Dynamic Inefficiencies in Contracts

May 24, 2015

1 Introduction

1.1 The Big Picture

A vast literature in contract theory considers a question that is fundamental to organizational economics: "what contracting imperfections might lead firms to act differently from the 'black box' production functions that we typically assume?" This is an important question if we want to understand how firms operate within an economy, which in turn is important if we want to advise managers or policy-makers.

These notes will focus on a particular departure from 'black box' production functions highlighted by contract theory. In dynamic settings, an optimal contract may not be *sequentially optimal:* there may be histories at which the *continuation* contract looks very different than a contract that would be written if the relationship started at that history. In other words, history matters: past actions and outcomes can affect firm performance, even if they do not affect the technology or options available to the firm. To borrow a phrase from Jin Li, Mike Powell, and Niko Matouschek, these dynamic contracts carry the "burden of past promises:" actions are chosen not to maximize surplus in the future, but to fulfill past obligations to the players.

1.2 Overview

These lecture notes will cover four papers. We will begin by discussing a classic paper in contract theory by Fudenberg, Holmstrom, and Milgrom (1990). We will then link this seminal analysis to three recent papers:

1. Fuchs (2007 AER), "Contracting with Repeated Moral Hazard with Private Evaluations,"

2. Halac (2012 AER), "Relational Contracts and the Value of Relationships,"

3. and Board (2011 AER), "Relational Contracts and the Value of Loyalty."

As time permits, I will also discuss Barron and Powell (2015 Working Paper), "Policies in Relational Contracts."

1.3 Fudenberg, Holmstrom, and Milgrom (1990) - FHM

1.3.1 Set-Up[1]

Consider a very general dynamic contracting problem. The principal and agent interact for $T < \infty$ periods with a common discount factor $\delta \in [0, 1]$. In each period $t \geq 0$, the sequence of events is:

1. The agent learns some information $\theta_t \in \Theta$.
2. The agent chooses effort $e_t \in \mathbb{R}$.
3. Output y_t is realized according to the distribution $F(\cdot|\{\theta_{t'}, e_{t'}\}_{t' \leq t})$.
4. The principal pays $w_t \in W$ to the agent, where $W \subseteq \mathbb{R}_+$.
5. The agent chooses to consume c_t and save $s_t = s_{t-1} + w_t - c_t$.

Denote $y = \{y_1, ..., y_T\}$ and similarly for θ, e, w, and s. Payoffs for the principal and agent are respectively

$$\begin{array}{c} \Pi(y, w) \\ U(\theta, c, e, s_T) \end{array} .$$

We assume that the principal can commit to a long-term contract. Define $x_t \in X_t$ as the set of contractible variables in period t, and define $X^t = \times_{t'=0}^{T} X_t$ as the contractible histories of length t. In general, the sequence of outcomes $\{y_{t'}\}_{t'=0}^{T}$ will be contractible, but x_t might contain other variables as well. A formal contract is a mapping $w : \bigcup_{t=0}^{T} X^t \to W$ that gives a payment in period t for each possible history of contractible variables $x^t \in X^t$. We leave unspecified how this contract is offered. The principal and agent simultaneously accept or reject the contract, with outside options $\bar{\pi}$ and $\bar{u}$, respectively.

1.3.2 Definitions: Incentive compatibility, individual rationality, and sequential efficiency

Let $\mathcal{H}^t$ be set of full histories in period t. The agent's effort and consumption plans map $e, c : \bigcup_{t=0}^{T} \mathcal{H}^t \times \Theta \to \mathbb{R}$. Given a contract w, (c_w, e_w) is *incentive compatible (IC)* if at each history h^t,

$$(c_w, e_w) \in \arg\max_{c,e} E\left[U(\theta, c, e, s_T)|h^t, w\right].$$

Define U_w and Π_w as the agent and principal's total payoffs from contract w and and the IC (c_w, e_w) that maximize the principal's payoff. Then w is *Individually Rational (IR)* if

$$\begin{array}{c} U_w \geq \bar{u} \\ \Pi_w \geq \bar{\pi} \end{array} .$$

A contract w is *efficient* if there exists no other contract w' such that

$$\begin{array}{c} U_w \leq U_{w'} \\ \Pi_w \leq \Pi_{w'} \end{array}$$

[1] The notation and timing I use here are adapted from Fudenberg, Holmstrom, and Milgrom (1990).

with at least one inequality strict.

Define $U_w(x^t)$ and $\Pi_w(x^t)$ as the expected payoffs for the agent and principal given contractible history x^t. Then w is *sequentially efficient* if for every x^t, there exists no other contract w' such that

$$U_w(x^t) \leq U_{w'}(x^t)$$
$$\Pi_w(x^t) \leq \Pi_{w'}(x^t)$$

with at least one inequality strict. That is, a contract w' is *sequentially efficient* if at the start of each period, the continuation contract, effort plan, and message plan are efficient with respect to the information partition given by X^t.

1.3.3 When is the efficient contract *not* sequentially efficient?

If an efficient contract is sequentially efficient, then the contract at the start of each period resembles a contract that could have been written if the game were just starting. The relationship between the principal and the agent does not "develop inefficiencies" over time, except to the extent that the production technology itself changes.

Why might dynamic inefficiencies arise in an efficient contract? Fudenberg, Holmstrom, and Milgrom (1990) outline (at least) three reasons why an efficient contract is not necessarily sequentially efficient:

1. The payoff frontier between the principal and agent is downward sloping: Given contractible history X^t, if $U_w(x^t) \leq U_{w'}(x^t)$, then $\Pi_w(x^t) \geq \Pi_{w'}(x^t)$.

    - If this does not hold, then the only way to punish the agent may be to also punish the principal. But simultaneously punishing both principal and agent is inefficient.

2. The principal learns information about the agent's past effort over time: y_t is not a sufficient statistic for e_t

    - Suppose that $y_{t+t'}$ contains information about e_t that y_t does not. Then the optimal contract would motivate e_t by making payments contingent on $y_{t+t'}$ in a way that might not be sequentially efficient.

3. At the start of each period, all variables that determine future preferences and production technology are contractible.

    - If future payoffs depend on non-contractible information, then adverse selection might lead to sequential inefficiencies.

The next sections consider a series of examples that illustrate why sequential inefficiencies might arise if any of these three conditions do not hold.

1.3.4 Example 1: The payoff frontier is not downward sloping

Consider the following two-period example: in each period $t \in \{0, 1\}$,

1. The agent chooses effort $e_t \in \{0, 1\}$ at cost ke_t.
2. Output $y_t \in \{0, H\}$ realized, with

$$\text{Prob}\{y_t = H\} = q + (p - q)e_t$$

   and $0 < q < p < 1$.

3. Payoffs are:

$$\begin{array}{c} \Pi = y_0 + y_1 - w_0 - w_1 \\ U = c_0 + c_1 - ke_0 - ke_1 \end{array}.$$

   The agent has limited liability: $w_0, w_1 \geq 0$. He receives a payoff of $-\infty$ if $s_T < 0$.

The principal can write a long-term contract as a function of realized output: $w_0(y_0)$ and $w_1(y_0, y_1)$. However, the agent has limited liability, so $w_0, w_1 \geq 0$. It is easy to show that saving plays no role in this contract, so $c_t = w_t$ without loss of generality.

Let $w_1 = w_1^{y_1}$ following output y_1. In period $t = 1$, the IC constraint is

$$\frac{k}{p - q} \leq w_1^H - w_1^0$$

which implies that $w_1^H \geq \frac{k}{p-q}$. Thus, the agent's utility if $e_1 = 1$ can be no less than

$$\frac{k}{p - q}p - k = \frac{qk}{p - q} > 0$$

if $q > 0$. In other words, the agent must *earn a rent* to be willing to work hard. Suppose that motivating high effort in $t = 1$ is efficient, or

$$H - \frac{p}{p - q}k > 0$$

Now consider period $t = 0$. In any sequentially efficient contract, the agent must choose $e_1 = 1$, regardless of y_0. Therefore, the principal can earn no more than

$$2\left(H - \frac{p}{p - q}k\right)$$

in a sequentially efficient contract. Consider the following alternative: if $y_0 = 0$, then $w_1^H = w_1^0 = 0$ and $e_1 = 0$. Intuitively, following low output in $t = 0$, the agent is not motivated in $t = 1$ (and earns no rent). This alternative is clearly sequentially inefficient. However, the agent is now more willing to work hard in $t = 0$, since he loses both a bonus *and* a future rent if output is low.

Therefore, the principal can motivate the agent to work hard in $t = 0$ if

$$w_0^H + \frac{q}{p - q}k \geq \frac{k}{p - q}$$

or

$$w_0^H \geq \frac{1-q}{p-q}k.$$

Relative to the sequentially efficient payoff, the alternative with low effort leads to a higher payoff for the principal if

$$H - \frac{p(1-q)}{p-q}k + p\left(H - \frac{p}{p-q}k\right) \geq 2\left(H - \frac{p}{p-q}k\right)$$

or

$$p(H-k) \geq H - \frac{p}{p-q}k$$

Both sides are strictly positive. This inequality holds if, e.g., $p \approx 1$ and $1 - \frac{k}{H} \approx q$. So in some circumstances, the principal would like to implement a contract that hurts *both* the agent *and* herself following low output in order to provide better incentives to the agent in period 0. But such a contract is clearly not sequentially efficient.

1.3.5 Example 2: Information about past performance revealed over time

Consider the following two-period example, which is adapted from Fudenberg and Tirole (1990):

1. In period 0, the agent chooses $e_0 \in \{0,1\}$ at cost ke_0. Output is $y_0 = 0$ with probability 1.

2. In period 1, the agent has no effort choice: $e_1 = 0$. Output is $y_1 \in \{0, H\}$, with $y_1 = H$ with probability pe_0.

3. Payoffs are

$$\begin{aligned} \Pi &= y_1 - w_0 - w_1 \\ U &= u(c_0) + u(c_1) - ke_0 \end{aligned}$$

with $u(\cdot)$ strictly concave. The agent receives a payoff of $-\infty$ if $s_T < 0$.

As in the first example, output y_1 is contractible. It is easy to see that $w_0 = c_0 = 0$ and $c_1 = w_1$ in this setting.

The optimal contract will condition only on y_1. In order to motivate the agent in $t = 0$, the payment w^{y_1} for output y_1 must satisfy

$$pu(w^H) + (1-p)u(w^0) - k \geq u(w^0)$$

or

$$u(w^H) - u(w^0) \geq \frac{k}{p}.$$

Now, consider the contract starting at the beginning of $t = 1$. The effort e_0 has already been chosen in this contract. Because u is strictly concave,

$$pu(w^H) + (1-p)u(w^0) < u(pw^H + (1-p)w^0).$$

The principal would earn strictly higher profits if she instead offered a contract with $w = pw^H + (1-p)w^0$. So if $e_0 = 1$, then any sequentially efficient contract must have a constant payment in $t = 1$. But then $e_0 = 1$ is not incentive compatible.

In this example, the efficient contract requires the payment in $t = 1$ to vary in output in order to motivate the agent to work hard in $t = 0$. After the agent has worked hard, however, the parties have an incentive to renegotiate in order to shield the agent from risk (since u is strictly concave). The efficient contract is not sequentially efficient because y_1 contains information about e_0 that is not also contained in y_0.

1.3.6 Example 3: Players have private information about the future

Consider the following two-period example, which is adapted from Fudenberg, Holmstrom, and Milgrom (1990). In each $t \in \{0, 1\}$,

1. In period $t = 0$, the agent does not exert effort ($e_0 = 0$) and produces no output ($y_0 = 0$).
2. In period $t = 1$, the agent chooses $e_1 \in \{0, 1\}$ at cost ke_1.
3. Output in $t = 1$ is $y_1 \in \{0, H\}$, with $\text{Prob}\{y_1 = H\} = q + e(p - q)$ for $0 < q < p < 1$.
4. Payoffs are

   $$\begin{aligned} \Pi &= y_1 - w_0 - w_1 \\ U &= u(c_0) + u(c_1) - ke_1 \end{aligned} .$$

   The agent receives a payoff of $-\infty$ if $s_T < 0$.

Importantly, note that the agent works only in period 1, but consumes in both periods 0 and 1. The agent is able to **borrow money** to consume early (so $c_0 > w_0$). The principal can write a formal contract on y_1, but **not** on consumption or savings.

Suppose the agent consumes c_0 in $t = 0$ and chooses $e_1 = 1$. The agent could always deviate by choosing some other consumption $\tilde{c}_0$ in $t = 0$ and then shirking: $e_1 = 0$. Therefore, the agent's IC constraint is:

$$u(c_0) + pu(w^H - c_0) + (1-p)u(w^0 - c_0) - k \geq u(\tilde{c}_0) + qu(w^H - \tilde{c}_0) + (1-q)u(w^0 - \tilde{c}_0).$$

Any IC contract must satisfy $w^H - w^0 > 0$. Because u is strictly concave, it must be that $c_0 > \tilde{c}_0$. That is, the agent consumes more in $t = 0$ if he expects to work hard in $t = 1$. He does so because he expects higher monetary compensation in $t = 1$.

Now, suppose the agent chooses to consume c_0 because he anticipates working hard. Once he has consumed that amount, his IC constraint becomes

$$pu(w^H - c_0) + (1-p)u(w^0 - c_0) - k \geq qu(w^H - c_0) + (1-q)u(w^0 - c_0).$$

This constraint is slack because $c_0 \neq \tilde{c}_0$. Therefore, the parties could renegotiate the original contract to expose the agent to less risk (by making w^H and w^0 closer to one another). So the efficient contract is sequentially inefficient.

The key for this inefficiency is that the contract cannot condition on consumption c_0. But consumption in period 0 determines the agent's willingness to work hard in $t = 1$. That is, c_0 is effectively "private information" about the agent's utility in $t = 1$.

2 Modern Papers that apply FHM

2.1 Upward-sloping payoff frontier: Board (2011)

In many real-world environments, a principal interacts with several agents. For example, Toyota allocates business among its suppliers. The government interacts with multiple companies in procurement auctions. Bosses oversee multiple workers. And so on.

This paper, along with Andrews and Barron (2014) and Barron and Powell (2015), focus on dynamics in *multilateral* relationships.

2.1.1 Set-Up

A single principal P interacts repeatedly with N agents A with discount factor δ. In each period of the interaction:

1. For each A $i \in \{1, ..., N\}$, the cost of investing in i, $c_{i,t}$, is publicly observed.
2. P invests in one agent i, pays $c_{i,t}$. Let $Q_{i,t}$ be probability of investing in agent i.
3. Chosen agent creates value v and chooses an amount $p_t \in [0, v]$ to keep. The remainder goes to P.

Payoffs are $u_{i,t} = p_t Q_{i,t}$ for agent i and $\pi_t = \sum_{i=1}^N Q_{i,t}(v - p_t - c_{i,t})Q_{i,t}$ for P.

A note about the model: this problem has a limited liability constraint, which is built into the requirement that $p_t \in [0, v]$.

2.1.2 Limited Liability leads to Sequential Inefficiencies

Suppose that the principal can **commit** to an investment scheme. That is, $Q_{i,t}$ can be made conditional on any past variables. However, the agent cannot commit to repay the principal.

If this game is played once, then $p_t = v$ and so P chooses not to invest. Suppose the game is played repeatedly, but P invests in agent i only once. Then again, i has no incentive to give money to the principal, $p_t = v$, and the principal prefers not to invest in i. So **repeated interaction with a single agent** is key to providing incentives.

Define

$$U_{i,t} = E\left[\sum_{s=t}^{\infty} \delta^{s-t} p_s Q_{i,s}\right]$$

as agent i's continuation surplus. Define

$$\Pi_t = \sum_{i=1}^{N} E\left[\sum_{s=t}^{\infty} \delta^{s-t}(v - p_s - c_{i,s})Q_{i,s}\right]$$

as P's continuation.

Dynamic Enforcement: Agent i is only willing to follow a strategy if

$$(U_{i,t} - v)Q_{i,t} \geq 0$$

for all t. Otherwise, if P invests in i, then i can run away with the money and earn v. P can always choose not to invest in i, so if i does run away then he earns 0 in the continuation game.

Principal's Problem: At time $t = 0$, P chooses $\{Q_{i,t}\}$ and $\{p_t\}$ to maximize his profit subject to the dynamic enforcement constraint. **Principal profit at time** 0 **equals total surplus minus agent payoff at time** 0**:**

$$\Pi_0 = E\left[\sum_{t=0}^{\infty} \delta^t (v - c_{i,t}) Q_{i,t}\right] - \sum_{i=1}^{N} U_{i,0}.$$

The key observation is that an agent can be motivated by **promises of future rent.** In particular, promising future rent motivates an agent **in every period before that rent is paid**. Therefore, the principal only really needs to give agent i rent **once.**

More precisely, define $\tau_i(t) \geq t$ as the period on or after period t in which $Q_{i,\tau_i(t)} > 0$. Agent i earns 0 surplus P does not invest in i, so

$$U_{i,t} = E\left[U_{i,\tau_i(t)} \delta^{\tau_i(t)-t}\right].$$

i's dynamic enforcement is only satisfied if

$$U_{i,\tau_i(t)} \geq v,$$

so $U_{i,t} \geq E[v\delta^{\tau_i(t)-t}]$.

Can we make this inequality bind? **Yes**. One way is to ask the agent to keep just enough so that he earns v continuation surplus. For example,

$$p_t = vE_t\left[1 - \delta^{\tau_i(t+1)}\right]$$

would work. To see why, suppose $U_{i,\tau_i(t+1)}) = v$. Then if $Q_{i,t} = 1$, agent i's continuation surplus is

$$p_t + \delta^{\tau_i(t+1)} v = v.$$

Biases in Investment Decision. We know that $U_{i,\tau_i(0)} = v$. So

$$\Pi_0 = E\left[\sum_{t=0}^{\infty} \delta^t (v - c_{i,t}) Q_{i,t}\right] - \sum_{i=1}^{N} E[v\delta^{\tau_i(0)}].$$

So P's objective is to maximize **total surplus minus a "rent cost"** v **that is incurred the first time P trades with a new agent**.

Main Result: Define $\mathcal{I}_t$ as the set of agents with whom P has already traded in period t. Then in each period:

1. If P invests in $i \in \mathcal{I}$, then $i \in \mathcal{I}$ has the lowest cost **among agents in $\mathcal{I}$.**

2. If $i \notin \mathcal{I}$ and there exists $j \in \mathcal{I}$ with

$$(c_{j,t} - c_{i,t}) \leq (1-\delta)v,$$

then P never invests in i.

If costs are iid across agents and over time, then there exists a unique integer n^* such that the optimal contract entails at most n^* insiders.

Principal Dynamic Enforcement: What if the principal cannot commit to an investment plan? If P is punished by reversion to static Nash following a deviation, then it is easy to sustain the contract above.

But what if agents have trouble coordinating their punishments? More precisely, suppose that P loses no more than

$$\Pi_{i,t} = \sum_{s=t}^{\infty} \delta^{s-t} Q_{i,s}(v - p_s - c_{i,s})$$

if he deviates by not investing in agent i. Intuitively, agent i stops repaying P, but the other agents keep on repaying as before.

Suppose $c \in [\underline{c}, \bar{c}]$ is iid across agents and periods. Suppose $v > \bar{c}$ or $\underline{c} > 0$. Then $\Pi_{i,t}$ increases in δ. This is not *a priori* obvious because the **number of insiders** increases in δ. However, as the number of insiders increases, it is increasingly likely that an insider has costs very close to $\underline{c}$. Therefore, P does not gain much by including an additional insider. This effect dominates as $\delta \to 1$.

The upshot: as $\delta \to 1$, the principal is willing to follow the optimal contract **even if punishment is bilateral.**

2.2 Information about past outcomes is revealed over time: Fuchs (2007)

2.2.1 Set-Up

The following is a simplified version of Fuchs (2007)

Consider a repeated game with a principal P and agent A who share a discount factor $\delta \in (0,1)$. At the start of the game, P offers a long-term contract to A that maps verifiable information into payments and termination decisions. If A rejects, the parties earn 0. Otherwise, the following stage game is repeatedly played:

1. A chooses effort $e_t \in \{0,1\}$ at cost ce_t.

2. P **privately** observes output $y_t \in \{L, H\}$. If $e_t = 1$, $y_t = H$ with probability p. If $e_t = 0$, $y_t = H$ with probability $q < p$.[2]

3. P sends a public message m_t. A public randomization device x_t is realized after P's message is sent.

4. The formal contract determines transfers: wage w_t, bonus b_t, and burnt money B_t.

5. The parties decide whether to continue the relationship or not. Outside options are 0.

Payoffs are $\pi_t = y_t - w_t - b_t - B_t$ for P and $u_t = w_t + b_t - ce_t$ for A, with discounted continuations $\sum \delta^t(1-\delta)\pi_t$ and $\sum \delta^t(1-\delta)u_t$. For the moment, assume that parties are locked into the contract and cannot choose to terminate the relationship.

What is verifiable? The wage w_t can depend only on past realizations of the public randomization device $x_{t'}$. The bonus b_t and burnt money B_t can depend both on past $x_{t'}$ and on past messages $m_{t'}$ (including the message from the current period).

Note that if parties are locked into the contract, then this is not really a "relational contract." Everything observable is verifiable. We'll be altering that assumption in a bit.

2.2.2 Intuition - Formal Contract

Set $\delta = 1$ for simplicity.

One-shot game: Suppose the game is played once, and suppose moreover that y is verifiable. Then the parties can easily attain first-best: P pays b_H following $y = H$ and b_L following $y = L$, where

$$pb_H + (1-p)b_L - c \geq qb_H + (1-q)b_L$$

or

$$b_H - b_L \geq \frac{c}{p-q}.$$

The wage is set to satisfy the agent's outside option.

Why doesn't this simple contract work if P privately observes y? The short answer is that **P would have an incentive to report $m = L$ regardless of y.** The principal must pay $b_H > b_L$ if he reports $m = H$, which *ex post* he would rather not do.

In order for P to have incentives to tell the truth, it must be that

$$b_L + B_L = b_H + B_H$$

which immediately implies that

$$B_L - B_H \geq \frac{c}{p-q}.$$

Following low effort, P must burn some money in order to "convince" A that he is telling the truth.

[2] In the full model, P also observes the outcome of a random variable ϕ_t.

Two-shot game: Now, suppose the game is played **twice**. What contracts motivate the agent to work hard?

One easy option is to simply repeated the one-shot contract twice. In this case, $\frac{c}{p-q}$ surplus is burnt **whenever** $y_t = L$. So this contract isn't very efficient.

An alternative would be to make the contract **history-dependent**. For example, suppose the contract allowed the agent to avoid punishment if he produces L in the first period but H in the second period. In that case, the money burnt is 0 following (H, H), $\frac{c}{p-q}$ following (H, L), 0 following (L, H), and $\frac{c}{p-q} + \frac{c}{(p-q)(1-p)}$ following (L, L). One can show that this alternative scheme also induces high effort (check it as an exercise!). It also leads to the **same expected efficiency loss**.

Both of these contracts assume that P reports A's output after **each period.** However, P could instead **keep silent** until the very end of the game. In that case, the agent does not know whether he produced high output or not in period 1, and so his second-period IC constraint is satisfied so long as

$$E[b|e_1 = 1, y_2 = H] - E[b|e_1 = 1, y_2 = L] \geq \frac{c}{p-q}.$$

This is clearly easier to satisfy than the outcome-by-outcome IC constraints if the principal reveals information. **So the principal will not reveal information until the very end of the game in the optimal contract.**

Let $b_{y_1 y_2}$ be the bonus following (y_1, y_2). One can show that $b_{HH} \geq b_{HL} = b_{LH} \geq b_{LL}$ in the optimal contract. Moreover, suppose $b_{HH} > b_{HL}$. Consider increasing b_{HL} by $\epsilon > 0$ and decreasing b_{LL} by $\frac{2p}{1-p}\epsilon$. This change is rigged to ensure that the same amount of money is burnt under the new scheme. A receives the same payoff if he works hard.

If A shirks in a single period, his payoff is now

$$pqb_{HH} + [(1-p)q + (1-q)p](b_{HL} + \epsilon) + (1-q)(1-p)(b_{LL} - \frac{2p}{1-p}\epsilon) - c.$$

The coefficient on ϵ in this expression is

$$(1-p)q - p(1-q) < 0.$$

Therefore, **holding on-path surplus fixed, A's surplus following a deviation is strictly lower under this alternative contract.** So deviations are easier to deter.

Hence, the optimal contract has $b_{HH} = b_{HL} > b_{LL}$. **P does not communicate until the end of the game, and A is only punished if he produces low output in both periods**.

2.2.3 What if the principal cannot commit?

The intuition outlined above extends to any number of periods. P does not communicate and A is punished only if he produces low output in **every** period.

However, this is not a terribly realistic solution. If the game is played repeatedly, then the "optimal contract" would entail an infinitely large punishment infinitely far in the future, accompanied by an infinitely large amount of burnt surplus. Instead, we might think that the amount of surplus that can be burnt equals the **future value of the relationship**.

In other words, the principal **cannot commit** to burn money, so the worst that could happen is that the agent leaves the relationships. This is a relational contract that "burns money" by termination: because termination is inefficient (it hurts both principal and agent), it can be used to induce the principal to truthfully report output.

The main result of the paper argues that any optimal relational contract is equivalent to a relational contract in which:

1. A is paid a constant wage w that is strictly above his outside option, until he is fired.

2. A exerts $e = 1$ until fired.

3. P sends no messages to A until A is fired.

The upshot? Efficiency wages are an optimal contract. Note that neither this result nor the paper pin down **when** firing occurs, although it must occur with positive probability on the equilibrium path.

2.3 There is private information about payoff-relevant variables: Halac (2012)

Consider Levin's (2003) relational contract. In this relational contract:

1. **Total** surplus is independent of how **rent is split**. Therefore, Levin has nothing to say about bargaining between players.

2. Surplus depends on **outside options**. Recall the dynamic enforcement constraint:

$$c(y^*) \leq \frac{\delta}{1-\delta}(S^* - \bar{u} - \bar{\pi}).$$

   The **larger** the outside options, the **lower** the output.

In Levin's world, the parties would love to **decrease** their outside options, which would increase total surplus on the equilibrium path. In particular, if the principal could **pretend to have a worse outside option, then he would.**

In the real world, the principal might be wary of small outside options because he is afraid he will be **held up by the other player**. Suppose relationship rents are split according to Nash bargaining. The principal has bargaining weight λ and so earns

$$(1-\lambda)\bar{\pi} + \lambda(S^* - \bar{u}).$$

If $\lambda = 1$, then the principal would like to misreport his type to be **smaller** in order to increase S^*. If $\lambda = 0$, then the principal would like to pretend his type is **larger** in order to capture more rent.

Halac (2012) formalizes this loose intuition by considering a model in which the principal has persistent private information about his outside option.

2.3.1 Set-Up

A principal P and agent A interact repeatedly with a common discount factor $\delta \in (0,1)$. At the beginning of the game, P learns her type $\theta \in \{l, h\}$, which determines her outside option r_θ with $r_h > r_l$. This type is private and constant over time. $\text{Prob}\{\theta = l\} = \mu_0$.

In each period of the game:

1. With probability λ, P makes an offer of a wage w_t and promised bonus $b_t(y_t)$. Otherwise, A makes the offer.

2. The party that did not make the offer accepts or rejects.

3. If accept, A chooses effort $e_t \in [0,1)$.

4. Output $y_t \in \{\underline{y}, \bar{y}\}$ is realized, with $\text{Prob}\{y_t = \bar{y}|e_t\} = e_t$.

5. Payments: the fixed wage w_t is enforced by a court. The bonus b_t is discretionary.[3]

Denote P and A's payoffs by π_t and u_t, respectively. If the offer is accepted, $u_t = w_t + b_t - c(e_t)$ and $\pi_t = y_t - w_t - b_t$. If the offer is rejected, $u_t = r_A$ and $\pi_t = r_\theta$.

To highlight the intuition outlined at the start of this section, the paper makes several restrictions to equilibrium. The paper looks for a Perfect Public Bayesian Equilibrium that is on the Pareto frontier. Moreover:

1. Once a posterior assigns probability 1 to a type, it forever after assigns probability 1 to that type.

2. If a party reneges on a **payment**, then the relationship either breaks down or remains on the Pareto frontier.

3. If a party deviates in **any other way**, then the relationship remains on the Pareto frontier.

Note that "Pareto frontier" is a little strange here, since parties have different beliefs about payoffs. What is assumed is that the equilibrium is Pareto efficient **given the (commonly known) beliefs of the agent**.

The paper also makes assumptions about the **symmetric-information** relational contract. Regardless of r_θ, the discount factor δ is such that

1. If both players know that $\theta = l$, then first-best is not attainable in a relational contract.

2. If both players know that $\theta = h$, then some positive effort is attainable in a relational contract.

2.3.2 Sketch of Results

Proposition 1: Suppose that the two types of P choose different actions in period t. Then it must be that either (i) one type rejects an offer by A, or (ii) one type reneges on a bonus.

[3]For precisely, parties simultaneously choose non-negative payments to make to one another.

Why? Types could separate in one of four ways: either (i) or (ii) above; or (iii) they accept different contracts from A; or (iv) they offer different contracts to A. **By assumption,** once types separate play is on the symmetric-information Pareto frontier. In particular, players never separate in the future on the equilibrium path.

Suppose (iii). Then there is no rejection in the current period. So P's on-path payoff doesn't depend on type. So on-path payoffs must be equal. But then $\theta = h$ can imitate $\theta = l$ and then take his outside option or renege to earn earn a strictly larger profit.

Suppose (iv). If the l-type is supposed to reject, then the h-type also wants to reject because the l-type has a better symmetric info contract and the h-type has a better outside option. If the h-type is supposed to reject, then the h-type can deviate and offer his symmetric-info contract immediately. This contract is an equilibrium for **any** agent beliefs. The deviation generates a strictly higher payoff because it doesn't involve any breakdown. Moreover, whenever the agent offers a payoff below r_h in future periods, the principal can simply reject.

If P has bargaining power, $\lambda = 1$: Separation occurs only if P defaults on a payment (because A never makes a contract offer). P is the residual claimant, so the h-type wants to imitate the l-type.

How do the parties separate? If P is supposed to pay a large bonus and threatened by breakdown, then h-type is less willing to pay. So the **cost of separation is the probability of breakdown.** The **benefit of breakdown is that the l-type can credibly induce higher effort than the h-type.**

As a result, separation is optimal if **l-types are sufficiently likely:**

Proposition 2: $\exists\, \hat{\mu}_0$ such that if $\mu_0 > \hat{\mu}_0$, the optimal contract entails separation. Otherwise, the optimal contract pools on the h-type symmetric-info contract.

Under further restrictions on equilibrium, the paper characterizes the **speed of separation**. Because the l-type must compensate A for the possibility of default, l-type's payoff is larger when separation is **slower**.

If A has bargaining power, $\lambda = 0$: Separation occurs only if P rejects a contract (because P can never credibly promise a positive payoff). P earns his outside option, so the l-type wants to imitate the h-type.

After separation, each type earns his outside option (on-path). However, l-type can imitate h-type and earn r_h in all future periods. So it must be that the h-type rejects the contract while the l-type accepts. Let v_l be **today's payoff from accepting the contract**.

	Today's Payoff	Tomorrow's Payoff
h playing h	$(1-\delta)r_h$	δr_h
h playing l	$(1-\delta)v_l$	δr_h
l playing l	$(1-\delta)v_l$	δr_l
l playing h	$(1-\delta)r_l$	δr_h

Players are only willing to follow their specified roles if there exists a v_l that satisfies both IC constraints:

$$(1-\delta)v_l + \delta r_l \geq (1-\delta)r_l + \delta r_h$$

$$r_h \geq (1-\delta)v_l + \delta r_h$$

So $r_h \geq v_l$ and hence $\delta \leq \frac{1}{2}$. So **separation is only feasible if players are impatient.** If P is patient, then all types are willing to mimic h-type to get a better continuation payoff.

The paper shows that separation occurs immediately if it occurs at all. Separation is only optimal if r_l is sufficiently likely.