. There is two-sided asymmetric information: "High" types prefer to cooperate as long as their partners also cooperate whereas "low" types have an incentive to take advantage of the other player's trust. There is an initial testing phase in which partners of high type build trust through their actions, and the stakes in the relationship increase gradually as players trust each other more. In Ghosh and Ray, testing phase lasts one period, and if at least one of the partners reveal himself to be low type then the match dissolves and players match with other partners from the pool of potential matches. In Watson, testing phase is longer, and information revelation is gradual -i.e., the low type is indifferent between defecting and cooperating and uses a mixed strategy; hence learning takes place gradually through default. In my paper, the dynamic buildup of the relationship is similar to that in Watson (1999, 2002). However, my model is distinct from the "starting small" literature from other aspects. I consider a principal-agent setting with relational contracts. Also, unlike the gradual learning in

Watson, the agent learns the true type of the principal with probability one by the end of the initial period although costly signaling sometimes takes much longer than one period. Hence, there is no betrayal on the equilibrium path.

Finally, Rubinstein (1985) analyzes a bargaining problem where one of the parties is privately informed about his discount factor. As is typical in most signaling games, there are many equilibria of the bargaining game with incomplete information, but unreasonable equilibria are eliminated by making additional requirements on beliefs. In the unique equilibrium that survives these requirements, the bargaining either ends in the first period or continues for a while, depending on the prior beliefs of the uninformed party.

2.2 Relational Contracting

If the performance measure is nonverifiable then a relational contract is used to sustain trade between a principal and an agent. Numerous relational contracting models focused on environments with symmetric information (see, for example, MacLeod and Malcolmson (1989), Levin (2003b)).[8] Asymmetric information has also been incorporated to relational contracting. Shapiro and Stiglitz (1984), and Baker, Gibbons and Murphy (2002) consider relational contracts with moral hazard, whereas Levin (2003a) analyzes two distinct scenarios. He assumes that either the agent effort or the agent's time-specific cost parameter cannot be observed by the principal. Finally, MacLeod (2003) and Fuchs (2007) consider asymmetric information about output realization. In these papers, asymmetric information has no persistence.

Halac (2012) is closely related to this paper. In her model, the outside option of the principal is her private information. The outside option of the principal is either high or low, and it remains fixed throughout the game. Both types of asymmetric information (about the outside option and the discount factor) represent the case in which the agent is uncertain about the commitment of the principal in the relationship. However, the two papers are distinct because my modeling assumptions give rise to different economic trade-offs than those in Halac. Consequently, I obtain different equilibrium outcomes.

[8] There are also dynamic models that analyze the joint use of relational incentives and explicit incentives in the optimal contract, such as Baker, Gibbons and Murphy (1994), and Pearce and Stacchetti (1998).

In Halac, default by the high type (the type with a higher outside option) is a necessary condition for separation, which is not the case in my model.[9] In such an equilibrium, the high type is indifferent between defaulting and imitating the low type and uses a mixed strategy; hence learning takes place gradually through default. In contrast, I find that the agent learns the type of the principal immediately through the contract choice; still, the high type principal has to engage in dynamic costly signaling and behavior is distorted for an extended period with some parameter values. Finally, in Halac the existence of a separating equilibrium depends on the prior belief that the principal is low type. On the other hand, all my results are independent of the prior belief.

3 The Model

Two risk-neutral parties, a principal (she) and a single agent (he), interact repeatedly in periods $t = 0, 1, \ldots$. The agent's discount factor is δ, which is fixed and known, whereas the principal's discount factor is δ_θ, where $\theta \in \{l, h\}$ is the principal's private information and $\delta_l < \delta_h$. The principal learns her type at the beginning of the initial period and this type remains the same in all subsequent periods.

At the beginning of period t, the principal makes a contract offer to the agent. The agent can either accept or reject this offer, where $d_t \in \{0, 1\}$ denotes the agent's decision. If the offer is accepted, the agent chooses effort $e_t \in [0, 1]$ incurring a cost $c(e_t)$, where $c'(\cdot) > 0$, $c''(\cdot) > 0$ and $c'(\bar{e}) = \infty$. The principal cannot observe the effort choice of the agent. The agent's effort generates stochastic output y_t, where $y_t \in \{L, H\}$ and L (low) $< H$ (high). The probability that $y_t = H$ given effort e_t is e_t.[10] [11]Output is observed by both the principal and the agent, but it is not verifiable by a third party. The contract offer at time t consists of a fixed wage w_t and a bonus transfer b_t that is contingent on y_t. The fixed wage is enforceable, but the bonus payment is not. At the end of period t, parties decide whether or not to honor the bonus payment. If $b_t > 0$, the decision belongs to the principal whereas the agent makes the decision if $b_t < 0$. Total payment

[9]Halac assumes that degenerate beliefs cannot change. However, her result would be the same even if degenerate beliefs were allowed to change.

[10]Instead, it can be assumed that the probability of H is given by a function $g : e \to [0, 1]$. This doesn't change the results as long as $g(e)$ is weakly concave.

[11]Due to these assumptions, first order approach is valid.

from the principal to the agent is denoted by P_t, where $P_t = b_t + w_t$ if the promised payment is honored, and $P_t = w_t$ if not. Thus, the agent's per-period payoff is $P_t - c(e_t)$, and the principal's is $y_t - P_t$. The term $s(e) \equiv E_y[y - c|e]$ denotes the expected joint surplus.

If the agent rejects the principal's offer, both the principal and the agent receive their outside options in the current period, $\bar{\pi}$ for the principal and $\bar{u}$ for the agent. I assume that there exists an incentive compatible effort level e such that $s(e) > \bar{\pi} + \bar{u} \geq s(0)$. Thus, the trade between the two parties is more desirable than the outside options so long as the agent can be motivated to work.

The parties care about their discounted payoff stream. As of period t, the respective payoffs for the principal (of type θ) and the agent can be written as

$$\pi_{\theta t} = \mathbb{E}\sum_{\tau=t}^{\infty}\delta_{\theta}^{\tau-t}[d_\tau(y_\tau - P_\tau) + (1 - d_\tau)\bar{\pi}],$$

$$u_t = \mathbb{E}\sum_{\tau=t}^{\infty}\delta^{\tau-t}[d_\tau(P_\tau - c(e_\tau)) + (1 - d_\tau)\bar{u}],$$

Finally, the expected surplus in period t is given by $s_{\theta t} = \pi_{\theta t} + u_t$.

3.1 Equilibrium Concept

Let the term $h_t = (w_t, d_t, y_t, P_t)$ denote the public outcome at time t, and let $h^t = (h_0, ..., h_t)$ denote the public history up to time t with $\mathcal{H}^t$ representing the set of all possible date t histories. A relational contract describes a complete plan for the relationship (Levin (2003a)). For each principal type θ, date t and history $h^t \in \mathcal{H}^t$ a relational contract must specify (i) the contract the principal offers; (ii) whether the agent accepts or rejects the offer; (iii) the effort choice in case the agent accepts the contract offer; (iv) the bonus payment decision given the output realization; and (v) the conditional beliefs for the agent. Such a contract is self-enforcing if it describes a perfect public Bayesian Nash equilibrium of the repeated game. A PPBE is a set of public strategies and posterior beliefs such that the strategies form a Bayesian Nash equilibrium in every continuation game given the posterior beliefs, and the beliefs are updated according to Bayes' rule whenever possible. A public strategy depends only on $\mathcal{H}^t$ and the player's (payoff-relevant) private information. Put differently, the agent conditions his strategy on $\mathcal{H}^t$ whereas the principal conditions her strategy on $\theta \times \mathcal{H}^t$.

First, I focus on information revelation and I characterize PPBE with a property that allows for costly signaling even after types fully separate. There are many contracts that can separate the two types. But I focus on the optimal separating contract (within a wide class of contracts); this can be thought of as an informal equilibrium selection procedure. I define the optimal contract as the contract that maximizes the profit of the principal (and is pareto-optimal in the usual sense). Finally, I analyze the optimal pooling contract and I compare the two types of contracts.

3.2 Benchmark Model: Symmetric Information Setting

The analysis for the symmetric information setting consists of two cases: $\delta < \delta_\theta$ and $\delta > \delta_\theta$.[12] If $\delta < \delta_\theta$ then the principal chooses a stationary contract to maximize the surplus in the symmetric information setup -i.e., given $\theta \in \{l, h\}$, $e_t = e(\theta)$, $b_t = b(y, \theta)$, and $w_t = w(\theta)$ for all t. Moreover, the optimal stationary contract is unique and there exists no optimal contract that is nonstationary on the equilibrium path if $\delta < \delta_\theta$. This is in contrast with Levin (2003a). In Levin, there are many nonstationary contracts that maximize the profit of the principal. The indeterminacy of the optimal contract stems from the knife-edge case in which the two parties have the same discount factor. Stationarity of the optimal contract follows from the following Lemma.

Lemma 1 *In the optimal symmetric information contract with $\delta < \delta_\theta$, $u(h^t) = \bar{u}$ must hold for any h^t with $t > 0$.*

By Lemma 1, the agent must receive an expected utility of $\bar{u}$ in every period, which implies that the difference in continuation payoffs is not used to discipline the agent. The principal captures all the surplus $s(\theta)$ (in excess of $\bar{u}$) in every period, and therefore the agent never makes a bonus transfer. The principal offers the same contract to the agent every period: She pays a fixed wage $w(\theta)$ and bonus $b(\theta) > 0$ contingent on high output, and no bonus transfer is made if the output realization is low. Then, the respective payoffs for the principal of type θ and the agent are given by

[12] The case with equal discount factors is analyzed in Levin (2003a). He shows that stationary contracts suffice for optimality.

$$\pi_\theta = \frac{y - w(\theta) - e(\theta)b(\theta)}{1-\delta_\theta},$$

and

$$u = \frac{w(\theta) + e(\theta)b(\theta) - c(e(\theta))}{1-\delta}.$$

In a self-enforcing contract, the principal does not default on a promised payment. Default is punished by ending the relationship, the worst possible punishment. This type of punishment generates the highest possible surplus. Moreover, as noted by Levin, it is renegotiation-proof. The dynamic enforcement constraint has the following form:

$$\frac{\delta_\theta}{1-\delta_\theta}(s(\theta) - \bar{u} - \bar{\pi}) \geq b(\theta).$$

The optimal symmetric-information contract maximizes the expected surplus subject to the incentive compatibility constraint for the agent's effort choice and the enforcement constraint for the principal. For $\theta \in \{l, h\}$,

$$max_{e(),b()} \quad s = \mathbb{E}_y[y - c|e(\theta)]$$

$$\text{subject to} \qquad e(\theta) \in \arg\max_e \; w + eb - c(e),$$

$$\frac{\delta_\theta}{1-\delta_\theta}(s - \bar{u} - \bar{\pi}) \geq b.$$

The terms e_θ and b_θ give the solution to the maximization problem and the maximized surplus is denoted by s_θ. To avoid trivialities, I assume that for both principal types, parameters are such that trade is feasible but the enforcement constraint binds. Following,

$$b_\theta = \frac{\delta_\theta}{1-\delta_\theta}(s_\theta - \bar{u} - \bar{\pi}),$$

whereas the fixed payment is given by $w_\theta = \bar{u} - \mathbb{E}_y[b_\theta - c|e(\theta)]$. The contract $C_\theta = \{w_\theta, b_\theta\}$ is called the symmetric-information contract of type θ.

If $\delta > \delta_\theta$, then the optimal contract is unique and nonstationary. However, it becomes stationary after the initial period. The principal receives a frontloaded payment at $t = 0$ through fixed wage, whereas the agent gets all the surplus in excess of $\bar{\pi}$ at all $t \geq 1$:

This generates the highest surplus due to the fact that the agent has a higher discount factor.[13] So the optimal symmetric-information contract maximizes expected surplus subject to the incentive compatibility constraint as well as the enforcement constraint for the agent. Hence, the following program is solved.

$$max_{e(),b()} \quad s = \mathbb{E}_y[y - c|e]$$

$$\text{subject to} \qquad e \in \arg\max_e \; w + (1-e)b - c(e),$$

$$\frac{\delta}{1-\delta}(s - \bar{u} - \bar{\pi}) \geq -b.$$

Notice that $b < 0$. The terms e^* and b^* give the solution to the maximization problem and the maximized surplus is denoted by s^*. Assuming that the enforcement constraint binds as in the previous case,

$$b^* = \frac{\delta}{1-\delta}(s^* - \bar{u} - \bar{\pi}),$$

and $w_1^* = \mathbb{E}_y[y - b|e] - \frac{s - \bar{u} - \delta\bar{\pi}}{1-\delta}$, whereas $w_t^* = \mathbb{E}_y[y - b|e] - \bar{\pi}$ for $t \geq 2$.

4 Results

4.1 Information Revelation and the Separating Equilibrium

The optimal symmetric information contract with $\delta > \delta_\theta$ is independent of the principal's discount factor, as the previous section has shown. Therefore, hidden information about the principal's discount factor does not create a problem if $\delta > \delta_h$, and I do not pursue this case further. If $\delta_h > \delta > \delta_l$, then all the results that are presented below are still valid, and I focus on the case $\delta < \delta_l$ throughout, without loss of generality.

The previous section shows that the relationship between a high type principal and an agent can generate a surplus of s_h (with a contract offer of C_h) provided that the principal's discount factor is common knowledge. So, a contract offer different from C_h

[13] Notice that in an optimal contract with unequal discount factors, the bonus is paid out only by the party with the higher discount factor as this generates the strongest incentives and results in higher effort.

is interpreted as costly signaling undertaken by the high type. I focus on pure strategy separating equilibria where the two principal types separate through their contract choices. I analyze separation through the contract choice because default is wasteful (as the agent anticipates default with some probability and reduces his effort), and consequently, separation through default is infeasible for a nontrivial set of parameters. In this section, I show that a pure-strategy separating equilibrium always exists and that it sometimes requires costly signaling even after types separate in the initial period.

Symmetric information analysis makes it transparent that it is always the low type who would like to imitate. Since $b_l < b_h$ the low type is better off being thought of as the high type, because imitation gives her a payoff that is strictly greater than $\frac{s_l - \bar{u}}{1-\delta_l}$, the symmetric information payoff for the low type. It follows that in a separating equilibrium, the high type must use a credible signaling mechanism whereas the low type offers her optimal symmetric information contract C_l. I assume that a default by the principal results in termination of the relationship, just as in the symmetric information setting. This assumption is without loss: For any separating contract that assigns a positive probability to the continuation of the relationship after default there exists an equivalent separating contract that terminates the relationship with probability one after default and gives the same expected payoffs.

It is instructive to present the incentive constraints in a hypothetical separating equilibrium. Let the sequence $\{C(h^t)\}$ be the equilibrium path of contracts offered by the high type in a separating equilibrium such that $\{C(h^t)\} = \{w(h^t), b(h^t)\}$. For expositional simplicity, the contract offer of the high type is assumed to depend on only the current period output so that $\{C(h^t)\} = \{C_t\}_{t=0}^{\infty}$ and the principal pays a bonus $b_t > 0$ if and only if $y_t = H$. For $\{C_t\}_{t=0}^{\infty}$ to be the equilibrium path of contracts for the high type principal, the low type should be deterred from imitating:

$$\mathbb{E}_y[y - P_0|e_0] + \max\{\delta_l V_{l,1}(1), e_0(b_0 + \frac{\delta_l}{1-\delta_l}\bar{\pi}) + (1-e_0)\delta_l V_{l,1}(1)\}. \leq \frac{(s_l - \bar{u})}{1-\delta_l},$$

where $V_{l,1}(1)$ is the expected payoff for the low type principal at $t = 1$ having imitated at $t = 0$.[14] If $y_0 = H$, the low type can choose one of two actions: She can either default

[14]Here $P_t = w_t + b_t$ if $y_t = H$.

or honor the bonus payment (so that she defaults at a later period). The decision depends on the benefit of the low type from each action. The term $V_{l,t}(1)$ can be defined recursively as

$$V_{l,t}(1) = \mathbb{E}_y[y - P_t|e_t] + \max\{\delta_l V_{l,t+1}(1), e_t(b_t + \frac{\delta_l}{1-\delta_l}\bar{\pi}) + (1-e_t)\delta_l V_{l,t+1}(1)\}.$$

On the other hand, the high type must be willing to engage in costly signaling at every period given $\{C_t\}_{t=0}^{\infty}$. I assume that an unexpected offer provides sufficient evidence to believe that the principal is certainly low type but it is not enough to terminate the relationship. Following, the high type must be willing to offer C_t at t rather than deviating and being taken as low type with probability one from period t onwards. Therefore, for all t

$$\sum_{k=t}^{\infty} \delta_h^{k-t} \mathbb{E}_y[y - P_k|e_k] \geq \frac{(s_l - \bar{u})}{1-\delta_h}$$

must hold.

First, I show that with some values for the parameters δ_h and δ_l, separation is impossible if behavior is distorted in only the initial period.

Proposition 2 *There exists an $\varepsilon > 0$ such that if $\delta_h - \delta_l < \varepsilon$, and $\delta_h > \delta_l > \delta$, then there exists no pure-strategy separating equilibrium where behavior is distorted in only the initial period.*

Proof. Suppose that both types separate by the end of the initial period, and that the continuation play following separation is the symmetric-information equilibrium. Let C_1 be the high type's contract and C_2 be the low type's contract at $t = 0$.[15] Also let e_1 and e_2 be the respective effort levels that the agent optimally chooses given C_1 and C_2 and his beliefs. I assume that the agent accepts both offers in the initial period, but this assumption is without loss. The same result is obtained if the agent rejects one of the offers or both. If the high type deviates to mimicking the low type, the bonus promise that she can credibly make from $t = 1$ is at most b_l. The high type does not renege on the bonus payment and gets an expected payoff of $s_l - \bar{u}$ in each period. Therefore, the

[15]We can simply assume $C_2 = C_l$. This is inessential for the results.

continuation payoff of the high type principal who mimicks the low type is

$$\frac{s_l - \bar{u}}{1 - \delta_h}.$$

If, on the other hand, the low type mimics the high type, then the low type reneges on the bonus payment either at $t = 0$ or in the first period after $t = 0$ when high output is realized. The term $V_l(1)$ is the expected payoff of the low type principal from $t = 1$ onwards given that there has been no default at $t = 0$ and the agent believes that the principal is high type with probability one with no additional costly signaling required after the initial period. In a separating equilibrium, the following inequality must be satisfied so that the low type is deterred from mimicking the high type:

$$\mathbb{E}_y[y - P_2|e_2] + \frac{\delta_l}{1 - \delta_l}(s_l - \bar{u}) \geq$$

$$\mathbb{E}_y[y - P_1|e_1] + \max\{\delta_l V_l(1), e_1(b_1 + \frac{\delta_l}{1 - \delta_l}\bar{\pi}) + (1 - e_1)\delta V_l(1)\}.$$

I will focus on the case

$$\delta_l V_l(1) \geq b_1 + \frac{\delta_l}{1 - \delta_l}\bar{\pi},$$

as the other case makes separation even more difficult. It can be shown that

$$V_l(1) = \frac{(s_h - \bar{u}) + e_h(b_h + \frac{\delta_l}{1-\delta_l}\bar{\pi})}{1 - \delta_l(1 - e_h))}.$$

Then, two necessary conditions for a separating equilibrium are

$$\mathbb{E}_y[y - P_2|e_2] + \frac{\delta_l}{1 - \delta_l}(s_l - \bar{u}) \geq \mathbb{E}_y[y - P_1|e_1] + \delta_l V_l(1), \tag{1}$$

$$\mathbb{E}_y[y - P_1|e_1] + \frac{\delta_h}{1 - \delta_h}(s_h - \bar{u}) \geq \mathbb{E}_y[y - P_2|e_2] + \frac{\delta_h}{1 - \delta_h}(s_l - \bar{u}). \tag{2}$$

From (1) and (2), it follows that

$$\frac{\delta_h}{1 - \delta_h}(s_h - \bar{u}) + \frac{\delta_l}{1 - \delta_l}(s_l - \bar{u}) \geq \frac{\delta_h}{1 - \delta_h}(s_l - \bar{u}) + \delta_l V_l(1) \tag{3}$$

must hold. Assume for simplicity (and without loss of generality) that $\bar{u} = \bar{\pi} = 0$. Then,

$$\frac{\delta_h}{1-\delta_h} s_h + \frac{\delta_l}{1-\delta_l} s_l \geq \frac{\delta_h}{1-\delta_h} s_l + \delta_l V_l(1),$$

must hold. However, this inequality cannot hold when δ_h and δ_l are sufficiently close. From the enforcement constraint for the high type principal,

$$b_h = \frac{\delta_h}{1-\delta_h} s_h.$$

Therefore,

$$V_l(1) = \frac{s_h + e_h b_h}{1-\delta_l(1-e_h))} = \frac{[1-\delta_h(1-e_h)]s_h}{(1-\delta_h)[1-\delta_l(1-e_h)]}.$$

From this it follows that

$$\delta_l V_l(1) - \frac{\delta_h}{1-\delta_h} s_h = \frac{s_h}{1-\delta_h}(\frac{\delta_l[1-\delta_h(1-e_h)]}{[1-\delta_l(1-e_h)]} - \delta_h).$$

Finally, it can be checked that

$$(\delta_l V_l(1) - \frac{\delta_h}{1-\delta_h} s_h) + (\frac{\delta_h}{1-\delta_h} - \frac{\delta_l}{1-\delta_l}) s_l = \frac{\delta_h - \delta_l}{1-\delta_h}[\frac{s_l}{1-\delta_l} - \frac{s_h}{1-\delta_l(1-e_h)}].$$

The term in parentheses on the right hand side of the equality is strictly positive for all δ_l strictly greater than zero, if δ_h and δ_l are sufficiently close assuming that e_h is positive. This shows that, with at least some parameter values, a contract-separating equilibrium is impossible if costly signaling lasts only one period.[16] ■

Let me summarize the idea of the proof in a couple of sentences. Costly signaling (that is, The gain of the low type from imitation accrues very early and is discounted little relative to what she would have gotten had she revealed her type. This is problematic, especially if the discount factors, δ_h and δ_l are very close. Therefore the high type must delay symmetric information contract to a sufficiently late period. This is the gist behind showing that there *always* exists a separating equilibrium provided that costly signaling can be spread over multiple periods. To that I aim, I analyze the PPBE with

[16] The result still holds even if default by the low type is allowed in the initial period as part of the separation mechanism.

the following property: I allow out of equilibrium degenerate beliefs to change. In other words, even if beliefs put zero probability on a type of player at time t, beliefs can still attach positive probability on that type after t off-the-equilibrium path. Due to this property, costly signaling can take place even after types separate in the initial period. Then, the reason why a separating equilibrium always exists becomes very intuitive. Initially, the high type settles for low enough profits and delays full cooperation (the contract offer C_h) to a sufficiently late period.[17] Since the two types differ in their discount factors, a long enough delay in full cooperation suffices to prevent the low type from imitating.

Suppose that the game parameters are such that Proposition 2 holds. Then, it is indeed necessary for the high type to offer contracts different from C_h for multiple periods in order to separate; otherwise, the low type would imitate the high type. Now, I construct a very simple separating equilibrium to illustrate how repeated signaling can work with such parameters. The idea is simple. The high type starts off with a contract that generates lower profit than C_l (the optimal symmetric-information contract for a low type principal), which is followed by C_l offers until a certain period, followed by C_h thereafter. Now I provide a more detailed description.

At $t = 0$, the high type's contract offer consists of a fixed wage w_0 and a bonus b_l contingent on high output, where b_l is the optimal symmetric-information bonus for the low type. For $t \in \{1, \ldots, T^* - 1\}$, the high type offers $C_{t,h} = C_l = \{w_l, b_l\}$, where T^* is to be determined endogenously in equilibrium, just like w_0. Starting in period $t = T^*$, the high type always offers C_h. On the other hand, the low type's contract offer is C_l for all $t \geq 0$. The respective incentive compatibility constraints for the low type and the high type are

$$(s_l - \bar{u}) + \delta_l(s_l - \bar{u}) + \delta_l^2(s_l - \bar{u}) + \ldots + \frac{\delta_l^{T^*}}{1-\delta_l}(s_l - \bar{u}) \geq \tag{4}$$

$$\mathbb{E}_y[y - P_0|e_1] + \delta_l(s_l - \bar{u}) + \delta_l^2(s_l - \bar{u}) + \ldots + \delta_l^{T^*} V_l(1),$$

$$\mathbb{E}_y[y - P_0|e_1] + \delta_h(s_l - \bar{u}) + \ldots + \frac{\delta_h^{T^*}}{1-\delta_h}(s_h - \bar{u}) \geq \tag{5}$$

[17] Recall that $C_h$is the optimal symmetric-information contract for the high type principal.

$$(s_l - \bar{u}) + \delta_h(s_l - \bar{u}) + \delta_h^2(s_l - \bar{u}) + \ldots + \frac{\delta_h^{T^*}}{1-\delta_h}(s_l - \bar{u})$$

where $P_0 = w_0$ if $y_0 = L$ and $P_0 = w_0 + b_l$ if $y_0 = H$. Note that

$$\frac{\delta_h^t}{1-\delta_h}(s_h - \bar{u}) - \frac{\delta_h^t}{1-\delta_h}(s_l - \bar{u}) > \delta_l^t V_l(1) - \frac{\delta_l^t}{1-\delta_l}(s_l - \bar{u}) \tag{6}$$

must hold for some t sufficiently large (but finite) since $(\delta_l/\delta_h)^t \to 0$ as $t \to \infty$. Let T^* be the minimum of t that makes the inequality (6) hold strictly. Next, let w_0 be given by

$$\mathbb{E}_y[y - P_0|e_1] - (s_l - \bar{u}) = \frac{\delta_l^{T^*}}{1-\delta_l}(s_l - \bar{u}) - \delta_l^{T^*} V_l(1). \tag{7}$$

By construction in (7),

$$(s_l - \bar{u}) + \delta_l(s_l - \bar{u}) + \delta_l^2(s_l - \bar{u}).... + \frac{\delta_l^{T^*}}{1-\delta_l}(s_l - \bar{u}) =$$

$$\mathbb{E}_y[y - P_0|e_1] + \delta_l(s_l - \bar{u}) + \delta_l^2(s_l - \bar{u}).... + \delta_l^{T^*} V_l(1),$$

which implies that (4) holds -i.e., the low type principal is indifferent between mimicking the high type and revealing her type. By definition of T^*,

$$\mathbb{E}_y[y - P_0|e_1] + \delta_h(s_l - \bar{u}) + \delta_h^2(s_l - \bar{u})...\frac{\delta_h^{T^*}}{1-\delta_h}(s_h - \bar{u}) >$$

$$(s_l - \bar{u}) + \delta_h(s_l - \bar{u}) + \delta_h^2(s_l - \bar{u})..\frac{\delta_h^{T^*}}{1-\delta_h}(s_l - \bar{u}),$$

which shows that (5) holds and the high type is willing to engage in costly signaling. Moreover, the high type cannot do better deviating to another bonus scheme at some $t \leq T^* - 1$ because this makes the agent believe she is low type with probability one. If $\delta_l < \delta < \delta_h$, then only a slight modification is needed to prove that there is always a separating equilibrium. In that case, one still apply C_l in the high type's separating contract *as if* $\delta < \delta_l$, The terms w_0 and T^* end up being slighty different than what I presented above; however, the difference is inessential for the argument. Hence, I have shown the following:

Proposition 3 *There always exists a separating equilibrium. With some parameter* *values, it requires dynamic costly signaling for an extended period.*

This simple equilibrium is very useful for conveying the idea behind repeated signaling in practice as it involves just two endogenous parameters, w_0 and T^*. Nevertheless, it is suboptimal. To see this, suppose that the bonus in period $T^* - 1$ increases to $b_l + \Delta b$. In the original contract, a low type who imitates would not default at $T^* - 1$ as it is worthwhile to wait until T^* when the bonus and surplus both get higher. The low type principal would also prefer honoring $b_l + \Delta b$ at $T^* - 1$ if Δb is sufficiently small because it it is *still* more profitable to default at $t \geq T^*$. The term

$$\delta_l^{T^*-1} \left(\begin{array}{c} \mathbb{E}_y[y - w - eb_l|e] + \max\{\delta_l V_l(1), (e(b_l + \Delta b)((b_l + \Delta b) + \\ \frac{\delta_l}{1-\delta_l}\overline{\pi}) + (1 - (e(b_l + \Delta b)))\delta_l V_l(1)\} \end{array} \right) \tag{8}$$

represents the period $T^* - 1$ component of the low type's imitation payoff. Since the low type prefers honoring $b_l + \Delta b$ at $T^* - 1$, this term must equal $\delta_l^{T^*-1} \left(\mathbb{E}_y[y - w - eb_l|e] + \delta_l V_l(1)\right)$ for Δb sufficiently small. Due to this and due to the difference in discount factors, the high type principal's payoff increases more than the imitation payoff of the low type principal with the new bonus $b_l + \Delta b$. Therefore, the high type principal can increase costly signaling (-i.e., fixed wage w_0) by the amount of change in the imitation payoff of the low type and increase her payoff strictly. This finding suggests that a gradual increase in bonus and suplus might be the optimal way to separate. In the next section, I will indeed show that gradualism is the optimal solution to the high type's separation problem within a wide class of contracts.

Halac (2012) is the only other work I am aware of that studies relational contracts with persistent private information, but her information asymmetry is about the value of outside options, and consequently, she finds different types of equilibria that do not display this sort of buildup of a relationship.

As I have already mentioned, Halac (2012) analyzes a dynamic relational contracting model with private information. The outside option of the principal can be either high or low, which remains fixed throughout the game and is the principal's private information. In contrast with my model, Halac assumes that degenerate beliefs cannot change. However, even if degenerate beliefs were allowed to change in her model, there would still exist no equilibrium with separation through choice of contract. In her setup, default is a necessary condition for information revelation, which is not the case in the current setup. Both types of information about the discount factor or the outside option both create private information about the value of future relationships, both types of private information relates to the commitment of the principal to the relationship. However, the resulting equilibria are very different especially from a learning aspect. My result contrasts with Halac (2010); there always exists a contract separating equilibrium in my framework.

4.2 The Optimal Separating Contract

4.2.1 Preliminary Findings

Before stating the results, let me note that having an initial pooling stage before separation is surely suboptimal, because, the high type is better off discarding the pooling stage. This also makes the low type weakly better-off. Hence, separation takes place in the initial period in the optimal contract.

At a sufficiently distant future, the difference in discount factors has a very strong impact, so offering Ch at a distant future has an impact on the low type's imitation payoff that is very small (relative to the impact on high type's payoff). Therefore, costly signaling must stop after a finite time in the optimal contract. has a much higher impact on the high type's payoff than it would low type's imitation payoff. Suppose that this was not the case. But due to the difference in discount factors, there is a finite (and possibly large) t such that terminating costly signaling and jumping to $\{w_h, b_h\}$ coupled with an appropriate increase in the initial fixed wage makes the high type better off and the low type worse off. Hence, costly signaling lasts finitely many periods in the optimal contract.

Lemma 4 *Costly signaling stops at a finite date -i.e.,* $(w(h^t), b(h^t)) = (w_h, b_h)$ *for all* t *sufficiently large.*

The next result shows that the principal does not use variations in future rewards as a discipline device in the optimal contract -i.e., the expected lifetime utility of the agent is always equal to his outside option after the initial period of the relationship. The reason is as follows. In order to motivate the agent, the high type principal may prefer using the difference in the continuation values that are contingent on the output realization because using the bonus promise may result in a temptation for the low type to imitate. However, the principal will not do that because the fact that $\delta < \delta_h$ makes the use of a future reward scheme too costly in comparison to the cost of signaling that allows for using a bonus scheme that is equivalent to the future reward scheme.

Lemma 5 *For all t and h^t with $t > 0$, $u(h^t) = \bar{u}$ must hold.*

This result implies that the agent never pays a bonus to the principal, which is stated in Corollary 5.

Corollary 6 *In the optimal contract, $b(h^t) \geq 0$.*

In the optimal contract, the principal does not pay a bonus to the agent if low output is realized. This is fairly intuitive because the principal can simply shift the bonus for the low output to the fixed wage and adjust the bonus for the high output accordingly, without increasing (and possibly decreasing) the temptation for the low type principal to cheat.

Lemma 7 *There is no bonus payment from the principal to the agent for a low output realization.*

4.2.2 Main Result

Finally, I characterize the (constrained) Pareto-optimal separating contract assuming that the bonus for the high type depends on calendar time but not on the past realizations of output. This doesn't seem to be an extreme assumption. Because as Lemma 5 shows, a contract that is contingent on the history is used to discipline only the low type principal, not the agent. So, if the bonus were contingent on the history of output, this would affect only the low type principal's temptation to cheat because continuation payoffs are not used to motivate the agent in the unconstrained Pareto-optimal contract as shown in Lemma 5. I show that the bonus of the high type gradually increases in the optimal contract. Together with Lemma 5, this implies that the surplus in the relationship increases, as well. Since $w_t = \bar{u} - \mathbb{E}_y[b_t - c|e]$ holds at every t, the bonus schedule, the surplus in the relationship, and the effort schedule increase in time, whereas the fixed wage decreases over time.

Proposition 8 *In the (constrained) optimal separating contract of the high type principal, the bonus schedule is strictly increasing until it reaches b_h, which takes place in finite time. Similarly, the effort schedule is strictly increasing until it reaches e_h. Finally, the fixed wage w_t is decreasing over time in the optimal contract.*

The intuition for this proposition is simple. The underlying mechanism is the ability of the high type principal to delay the benefits of higher cooperation to later periods due to the difference in discount factors. The high type principal should establish reputation at an early stage by accepting a stream of low payoffs. As the relationship moves forward the high type is able to offer higher and higher bonuses as she accumulates a stock of reputation from the previous stream of low payoffs. Moreover, the low type principal is less tempted by the higher bonuses that comes in later periods due to the difference in discount factors.

To make the idea transparent consider the hypothetical scenario where $T^* = 2$ (6). This implies that the high type principal manages to separate undertaking costly signaling for two periods. By construction, the bonuses in the initial two periods are b_l and starting from Period 2 onwards, the high type principal is able to offer the symmetric information bonus b_h. As I have discussed before, this is not the optimal contract for the high type. Because after the initial period, the high type principal can offer $b_l + \Delta b$ (coupled with an increase in w_0 in the initial period) which is enforceable and credible if Δb is not too high. Then the bonus schedule becomes $\{b_l, b_l + \Delta b, b_h, b_h, b_h, ...\}$. The bonus $b_l + \Delta b$ is higher than b_l but is still strictly lower than b_h. This increasing bonus schedule implies that as the relationship moves on, the high type gains credibility through the low payoff stream in the previous periods coupled with the impact of the difference in discount factors. *Waiting* before engaging in full cooperation makes imitating less and less tempting for the low type and the high type manages to offer higher bonuses as the relationship moves on.

Discuss how the result is different than the rest of the literature. Watson, Ghosh and Ray

It is easy to see that the optimal pooling contract is the symmetric information contract of the low type. Because in the optimal pooling contract there should be no revelation about types and parties either honor or default with the same probability. Mixing between defaulting and honoring cannot be optimal and pooling because what makes the high type indifferent makes the low type strictly prefer defaulting. Therefore the optimal pooling equilibrium is such that both types honor and the optimal pooling contract can only be the low type's symmetric information contract. However, in the efficient separating equilibrium the high type principal strictly prefers revealing her type over offering the symmetric information contract of the low type. Hence, the following proposition is proved.

Proposition 9 *The optimal separating equilibrium generates higher surplus than the optimal pooling equilibrium.*

4.3 An Extension: Multiple Types

If the principal has more than two types, there still exists a separating equilibrium, and just as in the case with two types, the behavior may be distorted for multiple periods. On the other hand, the precise characterization of the optimal contract is difficult and depends on the particular case because the relative magnitudes for discount factors matter in the analysis. To see this, consider the following scenario. If the principal has three possible types and the discount factors of the low and medium types are very close to each other, then the high type should typically separate fast, whereas there is delay in the separation of the middle type. If, on the other hand, the discount factors of the high and medium types are very close to each other then the separation of the high type takes longer. One can conjecture that even in the setting with multiple types the monotonicity of the bonus schedule, the effort schedule and the surplus in the relationship is preserved.

5 Conclusion

The relational contracting literature has usually relied on the assumption that the trading parties know each other's true motives at the beginning of the relationship. I drop this assumption and I address the issues resulting from the uncertainty of the agent regarding the principal's true motives. I characterize what actions a trustworthy principal should take in order to convince the agent that she is trustworthy. Put differently, I show how trust is built in this relational contracting context.

Discount factor is the proxy for trustworthiness. As we know from the theory of repeated games a high discount factor typically fosters cooperation whereas a low discount factor results in myopia. In this setup, I analyze both pooling equilibria and separating equilibria, in which the two principal types separate making different contract offers in the initial period. I show that there *always* exists a separating equilibrium and it sometimes requires dynamic costly signaling for an extended period of time. If, for example, the discount factors of the two types are very close the continuation play after full separation in period 1 is not the symmetric information equilibrium; the high type delays the fullest possible cooperation to a later period such that the low type is not tempted to imitate. Next I characterize the properties of the optimal contract-separating equilibrium. I show that in the optimal contract of the high type, the bonus for high output, the agent effort and the surplus in the relationship are all increasing; on the other hand, the low type offers immediately her full information contract because it is the low type who would like to imitate. So, the equilibrium for high type is characterized by gradual building of trust, which is intuitive because being more patient, the high type should forgo earlier payoffs for the benefits of later and higher cooperation. Next, I find that the optimal pooling contract is the optimal symmetric information contract of the low type. In a separating contract the high type principal prefers separation to the low type's symmetric information contract, which is the best pooling outcome. The agent is also better off under separation because some part of the costly signaling is a monetary transfer to the agent. That implies the following: The optimal separating contract generates more surplus than the optimal pooling contract.

In my model, promises are not broken on the equilibrium path. However, there are real world examples in which firms do break their implicit promises (albeit infrequently). A very interesting extension would be an analysis incorporating to the environment productivity shocks that may or may not be observable to the agent. In such an environment, broken promises and renegotiation may be part of an equilibrium. This is left for further research.

Appendix

Proof of Proposition 1 Suppose that following some history h^t, the optimal contract prescribes $u|_{h^t} > \bar{u}$ for type θ. Let $h^t(H) = \{h^{t-1}, h\}$ and $h^t(L) = \{h^{t-1}, h'\}$ where $h \equiv h_t$ with $H \in h$ and h' is the same as h except that the output realization at t differs. So, $h^t(H)$ and $h^t(L)$ differ only in their last period output realization. The effort chosen by the agent in period t following history h^{t-1} is given by the first order condition,

$$c'(e) = b(h^t(H)) - b(h^t(L)) + \delta(u|_{(h^t(H))} - u|_{(h^t(L))}).$$

Let's alter the contract as follows: $b(h^t(H))$ is increased by a small amount δx whereas $u|_{(h^t(H))}$ is reduced by x via a reduction in the fixed wage, w. Note that $u(h^{t-1})$ remains the same as in the initial contract and the incentives of the agent are unchanged as long as the increase in the bonus payment is incentive compatible -i.e., if $b(h^t(H)) \geq 0$ the principal must keep this bonus promise and if $b(h^t(H)) < 0$ the agent must keep his promise. The principal is willing to keep this promise if $b(h^t(H)) \geq 0$ because in the initial contract,

$$b(h^t(H)) + \frac{\delta_\theta}{1-\delta_\theta}\bar{\pi} \leq \delta_\theta \pi_\theta|_{(h^t(H))}.$$

Even though the bonus $(h^t(H))$ increases by the amount δx, $\pi_\theta|_{(h^t(H))}$ also increases by an amount x since the fixed wage at date $t+1$ following the history $(h^t(H))$ is reduced by x. Hence the enforcement constraint for the principal is relaxed. Moreover, the principal is better off with this change because his lifetime expected payoff increases by $P((h^t(H)))\delta_\theta^t(\delta_\theta x - \delta x)$, which is positive since $\delta_\theta > \delta$ by assumption.

If on the other hand $b(h^t(H)) < 0$ then the bonus $(h^t(H))$ is increased by a small amount δx whereas $u|_{(h^t(H))}$ is reduced by x via a reduction in the fixed wage, w. The agent must keep his promise. In the initial contract,

$$b(h^t(H)) + \frac{\delta}{1-\delta}\bar{u} \leq \delta u|_{(h^t(H))}$$

holds so the new contract leaves the enforcement constraint for the agent unchanged. Moreover the effort choice remains the same.

Similarly, suppose that $u|_{h^t} > \bar{u}$ for type θ and $h^t = (h^t(L))$. The effort chosen by the agent in period t following history h^{t-1} is given by the first order condition,

$$c'(e) = b(h^t(H)) - b(h^t(L)) + \delta(u|_{h^t(H)} - u|_{h^t(L)}).$$

Let's alter the contract as follows: $b(h^t(L))$ is increased by an amount δx whereas $u|_{(h^t(H)}$ is reduced by x via a reduction in the fixed wage, w. Note that $u(h^{t-1})$ remains the same and the incentives of the agent are unchanged as long as the increase in the bonus payment is incentive compatible -i.e., if $b(h^t(L)) \geq 0$ the principal must keep this bonus promise and if $b(h^t(L)) < 0$ the agent must keep his promise. If $b(h^t(L)) \geq 0$ the principal is better off eliminating the bonus, so that $b'(h^t(L)) = 0$ and $b'(h^t(H)) = b(h^t(H)) - b(h^t(L))$. Incentives for the agent remain the same and the enforcement constraint for the principal is relaxed. Then the principal can make a increase in $b'(h^t(L))$, which in turn increases the expected surplus. If, on the other hand, $b(h^t(L)) < 0$ then the bonus $b(h^t(L))$ is increased by a small amount δx whereas $u|_{(h^t(L))}$ is reduced by x via a reduction in the fixed wage, $w(h^t(L))$. The agent must keep his promise. In the initial contract,

$$b(h^t(L)) + \frac{\delta}{1-\delta}\bar{u} \leq \delta u(h^t(L))$$

holds so the new contract leaves the enforcement constraint for the agent unchanged. Moreover, the effort choice remains the same.

Proof of Lemma 4 Suppose not. Then take the optimal contract with the infinite costly signaling on some path of play. Let the contract be altered as follows: At time T which is to be determined endogenously, the high type principal switches to $(w(h^t), b(h^t)) = (w_h, b_h)$ for all $t > T$. This is, of course, enforceable for the high type provided that the contract remains separating. If T is at a sufficiently distant future, then the gain of the high type principal from proposing $(w(h^t), b(h^t)) = (w_h, b_h)$ for $t > T$ is greater than the maximum gain that the low type can make from imitation because of the difference in discount factors. So, the high type is still better off increasing costly signaling (i.e., the first period fixed wage) by a suitable amount. Hence, costly signaling must end at a finite time in the optimal contract.

Proof of Lemma 5 Suppose $u(h^t) > \bar{u}$ for some h^t where $t > 0$. Consider the following change in the high type's contract. The principal reduces the continuation payoff $u(h^t)$ following the history h^t by a small amount x (via reducing the fixed wage in period $t+1$

by x). She also increases $b(h^t)$ by δx. First of all, one can check that the enforcement constraint for the high type principal is relaxed if $b(h^t) \geq 0$, and in case $b(h^t) < 0$ the enforcement constraint for the agent remains unchanged. Second, to make the new contract separating, the costly signaling (i.e., the fixed wage in the initial period) must increase by the amount of increase in the low type's imitation payoff. To show that the high type is better off even with the increase in costly signaling, I use the following claim.

Claim *In the optimal contract, the low type principal who imitates is indifferent between defaulting and honoring at h^t if $b(h^t) \in (0, b_h)$.*

Proof of Claim. Suppose not. So,

$$\delta_l(\pi_l(h^t) - \frac{\bar{\pi}}{1-\delta_l}) < b(h^t)$$

holds for some h^t with $b_h > b(h^t) \geq 0$. But then the high type can increase her payoff making a (sufficiently) small increase in one of the future bonuses that reward high output following history h^t. This does not affect the low type's incentives to imitate since the low type strictly prefers defaulting if h^t realizes. Consider, for example, the bonus for a history following h^t that is rewarded for high output realization, *right before* costly signaling ends. This bonus is surely lower than b_h and it can be increased slightly. There must exist such a history, since by Lemma 4 costly signaling terminates at a finite date that is later than t, by the hypothesis that $u(h^t) > \bar{u}$. Hence, the low type must be indifferent between indifferent between defaulting and honoring at h^t if $b(h^t) > 0$. So, the claim is proved.

If $b(h^t) < 0$, then the low type's imitation payoff increases by at most $p(h^t)\delta_l^{t-1}(\delta_l x - \delta x)$, where $p(h^t)$ is the probability that the high type's play reaches history h^t. This is still true if $b(h^t) \geq 0$. By the claim above, the low type strictly prefers honoring at h^t with the small increase in $b(h^t)$ (the enforcement constraint for the low type is relaxed), and the low type's imitation payoff increases by at most $p(h^t)\delta_l^{t-1}(\delta_l x - \delta x)$. On the other hand, the high type's payoff increases by $p(h^t)\delta_h^{t-1}\delta_h^t(\delta_h x - \delta x)$. So, increasing the fixed wage in the initial period by (at most) $p(h^t)\delta_l^{t-1}(\delta_l x - \delta x)$ will make the high type strictly better off, which gives a contradiction.

Proof of Corollary 6 This follows from the enforcement constraint for the agent and

Lemma 4

Proof of Lemma 7 Let $h_H^t = \{h^{t-1}, H\}$ and $h_L^t = \{h^{t-1}, L\}$ be two histories until time t that differ only in their last output realization at time t. Suppose that $b(h_L^t) > 0$. If the fixed wage is arranged as $w + b(h_L^t)$ and the bonus paid for high output is reduced to $b(h_H^t) - b(h_L^t)$ (note that $b(h_H^t) - b(h_L^t) \leq 0$ cannot hold in the initial contract). Therefore, the incentives of the agent are not affected. Since the bonus payment following high output reduces, the required costly signaling may even fall.

Proof of Proposition 8 After the initial period the only costly signaling device is the offer of a sufficiently low bonus (costly signaling in the form of a high fixed wage will be used only in the initial period since by restriction to time-dependent contracts $u_t = u(h^t) = \bar{u}$ must hold for all $h^t \in \mathcal{H}^t$ for fixed $t > 1$.). Let's start from the last period with costly signaling. Let $T - 1$ denote the last time period such that $b_t \neq b_h$ holds and therefore, $b_T = b_h$. Just as in the proof of Lemma 4, it can be shown that there exists such a T. First note that $b^*_{T-1} < b_T$ by the definition of T. Otherwise $b^*_{T-1} > b_T$ in which case the high type principal defaults. Next I will show that $b^*_{T-2} < b^*_{T-1}$. Suppose towards a contradiction that $b^*_{T-2} \geq b^*_{T-1}$. Since $b^*_{T-1} < b_T = b_h$ it follows that $\pi_{l,T-1} < \pi_{l,T}$ where $\pi_{l,t}$ stands for the maximum lifetime utility of the low type at time t having mimicked the high type until then.
In case (i), assume that at time $T - 2$

$$b^*_{T-2} \leq \delta_l \pi_{l,T-1} - \frac{\delta_l}{1-\delta_l}\bar{\pi},$$

holds. If

$$\delta_l \pi_{l,T-1} - \frac{\delta_l}{1-\delta_l}\bar{\pi} \leq b^*_{T-1},$$

then from

$$b^*_{T-2} \leq \delta_l \pi_{l,T-1} - \frac{\delta_l}{1-\delta_l}\bar{\pi} < \delta_l \pi_{l,T} - \frac{\delta_l}{1-\delta_l}\bar{\pi} \leq b^*_{T-1}$$

it follows that $b^*_{T-2} < b^*_{T-1}$. But it cannot be the case that

$$\delta_l \pi_{l,T} - \frac{\delta_l}{1-\delta_l}\bar{\pi} > b^*_{T-1}$$

holds because then the high type would increase b^*_{T-1} slightly (coupled with an increase in the initial period fixed wage as costly signaling). If the strict inequlaity holds, the low type does not default with a bonus level of b^*_{T-1}, which is still true with the slight in increase in b^*_{T-1}. So, the payoff of the high type principal increases more than the imitation payoff of the low type with the small increase in b^*_{T-1} as

$$\frac{\partial \pi_h}{\partial b_{T-1}}|_{b^*_{T-1}} > \frac{\partial \pi_l}{\partial b_{T-1}}|_{b^*_{T-1}}.$$

Increasing costly signaling by the amount of increase in the low type's imitation payoff, the high type is still better off. Next, consider the case (ii)

$$b^*_{T-2} > \delta_l \pi_{l,T-1} - \frac{\delta_l}{1-\delta_l}\bar{\pi}.$$

Consider towards a contradiction that $b^*_{T-2} \geq b^*_{T-1}$. Then, I alter the equilibrium contract as follows: b^*_{T-2} is decreased slightly whereas b^*_{T-1} is increased slightly in a way that

$$\frac{\partial \pi_l}{\partial b_{T-2}}|_{b_{T-2}=b^*_{T-2}} + \frac{\partial \pi_l}{\partial b_{T-1}}|_{b_{T-1}=b^*_{T-1}} \frac{\partial b_{T-1}}{\partial b_{T-2}} = 0 \tag{9}$$

holds. Note that this small increase in b^*_{T-1} is still enforceable for the high type. I want to show that this results in

$$\frac{\partial \pi_h}{\partial b_{T-2}}|_{b_{T-2}=b^*_{T-2}} + \frac{\partial \pi_h}{\partial b_{T-1}}|_{b_{T-1}=b^*_{T-1}} \frac{\partial b_{T-1}}{\partial b_{T-2}} < 0 \tag{10}$$

which establishes a contradiction because this implies that the high type principal can decrease b^*_{T-2} slightly and increase b^*_{T-1} slightly which would benefit the high type strictly more than the low type. From (9) it follows that

$$\frac{\partial b_{T-1}}{\partial b_{T-2}} = -\frac{\frac{\partial \pi_l}{\partial b_{T-2}}|_{b_{T-2}=b^*_{T-2}}}{\frac{\partial \pi_l}{\partial b_{T-1}}|_{b_{T-1}=b^*_{T-1}}} = -\frac{\frac{\partial s_{T-2}}{\partial b_{T-2}} + e^*_{T-2} + b_{T-2}\frac{\partial e}{\partial b_{T-2}} - \delta_l \pi_{l,T-1}\frac{\partial e}{\partial b_{T-2}}}{\delta_l(1-e^*_{T-2})(\frac{\partial s_{T-1}}{\partial b_{T-1}} + e^*_{T-1} + b_{T-1}\frac{\partial e}{\partial b_{T-1}} - \delta_l \pi_{l,T}\frac{\partial e}{\partial b_{T-1}})} \tag{11}$$

$$= -\frac{c''(e^*_{T-1})(H-L) + e^*_{T-2}c''(e^*_{T-2}) - \delta_l \pi_{l,T-1}}{\delta_l c''(e^*_{T-2})(1-e^*_{T-2})((H-L) + e^*_{T-1}c''(e^*_{T-1}) - \delta_l \pi_{l,T})} \tag{12}$$

Note that $\frac{\partial b_{T-1}}{\partial b_{T-2}} < -\frac{c''(e^*_{T-1})}{\delta_l c''(e^*_{T-2})}$. Also

$$-\frac{\frac{\partial \pi_h}{\partial b_{T-2}}|_{b_{T-2}=b^*_{T-2}}}{\frac{\partial \pi_h}{\partial b_{T-1}}|_{b_{T-1}=b^*_{T-1}}} = -\frac{(H-L-c'(e^*_{T-2}))c''(e^*_{T-1})}{\delta_h(H-L-c'(e^*_{T-1}))c''(e^*_{T-2})} > -\frac{c''(e^*_{T-1})}{\delta_h c''(e^*_{T-2})} > -\frac{c''(e^*_{T-1})}{\delta_l c''(e^*_{T-2})}$$

Therefore, (10) must hold and $b^*_{T-2} \geq b^*_{T-1}$.cannot hold in the optimal contract. Hence $b^*_{T-2} < b^*_{T-1}$.

Next I will show that for $t \in \{1, \ldots, T-2\}$, $b^*_{t-1} < b^*_t$ must hold provided that b_τ is monotone increasing for $\tau \geq t$. The proof is analogous to the proof presented above for showing $b^*_{T-2} < b^*_{T-1}$. First, one needs to verify that $\pi_{l,t} < \pi_{l,t+1}$ and $\pi_{h,t} < \pi_{h,t+1}$. But, this is true by the hypothesis that b_τ is monotone increasing for $\tau \geq t$.

First, consider the case where

$$b^*_{t-1} \leq \delta_l \pi_{l,t} - \frac{\delta_l}{1-\delta_l}\bar{\pi}.$$

Assume towards a contradiction that $b^*_{t-1} \geq b^*_t$. So, from

$$b^*_t \leq b^*_{t-1} \leq \delta_l \pi_{l,t} - \frac{\delta_l}{1-\delta_l}\bar{\pi} < \delta_l \pi_{l,t+1} - \frac{\delta_l}{1-\delta_l}\bar{\pi}$$

we obtain

$$b^*_t < \delta_l \pi_{l,t+1} - \frac{\delta_l}{1-\delta_l}\bar{\pi}.$$

But this implies that

$$b^*_t = \delta_h \pi_{h,t+1} - \frac{\delta_h}{1-\delta_h}\bar{\pi}.$$

Otherwise, the high type would have a higher b^*_t in the optimal contract. But then,

$$b^*_t = \delta_h \pi_{h,t+1} - \frac{\delta_h}{1-\delta_h}\bar{\pi} > \delta_h \pi_{h,t} - \frac{\delta_h}{1-\delta_h}\bar{\pi} \geq b^*_{t-1}$$

implies that $b^*_{t-1} < b^*_t$, a contradiction.

Next, consider the case where

$$b^*_{t-1} > \delta_l \pi_{l,t} - \frac{\delta_l}{1-\delta_l}\bar{\pi}.$$

Assume towards a contradiction that $b^*_{t-1} \geq b^*_t$. This implies that

$$b^*_t \leq b^*_{t-1} \leq \delta_l \pi_{h,t} - \frac{\delta_h}{1-\delta_h}\bar{\pi} < \delta_h \pi_{h,t+1} - \frac{\delta_h}{1-\delta_h}\bar{\pi}.$$

Hence, a small increase in b^*_t would be enforceable for the high type. Now, let's change the contract as follows: b^*_{t-1} is decreased and b^*_t is increased slighlty in a way that

$$\frac{\partial \pi_l}{\partial b_{t-1}}|_{b_{t-1}=b^*_{t-1}} + \frac{\partial \pi_l}{\partial b_t}|_{b_t=b^*_{t-1}} \frac{\partial b_t}{\partial b_{t-1}} = 0 \tag{13}$$

holds. This results in

$$\frac{\partial \pi_h}{\partial b_{t-1}}|_{b_{t-1}=b^*_{t-1}} + \frac{\partial \pi_h}{\partial b_t}|_{b_t=b^*_t} \frac{\partial b_t}{\partial b_{t-1}} < 0 \tag{14}$$

giving a contradiction. The proof for showing this follows exactly the same steps as I used in case (ii) above.

Finally, to show that the fixed wage is decreasing in the optimal contract, let $u(h^t)$ denote the expected stage game payoff of the agent at period t following history $h^t \in \mathcal{H}^t$ where we understand $\{\emptyset\} \in \mathcal{H}^t$. Then, either $u(\emptyset) > \bar{u}$ or $b_1 < b_l$.Suppose $u(\emptyset) = \bar{u}$ and $b_1 \geq b_l$. The bonus schedule is increasing from Proposition 7. Then, the total expected payoff of the low type principal from imitating exceeds $\frac{(s_l - \bar{u})}{1-\delta_l}$ and separation is impossible, a contradiction. Hence either $u(\emptyset) > \bar{u}$ or $b_1 < b_l$ must hold.

REFERENCES

Abreu, Dilip. (1988). "On the Theory of Infinitely Repeated Games with Discounting," Econometrica, 56(6), 383-396.

Admati Anat R. and Motty Perry. (1987) "Strategic delay in bargaining," Review of Economic Studies, 54(3), 345-364.

Baker G., Gibbons R., and K. J. Murphy. (1994). "Subjective Performance Measures in Optimal Incentive Contracts," Quarterly Journal of Economics, 109(4), 11215-1156.

Baker G., Gibbons R., and K.J. Murphy (2003) "Relational Contracts and the Theory of the Firm," Quarterly Journal of Economics, 117(1), 39–84.

Bull, Clive. (1988) "The Existence of Self-Enforcing Implicit Contracts," Quarterly Journal of Economics, February 1987, 102(1), pp. 147-159.

Cho I. K. and D. Kreps. (1987) "Signaling games and stable equilibria," Quarterly Journal of Economics, 56, 601-612.

Cramton, Peter. (1992) "Strategic delay in bargaining with two-sided uncertainty,". Review of Economic Studies, 59, 205-225.

Ghosh, P., and D. Ray (1996). "Cooperation in Community Interaction without Information Flows, Review of Economic Studies, 63, 491–519.

Halac, Marina. (2010) "Relational Contracts and Value of Relationships," forthcoming in American Economic Review.

Kaya, Ayca. (2009) "Repeated Signaling," Games and Economic Behavior, 66(2), 841-854.

Levin, Jonathan. (2003a) "Relational Incentive Contracts," American Economic Review, 93(3), 835–57.

Levin, Jonathan. (2003b) "Multilateral Contracting and the Employment Relationship," Quarterly Journal of Economics, 117(3), 1075-1103.

MacLeod, W. B., and J. M. Malcomson (1989) "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," Econometrica, 57(2), 447-480.

Noldeke, George and Eric van Damme. (1990) “Signaling in a dynamic labor market,” Review of Economic Studies, 57(1), 1-23.

Riley, John G. (2001) “Silver signals: Twenty-¯ve years of screening and signaling,” Journal of Economic Literature, 39, 432-478.

Rubinstein, Ariel. (1985). “A bargaining model with incomplete information about preferences,” Econometrica, 1151-1172.

Watson, J. (1999). “Starting Small and Renegotiation,” Journal of Economic Theory, 85, 52–90.

Watson, J. (2002): “Starting Small and Commitment,” Games and Economic Behavior, 38, 176–99.