Melis Kartal

Honest Equilibria in Reputation Games: The Role of Time Preferences

American Economic Journal: Microeconomics 2018, 10(1): 278–314
https://doi.org/10.1257/mic.20160178

Honest Equilibria in Reputation Games: The Role of Time Preferences†

By Melis Kartal*

New relationships are often plagued with uncertainty because one of the players has some private information about her "type." The reputation literature has shown that equilibria that reveal this private information typically involve breach of trust and conflict. But are these inevitable for equilibrium learning? I analyze self-enforcing relationships where one party is privately informed about her time preferences. I show that there always exist honest reputation equilibria, which fully reveal information and support cooperation without breach or conflict. I compare these to dishonest reputation equilibria from several perspectives. My results are applicable to a broad class of repeated games. (*JEL* C73, D82, D83, D86, Z13)

How can two agents build cooperation when one of them does not fully know the other's motives and preferences? Consider a repeated relationship (e.g., a buyer-seller, principal-agent or lender-borrower relationship), where incentives are informal because actions and outcomes are not verifiable or contractible. In such a relationship that lacks formal enforcement, trust is essential. However, a buyer may be uncertain about the trustworthiness of a seller, a lender may have doubts about the repayment incentives of a borrower, and an agent may not know the quality of a prospective principal's work environment.

Whenever there are "better" types in a market, the "worse" types have an opportunistic incentive to imitate them. Starting with the seminal works by Kreps and Wilson (1982), Milgrom and Roberts (1982), and Sobel (1985), the reputation literature has shown that the "bad" type exploits the informational asymmetry and accepts short-run losses in order to build a "good" reputation and obtain a higher long-run profit. This profit is eventually attained by breaching trust: a bad type borrower defaults, a bad type principal withholds the agent's bonus, and a bad type seller delivers a defective product. Such a deviation from the implicit agreement generates conflict, and the relationship likely breaks down. So, a natural question follows. Are imitation and breach of trust inevitable for equilibrium learning? For example, young firms are often advised to invest in a "name" via advertising, hiring reputable executives or board members, and engaging in social responsibility acts

* Department of Economics, Vienna University of Economics and Business, Welthandelsplatz 1, 1020, Vienna, Austria (email: melis.kartal@wu.ac.at). I would like to thank two anonymous referees, Alessandro Lizzeri, Debraj Ray, Ariel Rubinstein, Karl Schlag, Matan Tsur, and Kemal Yildiz for helpful comments and suggestions.

† Go to https://doi.org/10.1257/mic.20160178 to visit the article page for additional materials and author disclosure statement or to comment in the online discussion forum.

in order to signal a good type.[1] How effective are "costly signals" when incentives are only informal? Is there still an effective signaling mechanism that reveals information and supports cooperation without breach and conflict, and if so, under what conditions?

These questions are important because a risk of breach may be inefficient as it is correctly anticipated by the other party in equilibrium, and costlier than is presumed in standard models because individuals may exhibit an aversion to disappointment or betrayal.[2] I study these questions in the context of relational contracts with persistent private information—my results also generalize to other settings. In the model, the principal is privately informed about her time preferences. This represents a situation in which the agent does not fully know the preferences and the commitment of the principal at the beginning of the relationship. I characterize the "honest" and "dishonest" information revelation regimes of the game and compare them from several perspectives.[3] I show that there *always* exists an equilibrium that fully *separates* the principal types in an honest manner via costly signaling. Differential time preferences of types have a unique role in this result because a separating equilibrium does not generally exist with other types of private information. Hence, differential discounting is analogous to a "monotonicity condition" for repeated private information games. Moreover, I find that there exist parameters under which separating equilibria dominate other types of equilibria. These results have a simple, intuitive structure and generalize to other settings in which reputations matter.

I now outline the model and its assumptions. I develop a relational contracting model in which the agent's discount factor is fixed and known, whereas the principal's discount factor is her private information. The principal is one of two types: high or low. The high type has a higher discount factor than the low type. 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 opportunism. Similarly, in the context of relational contracts, the higher the discount factor, the easier it is to honor promises because it increases the value of future trade between the two parties. Thus, the discount factor is a simple proxy for trustworthiness and commitment in my model. Traits such as commitment and trustworthiness can be difficult to observe. This is why, for example, Forbes announces "America's most trustworthy companies" every year. Indeed, some high-profile cases that lie at the opposite ends of the trustworthiness scale indicate that some firms are more trustworthy and more committed to their employees than others. Southwest Airlines famously declined to abandon its no-layoff policy and kept employee morale high

[1] See Fombrun and Shanley (1990) and Deutsch and Ross (2003).

[2] One important exception in contract theory is the study by Hart and Moore (2008), which assumes that people care about what they receive relative to what they are "entitled" to, and getting less than what one is entitled to causes a disutility in the form of "aggrievement" and even leads to retaliation. Also, see Bohnet et al. (2008), Abeler et al. (2011), and Gill and Prowse (2012) for experimental studies on betrayal aversion and disappointment aversion, and Gul (1991) and Kőszegi and Rabin (2006, 2007) for theoretical models of disappointment aversion and loss aversion. In these models, agents are sensitive to downside deviations from their expectations and incur a psychological disutility when they receive less than what they expect or deem "fair."

[3] I refer to an equilibrium as "dishonest" if the equilibrium path involves breach of trust with a positive probability and as "honest" otherwise.

even in difficult times, whereas IBM and Credit Suisse First Boston chose to renege on their promises regarding bonus payments or layoff policies in order to cut costs (Stewart 1993; Conlin 2001).

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 a legally enforceable fixed wage and a performance-contingent bonus transfer, which is not legally enforceable. If the agent accepts the offer, he then chooses an effort level. Output is strictly increasing in effort, and exerting effort is costly. Output is observed by both the principal and the agent but 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.

Nonverifiability of output limits the scope of cooperation even in the absence of private information because incentives are only informal. Introducing private information to the standard relational setting exacerbates the incentive problem further. How does this informational asymmetry make things worse? Since a bonus promise is not legally enforceable, the willingness of the principal to honor a bonus promise depends on her future profit from the relationship, which is higher for the high-type principal because she is more patient. Thus, the high-type principal can fulfill higher bonus promises than the low type. The low type then wants to mimic the high type so that she can obtain high effort from the agent with a high bonus promise on which she subsequently defaults. Therefore, the high-type principal needs a credible signaling mechanism in order to separate herself from the low type.

I show that there always exists a separating equilibrium; the two types separate immediately by offering different contracts.[4] Although learning is immediate in such an equilibrium, the high type must distort her behavior for an extended period of time, and there is a delay in full cooperation with many parameter values.

A separating equilibrium is honest and free of conflict as it involves no imitation, breach of trust or punishment. This result contrasts with the standard reputation literature, where equilibrium information revelation involves dishonesty and conflict; in these "dishonest reputation equilibria," there is a strictly positive probability that a bonus is unpaid, a rip-off price is charged for a low-quality product, a loan is not repaid, etc.[5] My framework identifies the conditions under which equilibrium information revelation and cooperation do not involve dishonesty or conflict. In particular, it shows that honest information revelation is possible in any repeated private information game as long as types differ in their time preferences; and the game allows for costly signals, such as monetary transfers between parties or "money-burning" activities.[6]

How can one reconcile the result above with the defaults at IBM and Credit Suisse First Boston discussed earlier? Earle and Sabirianova Peter (2009) note that

[4] I define a separating equilibrium as a perfect Bayesian equilibrium in which the agent's belief about the principal's type is degenerate after the initial contract offer as well as any equilibrium path history.

[5] In addition to the classic references mentioned before, see Diamond (1989), Watson (1999; 2002), Mailath and Samuelson (2001), Halac (2012), and Jullien and Park (2014). Private information is not about time preferences in these models.

[6] I show explicitly that the main results hold in a buyer-seller relationship with moral hazard and asymmetric information à la Mailath and Samuelson (2001).

"there is no systematic data collection about breaches of the wage contract in most economies—perhaps because they are rare." The rarity of contract breach is justifiable in my framework because dishonest reputation equilibria can be inefficient, and they are generally not robust to a dynamic version of the Intuitive Criterion, as I discuss below. Moreover, the inefficiency of dishonest equilibria can be amplified in practice, in light of the recent findings regarding disappointment and betrayal aversion.[7] Many people may prefer rejecting a contract that exposes them to a nontrivial risk of exploitation, and in turn, rejected contracts will presumably adapt in order to become acceptable. Thus, in both theory and practice, an honest and conflict-free arrangement may be a plausible norm in economic relationships even if individuals enter a relationship with private information.

Next, I investigate the properties of the optimal separating contract. The low-type principal offers her optimal symmetric information contract immediately, and the game continues as in the symmetric information setting for her. The high type, however, engages in dynamic signaling that evolves until it reaches 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. This is intuitive because discounting affects the prospective payoffs of the two types differently: the high type must forgo earlier profits and delay full cooperation in order to make imitation less tempting for the low type. This *gradualism* result also emerged in previous studies with private information, albeit in the context of dishonest reputation equilibria. As it turns out, a gradualist pattern of cooperation is superior even in settings where learning is immediate because it optimally spreads costly signaling over time.[8]

I then characterize other equilibria of the game: namely, pooling equilibria, in which no information is revealed, and hybrid equilibria, which are dishonest as they involve imitation and breach of trust by the low type. I find that the optimal separating contract always dominates pooling equilibria, and I characterize the conditions under which it also dominates hybrid equilibria. Next, I develop a dynamic version of the Intuitive Criterion (Cho and Kreps 1987) and show that it typically selects the optimal separating equilibrium as the unique reasonable equilibrium.

To reiterate, differential time preferences of types have a unique role in repeated private information games. A separating equilibrium does not generally exist if the principal's private information is uncorrelated with her time preferences. This role of discounting connects my paper to the model of Becker and Mulligan (1997), in which an individual is endowed with a baseline level of discount rate and can choose to increase this level at some cost. One of the main insights of their model is that an increase in future payoffs boosts the incentive to invest in one's patience; therefore, they argue, higher income individuals are more patient. If a "good" type expects a higher future payoff than a "bad" type in a private information game (as is often the case), there is yet another incentive to invest in a higher discount factor: separation

[7] There is empirical evidence for this from labor markets, in addition to the experimental evidence from the studies mentioned in Footnote 2. For example, Mas (2006) shows that police performance declines sharply when police officers lose arbitrations and argues that considerations of fairness and disappointment affect workplace behavior.

[8] Even after information is fully revealed, costly signaling may continue (see the discussion in Section IA).

from the bad type is surely possible and also cheaper. The upshot is that if individuals can invest in their time preferences, and if the good type has a higher lifetime payoff, then a separating equilibrium always exists.

I. Related Literature

This paper draws from the study of relational contracts and the literature on reputation and dynamic signaling. Below, I discuss the two literatures separately.

A. *Reputation and Dynamic Signaling*

My work relates to the literature on reputation building, and in particular to Sobel (1985), Diamond (1989), Ghosh and Ray (1996), Kranton (1996), Watson (1999, 2002), Tadelis (1999), Mailath and Samuelson (2001), Halac (2012), and Jullien and Parks (2014). In these models, separating equilibria do not exist, and information revelation is dishonest: in equilibrium, there is a nontrivial probability that a bonus is unpaid, a rip-off price is charged for a low-quality product, a loan is not repaid, etc. Furthermore, such events often take place at *high stakes* if relationships can build up over time (see, for example, Sobel 1985, Watson 1999, 2002 and Halac 2012). In contrast with these models, I identify conditions under which separating equilibria always exist, and thus, equilibrium learning is honest and partnerships are long-lasting. I analyze the efficiency properties of separating equilibria and compare them to other types of equilibria. I also show that in the optimal separating equilibrium, the relationship becomes more cooperative and valuable over time provided that the informed party is a high type. While I obtain these results in a principal-agent model, they generalize to other settings. Section VIA presents an application of the results to a buyer-seller relationship à la Mailath and Samuelson (2001).

My work also relates to a subset of the dynamic signaling literature in which out-of-equilibrium degenerate beliefs are allowed to change. This property implies that costly signaling can continue even after beliefs have become degenerate.[9] This property has been applied in numerous models such as Admati and Perry (1987), Noldeke and van Damme (1990), Cramton (1992), Kremer and Skrzypacz (2007), and Kaya (2009). Admati and Perry (1987) and Cramton (1992) analyze a bargaining game, whereas Noldeke and van Damme (1990) and Kremer and Skrzypacz (2007) extend Spence's job market signaling model to a dynamic environment.[10] With the exception of Kaya (2009), these games are dynamic but not repeated; and the range of the signaling variable is too small, and therefore, distorting behavior

[9] It may be argued that once the agent attaches probability one to the principal being a certain type, the subsequent game should essentially be one of symmetric information. Yet, this requires a restriction on beliefs because it is not implied by the definition of perfect Bayesian equilibrium. Moreover, this type of restriction sometimes leads to the nonexistence of a perfect Bayesian equilibrium.

[10] In the basic job market signaling model, 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 equilibria, only the Riley outcome (i.e., the Pareto-efficient separating equilibrium) survives the Intuitive Criterion. Could the Riley outcome survive in a continuous time setting? If it is to survive, then costly education *must* continue after beliefs have become degenerate. This example, presented in Noldeke and van Damme (1990), illustrates why it is plausible to have costly signaling after beliefs have become degenerate.

in only one period is not sufficient to achieve separation. Kaya (2009) extends the concept of dynamic signaling to a class of repeated games where the stage game has a separating equilibrium by construction. She then characterizes the least cost separating equilibrium of the repeated game. My model differs from the aforementioned papers in two ways. Among them, only Kaya (2009) analyzes a repeated game. But the stage game of her model has a separating equilibrium by assumption, whereas different types would exhibit identical behavior if the stage game in my setting is repeated only once because incentives are only *informal*. Thus, these papers do not address the issues that I study. Moreover, the signal space is essentially endogenous in my model (as it depends on the surplus generated within the relationship), and it is not clear whether or not a separating equilibrium exists.

B. *Relational Contracting*

A relational contract sustains trade between a principal and an agent if the performance measure is nonverifiable. Numerous relational contracting models focus on environments with symmetric information (see, for example, MacLeod and Malcomson 1989, 1998, and Board and Meyer-Ter-Vehn 2015). Asymmetric information has also been incorporated into relational contracting. Shapiro and Stiglitz (1984), and Baker, Gibbons, and Murphy (2002) consider relational contracts with moral hazard, whereas Levin (2003) analyzes two distinct scenarios. He assumes that either the agent's effort or the agent's time-specific cost parameter cannot be observed by the principal. MacLeod (2003) and Fuchs (2007) consider asymmetric information about output realization, whereas Li and Matouschek (2013) assume that there is asymmetric information about the state of the world which affects the opportunity cost of paying a bonus to the agent. Asymmetric information has no persistence in these papers.

To my knowledge, there are only two papers that analyze relational contracts with persistent asymmetric information. Halac (2012) analyzes relational contracts in a setting where the principal has persistent private information regarding her outside option. As discussed above, separating equilibria do not exist, and information revelation is dishonest in Halac (2012): if the principal is the less cooperative type, then the relationship breaks down after a while because the principal defaults on a payment promise. Yang (2013) assumes that each agent has persistent private information about his ability but his setting differs from mine, as he considers relational contracting in a repeated matching market where matches are constantly reshuffled, and the informed party (i.e., the agent) is protected by limited liability.

In Section VIA, I reanalyze the model of Mailath and Samuelson (2001) in a relational environment and show that my main results apply to this setting. In Mailath and Samuelson, a seller can repeatedly sell a product to a long-lived buyer or to a sequence of short-lived buyers. In each period, the quality of the product is stochastic. The quality depends on the effort chosen by the seller in that period, and exerting effort is costly. Unlike Mailath and Samuelson (2001), I allow the seller to send a nonbinding message regarding the realized quality (see Jullien and Park 2014 for a similar approach). The quality announcement is cheap-talk, like a principal's bonus announcement, but truthful announcements may be enforceable, and the seller may be induced to exert effort perpetually if the future profit from the relationship is high enough.[11] If the seller types differ in their discount factors, then my main results hold, whereas separation is not generally possible if the private information relates to another parameter, such as the seller's cost of effort or ability.[12]

II. The Model

Two risk-neutral parties, a principal (she) and an 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 $t \geq 0$, the principal makes a contract offer to the agent. The agent either accepts this offer or rejects it: $d_t \in \{0, 1\}$ denotes the agent's decision, where $d_t = 1$ if the agent accepts the offer, and $d_t = 0$ otherwise. If the agent accepts the offer, then he chooses effort $e_t \in [0, \overline{e}]$ and incurs a cost $c(e_t)$, where c is strictly increasing, differentiable, and convex with $c(0) = 0$ and $c'(\overline{e}) = \infty$. The agent's effort e_t generates the output $y_t = y(e_t)$, where y is strictly increasing, differentiable, and concave. The term $s(e) \equiv y(e) - c(e)$ represents the expected surplus given the effort level e. The output is observed by both the principal and the agent but cannot be verified by a third party.[13] The contract offer at the beginning of period t consists of a fixed wage w_t and a bonus transfer b_t contingent on performance. This contract offer is denoted by $C_t = \{w_t, b_t\}$. The fixed wage w_t is legally enforceable, whereas the bonus payment b_t is not. After the output realization in period t, the party that is responsible for making the bonus payment b_t decides whether or not to honor the payment. If $b_t > 0$, then the decision belongs to the principal, whereas the agent makes the decision if $b_t < 0$. Total payment from the principal to the agent is denoted by P_t, where $P_t = w_t + b_t$ if the promised payment is honored, and $P_t = w_t$ otherwise. Thus, the agent's per period payoff is $P_t - c(e_t)$ and the principal's is $y_t - P_t$.

If the agent rejects the principal's offer, both parties receive their outside options in the current period: $\overline{\pi}$ for the principal and $\overline{u}$ for the agent. There exists an effort level e, such that $s(e) > \overline{\pi} + \overline{u} \geq s(0)$.

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

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

[11] Such cheap talk allows for equilibria in which the seller can be motivated to exert effort even in a symmetric information setting, unlike the case in Mailath and Samuelson (2001). A similar idea still applies if the seller observes the realized quality only after the sale. If the value of the repeated relationship is sufficiently high, the seller will agree to give a refund to the buyer after observing that the sold product is of low quality.

[12] My results also extend to the case in which sellers enter and exit the market in every period in a stochastic fashion, and names are tradable assets, as in Tadelis (1999) and Mailath and Samuelson (2001).

[13] Since output is a perfect measure of effort, whether or not effort is observable is inconsequential. Therefore, I assume without loss of generality that effort is unobservable.

and

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

Let μ_0 denote the prior probability that the principal is a high type. Within period $t \geq 0$, the agent updates his beliefs twice: first, after C_t (but before P_t) is observed, and second, after P_t is observed. The term μ_t^1 denotes the posterior belief of the agent at t after only C_t is observed, and μ_t^2 denotes the posterior belief of the agent at t after P_t is observed.

A. *Equilibrium Concept*

Let $h_t = (C_{t-1}, d_{t-1}, y_{t-1}, P_{t-1})$ denote the public outcome at the end of period $t-1$, and let $h^t = (h_0, \ldots, h_t) \in \mathcal{H}^t$ denote the history up to the beginning of t, where $\mathcal{H}^t$ represents the set of all possible h^t realizations with $h^0 = \mathcal{H}^0 = \varnothing$. A relational contract describes a complete plan for the relationship (Levin 2003). That is, for every $h^t \in \mathcal{H}^t$, a relational contract must specify (i) the contract that a principal of type $\theta \in \{h, l\}$ offers at t; (ii) the posterior belief μ_t^1 given the contract offer at t; (iii) whether the agent accepts or rejects the offer; (iv) the effort choice if the agent accepts the contract offer; (v) the bonus payment decision given the output realization; and (vi) the posterior belief μ_t^2 given the total payment P_t. Such a contract is self-enforcing if it describes a perfect public Bayesian equilibrium (PPBE) of the repeated game. A PPBE is a set of public strategies and posterior beliefs, such that strategies form a Bayesian Nash equilibrium in every continuation game given the posterior beliefs, and beliefs are updated according to Bayes' rule whenever possible. A public strategy depends only on the publicly observed history of play and the player's payoff-relevant private information. More specifically, the agent conditions his strategy only on the public history, whereas the principal conditions her strategy on her discount factor and the public history.

First, I study separating equilibria. A separating equilibrium is a PPBE in which the agent's belief about the principal's type is degenerate after the first contract offer at $t = 0$ as well as any equilibrium path history; that is, $\mu_t^i \in \{0, 1\}$, $i \in \{1, 2\}$ for every $t \geq 0$ on the equilibrium path. Second, I analyze pooling contracts.[14] Third, I analyze hybrid equilibria, in which both types start by offering the same contract, and separation is induced through default (or the absence thereof).[15] Finally, I refer to an equilibrium as "dishonest" if the equilibrium path involves breach of contract with a positive probability and as "honest" otherwise. Thus, separating and pooling equilibria are honest, whereas hybrid equilibria are dishonest.

[14] There can be equilibria, which are separating after an initial pooling phase during which both types' contract offers and actions are identical. Eventually, such contracts will prove to be irrelevant for my findings because, as Proposition 7 shows, pooling contracts are dominated contracts, which implies that a contract that is separating after an initial pooling phase is also dominated (see the discussion in Section IV).

[15] The only remaining type of equilibrium is an equilibrium that is hybrid with probability $\alpha \geq 0$, separating with probability $\beta \geq 0$, and pooling with probability $(1 - \alpha - \beta)$ at $t = 0$. See Footnote 27 for details.

B. *Symmetric Information Model*

First, I analyze the benchmark setting in which the agent knows the discount factor of the principal. The analysis of the symmetric information model consists of two cases: $\delta < \delta_\theta$ and $\delta > \delta_\theta$. The case with $\delta = \delta_\theta$ is analyzed in Levin (2003).

I define an optimal contract as the contract that maximizes the profit of the principal and is not pareto-dominated by another contract. Lemma 1 implies that the optimal symmetric information contract of a type-θ principal is unique if $\delta \neq \delta_\theta$. In particular, if $\delta < \delta_\theta$, then the optimal contract is unique and stationary; i.e., $e_t = e$, $b_t = b$, and $w_t = w$ for all t. Thus, there exists no optimal contract that is nonstationary on the equilibrium path. These contrast with Levin (2003). In Levin, there are both stationary and nonstationary optimal contracts. It turns out that the indeterminacy of the optimal contract stems from the knife-edge case in which the principal and the agent have the same discount factor.[16]

LEMMA 1: *In the optimal symmetric information contract with $\delta < \delta_\theta$, $u_t = \frac{\overline{u}}{1-\delta}$ for every $t \geq 0$ on the equilibrium path. In the optimal contract with $\delta > \delta_\theta$, $\pi_{\theta,t} = \frac{\overline{\pi}}{1-\delta}$ for every $t \geq 1$ on the equilibrium path.*

Lemma 1 implies that if $\delta < \delta_\theta$, then variations in continuation payoffs are never used to discipline the agent in the optimal contract. As a result, it is optimal for the principal to offer the same contract to the agent in every period: she offers a fixed wage w and a bonus reward $b > 0$ that is contingent on output y. The respective lifetime payoffs for the type-θ principal and the agent are given by

$$\pi_\theta = \frac{y - w - b}{1 - \delta_\theta} = \frac{s - \overline{u}}{1 - \delta_\theta},$$

and

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

In a self-enforcing contract, the principal does not default on a bonus promise (obviously, the optimal contract must be self-enforcing). So, there is no loss in assuming that default is punished by ending the relationship, which is the worst possible punishment. Hence, I obtain the enforcement constraint for the type-θ principal:

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

The optimal symmetric information contract maximizes the expected surplus $s(e) = y(e) - c(e)$ subject to the incentive compatibility constraint for the

[16] My definition of an optimal contract is more restrictive than that of Levin. However, this is inconsequential for my claims. If $\delta = \delta_\theta$, then there are both stationary and nonstationary optimal contracts according to my definition, as in Levin.

agent's effort choice ($b \geq c(e)$ by Lemma 1), the agent's participation constraint ($w + b - c(e) \geq \overline{u}$ by Lemma 1), and the enforcement constraint for the type-θ principal. The terms e_θ and b_θ give the solution to the maximization problem, and the maximized surplus is denoted by s_θ. If δ_θ is too low, then no contract can be self-enforcing. If, however, δ_θ is sufficiently close to one, then even a relational contract achieves the first-best effort. Consistent with the previous literature on self-enforcing contracts, I focus on environments in which trade is feasible but the enforcement constraint is binding, and the first-best outcome cannot be attained. Therefore,

$$b_\theta = \frac{\delta_\theta}{1 - \delta_\theta}(s_\theta - \overline{u} - \overline{\pi}) = c(e_\theta)$$

in the optimal contract, and the fixed wage is given by $w_\theta = \overline{u}$. The contract $C_\theta = \{w_\theta, b_\theta\}$ implements e_θ and is called the optimal symmetric information contract of type θ. Note that since $\delta_h > \delta_l$, it follows that $b_h > b_l$, and $s_h > s_l$.

If $\delta > \delta_\theta$, then the optimal contract is independent of θ, unique, and nonstationary as a result of Lemma 1. However, surplus and effort are stationary. Moreover, the contract becomes fully stationary after the initial period. The principal receives a frontloaded transfer at $t = 0$ via w_0 (the fixed wage at $t = 0$), the agent gets the excess of the surplus over $\overline{\pi}$ in every $t \geq 1$ via w_t, and a bonus is never paid out; that is, $b = 0$ in equilibrium. Thus, in the optimal contract parties trade payoffs across time if the agent is more patient than the principal. To see why $b = 0$, first note that $b \leq 0$, since $\pi_{\theta,t} = \overline{\pi}/(1 - \delta_\theta)$ at every $t > 0$ by Lemma 1. But if it were the case that $b < 0$ in equilibrium, then the enforcement constraint for the agent would become

$$\frac{\delta}{1 - \delta}(s(e) - \overline{u} - \overline{\pi}) \geq c(e) - b,$$

rather than

$$\frac{\delta}{1 - \delta}(s(e) - \overline{u} - \overline{\pi}) \geq c(e).$$

As before, I focus on environments in which trade is feasible but the enforcement constraint is binding. Note that if $b < 0$, then the former inequality poses a tighter constraint and generates lower equilibrium surplus than the latter inequality. Therefore, the principal chooses $e \in (0,1)$ to maximize $s(e) = y(e) - c(e)$ subject to the latter enforcement constraint. The term e^* gives the solution to this maximization problem, and s^* denotes the maximized surplus. In the optimal contract, $w_0 = (y(e^*) - \overline{\pi}) - (s^* - \overline{u} - \overline{\pi})/(1 - \delta)$ and $w_t = w = y(e^*) - \overline{\pi}$ for every $t \geq 1$. The contracts $C_0 = \{w_0, 0\}$ and $C_t = C = \{w, 0\}$ at $t > 0$ implement e^* and denote the optimal symmetric information equilibrium with either principal type.

The symmetric information benchmark relates to the work by Lehrer and Pauzner (1999). Lehrer and Pauzner analyze a class of repeated games in which players have different discount factors. They show that players can mutually benefit from trading payoffs across time; i.e., it is efficient to reward the patient player later and the impatient player earlier. The outcome in Lemma 1 is consistent with their result.

III. Separating Contracts

A one-shot game has a separating equilibrium provided that it satisfies some monotonicity condition, such as the single-crossing property. In infinitely repeated games, it may not be straightforward to find a corresponding property, and the problem is further mitigated in games where incentives are only informal (as in my model) because different types would exhibit identical behavior if the game were played only once. This section shows that a separating equilibrium always exists, and that differential time preferences of types are essential for this result. I characterize the optimal separating contract in Section IIIA and show when it dominates other equilibria and when it is dominated in Section IV.

On a related note, several studies have documented that individuals may exhibit an aversion to disappointment and betrayal. Several others have explicitly modeled disappointment aversion and loss aversion, and the insights from these studies have also been incorporated into contract theory.[17] I discuss the implications of a similar approach for my model in Section VIC.

I focus on the case in which $\delta < \delta_l$ throughout the remainder of the paper. This is due to the following:

- If $\delta > \delta_\theta$, then the optimal symmetric information contract of the type-θ principal depends on δ but not on δ_θ, as discussed in the previous section. Hence, if $\delta \geq \delta_h$, then private information does not distort incentives, and symmetric information contracts can still be implemented. I thus do not pursue this case further.
- If $\delta_h > \delta \geq \delta_l$, then all the forthcoming results that I derive under the assumption that $\delta < \delta_l$ are still valid.

The previous section has shown that the trade between a high-type principal and the agent generates a surplus of s_h via the optimal symmetric information contract $C_h = \{w_h, b_h\}$ if the principal's discount factor is common knowledge. However, C_h is not a credible offer if the principal is privately informed about her discount factor because a low-type principal would like to imitate a high type offering C_h. To see why, note that $b_h > b_l$ and $s_h > s_l$ since $\delta_h > \delta_l$. Thus, a low-type principal who successfully imitates the high type can obtain at least $(s_h - \overline{u})/(1-\delta_l)$, which is obviously greater than her payoff from C_l. Moreover, the low type chooses to default immediately because

$$b_h > \frac{\delta_l}{1-\delta_l}(s_h - \overline{u} - \overline{\pi}),$$

[17] See the references in Footnotes 2 and 7 and Section VIC.

which further increases her imitation payoff. It follows that the high-type principal must use a credible signaling mechanism in order to separate herself from the low type.

The characterization of separating equilibria involves two sets of incentive compatibility constraints. One set of constraints ensures that the low type is deterred from imitating the high type, whereas the second set ensures that the high type is willing to separate and signal her type. In order to simplify the exposition of the incentive constraints below, I assume that the low type offers her optimal symmetric information contract C_l in every period and obtains $(s_l - \bar{u})/(1-\delta_l)$ in every separating equilibrium. This assumption is without loss of generality for Proposition 2 and Proposition 3. Moreover, optimality dictates that the low type will offer C_l in every period of a separating equilibrium (see the discussion in Section IIIA).[18]

The low-type principal who decides to imitate the high type deviates from the separating contract at some t, either refusing to pay $b_t > 0$ or offering a contract different than the equilibrium prescription. I assume that $b_t \geq 0$ in order to simplify the exposition of the incentive compatibility constraints. None of the upcoming results rely on this simplification. Moreover, $b_t > 0$ must hold in the optimal separating contract, as implied by Lemma 4. The sequence $\{w_t, b_t\}_{t=0}^{\infty}$ represents a separating equilibrium provided that it satisfies conditions (1)–(4); if these conditions are satisfied, then the low-type principal is deterred from imitation and offers C_l in every period, the high-type principal is willing to signal her type, and $\{w_t, b_t\}_{t=0}^{\infty}$ is enforceable for the high type.[19] Conditions (1)–(4) are as follows:

$$(1) \qquad \forall t \geq 0\colon \sum_{k=0}^{t-1} \delta_l^k\big(y(e_k) - W_k\big) + \delta_l^t\big(y(e_t) - w_t\big) + \frac{\delta_l^{t+1}}{1-\delta_l}\bar{\pi} \leq \frac{s_l - \bar{u}}{1-\delta_l},$$

$$(2) \quad \sum_{k=0}^{\infty} \delta_h^k\big(y(e_k) - W_k\big) \geq \frac{s_l - \bar{u}}{1-\delta_h} \quad \text{and} \quad \forall t \geq 0\colon \sum_{k=t}^{\infty} \delta_h^{k-t}\big(y(e_k) - W_k\big) \geq \frac{\bar{\pi}}{1-\delta_h},^{20}$$

where $W_k = w_k + b_k$, and e_k is implemented by the contract $\{w_k, b_k\}$ and satisfies the agent's incentive compatibility constraint for effort as well as his participation constraint at every $k \geq 0$:

$$(3) \ \ \forall k \geq 0\colon \ W_k - c(e_k) + \sum_{t=k+1}^{\infty} \delta^{t-k}\big(W_t - c(e_t)\big) \geq \max\left\{w_k + \frac{\delta}{1-\delta}\bar{u}, \frac{\bar{u}}{1-\delta}\right\},$$

[18] I define the optimal contract as the equilibrium contract that maximizes a weighted average of the two principal types' equilibrium payoffs. Section IIIA provides a detailed description.

[19] Assuming that there is no default in a separating equilibrium and that the worst punishment is imposed in case of a default is without loss of generality for the upcoming results. Of course, the optimal separating contract cannot involve default on the equilibrium path.

[20] The second incentive compatibility constraint for the high-type principal can also be written as

$$\forall t \geq 0\colon \sum_{k=t}^{\infty} \delta_h^{k-t}\big(y(e_k) - W_k\big) \geq \frac{s_l - \bar{u}}{1-\delta_h}$$

if she wants to stop costly signaling at t and offer C_l from then on—assuming that she has never defaulted on a bonus payment and a future payoff greater than $\bar{u}/(1-\delta)$ has not been promised to the agent. It might be reasonable to assume that the agent does not punish this because C_l is enforceable with both types. However, this incentive constraint formulation does not affect any of my results.

where $W_k = w_k + b_k$, as defined above. Finally, the high type's enforcement constraint holds at every $t \geq 0$ if the following is satisfied:

$$(4) \qquad \forall t \geq 0\colon\ b_t + \frac{\delta_h}{1-\delta_h}\overline{\pi} \leq \sum_{k=t+1}^{\infty} \delta_h^{k-t}\big(y(e_k) - w_k - b_k\big).$$

Since C_h is the optimal symmetric information contract of the high-type principal, a contract offer different than C_h generates a surplus that is lower than s_h and reduces the profit of the principal. Thus, the high type's behavior is "distorted" (put differently, the high type engages in costly signaling) in a separating equilibrium whenever her contract offer differs from C_h. It is sometimes *necessary* that costly signaling lasts multiple periods because, as Proposition 2 shows, separation may be impossible otherwise. To be more precise, there exists a nontrivial set of parameters, such that if the high type offers a contract different from C_h only at $t = 0$, then separation is impossible.

PROPOSITION 2: *Separation is not generally possible if the high type's behavior is distorted only at* $t = 0$. *To be more precise, there exists a* $\Lambda > 0$, *such that if* $\delta_h - \delta_l < \Lambda$, *and* $\delta_h > \delta_l > \delta$, (*or, if* $\delta_h - \delta < \Lambda$, *and* $\delta_h > \delta \geq \delta_l$), *then there exists no separating equilibrium in which the high type's behavior is distorted in only the initial period.*

The formal proof of Proposition 2 is relegated to the Appendix. Note that Λ need not be small. As the discussion below shows, if, for example, $\delta_l = 2/3$ and $\overline{\pi} = \overline{u} = 0$, then δ_h values that satisfy $3s_l > s_h$ also satisfy $\delta_h - \delta_l < \Lambda$.

I will now explain the intuition behind this proposition, since it is instrumental to understanding why there always exists a separating equilibrium if the private information is on time preferences, and why a separating equilibrium may not exist with other types of private information. To see the intuition behind Proposition 2, first note that if the high type's behavior is not distorted *at all*, then the benefit of separation for the high-type principal is strictly lower than the benefit of imitation for the low type. For simplicity, assume that $\overline{\pi} = \overline{u} = 0$, and note that the benefit of separation for the high type equals

$$\beta_h = \frac{s_h}{1-\delta_h} - \frac{s_l}{1-\delta_h} = \frac{s_h - s_l}{1-\delta_h},$$

whereas the benefit of imitation for the low type equals

$$\beta_l = s_h + b_h - \frac{s_l}{1-\delta_l} = \frac{s_h}{1-\delta_h} - \frac{s_l}{1-\delta_l}.$$

Since $\delta_h > \delta_l$, it follows that $\beta_l > \beta_h$. The left-hand side of the second equality above follows because an imitator promises b_h and obtains output $y(e_h)$ but defaults on b_h subsequently. The right-hand side follows because the enforcement constraint of the high type is binding in the contract C_h, and thus $b_h = \delta_h \frac{s_h}{1-\delta_h}$ with $\overline{\pi} = \overline{u} = 0$.

Now, assume that the behavior of the high-type principal is distorted *only* at $t = 0$ in a separating equilibrium. This implies that a low-type principal who chooses to imitate the high type optimally defaults either at $t = 0$ or at $t = 1$.[21] A key observation is that the equilibrium benefit of separation for the high type must exceed the equilibrium benefit of imitation for the low type in a separating contract. For simplicity, I assume that the imitator chooses to imitate the high type fully at $t = 0$ honoring the bonus promise and offers C_h at $t = 1$.[22] Thus, $t = 0$ can be ignored in the comparison of the equilibrium benefit of separation and the equilibrium benefit of imitation because there is no discounting at $t = 0$, and the payoff prospects are identical for both types. From $t = 1$ onward, the low type obtains the benefit of imitation β_l as described above, and the high type obtains the benefit of separation β_h. Recall that $\beta_l > \beta_h$. The "discounted" benefit of imitation $\delta_l \beta_l$ is also strictly higher than the "discounted" benefit of separation $\delta_h \beta_h$ if $s_l/(1 - \delta_l) > s_h$ holds. But this condition can easily hold if δ_l and δ_h are not far from each other (as an example, if $\delta_l = 2/3$, then the condition is satisfied with δ_h values, such that $3s_l > s_h$). In that case, separating contracts in which the high type's behavior is distorted at only $t = 0$ requires excessively costly signaling (because the low type benefits more from imitation than does the high type from separation), and therefore, the high type prefers imitating the low type and offering C_l instead.

Importantly, this problem is fully resolved if the high type continues to distort her behavior at $t > 0$ and delays C_h as well as other high-surplus contracts to a sufficiently distant future. Such a delay is effective in separating the principal types because they differ in their patience. Note that Proposition 2 is not specific to the case where the principal's private information relates to time preferences. Similar or stronger versions of this result obtain with other types of private information, such as productivity or outside option; that is, separation is not generally possible if the high type's behavior is distorted only at $t = 0$. But, as discussed in more detail below, delaying high cooperation is typically ineffective if the principal types have the same discount factor, and, therefore, a result similar to Proposition 2 has general implications when the private information relates to another parameter; in particular, a separating equilibrium may not exist.

Delaying C_h and other high-surplus contracts enables separation in my model because the cost of waiting is "monotonic" in type due to the difference in the time preferences of types. More precisely, the equilibrium benefit of separation for the high type exceeds the equilibrium benefit of imitation for the low type with a sufficiently long delay because $\delta_h > \delta_l$, and thus, $\delta_h^T \beta_h > \delta_l^T \beta_l$ must hold with a sufficiently large (but finite) T. That is,

$$(5) \qquad \delta_h^T \left(\frac{s_h - s_l}{1 - \delta_h} \right) > \delta_l^T \left(\frac{s_h}{1 - \delta_h} - \frac{s_l}{1 - \delta_l} \right).$$

[21] Since the bonus equals b_h at every $t \geq 1$ by hypothesis, defaulting for the first time at $t > 1$ is strictly dominated.

[22] The argument becomes even stronger if the low type strictly prefers defaulting at $t = 0$.

As a result, the high-type principal who delays C_h until period T, limits the surplus of the relationship to a low level (for example, s_l) in the first $T-1$ periods, and spends the amount

$$\delta_l^T \beta_l = \delta_l^T \left(\frac{s_h}{1-\delta_h} - \frac{s_l}{1-\delta_l} \right)$$

(being precisely the discounted benefit of imitation for the low type) on costly signaling at $t = 0$ and is able to separate herself from the low type and strictly prefers doing so. This is the key insight behind the following result.[23]

PROPOSITION 3: *There always exists a separating equilibrium.*

PROOF:

I outline a separating mechanism based on the intuition above. The low-type principal's separating equilibrium contract is always C_l. At $t = 0$, the high-type principal offers the contract $\{w_l + \Delta, b_l\}$, where w_l is the fixed wage and b_l is the bonus payment contingent on output $y(e_l)$ as specified in C_l. Thus, the high type's contract at $t = 0$ differs from C_l in that the fixed wage offer exceeds w_l by Δ. The amount Δ represents the initial cost of signaling and will be determined below endogenously (Δ can be interpreted as a lump-sum signing bonus for the agent). At every $t \in \{1, \ldots, T-1\}$, the high type offers $C_l = \{w_l, b_l\}$, where T is to be determined endogenously, just like Δ.[24] From period T onward, the high type offers C_h. In other words, the high type delays C_h until period T and limits the surplus of the relationship to s_l in the first $T-1$ periods. A low-type principal who chooses to imitate the high type will honor b_l at every $t < T$ and optimally default on b_h in period $t = T$, by construction. Therefore, it suffices to focus on the following incentive compatibility constraint for the low type:

$$\frac{s_l - \overline{u}}{1-\delta_l} \geq y(e_l) - (w_l + \Delta) - b_l + \sum_{t=1}^{T-1} \delta_l^t (s_l - \overline{u}) + \delta_l^T V_l, \tag{6}$$

where V_l denotes the undiscounted imitation payoff of the low type from T onward; i.e., $V_l = s_h - \overline{u} + b_h + \frac{\delta_l}{1-\delta_l}\overline{\pi}$. The only relevant incentive compatibility constraint for the high type is

$$y(e_l) - (w_l + \Delta) - b_l + \sum_{t=1}^{T-1} \delta_h^t (s_l - \overline{u}) + \frac{\delta_h^T}{1-\delta_h}(s_h - \overline{u}) \geq \frac{s_l - \overline{u}}{1-\delta_h}, \tag{7}$$

because the high type would never deviate from costly signaling at $t > 0$, and prefers honoring the bonus promise at every $t \geq 0$ by construction. Thus, (6) and (7) are necessary and sufficient conditions for this construction to result in a separating

[23] I allow out-of-equilibrium degenerate beliefs to change in order to allow for a delay in the contract offer C_h in separating equilibria. In other words, even if beliefs put zero probability on a type of principal at some $t \geq 0$, off-the-equilibrium path beliefs can still attach positive probability on that type after t. As a result, it is possible for the high type to delay a C_h offer for multiple periods and deter the low type from imitation.

[24] If it turns out that $T = 1$, then the high type offers C_h from period 1 onward.

equilibrium. Period T is determined in a way that the discounted benefit of separation for the high type exceeds the discounted benefit of imitation for the low type, just as in (5). Let T be the lowest t, such that

$$\delta_h^t\left(\frac{s_h - \overline{u}}{1-\delta_h} - \frac{s_l - \overline{u}}{1-\delta_h}\right) > \delta_l^t\left(V_l - \frac{s_l - \overline{u}}{1-\delta_l}\right)$$

holds. Once T is found, Δ is determined in a way that the low type is deterred from imitation in the least costly way; that is, Δ is equal to $\delta_l^T\big(V_l - (s_l - \overline{u})/(1-\delta_l)\big)$, which is the discounted benefit of imitation for the low type. Using this, it can easily be verified that (6) holds and, hence, the low type is indifferent between mimicking and revealing her type. Finally, it can be checked that (7) holds with strict inequality, and thus, the high type strictly prefers separating. As a result, a separating equilibrium always exists.[25] ∎

The separation mechanism described above is fairly simple as it involves only two costly signals, Δ and T. The amount Δ is the initial cost of separation (to reiterate, this can be interpreted as a signing bonus paid upfront to the agent), and T refers to the delay in offering C_h. Although this separating equilibrium need not be optimal, its simplicity is instructive. In general, the key features of any separating equilibrium are an initial transfer to the agent (or some form of money-burning); and delaying high-surplus contracts to a sufficiently distant future. These two deter the low-type principal from imitating the high type in a cost-efficient way due to the following:

- Delaying high-surplus contracts implies that imitating the high type and defaulting early on does not pay off, given the initial cost of signaling Δ.
- Waiting for high-surplus contracts (such as C_h) is costlier for the low type—recall that the cost of delay is *monotonic* in type, which allows me to obtain an inequality such as (5). If the delay is sufficiently long, then the benefit of separation for the high type exceeds the benefit of imitation for the low type in discounted terms. The initial costly signaling Δ is then simply set equal to the discounted benefit of imitation for the low type.

Thus, a separating equilibrium always exists: the low type is deterred from imitation, and the high type strictly prefers revealing her type. This simple idea can be applied to any infinitely repeated game where the private information relates to time preferences.

To reiterate, differential time preferences of types is essential for a general separation result. As I stated before, similar or stronger versions of Proposition 2 hold with other types of private information; that is, separation is not generally possible

[25] If $\delta_l < \delta < \delta_h$, then only a slight modification is needed in the construction above. The idea of the proof remains the same but the incentive compatibility constraints end up being slightly different than those presented above. This is because if $\delta_l < \delta$, then the symmetric information contract of the low type depends on δ, not on δ_l. However, this difference is inessential for the construction above.

if the high type's behavior is distorted only at $t = 0$. But such a result has broader implications with other types of private information because if the principal types have the same discount factor, then delaying better contracts is not effective, unlike in my model. For example, it is not possible to come up with an inequality similar to (5) and ensure that there exists an equilibrium in which the benefit of separation for the high type exceeds the benefit of imitation for the low type. As a result, a separating equilibrium may not exist with other types of private information. As Halac (2012) has shown, there exists no separating contract if the principal types differ in their outside options. Consider next the case where the high type and the low-type principals differ in (and are privately informed regarding) their productivity and are identical in other respects. Let y_θ represent the production function of the type-θ principal, where $\theta \in \{h, l\}$. If the high-type principal has a productivity advantage of $\eta > 0$ over the low type, such that $\eta \equiv y_h(e) - y_l(e) > 0$ given effort e, then a separating contract does not exist. Furthermore, it is not generally sufficient for separation if $y_h(e) > y_l(e)$ and $y_h(e) - y_l(e)$ is increasing in e. I formalize these claims in the online Appendix B. For an intuition, consider the simple case where $y_h(e) - y_l(e) = \eta > 0$ for every e. A separating equilibrium does not exist in this case because a strong version of Proposition 2 obtains: separation is impossible if the high type's behavior is distorted only at $t = 0$. Moreover, delaying high-surplus contracts is ineffective. It follows that a separating equilibrium does not exist, and there is no analogue of Proposition 3. These also hold in Halac (2012): separation is impossible if the high type's behavior is distorted only at $t = 0$, and delay is also ineffective. As a result, there exists no separating equilibrium.

In Section VIA, I show that my analysis and results are applicable to a reputation model à la Mailath and Samuelson (2001). In particular, I show that a separating equilibrium always exists if the seller's private information is about her time preferences, whereas a separating contract does not generally exist if the seller is privately informed regarding another parameter, such as her ability or cost of effort.

The Optimal Separating Contract.—I define the optimal contract as the equilibrium contract that maximizes a weighted average of the two principal type's equilibrium payoffs, where the weight for the high type's payoff, denoted by γ, is arbitrary, and $\gamma \in (0, 1)$. In the optimal separating contract, the low-type principal always offers C_l. A separating equilibrium in which the low type offers a different contract is suboptimal since changing the low type's contract offer to C_l does not affect the high type at the optimum (the construction in the proof of Proposition 3 shows that the high type's separation payoff must exceed $(s_l - \overline{u})/(1 - \delta_h)$ at the optimum) but improves the low type's separation payoff. Thus, the optimal separating contract boils down to the contract that maximizes the payoff of the high-type principal among all separating contracts in which the low type's contract offer is C_l. Formally, the optimal separating contract $\{w_t, b_t\}_{t=0}^{\infty}$ maximizes $\sum_{t=0}^{\infty} \delta_h^t\big(y(e_t) - w_t - b_t\big)$ subject to (1)–(4).[26]

[26] The optimal contract is also pareto-optimal in the usual sense: if there are two contracts that maximize this weighted average, then the one that gives the agent a higher payoff is optimal.

The first step toward characterizing the optimal separating contract is to show that the high-type principal does not use future rewards (i.e., variations in continuation payoffs) as a discipline device. The reason for this is as follows. In order to motivate the agent, the high-type principal can use either a bonus payment (i.e., an immediate reward) or a future reward (or both). However, using a future reward scheme is costlier than a bonus scheme because the agent is less patient than the high-type principal, and a future reward is not legally enforceable like a bonus reward. Thus, the high-type principal strictly prefers using a bonus scheme in order to motivate the agent, and the agent is rewarded immediately conditional on performance.

LEMMA 4: *In the optimal separating contract of the high type,* $u_t = \frac{\overline{u}}{1-\delta}$ *for every* $t > 0$.

The next lemma shows that costly signaling stops at a finite period in the optimal separating contract of the high type; that is, the high type offers C_h at all sufficiently large t. This is because the optimal separating contract is such that a low-type principal who were to imitate the high type would strictly prefer defaulting at some finite period T. Otherwise, the high type's optimal separating contract is a contract that can be implemented with a low type in a symmetric information setting, which is a contradiction. This implies that the high type will start offering C_h from a sufficiently large $t \geq T$ onward, rather than distort her behavior indefinitely.

LEMMA 5: *Costly signaling stops at a finite date; i.e., the high type offers* C_h *at all sufficiently large* t.

Finally, I characterize the optimal separating contract. Proposition 6 shows that in the optimal separating contract of the high type, $b_t > b_l$ and $e_t > e_l$ at every $t \geq 0$. This implies that from the very beginning the relationship with a high-type principal generates a higher surplus than the low type's contract. In turn, the high type must pay a relatively high fixed wage at $t = 0$ so that the low type does not imitate. In other words, a high-type principal offers an initial lump-sum *signing bonus* and provides stronger performance incentives than a low-type principal in every period. Moreover, b_t is strictly increasing until it reaches b_h, and $b_t = c(e_t)$ at every $t \geq 0$ in the optimal contract. Thus, the strength of the performance incentives increases over time, and the percentage of the agent compensation that comes from the performance bonus increases progressively. As incentives become stronger, the surplus increases, and the relationship becomes more valuable over time. This "gradualism" result is reminiscent of the results obtained in Sobel (1985), Diamond (1989), Ghosh and Ray (1996), Kranton (1996), Watson (1999; 2002), and Halac (2012). However, I obtain gradualism in honest separating equilibria, which is a novel result. As it turns out, a gradualist pattern of cooperation is optimal even in settings where learning is immediate because separation is cheaper for the high type if costly signaling is spread over time in a gradualist manner.

PROPOSITION 6: *In the optimal separating contract of the high type,* $b_t < b_{t+1}$ *for every* $t \geq 0$ *until* $b_t = b_h$*, which takes place in finite time. Similarly, the effort* *schedule and the surplus are strictly increasing until they reach* e_h *and* s_h, *respectively. Moreover,* $b_t > b_l$ *and* $s_t > s_l$ *at every* $t \geq 0$. *Finally,* $w_0 > \overline{u}$ *and* $w_t = \overline{u}$ *at* $t > 0$.

The intuition for Proposition 6 is as follows. Offering a lump-sum signing bonus, delaying the contract offer C_h, and using relatively weak incentives at the early stages of the relationship are very effective as costly signals. These generate low surplus and low principal profits early in the relationship, and thus a low imitation payoff because the low type is relatively impatient. As the relationship builds up over time, incentives become stronger, and the profit level increases progressively. Such an arrangement is optimal because the initial signing bonus together with the gradually increasing profit levels deter an impatient low type from imitation in the least costly way.

To make the intuition regarding the monotonicity of the bonus schedule more transparent, consider a scenario in which $T = 2$ according to the construction that I used in order to prove Proposition 3. Since $T = 2$, the high-type principal offers b_l in the first two periods followed by b_h, thereafter. The corresponding effort levels are $e_0 = e_1 = e_l$, and $e_t = e_h$ for $t \geq 2$. The high type can do better than this contract by using a *strictly increasing* bonus schedule. She can (i) change the bonus at $t = 1$ so that $b_1' = b_l + \epsilon$ instead of b_l, where $\epsilon > 0$ is arbitrarily small; (ii) make $b_l + \epsilon$ contingent on $y(e_1')$, such that $c(e_1') = b_l + \epsilon$ (thus, $e_1' > e_l$); and (iii) increase the initial fixed-wage w_0 by $\delta_l[s(e_1') - s(e_l)]$. Since ϵ is arbitrarily small, a low type who were to imitate a high type would still optimally default at $T = 2$. Thus, the changes in e_1 and w_0 are such that the low type's imitation payoff is exactly the same as before, and thus, she is still deterred from imitation. Moreover, by virtue of her higher discount factor, the high type is strictly better off. This illustrates the key mechanism behind Proposition 6. If a separating contract involves bonuses such that $b_{t+1} \leq b_t < b_h$ for some t, then one can always construct a separating contract that strictly dominates it.

IV. Other Contracts

In this section, I investigate the remaining types of equilibria, namely, pooling equilibria and hybrid equilibria.[27] There is no information revelation in a pooling equilibrium. In a hybrid equilibrium, both types start by offering the same contract, and information is revealed over time through default (or the absence thereof). To be more precise, the high type honors implicit bonus or future contract promises in the hybrid contract, whereas the low type eventually reneges on a promise.

The optimal pooling equilibrium is stationary and implements C_l in every period. To see why, first note that in every period of a pooling equilibrium, the principal types must honor the bonus promise or default with identical probability so that

[27] As discussed in Footnote 15, there can be equilibria that are separating after an initial pooling phase, but such contracts are not relevant for my main results because they are inefficient, as implied by Proposition 7. There is one final type of contract that is possible: a contract that is hybrid with probability $\alpha \geq 0$, separating with probability $\beta \geq 0$, and pooling with probability $(1 - \alpha - \beta)$ at $t = 0$. However, none of the following results are affected by this possibility.

there is no revelation of information regarding the principal's type. However, a pooling equilibrium that involves default in some period is inefficient. Therefore, the principal always fulfills her promises in the optimal pooling contract, regardless of her type. Put differently, the optimal pooling contract is enforceable with either type in a symmetric information setting.[28] But a contract that is enforceable with the low-type principal is surely enforceable with the high type since the latter is more patient. Thus, the optimal pooling equilibrium implements C_l in every period because this is the best contract which is enforceable with the low type.

In the optimal separating equilibrium, the low-type principal always offers C_l, whereas the high-type principal strictly prefers revealing her type to imitating the low type, as the equilibrium construction in the proof of Proposition 3 indicates. Hence, the following result obtains.

PROPOSITION 7: *The optimal separating equilibrium strictly dominates pooling equilibria.*

Next, I analyze the optimal hybrid contract. One factor that complicates the analysis is that the low-type principal defaults on the equilibrium path, unlike the case in the separating contract. Thus, the form of the punishment is an important feature of the equilibrium. One useful observation is that after the posterior belief becomes degenerate, imposing the worst punishment on a principal who defaults is without loss of generality. Put differently, if there exists a hybrid equilibrium, such that a default that takes place after the posterior belief has become degenerate is not punished in the worst possible way, then there exists a payoff-equivalent hybrid equilibrium that imposes the worst punishment instead.

While it is possible to obtain some intermediate results regarding the optimal hybrid contract without making any assumption about the form of the punishment, my main interest lies in comparing different types of equilibria, and this requires some restriction on punishments, given the complexity of the problem. I impose Assumption 1 (A1) in order to be able to obtain a precise characterization of the optimal hybrid contract and compare efficiency under different types of contracts. After I present Assumption 1 (A1), I discuss in detail this assumption and the conditions under which it is optimal (or at least plausible), and relate it to the recent literature.

ASSUMPTION 1 (A1): *If the principal defaults on a payment promise*, *then the worst punishment is imposed and the agent terminates the relationship.*

As discussed above, A1 is without loss of generality after the posterior belief becomes degenerate. With nondegenerate posteriors, A1 may in principal be

[28] In a pooling equilibrium that involves default (that is, both types default in some period with identical probability), incentive provision must involve a default risk premium so that the agent is willing to participate and exert effort despite the default risk. But this is entirely wasteful, as it would be in a symmetric information setting: given an equilibrium in which the two types default in some period with the same probability, there exists a strictly better pooling equilibrium in which both types honor their promises with probability one, the output implemented is the same as in the original pooling contract, and the agent compensation is strictly lower because there is no need to pay the default risk premium.

restrictive. However, there are various conditions under which A1 is optimal even with nondegenerate beliefs. For example, if C_h is sufficiently far from the first-best contract (that is, if $y'(e_h)/c'(e_h)$ is sufficiently larger than 1), then it is optimal to impose the worst punishment following a default. Imposing the worst punishment is also optimal if δ_l and δ_h are not far from each other or if μ_0 is sufficiently high.[29] A1$'$ below is one possible alternative for A1.

ASSUMPTION 1$'$ (A1$'$): *Parameters are such that $y'(e_h)/c'(e_h)$ is sufficiently larger than* 1, *and therefore, imposing the worst punishment following a default is optimal.*

Apart from optimality, A1 may depict those cases where defaulting on payment promises may require the firm to shut down the business or change its name and location. This is reportedly a common form of wage violation.[30] Moreover, adopting an assumption such as A1 or A1$'$ has a precedent in the literature. Halac (2012) focuses on parameters, such that imposing the worst punishment after default is optimal, and this allows for characterizing the evolution of equilibrium bonuses and effort levels in the optimal hybrid contract.

While the optimal separating equilibrium strictly dominates the optimal pooling equilibrium, the comparison of the optimal separating equilibrium and the optimal hybrid equilibrium depends on the magnitude of μ_0 and γ, as Proposition 8 shows.

PROPOSITION 8: *Assume* A1 *or* A1$'$, *and fix* $\gamma \in (0,1)$*:*

(*i*) *There exists an* $\varepsilon > 0$, *such that if* $\mu_0 \leq \frac{b_l}{b_h} + \varepsilon$, *then hybrid contracts are strictly dominated, whereas if* μ_0 *is sufficiently close to* 1, *then separating contracts are strictly dominated. More generally, given* $\gamma \in (0,1)$ *there exists a* $\mu_\gamma > \frac{b_l}{b_h}$, *such that the optimal contract is separating if* $\mu_0 \leq \mu_\gamma$, *and hybrid otherwise. Moreover,* μ_γ *is increasing in* γ.[31]

(ii) *If the optimal contract is hybrid, then it fully reveals information in finite time and exhibits gradualism; that is, there exists a* $T < \infty$ *such that as long as the principal honors the promised payments,* b_t *and* s_t *are strictly increasing until period* T *when they reach* b_h *and* s_h, *respectively.*

The proof of Proposition 8 is provided in the online Appendix A. For an intuition of part (i), first note that if the prior probability μ_0 is sufficiently high and close to 1, then the optimal contract is hybrid regardless of γ because the chances that the

[29] I prove in online Appendix A that if $y'(e)/c'(e)$ is sufficiently far from 1, then the optimal contract implements the worst punishment following default. I also prove that the same is true if δ_l and δ_h are not far from each other or if μ_0 is sufficiently high. Note that, in order to obtain sharper results regarding the optimality of A1, I disregard certain hybrid equilibria that are strictly dominated regardless of the form of the punishment; this is without loss of generality as it has no consequence for the results stated in Proposition 8. Details are relegated to online Appendix A (see in particular the discussion preceding Claim 6).

[30] http://www.labor.ucla.edu/wage-theft/.

[31] Note that it is possible to fix $\gamma = \mu_0$. Then, the result would be as follows: there exists a $\tilde{\mu}_0 > b_l/b_h$ such that the optimal contract is hybrid if $\mu_0 > \tilde{\mu}_0$ and separating otherwise.

principal is a low type are sufficiently small, and it is possible to construct a hybrid equilibrium that approximates C_h at $t = 0$ and is identical to C_h at every $t > 0$, provided that the bonus promise is honored at every $\tau < t$. Such an equilibrium makes both types strictly better off than they would be in a separating equilibrium. If however μ_0 is sufficiently low (for example, lower than b_l / b_h), then hybrid equilibria are wasteful regardless of γ because equilibrium information revelation involves a relatively high probability of default for several periods, and thus, incentivizing the agent is too costly relative to other types of equilibria. To be more precise, incentive provision in a hybrid equilibrium requires insuring the agent against default to some extent so that the agent is still willing to participate and exert effort, and this requirement becomes too costly when the prior is relatively low (equivalently, the probability of default is relatively high). Note that the lower bound on μ_0 stated in Proposition 8 may be quite stringent. For example, if δ_l and δ_h are close, then b_l and b_h are also close, and hybrid equilibria are inefficient unless the prior μ_0 is close to 1.

If μ_0 is an intermediate value bounded away from both b_l / b_h and 1, then whether the optimal contract is hybrid or separating depends on both the magnitude of μ_0 and the welfare weights of the principal types. While the payoff of the optimal separating equilibrium does not depend on μ_0, the payoff of the optimal hybrid contract is strictly increasing in μ_0. Therefore, given fixed $\gamma \in (0, 1)$, there exists a unique threshold value of μ_0 such that the payoff of the optimal separating equilibrium and the payoff of the optimal hybrid equilibrium are identical. Hence, the optimal contract is hybrid if the prior exceeds this threshold and separating otherwise. This threshold value, denoted by μ_γ, is monotone increasing in γ since the high type prefers separating equilibria over hybrid equilibria for a wider range of μ_0 values than the low type.

I now focus on part (ii) and discuss the structure of the optimal hybrid equilibrium assuming that the optimal contract is hybrid. The low type is initially indifferent between defaulting and honoring the bonus promise but eventually defaults at some $t \geq 0$ with probability one. Once the principal defaults, the agent learns that the principal is a low type with probability one. As long as the principal keeps honoring her bonus promises, trust is gradually established, higher bonus promises become more credible, and the surplus of the relationship increases progressively until it becomes stationary at the high type's symmetric information surplus level.

V. Equilibrium Selection: Dynamic Intuitive Criterion

As discussed in the previous section, hybrid equilibria are dominated unless μ_0 is above a certain threshold. I now appeal to a dynamic version of the Intuitive Criterion and show that most hybrid equilibria are "unreasonable."

In what follows, let $\{C_t\}_{t=0}^{\infty}$ denote an equilibrium set of contracts. Assume that the principal deviates and announces an out-of-equilibrium set of contracts $\{D_t\}_{t \geq k}$ at an arbitrary period $k \geq 0$. Define $\{D_t\}_{t \geq k}$ to be *equilibrium-dominated* for type-θ if the equilibrium payoff of a type-θ principal is strictly higher than the highest possible payoff she gets if she honors all the payments until period k and deviates to $\{D_t\}_{t \geq k}$ (the highest possible payoff may of course involve default, for example reneging on a bonus promise of an arbitrary D_t).[32]

I define the *Dynamic Intuitive Criterion* as follows. An equilibrium $\{C_t\}_{t=0}^{\infty}$ fails to satisfy the Dynamic Intuitive Criterion (DIC) if there exists a $k \geq 0$ and an out-of-equilibrium set of contracts $\{D_t\}_{t \geq k}$ such that:

(i) $\{D_t\}_{t \geq k}$ is *equilibrium-dominated* for the low-type principal; and

(ii) $\{D_t\}_{t \geq k}$ is *enforceable* for the high type (i.e., the high type would never default) and *strictly profitable* if the agent best-responds according to the belief $\mu_t = 1$ at every $t \geq k$.

The DIC eliminates every separating equilibrium other than the optimal separating equilibrium, which is always robust to the DIC. The DIC also eliminates every pooling or hybrid equilibrium—with only one possible exception. There is only one hybrid equilibrium that *might* be immune to the DIC, the one in which period-0 hybrid contract $\{w_0, b_0\}$ is such that $w_0 = \overline{u}$,

$$b_0 = \delta_l\left(s_h - \overline{u} + b_h + \frac{\delta_l}{1 - \delta_l}\overline{\pi}\right),$$

the low type defaults with probability one at $t = 0$, and the high type offers C_h from $t = 1$ onward. Note that such an equilibrium can be robust to the DIC only if $\mu_0 > b_l/b_0$. Also, note that this type of lower bound is more stringent than b_l/b_h, the lower bound discussed in Proposition 8, because $b_0 < b_h$. However, even when $\mu_0 > b_l/b_0$ holds, it is still not obvious whether this equilibrium is robust to the DIC: the high type may find a deviation $\{D_t\}_{t \geq 0}$ that is equilibrium dominated for the low type and makes her strictly better off unless μ_0 is sufficiently close to 1. An extensive discussion regarding the DIC and the results of this part are provided in the online Appendix C.

VI. Extensions

A. *Buyer-Seller Relationships*

In this section, I analyze an infinitely repeated game à la Mailath and Samuelson (2001) and show that my main results extend to this setting. In this model, a long-lived seller sells one product to a long-lived buyer or to a sequence of short-lived buyers in each period $t = 0, 1, \ldots$. The quality of the product for sale, denoted by q, is either low (i.e., $q = q_L$), medium (i.e., $q = q_M$), or high (i.e., $q = q_H$).[33] The terms u_H, u_M, and u_L denote the buyers utility from consuming a product with $q = q_H$, $q = q_M$, and $q = q_L$, respectively. I assume that $u_H > u_M > u_L \geq 0$.

[32] For $\{D_t\}_{t \geq k}$ to be meaningful as out-of-equilibrium behavior, the game must have proceeded to period $k \geq 0$ without default. Otherwise, the agent knows with certainty that the principal is a low type.

[33] I extend the setting in Mailath and Samuelson (2001) and allow for more than two quality levels. I also assume that both seller types can be strategic, as in my main model.

In each period, q is stochastically determined according to the effort the seller chooses. There are three possible effort levels: low (i.e., $e = L$), medium (i.e., $e = M$), and high effort (i.e., $e = H$). The disutility of effort increases in the effort level. I normalize the cost of L to 0. Since higher effort levels are costlier, $c_H > c_M$, where c_M and c_H denote the cost of M and H, respectively. If the seller chooses $e = H$, then the product is high quality ($q = q_H$) with probability Φ and medium quality ($q = q_M$) with probability $(1 - \Phi)$:

$$q\,|_{e=H} = \begin{cases} q_H & \text{with probability } \Phi \\ q_M & \text{with probability } 1 - \Phi \end{cases}.$$

If the seller chooses $e = M$, then the product is medium quality ($q = q_M$) with probability Φ and low quality ($q = q_L$) with probability $(1 - \Phi)$:

$$q\,|_{e=M} = \begin{cases} q_M & \text{with probability } \Phi \\ q_L & \text{with probability } 1 - \Phi \end{cases}.$$

Finally, if the seller chooses $e = L$, then the product is low quality ($q = q_L$) with probability one. High effort is socially efficient, whereas low effort generates the lowest social surplus, i.e.,

$$\Phi\, u_H + (1 - \Phi)\, u_M - c_H > \Phi\, u_M + (1 - \Phi)\, u_L - c_M > u_L.$$

I now assume that there are two seller types. I consider three scenarios. In the first scenario, I consider a "good" type (type-g) and a "bad" type (type-b) such that type-g is more able and has a lower cost of effort than type-b; that is, $c_e^b > c_e^g$, where c_e^i denotes the cost of effort $e \in \{M, H\}$ for the type-i seller, $i \in \{g, b\}$. In the second scenario, type-g has a higher Φ value than type-b. Finally, in the third scenario type-g has a higher discount factor than type-b. In the first two scenarios, I assume that the seller types have a common discount factor δ.

I focus on parameter values such that type-g has an incentive to separate himself from type-b in order to avoid uninteresting cases (equivalently, type-b has an incentive to imitate type-g). In particular, I focus on the case in which type-g can be induced to exert high effort in a symmetric information setting, whereas type-b can only be induced to exert medium effort (thus, type-g is analogous to the high type, and type-b is analogous to the low type in my main model).[34]

[34] As discussed in Section IB, I analyze this model in a relational setup. As a result, the seller may be motivated to exert effort in a symmetric information setting. If the seller can observe the quality realization before sale, she may send a message regarding the realized quality of the product. Similar to the bonus promise of a principal, the message regarding the quality is only cheap-talk and nonbinding; however, it can be truthful and credible (like a self-enforcing bonus scheme in a relational contract) provided that the future profit of the seller from "honest" trade is high enough. Details are provided in online Appendix D. If the seller can observe the quality realization only *after* the sale, then the seller may agree to give a refund to the buyer if the product quality is observed to be lower than some benchmark level. Again, this is sustainable in equilibrium if the future profit from honest trade is sufficiently high.

A seller may choose to burn money (via, for example, advertising) in order to signal his type. It is reasonable to assume that the cost of such money burning is identical for the two types.

I first show that a separating equilibrium does not exist in the first case if $c_H^b - c_H^g \leq c_M^b - c_M^g$. Intuitively, this is because the benefit of separation for a type-g seller is strictly lower than the benefit of imitation for a type-b seller. Thus, equilibrium information revelation must involve dishonest behavior by the type-b seller if $c_H^b - c_H^g \leq c_M^b - c_M^g$. Moreover, I show that the condition $c_H^b - c_H^g > c_M^b - c_M^g$ is not sufficient for the existence of a separating equilibrium. Details of this model and the analysis can be found in the online Appendix D. In a similar vein, separation is not generally possible in the second case, where type-g has a higher Φ value than type-b.[35]

However, a separating equilibrium *always* exists in the third case, where the seller types differ in their time preferences. The construction in the proof of Proposition 3 can be directly applied to show this result (see the online Appendix D). Thus, there exist equilibria which reveal information fully and without breach of trust: consumers know the type of the seller and are never deceived about the quality of the product that they purchase in a separating equilibrium. Other main results also apply in the third case: the optimal separating equilibrium is gradualist (i.e., it involves increasing profits for the type-g seller), and hybrid equilibria are inefficient unless the prior probability that the seller is type-g is sufficiently high.

Finally, my results extend to the market setting where sellers enter and exit the economy stochastically, and names can be traded, as modeled in Tadelis (1999) and Mailath and Samuelson (2001). In addition to the trading of names, I allow for name changes; for example, an existing type-b seller that has a bad reputation can try to erase the public memory about his type by choosing a new name. I also maintain the assumption in Tadelis (1999) and Mailath and Samuelson (2001) that changes in namesownership are unobservable. I show that my main result also holds in this setting provided that type-g sellers have a sufficiently high discount factor. In particular, there exists a separating equilibrium such that a good name never becomes bad because good names are too expensive for type-b sellers. This result contrasts with Tadelis (1999) and Mailath and Samuelson (2001), in which the seller types have identical discount factors.

B. *Multiple Types*

If there are more than two principal types, there still exists a separating equilibrium—the mechanism in Proposition 3 can be extended to show the existence of a separating equilibrium with multiple types. The precise characterization of the optimal separating contract is tedious, but I conjecture that for the better types, the monotonicity of the bonus schedule and the effort schedule is preserved.

In the presence of multiple types, hybrid equilibria with full information revelation will be costlier. As the number of types increases, the number of types with

[35] To be more precise, if $q_H - q_M \leq q_M - q_L$, then a separating equilibrium does not exist. Moreover, the condition $q_H - q_M > q_M - q_L$ is not sufficient for separation.

an incentive to imitate naturally rises. As a result, the measure of the best type typically decreases. This in turn restricts the surplus from hybrid equilibria that fully reveal information because the uninformed party will anticipate default with a high probability for a lengthy period and behave accordingly.[36] As a result, the optimal contract is more likely to be separating in the presence of multiple types.

C. *A Behavioral Approach*

As discussed before, individuals may be sensitive to downside deviations from their expectations and incur a psychological cost when they receive less than what they expect or are entitled to. This is the main theme in models of disappointment aversion and loss aversion, such as Gul (1991), and Kőszegi and Rabin (2006, 2007). This idea has also been incorporated into contract theory in several papers, such as Hart and Moore (2008) and Herweg, Müller, and Weinschenk (2010).

Disappointment aversion and loss aversion do not affect separating contracts, which are honest. However, hybrid contracts become costlier; in particular, the set of priors with which the optimal contract is hybrid becomes smaller. This is because hybrid equilibria involve a breach of contract with positive probability, and the agent faces a lottery in each period of a hybrid contract: until information is fully revealed, there is a risk that the agent is not compensated for his effort, which entails a psychological cost. Let λ_t denote the equilibrium probability with which the bonus is paid at t in a hybrid equilibrium provided that there was no default until t. In the optimal hybrid contract with standard, risk-neutral agent preferences, w_t, λ_t, b_t, and e_t are such that if $\mu_t^1 \in (0,1)$, then $\lambda_t < 1$ and $w_t + \lambda_t b_t - c(e_t) = \bar{u}$ for at least some $t \geq 0$. I now adopt a simple formulation of disappointment-averse preferences and assume that the disappointment-averse expected utility of the agent for time t is given by

$$w_t + \lambda_t b_t - c(e_t) - \theta(1 - \lambda_t)\lambda_t b_t,$$

where θ is the disappointment parameter, and material utility is linear for simplicity. This formulation is now standard in the literature with two-outcome games or lotteries.[37] With probability $1 - \lambda_t$, the agent does not receive b_t, and suffers a disutility of $\theta\lambda_t b_t$ since the material utility falls below the expected material utility by $\lambda_t b_t$. If $w_t + \lambda_t b_t - c(e_t) = \bar{u}$, then the agent will *not* agree to participate as $\theta > 0$. Let $\theta = 1$ for simplicity. If the bonus is b_t/λ_t in every $t \geq 0$ instead of b_t, and everything else is kept the same in the contract, then the agents per period expected utility is $\bar{u}$ and the agent agrees to participate, as before. What about the agents effort incentives? Conditional on the agents participation and the same λ_t as before, the required bonus b_t' for the agent to continue to exert effort e_t is such that $c(e_t) = \lambda_t(b_t' - (1 - \lambda_t)b_t') = \lambda_t b_t$ holds. Thus, if $b_t' = b_t/\lambda_t$, the agent will agree to participate and exert the same effort as before, but the contract is now costlier, as bonus payments must be higher. Put differently, the optimal hybrid contract is

[36] Hybrid equilibria with full information revelation require all principal types but the best type to default.

[37] See Gill and Prowse (2012) and the discussion therein.

no longer feasible, and the profit of the principal (of either type) will be lower. Obviously, increasing λ_t cannot alleviate the problem: if λ_t is increased, the bonus will still have to be higher than b_t (assuming that $\lambda_t < 1$ still holds) and λ_{t+k} will have to be lower for some $k > 0$ in order to make up for the increase in λ_t. But b'_{t+k} must then increase $\left(\text{i.e., } b'_t > b_{t+k}/\lambda_{t+k}\right)$ in order to compensate the agent for betrayal risk. To conclude, disappointment-aversion extends the set of priors such that hybrid equilibria are inefficient.

D. *Stochastic Output*

Many of my main results in Sections III and IV also apply in a richer model, where the agents unobservable effort gives rise to a stochastic output. More specifically, assume that the agents effort in period t generates stochastic output y_t, where y_t is either "high" or "low" (i.e., $y_t \in \{L, H\}$). The probability that $y_t = H$ is equal to e_t. In this case, Propositions 1–3, 7, and Lemmas 4 and 5 go through without additional assumptions. A version of Proposition 6 also holds; however, it requires assuming that the bonus reward is contingent on calendar time, the current output, and the agents beliefs, but not on the whole history of output. This constraint does not seem to be extreme, thanks to (a version of) Lemma 4: a history-dependent bonus scheme is never used to motivate the agent in the optimal contract. Such a scheme could still be the optimal arrangement in order to deter the low-type principal from imitating. I cannot rule out this possibility, yet history-dependent bonus schemes are very complicated, and therefore, I focus on a constrained set of contracts in which the bonus reward is contingent on calendar time, the current output, and the agents beliefs. Let b_t^* denote the bonus at t in the constrained optimal contract. Then, Proposition 9 follows.

PROPOSITION 9: *Assume that $ec''(e)$ is weakly increasing. In the constrained optimal separating contract of the high-type principal, $b_t^* < b_{t+1}^*$ for every $t \geq 0$ until $b_t^* = b_h$, which takes place in finite time. Similarly, the effort schedule is strictly increasing until it reaches e_h. Finally, the fixed wage is strictly decreasing over time until it reaches w_h.*

The proof of Proposition 9 is relegated to the online Appendix E.[38] Finally, note that a version of Proposition 8 also holds in this setup; however, it requires additional assumptions, such as imposing a Markovian restriction on the principals strategy space.

VII. Conclusion

This paper sheds light on the dynamics of cooperation and information revelation in self-enforcing relationships with private information. In the previous reputation literature, information revelation has typically involved dishonesty and breach of trust. This paper identifies the conditions under which breach of trust is

[38] The rest of the proofs for this section are available upon request. They are omitted since they are similar to the current proofs (albeit, more tedious).

not necessary for information revelation. In particular, I show that costly signals, such as signing bonuses, advertising, hiring reputable executives, are effective, and a separating equilibrium always exists if the private information relates to time preferences. In fact, this type of private information is essential for a "general" separation result; with other types of private information, separation via costly signaling is not generally possible. Finally, note that if the "good" type expects a higher future payoff than the "bad" type (as is often the case), then the good type can "endogenously" be more patient than the bad type in view of the Becker-Mulligan theory of discounting. But if these hold, then a separating equilibrium always exists with any type of private information.

Appendix

PROOF OF LEMMA 1: First, note that $u_0 = \frac{\overline{u}}{1-\delta}$ holds in the optimal contract. Otherwise, a small reduction in w_0 is strictly profitable for the principal because this change doesn't affect the participation and the incentive constraints of the agent in any period. Next, assume that $\delta_\theta > \delta$. Let $\{w_t, b_t\}_{t=0}^{\infty}$ denote the optimal equilibrium set of contracts, and let e_t denote the effort implemented at t in this equilibrium. Assume towards a contradiction that the optimal contract specifies $u_{t+1} > \frac{\overline{u}}{1-\delta}$ on-the-equilibrium path for some $t \geq 0$. Consider the following changes in the contract at t and $t+1$: b_t is increased by a small amount $\delta\varepsilon$, whereas w_{t+1} is decreased by ε (this implies that u_{t+1} decreases and $\pi_{\theta,t+1}$ increases by ε). The bonus reward $b_t + \delta\varepsilon$ is still contingent on the effort level e_t (that is, output level $y(e_t)$) as in the original contract. This modified contract strictly increases the payoff of the principal, and the agent is unaffected. To see why, first consider the case in which $b_t \geq 0$. The increase in b_t coupled with the increase in $\pi_{\theta,t+1}$ is enforceable since

$$b_t + \frac{\delta_\theta}{1-\delta_\theta}\overline{\pi} \leq \delta_\theta \pi_{\theta,t+1}$$

and $\delta_\theta > \delta$ imply that

$$b_t + \delta\varepsilon + \frac{\delta_\theta}{1-\delta_\theta}\overline{\pi} < \delta_\theta\left(\pi_{\theta,t+1} + \varepsilon\right).$$

Thus, $b_t + \delta\varepsilon$ is enforceable. Moreover, the agent's participation constraint is satisfied at every t (since ε is small), and the agent's incentive-compatibility constraint for choosing effort level e_t is still satisfied by construction; that is

$$b_t + \delta\varepsilon - c(e_t) + \delta(u_{t+1} - \varepsilon) \geq \frac{\delta}{1-\delta}\overline{u}.$$

It can easily be checked that the incentive compatibility constraint for effort is still satisfied in every other period as well. Thus, the principal gains $(\delta_\theta - \delta)\varepsilon > 0$, a contradiction. Second, consider the case in which $b_t < 0$. Again, increasing b_t by a small amount $\delta\varepsilon$ and decreasing w_{t+1} by ε while keeping the effort requirement the same satisfies the participation constraint of the agent at every t (since ε is small). Moreover, these changes satisfy the enforcement and the incentive-compatibility constraints of the agent since

$$b_t - c(e_t) + \delta u_{t+1} \geq \frac{\delta}{1-\delta}\overline{u}$$

implies that

$$(b_t + \delta\varepsilon) - c(e_t) + \delta(u_{t+1} - \varepsilon) \geq \frac{\delta}{1-\delta}\overline{u}.$$

This ensures that the agent optimally chooses effort level e_t, and the enforcement constraint is satisfied as well. As a result, the principal gains $(\delta_\theta - \delta)\varepsilon > 0$. Hence, $u_t = \frac{\overline{u}}{1-\delta}$ must hold in the optimal contract at every $t \geq 0$ if $\delta_\theta > \delta$.

Now, assume that $\delta_\theta < \delta$. Assume toward a contradiction that $\pi_{\theta,t+1} > \frac{\delta_\theta}{1-\delta_\theta}\overline{\pi}$ for some $t \geq 0$. First, consider the case in which $b_t > 0$. Consider the following changes in the contract: b_t is decreased by a small amount $\delta\varepsilon$, whereas w_{t+1} is increased by ε (this implies that u_{t+1} increases and $\pi_{\theta,t+1}$ decreases by ε). The bonus reward $b_t - \delta\varepsilon$ is contingent on the effort level e_t as in the original contract. The decrease in b_t coupled with the decrease in $\pi_{\theta,t+1}$ is enforceable since

$$b_t + \frac{\delta_\theta}{1-\delta_\theta}\overline{\pi} \leq \delta_\theta \pi_{\theta,t+1}$$

and $\delta_\theta < \delta$ imply that

$$b_t - \delta\varepsilon + \frac{\delta_\theta}{1-\delta_\theta}\overline{\pi} < \delta_\theta(\pi_{\theta,t+1} - \varepsilon).$$

Thus, $b_t - \delta\varepsilon$ is enforceable. Moreover, the agent's participation constraint as well as the incentive compatibility constraint for effort is still satisfied at every t by construction. It follows that the principal gains $(\delta - \delta_\theta)\varepsilon > 0$, a contradiction. Thus, $\pi_{\theta,t} = \frac{\delta_\theta}{1-\delta_\theta}\overline{\pi}$ for every $t \geq 0$. The proof for the case in which $b_t \leq 0$ is similar, and therefore, omitted. ∎

PROOF OF PROPOSITION 2:

I provide the proof only for the case in which $\delta_h > \delta_l > \delta$. The proof for the case in which $\delta_h > \delta \geq \delta_l$ is almost identical, and therefore, omitted. Now, assume that the two types separate offering different contracts at $t = 0$, and that the continuation play following separation consists of the optimal symmetric information contract. Let $C_0 = \{w_0, b_0\}$ be the contract offer of the high type at $t = 0$, where the bonus reward b_0 is contingent on the effort level e_0 (that is, output level $y(e_0)$). Of course, $w_0 + b_0 - c(e_0) \geq \overline{u}$ must hold.[39] Note that the agent must accept the offer C_0; it

[39] I assume without loss of generality that the high type always honors b_0, and that $w_0 + b_0 - c(e_0) \geq \overline{u}$. Refusing to pay b_0 or offering a contract such that $w_0 + b_0 - c(e_0) < \overline{u}$ cannot be credible costly signals for the high type.

can easily be checked that a contract offer at $t = 0$ which the agent rejects cannot be part of this separating equilibrium given that the continuation play consists of C_h. Given these, the relevant incentive compatibility constraints boil down to

$$\text{(A1)} \qquad y(e_0) - w_0 - b_0 + \frac{\delta_h}{1-\delta_h}(s_h - \overline{u}) \geq \frac{s_l - \overline{u}}{1-\delta_h}$$

for the high type and

$$\text{(A2)} \qquad \frac{s_l - \overline{u}}{1-\delta_l} \geq y(e_0) - w_0 - b_0 + \max\left\{\delta_l V_l, b_0 + \frac{\delta_l}{1-\delta_l}\overline{\pi}\right\}$$

for the low type, where

$$V_l = s_h - \overline{u} + b_h + \frac{\delta_l}{1-\delta_l}\overline{\pi}.$$

The term V_l denotes the continuation payoff for an imitator low-type principal at $t \geq 1$ given that she has honored b_0 and the continuation play is the symmetric information contract of the high type. Note that (A2) follows because an imitator reneges either on b_0 at $t = 0$, or on b_h at $t = 1$ depending on the magnitude of b_0 relative to b_h. Combining inequalities (A1) and (A2) implies that

$$\text{(A3)} \qquad \frac{\delta_h}{1-\delta_h}(s_h - \overline{u}) + \frac{\delta_l}{1-\delta_l}(s_l - \overline{u}) \geq \frac{\delta_h}{1-\delta_h}(s_l - \overline{u}) + \delta_l V_l$$

is a *necessary* condition for separation. Recall that

$$b_h = \frac{\delta_h}{1-\delta_h}(s_h - \overline{u} - \overline{\pi}),$$

due to the binding enforcement constraint of the high type. Thus,

$$V_l = \frac{s_h - \overline{u}}{1-\delta_h} - \frac{\delta_h - \delta_l}{(1-\delta_h)(1-\delta_l)}\overline{\pi}.$$

It follows that

$$\delta_l V_l - \frac{\delta_h}{1-\delta_h}(s_h - \overline{u}) = \frac{(\delta_l - \delta_h)}{1-\delta_h}\left[(s_h - \overline{u}) + \delta_l \frac{\overline{\pi}}{(1-\delta_l)}\right].$$

As a result,

$$\begin{aligned}\left(\delta_l V_l - \frac{\delta_h(s_h - \overline{u})}{1-\delta_h}\right) + \left(\frac{\delta_h}{1-\delta_h} - \frac{\delta_l}{1-\delta_l}\right)(s_l - \overline{u}) &= \frac{\delta_h - \delta_l}{1-\delta_h}\left[\frac{s_l - \overline{u}}{1-\delta_l} - (s_h - \overline{u}) - \delta_l \frac{\overline{\pi}}{(1-\delta_l)}\right] \\ &= \frac{\delta_h - \delta_l}{1-\delta_h}\left[\frac{s_l - \overline{u} - \overline{\pi}}{1-\delta_l} - (s_h - \overline{u} - \overline{\pi})\right].\end{aligned}$$

The term inside the square brackets on the right-hand side of the equality is strictly positive if $s_l - \overline{u} - \overline{\pi} > (1 - \delta_l)(s_h - \overline{u} - \overline{\pi})$. But this will be the case if, for example, δ_h and δ_l are sufficiently close (note that s_h and s_l are close if δ_h and δ_l are close, by the Theorem of the Maximum). Thus, (A3) cannot hold if δ_h and δ_l are close. This implies that if δ_h and δ_l are sufficiently close, then a separating equilibrium in which the high type's behavior is distorted at only $t = 0$ is impossible. ∎

PROOF OF LEMMA 4:

Let $\{w_t, b_t\}_{t=0}^{\infty}$ denote the optimal separating set of contracts, and let e_t denote the effort implemented at t. Assume toward a contradiction that the optimal separating contract specifies $u_{t+1} > \frac{\overline{u}}{1-\delta}$ for some $t \geq 0$ on-the-equilibrium path. In period t of a separating equilibrium, either (i) $b_t > \delta_l \pi^i_{l,t+1} - \frac{\delta_l}{1-\delta_l}\overline{\pi}$; or (ii) $b_t \leq \delta_l \pi^i_{l,t+1} - \frac{\delta_l}{1-\delta_l}\overline{\pi}$, where $\pi^i_{l,t+1}$ represents the imitation payoff of a low-type principal who has not defaulted until $t+1$ from $t+1$ onwards. First, I show that if (i) holds, then $u_{t+1} > \frac{\overline{u}}{1-\delta}$ cannot be optimal. The reason is as follows. Consider the following changes in the separating contract at t and $t+1$: b_t is increased by a small amount $\delta\varepsilon$, whereas w_{t+1} is reduced by ε ; this implies that u_{t+1} decreases by ε, whereas $\pi_{h,t+1}$ and $\pi_{l,t+1}$ increase by ε. The bonus reward $b_t + \delta\varepsilon$ is still contingent on effort level e_t (that is, output level $y(e_t)$) as in the original contract. This modified separating contract strictly increases the payoff of the high type, the low type is still deterred from imitation, and the agent is unaffected. To see why, first note that the agent's participation constraint is still satisfied at every t (since ε is small), and the agent's incentive-compatibility constraint for choosing effort level e_t is still satisfied at every $t \geq 0$ by construction provided that the increased bonus payment is enforceable with the high-type principal. But this is indeed the case since $\pi_{h,t+1}$ increases by ε, and

$$b_t + \frac{\delta_h}{1-\delta_h}\overline{\pi} \leq \delta_h \pi_{h,t+1}$$

together with $\delta_h > \delta$ implies that

$$b_t + \delta\varepsilon + \frac{\delta_h}{1-\delta_h}\overline{\pi} < \delta_h(\pi_{h,t+1} + \varepsilon).$$

Thus, $b_t + \delta\varepsilon$ is enforceable with the high type. Assuming that ε is sufficiently small, $b_t + \delta\varepsilon > \delta_l(\pi^i_{l,t+1} + \varepsilon) - \frac{\delta_l}{1-\delta_l}\overline{\pi}$ also holds. Thus, the low type's strategy is unaffected, and furthermore, there is no change in the low type's imitation payoff. But the high type's payoff increases by $\delta_h^{t-1}(\delta_h - \delta)\varepsilon > 0$ in the modified separating contract, a contradiction.

If $b_t \leq \delta_l \pi^i_{l,t+1} - \frac{\delta_l}{1-\delta_l}\overline{\pi}$, then again, $u_{t+1} > \frac{\overline{u}}{1-\delta}$ is suboptimal. Consider the following changes to the contract: b_t is increased by a small amount $\delta\varepsilon$, w_{t+1} is reduced by ε (i.e., u_{t+1} is reduced by ε), and w_0 is increased by $\delta_l^t(\delta_l - \delta)\varepsilon$. This modified separating contract strictly increases the payoff of the high type, the low type is still deterred from imitation, and the agent is unaffected. To see why, first assume that $b_t \geq 0$. Note that the agent's participation constraint is still satisfied at every t (since ε is small), and the agent's incentive-compatibility constraint for choosing effort level e_t is still satisfied at every t by construction because the increased bonus payment is enforceable with the high-type principal (note that the decrease in w_{t+1} increases $\pi_{h,t+1}$ by ε). Since $b_t \leq \delta_l \pi^i_{l,t+1} - \frac{\delta_l}{1-\delta_l}\overline{\pi}$, it follows that

$$b_t + \delta\varepsilon < \delta_l(\pi^i_{l,t+1} + \varepsilon) - \frac{\delta_l}{1-\delta_l}\overline{\pi}.$$

This implies that $\pi^i_{l,t}$ increases by $(\delta_l - \delta)\varepsilon$. But since w_0 is increased by $\delta^t_l(\delta_l - \delta)\varepsilon$, the imitation payoff of the low type $\pi^i_{l,0}$ is either the same as before or even lower than before, depending on when it would be optimal to default for an imitator. However, the high type's payoff increases by $\left[\delta^t_h(\delta_h - \delta) - \delta^t_l(\delta_l - \delta)\right]\varepsilon > 0$, a contradiction. The argument is similar if $b_t < 0$ —this time, the agent's enforcement constraint matters. Hence, $u_{t+1} = \frac{\overline{u}}{1-\delta}$ must hold in the optimal contract for every $t \geq 0$. ■

PROOF OF LEMMA 5:

First, note that the optimal separating contract $\{w_t, b_t\}^{\infty}_{t=0}$ must be such that a low-type principal who imitates the high type strictly prefers defaulting at some $t \geq 0$. Suppose not. Then, $\{w_t, b_t\}^{\infty}_{t=0}$ is an equilibrium that can be implemented with the low type in a symmetric information setting. Thus, either $\{w_t, b_t\}^{\infty}_{t=0}$ is such that $\{w_t, b_t\} = C_l$ and e_l is implemented at every t, which is a contradiction, or the imitation payoff of the low type from $\{w_t, b_t\}^{\infty}_{t=0}$ is strictly lower than $\frac{y_l(e_l) - w_l - b_l}{1-\delta}$. This latter is also a contradiction because then $\{w_t, b_t\}^{\infty}_{t=0}$ can be strictly improved upon as follows. There exists a large enough (but finite) T such that if the high type starts offering C_h from T onwards (without any change prior to T), then the imitation payoff of the low type increases by a very small amount—i.e., the imitation payoff is still lower than $\frac{y_l(e_l) - w_l - b_l}{1-\delta}$, and the low type is still deterred from imitation. Moreover, the agent's incentives are unaffected given Lemma 4. But the high type is strictly better off, a contradiction. Thus, the optimal separating contract $\{w_t, b_t\}^{\infty}_{t=0}$ is such that the low type strictly prefers defaulting at some $t \geq 0$. Let T be the first period, such that

$$b_T > \delta_l \pi^i_{l,T+1} - \frac{\delta_l}{1-\delta_l}\overline{\pi}$$

in the optimal separating contract (recall that $\pi^i_{l,t}$ represents the imitation payoff of a low-type principal who has not defaulted until t from t onwards). Either $b_T = b_h$, or $b_{T+1} = b_h$, or there exists a $t > T$ such that the high type will offer C_h from t onward. In the first two cases, the desired result is obtained. Next, assume that $b_T < b_h$, and that $b_{T+1} < b_h$. Given that the low type strictly prefers defaulting at T, there must exist a $t > T$ such that the high type starts offering C_h from t onward. Suppose not. But there exists a large enough $T' > T$, such that if the high type offers C_h from T' onward, then $\pi^i_{l,T+1}$ increases by a very small amount, and the imitation payoff of the low type is unaffected as the imitator still strictly prefers defaulting at T. But then the high type is strictly better off, a contradiction. ∎

PROOF OF PROPOSITION 6:

First, I show that $b_t < b_{t+1}$ for every $t \geq 0$ until $b_t = b_h$, which takes place in finite time. Let $T = \min\{t \in \mathbb{N} \,|\, b_t = b_h\}$. Hence, $b_t = b_h$ for all $t \geq T$. Given Lemma 5, costly signaling ends at a finite period, and thus, $T < \infty$. First, note that $b_{T-1} < b_T = b_h$ by the definition of T. Otherwise, either $b_{T-1} > b_T$, in which case the high-type principal defaults or $b_{T-1} = b_h$, which is impossible due to the definition of T. Next, I show that $b_{T-2} < b_{T-1}$. Suppose toward a contradiction that $b_{T-2} \geq b_{T-1}$. I will make use of the following lemma.

LEMMA 10: *Let* $T = \min\{t \in N \,|\, b_t = b_h\}$. *In the optimal contract,*

$$b_t + \frac{\delta_l}{1-\delta_l}\overline{\pi} \leq \delta_l \pi^i_{l,t+1}$$

for every $t < T-1$.

PROOF:

Suppose not, so that $b_\tau + (\delta_l \overline{\pi})/(1-\delta_l) > \delta_l \pi^i_{l,\tau+1}$ for some $\tau < T-1$. Thus, a low type who imitates the high type would strictly prefer defaulting at τ. Since $b_{T-1} < b_h$, it follows that $b_{T-1} < \frac{\delta_h}{1-\delta_h}(s_h - \overline{u} - \overline{\pi})$. So, the high type can increase b_{T-1} by a small amount ε, which is still enforceable for the high type. This enables the high type to demand an output level $y(e'_{T-1})$ such that $c(e'_{T-1}) = b_{T-1} + \varepsilon$, which strictly increases the surplus and the high type's payoff at $T-1$.[40] Observe that this change in the contract does not affect the imitation payoff of the low type provided that ε is small enough. But the high type is strictly better off, a contradiction. Thus, in the optimal contract $b_t + (\delta_l\overline{\pi})/(1-\delta_l) \leq \delta_l \pi^i_{l,t+1}$ must hold for every $t < T-1$, where $T = \min\{t \in \mathbb{N} \,|\, b_t = b_h\}$. ∎

Since $b_{T-1} < b_T = b_h$ it follows that $\pi^i_{l,T-1} < \pi^i_{l,T} = V_l$. From Lemma 10,

$$b_{T-2} \leq \delta_l \pi^i_{l,T-1} - \frac{\delta_l}{1-\delta_l}\overline{\pi}.$$

Since $\pi^i_{l,T-1} < \pi^i_{l,T}$, it follows that

$$b_{T-1} \leq b_{T-2} < \delta_l \pi^i_{l,T} - \frac{\delta_l}{1-\delta_l}\overline{\pi}.$$

[40] In the optimal contract, $b_t = c(e_t)$ for every $t \geq 0$. To see why, first note that future rewards are not used in the optimal separating contract by Lemma 4, and thus, $b_t \geq c(e_t)$. To see why $b_t = c(e_t)$, suppose toward a contradiction that $b_t > c(e_t)$ for some $t \geq 0$. Then, e_t and w_t can be changed to e'_t and w'_t, respectively, such that $c(e'_t) = b_t$, and $w'_t = \overline{u}$, and w_0 would be increased to w'_0 in a way that $w'_0 = w_0 + \delta^t_l(s(e'_t) - s(e_t))$ holds. These changes leave the imitation payoff of the low type unaffected. If $t > 0$, then the high type is strictly better off due to her higher discount factor, whereas if $t = 0$, then the arrangement makes the agent strictly better off.

However,

$$b_{T-1} < \delta_l \pi^i_{l,T} - \frac{\delta_l}{1-\delta_l}\overline{\pi} \tag{A4}$$

cannot hold in the optimal contract. To see why, first note that an arbitrarily small increase in b_{T-1} is enforceable for the high type since the surplus becomes s_h after $T-1$, and $b_{T-1} < b_h$. Second, note that since (A4) holds, a very small increase in b_{T-1} increases the payoff of the high type by strictly more than the imitation payoff of the low type. This is because the discount factors of the two types are different, and a low type who imitates until $T-1$ still prefers honoring at $T-1$ as the increase in b_{T-1} is arbitrarily small and (A4) holds. Thus, the high type can (i) increase b_{T-1} by a small $\epsilon > 0$, (ii) increase the required output level to $y(e'_{T-1})$ in a way that $c(e'_{T-1}) = b_{T-1} + \epsilon$ so that the surplus at $T-1$ increases by $(s(e'_{T-1}) - s(e_{T-1}))$, and (iii) increase the initial fixed wage w_0 by $\delta_l^{T-1}(s(e'_{T-1}) - s(e_{T-1}))$ in order to deter the low type from imitating.[41] This modified contract is still separating, and the high types payoff increases by $(\delta_h^{T-1} - \delta_l^{T-1})(s(e'_{T-1}) - s(e_{T-1})) > 0$. Hence, a contradiction. As a result, (A4) cannot hold, and $b_{T-2} < b_{T-1}$.

Next, assume that $b_\tau < b_{\tau+1}$ holds for all $\tau \in \{t, t+1, \ldots, T-1\}$ by the induction hypothesis. I now show that $b_{t-1} < b_t$ must also hold. The proof of this is very similar to the proof above for the claim that $b_{T-2} < b_{T-1}$. First, one needs to verify that $\pi^i_{l,\tau} < \pi^i_{l,\tau+1}$ and $\pi_{h,\tau} < \pi_{h,\tau+1}$ for all $T-1 \geq \tau \geq t$. But this is true due to the hypothesis that b_τ is monotone increasing for $\tau \geq t$, $b_t = c(e_t)$ for every $t \geq 0$ and due to the fact that continuation payoffs are not used to motivate the agent by Lemma 4. From Lemma 10, it follows that

$$b_{t-1} \leq \delta_l \pi^i_{l,t} - \frac{\delta_l}{1-\delta_l}\overline{\pi}.$$

Now, assume toward a contradiction that $b_{t-1} \geq b_t$. From

$$b_t \leq b_{t-1} \leq \delta_l \pi^i_{l,t} - \frac{\delta_l}{1-\delta_l}\overline{\pi} < \delta_l \pi^i_{l,t+1} - \frac{\delta_l}{1-\delta_l}\overline{\pi},$$

it follows that

$$b_t < \delta_l \pi^i_{l,t+1} - \frac{\delta_l}{1-\delta_l}\overline{\pi}.$$

But this implies that

$$b_t = \delta_h \pi_{h,t+1} - \frac{\delta_h}{1-\delta_h}\overline{\pi}.$$

[41] Recall that in the optimal contract $b_t = c(e_t)$ for every $t \geq 0$ (see Footnote 40).

Otherwise, the high type could (i) increase b_t by a small $\epsilon > 0$, (ii) increase the required output level to $y(e_t')$ in a way that $c(e_t') = b_t + \epsilon$ so that surplus increases by $(s(e_t') - s(e_t))$, and (iii) increase the initial fixed wage w_0 by $\delta_l^t(s(e_t') - s(e_t))$, and make a positive gain, as argued before. But then,

$$b_t = \delta_h \pi_{h,t+1} - \frac{\delta_h}{1-\delta_h}\overline{\pi} > \delta_h \pi_{h,t} - \frac{\delta_h}{1-\delta_h}\overline{\pi} \geq b_{t-1}$$

implies that $b_{t-1} < b_t$, a contradiction. Since $b_t < b_{t+1}$ and $b_t = c(e_t)$ for every $t \leq T-1$ it follows that $e_t < e_{t+1}$ and $s_t < s_{t+1}$ for every $t \leq T-1$. Next, I will show that $b_t > b_l$ at every $t \geq 0$ in the optimal separating contract. Suppose not. Then, there exists a $t \geq 0$, such that $b_t \leq b_l < b_{t+1}$, given what I already showed above. Note that in that case,

$$b_t < \delta_l \pi_{l,t+1}^i - \frac{\delta_l}{1-\delta_l}\overline{\pi}.$$

Otherwise, a contradiction would follow because $\pi_{l,t+1}^i > \frac{s_l - \overline{u}}{1-\delta_l}$ must hold given that bonuses are monotone increasing and $b_\tau > b_l$ for all $\tau \geq t+1$ given the way t is determined. This implies that

$$b_t = \delta_h \pi_{h,t+1} - \frac{\delta_h}{1-\delta_h}\overline{\pi}$$

must hold, similar to what I argued above. Yet, this is a contradiction since $b_t \leq b_l$ and $\pi_{h,t+1} > \frac{s_l - \overline{u}}{1-\delta_h}$ (this latter holds since $b_\tau > b_l$ at every $\tau \geq t+1$, and $b_t = c(e_t)$ at every $t \geq 0$ in the optimal separating contract). As a result, $b_t > b_l$ at every $t \geq 0$ in the optimal separating contract. It follows that $e_t > e_l$ and $s_t > s_l$ at every $t \geq 0$ since $b_t = c(e_t)$ at every $t \geq 0$. Recall that in the optimal contract, $u_{t+1} = \frac{\overline{u}}{1-\delta}$ for $t \geq 0$. This implies that $w_t = \overline{u}$ at every $t > 0$ because $b_t = c(e_t)$ at every $t \geq 0$ in the optimal separating contract. Moreover, $w_0 > \overline{u}$: since $b_t > b_l$, $s_t > s_l$, and $u_{t+1} = \frac{\overline{u}}{1-\delta}$ at every $t \geq 0$, the high type must offer a high enough w_0 so that the low type is deterred from imitating the high type.

REFERENCES

Abeler, Johannes, Armin Falk, Lorenz Goette, and David Huffman. 2011. "Reference Points and Effort Provision." *American Economic Review* 101 (2): 470–92.

Admati, Anat R., and Motty Perry. 1987. "Strategic Delay in Bargaining." *Review of Economic Studies* 54 (3): 345–64.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 2002. "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics* 117 (1): 39–84.

Becker, Gary S., and Casey B. Mulligan. 1997. "The Endogenous Determination of Time Preference." *Quarterly Journal of Economics* 112 (3): 729–58.

Board, Simon, and Moritz Meyer-Ter-Vehn. 2015. "Relational Contracts in Competitive Labor Markets." *Review of Economic Studies* 82 (2): 490–534.

Bohnet, Iris, Fiona Greig, Benedikt Herrmann, and Richard Zeckhauser. 2008. "Betrayal Aversion: Evidence from Brazil, China, Oman, Switzerland, Turkey, and the United States." *American Economic Review* 98 (1): 294–310.

Bull, Clive. 1988. "The Existence of Self-Enforcing Implicit Contracts." *Quarterly Journal of Economics* 102 (1): 147–59.

Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal of Economics* 102 (2): 179–221.

Conlin, Michelle. 2001. "Where Layoffs Are a Last Resort." *Bloomberg Business Week*, October 8. https://www.bloomberg.com/news/articles/2001-10-07/where-layoffs-are-a-last-resort.

Cramton, Peter. 1992. "Strategic Delay in Bargaining with Two-Sided Uncertainty." *Review of Economic Studies* 59 (1): 205–25.

Deutsch, Yuval, and Thomas W. Ross. 2003. "You Are Known by the Directors You Keep: Reputable Directors as a Signaling Mechanism for Young Firms." *Management Science* 49 (8): 1003–17.

Diamond, Douglas W. 1989. "Reputation Acquisition in Debt Markets." *Journal of Political Economy* 97 (4): 828–62.

Earle, John S., and Klara Sabirianova Peter. 2009. "Complementarity and Custom in Wage Contract Violation." *Review of Economics and Statistics* 91 (4): 832–49.

Fombrun, Charles, and Mark Shanley. 1990. "What's in a Name? Reputation Building and Corporate Strategy." *Academy of Management Journal* 33 (2): 233–58.

Fuchs, William. 2007. "Contracting with Repeated Moral Hazard and Private Evaluations." *American Economic Review* 97 (4): 1432–48.

Gill, David, and Victoria Prowse. 2012. "A Structural Analysis of Disappointment Aversion in a Real Effort Competition." *American Economic Review* 102 (1): 469–503.

Ghosh, Parikshit, and Debraj Ray. 1996. "Cooperation in Community Interaction without Information Flows." *Review of Economic Studies* 63 (3): 491–519.

Gul, Faruk. 1991. "A Theory of Disappointment Aversion." *Econometrica* 59 (3): 667–86.

Halac, Marina. 2012. "Relational Contracts and Value of Relationships." *American Economic Review* 102 (2): 750–79.

Hart, Oliver, and John Moore. 2008. "Contracts as Reference Points." *Quarterly Journal of Economics* 123 (1): 1–48.

Herweg, Fabian, Daniel Müller, and Philipp Weinschenk. 2010. "Binary Payment Schemes: Moral Hazard and Loss Aversion." *American Economic Review* 100 (5): 2451–77.

Jullien, Bruno, and In-Uck Park. 2014. "New, Like New, or Very Good? Reputation and Credibility." *Review of Economic Studies* 81 (4): 1543–74.

Kaya, Ayça. 2009. "Repeated signaling games." *Games and Economic Behavior* 66 (2): 841–54.

Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4): 1133–65.

Kőszegi, Botond, and Matthew Rabin. 2007. "Reference-Dependent Risk Attitudes." *American Economic Review* 97 (4): 1047–73.

Kranton, Rachel E. 1996. "The Formation of Cooperative Relationships." *Journal of Law, Economics, and Organization* 12 (1): 214–33.

Kremer, Ilan, and Andrzej Skrzypacz. 2007. "Dynamic signaling and market breakdown." *Journal of Economic Theory* 133 (1): 58–82.

Kreps, David, and Robert Wilson. 1982. "Reputation and imperfect information." *Journal of Economic Theory* 27 (2): 253–79.

Lehrer, Ehud, and Ady Pauzner. 1999. "Repeated Games with Differential Discount Factors." *Econometrica* 67 (2): 393–412.

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

Li, Jin, and Niko Matouschek. 2013. "Managing Conflicts in Relational Contracts." *American Economic Review* 103 (6): 2328–51.

MacLeod, W. Bentley. 2003. "Optimal Contracting with Subjective Evaluation." *American Economic Review* 93 (1): 216–40.

MacLeod, W. Bentley, and James M. Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57 (2): 447–80.

MacLeod, W. Bentley, and James M. Malcomson. 1998. "Motivation and Markets." *American Economic Review* 88 (3): 388–411.

Mailath, George J., and Larry Samuelson. 2001. "Who Wants a Good Reputation?" *Review of Economic Studies* 68 (2): 415–41.

Mas, Alexandre. 2006. "Pay, Reference Points, and Police Performance." *Quarterly Journal of Economics* 121 (3): 783–821.

Milgrom, Paul, and John Roberts. 1982. "Predation, Reputation, and Entry Deterrence." *Journal of Economic Theory* 27 (2): 280–312.

Noldeke, Georg, and Eric van Damme. 1990. "Signalling in a Dynamic Labor Market." *Review of Economic Studies* 57 (1): 1–23.

Shapiro, Carl, and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74 (3): 433–44.

Sobel, Joel. 1985. "A theory of credibility." *Review of Economic Studies* 52 (4): 557–73.

Stewart, James. 1993. "Taking the Dare." *New Yorker*, July 26, 34.

Tadelis, Steven. 1999. "What's in a Name? Reputation as a Tradeable Asset." *American Economic Review* 89 (3): 548–63.

Watson, Joel. 1999. "Starting Small and Renegotiation." *Journal of Economic Theory* 85 (1): 52–90.

Watson, Joel. 2002. "Starting Small and Commitment." *Games and Economic Behavior* 38 (1): 176–99.

Yang, Huanxing. 2013. "Nonstationary Relational Contracts with Adverse Selection." *International Economic Review* 54 (2): 525–47.