

No Contracts (Updated: Sep 26 2017)

In many environments, contractible measures of performance may be so bad as to render them useless. Yet, aspects of performance that are relevant for the firm's objectives may be observable, but for whatever reason, they cannot be written into a formal contract that the firm can commit to. These aspects of performance may then form the basis for informal reward schemes. We will discuss two classes of models that build off this insight.

Career Concerns (Updated: Sep 26 2017)

An Agent's performance within a firm may be observable to outside market participants—for example, fund managers' returns are published in prospectuses, academics post their papers online publicly, a CEO's performance is partly announced in quarterly earnings reports. Holmstrom (1982/1999) developed a model to show that in such an environment, even when formal performance-contingent contracts are impossible to write, workers may be motivated to work hard out of a desire to convince "the market" that they are intrinsically productive in the hopes that doing so will attract favorable outside offers in the future—that is, they are motivated by their own career concerns.

Description There are two risk-neutral Principals, whom we will denote by P_1 and P_2 , and a risk-neutral Agent (A) who interact in periods $t = 1, 2$. The Agent has ability θ , which is drawn from a normal distribution, $\theta \sim N(m_0, h_0^{-1})$. θ is unobservable by all players, but all players know the distribution from which it is drawn. In each period, the Agent chooses an effort level $e_t \in E$ at cost $c(e_t)$ (with $c(0) = c'(0) = 0 < c', c''$) that, together with his

ability and luck (denoted by ε_t), determine his output $y_t \in Y$ as follows:

$$y_t = \theta + e_t + \varepsilon_t.$$

Luck is also normally distributed, $\varepsilon_t \sim N(0, h_\varepsilon^{-1})$ and is independent across periods and independent from θ . This output accrues to whichever Principal employs the Agent in period t . At the beginning of each period, each Principal i offers the Agent a short-term contract $w_i \in W \subset \{w_i : M \rightarrow \mathbb{R}\}$, where M is the set of outcomes of a performance measure. The Agent has to accept one of the contracts, and if he accepts Principal i 's contract in period t , then Principal $j \neq i$ receives 0 in period t . For now, we will assume that there are no available performance measures, so short-term contracts can only take the form of a constant wage.

Comment on Assumption. Do you think the assumption that the Agent does not know more about his own productivity than the Principals do is sensible?

If Principal P_i employs the Agent in period t , the agent chooses effort e_t , and output y_t is realized, payoffs are given by

$$\begin{aligned}\pi_i(w_{it}, e_t, y_t) &= py_t - w_{it} \\ \pi_j(w_{it}, e_t, y_t) &= 0 \\ u_i(w_{it}, e_t, y_t) &= w_{it} - c(e_t).\end{aligned}$$

Players share a common discount factor of $\delta < 1$.

Timing There are two periods $t = 1, 2$. In each period, the following stage game is played:

1. P_1 and P_2 propose contracts w_{1t} and w_{2t} . These contracts are commonly observed.
2. A chooses one of the two contracts. The Agent's choice is commonly observed. If A chooses the contract offered by P_i , denote his choice by $d_t = i$. The set of choices is

denoted by $D = \{1, 2\}$.

3. A receives transfer w_{it} . This transfer is commonly observed.
4. A chooses effort e_t and incurs cost $c(e_t)$. e_t is only observed by A .
5. Output y_t is realized and accrues to P_i . y_t is commonly observed.

Equilibrium The solution concept is pure-strategy Perfect-Bayesian Equilibrium. A **Perfect-Bayesian Equilibrium** of this game consists of a strategy profile $\sigma^* = (\sigma_{P_1}^*, \sigma_{P_2}^*, \sigma_A^*)$ and a belief profile μ^* (defining beliefs of each player about the distribution of θ at each information set) such that σ^* is sequentially rational for each player given his beliefs (i.e., each player plays the best response at each information set given his beliefs) and μ^* is derived from σ^* using Bayes's rule whenever possible.

It is worth spelling out in more detail what the strategy space is. By doing so, we can get an appreciation for how complicated this seemingly simple environment is, and how different assumptions of the model contribute to simplifying the solution. Further, by understanding the role of the different assumptions, we will be able to get a sense for what directions the model could be extended without introducing great complexity.

Each Principal i chooses a pair of contract-offer strategies $w_{i1}^* : \Delta(\Theta) \rightarrow \mathbb{R}$ and $w_{i2}^* : W^2 \times D \times Y \times \Delta(\Theta) \rightarrow \mathbb{R}$. The first-period offers depend only on each Principal's beliefs about the Agent's type (as well as their equilibrium conjectures about what the Agent will do). The second-period offer can also be conditioned on the first-period contract offerings, the Agent's first-period contract choice, and the Agent's first-period output. In equilibrium, it will be the case that these variables determine the second-period contract offers only inasmuch as they determine each Principal's beliefs about the Agent's type.

The Agent chooses a set of acceptance strategies in each period, $d_1 : W^2 \times \Delta(\Theta) \rightarrow \{1, 2\}$ and $d_2 : W^4 \times D \times E \times Y \times \Delta(\Theta) \rightarrow \{1, 2\}$ and a set of effort strategies $e_1 : W^2 \times D \times \Delta(\Theta) \rightarrow \mathbb{R}_+$ and $e_2 : W^4 \times D \times E \times Y \times \Delta(\Theta) \rightarrow \mathbb{R}_+$. In the first period, the agent chooses which

contract to accept based on which ones are offered as well as his beliefs about his own type. In the present model, the contract space is not very rich (since it is only the set of scalars), so it will turn out that the Agent does not want to condition his acceptance decision on his beliefs about his own ability. This is not necessarily the case in richer models in which Principals are allowed to offer contracts involving performance-contingent payments. The Agent then chooses effort on the basis of which contracts were available, which one he chose, and his beliefs about his type. In the second period, his acceptance decision and effort choice can also be conditioned on events that occurred in the first period.

It will in fact be the case that this game has a unique pure-strategy Perfect-Bayesian Equilibrium, and in this Perfect-Bayesian equilibrium, both the Principals and the Agent will use **public** strategies in which $w_{i1}^* : \Delta(\Theta) \rightarrow \mathbb{R}$, $w_{i2}^* : \Delta(\Theta) \rightarrow \mathbb{R}$, $d_1 : W^2 \rightarrow \{1, 2\}$, $d_2 : W^2 \rightarrow \{1, 2\}$, $e_1 \in \mathbb{R}_+$ and $e_2 \in \mathbb{R}_+$.

The Program Sequential rationality implies that the Agent will choose $e_2^* = 0$ in the second period, no matter what happened in previous periods. This is because no further actions or payments that the Agent will receive are affected by the Agent's effort choice in the second period. Given that the agent knows his effort choice will be the same no matter which contract he chooses, he will choose whichever contract offers him a higher payment.

In turn, the Principals will each offer a contract in which they earn zero expected profits. This is because they have the same beliefs about the Agent's ability. This is the case since they have the same prior and have seen the same public history, and in equilibrium, they have the same conjectures about the Agent's strategy and therefore infer the same information about the Agent's ability. As a result, if one Principal offers a contract that will yield him positive expected profits, the other Principal will offer a contract that pays the Agent slightly more, and the Agent will accept the latter contract. The second-period contracts offered will therefore be

$$w_{12}^* \left(\hat{\theta}(y_1) \right) = w_{22}^* \left(\hat{\theta}(y_1) \right) = w_2^* \left(\hat{\theta}(y_1) \right) = pE[y_2 | y_1, \sigma^*] = pE[\theta | y_1, \sigma^*],$$

where $\hat{\theta}(y_1)$ is the equilibrium conditional distribution of θ given realized output y_1 .

If the agent chooses e_1 in period 1, first-period output will be $y_1 = \theta + e_1 + \varepsilon_1$. Given conjectured effort e_1^* , the Principals' beliefs about the Agent's ability will be based on two signals: their prior, and the signal $y_1 - e_1^* = \theta + \varepsilon_1$, which is also normally distributed with mean m_0 and variance $h_0^{-1} + h_\varepsilon^{-1}$. The joint distribution is therefore

$$\begin{bmatrix} \theta \\ \theta + \varepsilon_1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} \theta \\ \varepsilon_1 \end{bmatrix} \sim N \left(\begin{bmatrix} m_0 \\ m_0 \end{bmatrix}, \begin{bmatrix} h_0^{-1} & h_0^{-1} \\ h_0^{-1} & h_0^{-1} + h_\varepsilon^{-1} \end{bmatrix} \right)$$

Their beliefs about θ conditional on these signals will therefore be normally distributed:

$$\theta | y_1 \sim N \left(\varphi y_1 + (1 - \varphi) m_0, \frac{1}{h_\varepsilon + h_0} \right),$$

where $\varphi = \frac{h_\varepsilon}{h_0 + h_\varepsilon}$ is the signal-to-noise ratio. Here, we used the normal updating formula, which just to jog your memory is stated as follows. If X is a $K \times 1$ random vector and Y is an $N - K$ random vector, then if

$$\begin{bmatrix} X \\ Y \end{bmatrix} \sim N \left(\begin{bmatrix} \mu_X \\ \mu_Y \end{bmatrix}, \begin{bmatrix} \Sigma_{XX} & \Sigma_{XY} \\ \Sigma'_{XY} & \Sigma_{YY} \end{bmatrix} \right),$$

then

$$X | Y = y \sim N \left(\mu_X + \Sigma_{XY} \Sigma_{YY}^{-1} (y - \mu_Y), \Sigma_{XX} - \Sigma_{XY} \Sigma_{YY}^{-1} \Sigma'_{XY} \right).$$

Therefore, given output y_1 , the Agent's second-period wage will be

$$w_2^* \left(\hat{\theta}(y_1) \right) = p [\varphi (y_1 - e_1^*) + (1 - \varphi) m_0] = p [\varphi (\theta + e_1 + \varepsilon_1 - e_1^*) + (1 - \varphi) m_0].$$

In the first period, the Agent chooses a non-zero effort level, even though his first-period contract does not provide him with performance-based compensation. He chooses a non-zero effort level, because doing so affects the distribution of output, which the Principals use in

the second period to infer his ability. In equilibrium, of course, they are not fooled by his effort choice.

Given an arbitrary belief about his effort choice, \hat{e}_1 , the signal the Principals use to update their beliefs about the Agent's type is $y_1 - \hat{e}_1 = \theta + \varepsilon_1 + e_1 - \hat{e}_1$. The agent's incentives to exert effort in the first period to shift the distribution of output are therefore the same no matter what the Principals conjecture his effort choice to be. He will therefore choose effort e_1^* in the first period to solve

$$\max_{e_1} -c(e_1) + \delta E_{y_1} \left[w_2^* \left(\hat{\theta}(y_1) \right) \middle| e_1 \right] = \max_{e_1} -c(e_1) + \delta p (\varphi (\theta + e_1 - e_1^*) + (1 - \varphi) m_0),$$

so that he will choose

$$c'(e_1^*) = p\delta \frac{h_\varepsilon}{h_0 + h_\varepsilon},$$

and if $c(e) = \frac{c}{2}e^2$,

$$e_1^* = \frac{p}{c} \delta \frac{h_\varepsilon}{h_0 + h_\varepsilon}.$$

This second-period effort choice is, of course, less than first-best, since first-best effort satisfies $c'(e_1^{FB}) = 1$ or $e_1^{FB} = p/c$. He will choose a higher effort level in the first period the less he discounts the future (δ larger), the more prior uncertainty there is about his type (h_0 small), and the more informative output is about his ability (h_ε large). Finally, given that the Agent will choose e_1^* , the first-period wages will be

$$w_{11}^* = w_{21}^* = pE[y_1] = p(m_0 + e_1^*).$$

This model has a number of nice features. First, despite the fact that the Agent receives no formal incentives, he still chooses a positive effort level, at least in the first period. Second, he does not choose first-best effort (indeed, in versions of the model with three or more periods, he may initially choose excessively high effort), even though there is perfect competition in the labor market for his services. When he accepts an offer, he cannot commit

to choose a particular effort level, so competition does not necessarily generate efficiency when there are contracting frictions.

The model is remarkably tractable, despite being quite complicated. This is largely due to the fact that this is a symmetric-information game, so players neither infer nor communicate information about the agent's type when making choices. The functional-form choices are also aimed at ensuring that it not only starts out as a symmetric information game, but it also remains one as it progresses. At the end of the first period, if one of the Principals (say the one that the Agent worked for in the first period) learned more about the Agent's type than the other Principal did, then there would be asymmetric information at the wage-offering stage in the second period.

This model extends nicely to three or more periods. In such an extension, however, if the Agent's effort affected the variance of output, he would have more information about his type at the beginning of the second period than the Principals would. This is because he would have more information about the conditional variance of his own ability, because he knows what effort he chose. In turn, his choice of contract in the second period would be informative about what effort level he would be likely to choose in the second period, which would in turn influence the contract offerings. If ability and effort interact, and their interaction cannot be separated out from the noise with a simple transformation (e.g., if $y_t = \theta e_t + \varepsilon_t$), then the Agent would acquire private information about his marginal returns to effort, which would have a similar effect. For these reasons, the model has seen very little application to environments with more than two periods, except in a couple special cases (see Bonatti and Horner (2016) for a recent example with public all-or-nothing learning).

Finally, if the Agent's effort choice affects the informativeness of the public signal (e.g., $\varepsilon_t \sim N(0, h_\varepsilon(e_t)^{-1})$), then the model may generate multiple equilibria. In particular, the equilibrium condition for effort in the first period will be

$$c'(e_1^*) = p\delta \frac{h_\varepsilon(e_1^*)}{h_0 + h_\varepsilon(e_1^*)},$$

which may have multiple solutions if $h'_\epsilon(e_t) > 0$. Intuitively, if the Principals believe that the Agent will not put in effort in $t = 1$, then they think the signal is not very informative, which means that they will not put much weight on it in their belief formation. As a result, the Agent indeed has little incentive to put in effort in period 1. In contrast, if the Principals believe the Agent will put in lots of effort in $t = 1$, then they think the signal will be informative, so they will put a lot of weight on it, and the Agent will therefore have strong incentives to exert effort.

Exercise. *Can the above model be extended in a straightforward way to environments with more than 3 periods if the Agent has imperfect recall regarding the effort level he chose in past periods?*

An important source of conflicting objectives within firms is often the tension between the firm's desire to maximize profits and its workers' concerns for their careers. And importantly, as this model shows, these incentives are not chosen by the firm but rather, they are determined by the market and institutional context in which the firm operates. That is, career concerns provide *incidental*, rather than *designed*, incentives.

In this model, these incidental incentives motivate productive effort. Of course, these incentives may be excessively strong for young workers (see Landers, Rebitzer, and Taylor (1996) for evidence of this effect in law firms), and they may be especially weak for older workers (see Gibbons and Murphy (1992) for evidence of this effect among executives). More generally, however, career concerns incentives may motivate employees to make decisions that are counterproductive for the firm. If an employee is risk-averse, and he can choose between a safe project with outcomes that are independent of his ability and a risky project with outcomes that are more favorable if he is high-ability, he may opt for the safe project, even if the safe project is bad for the firm. In particular, if his expected future wage is linear in his expected ability, then since the market's beliefs about his ability are a martingale, he prefers the market's beliefs to remain constant. If a professional adviser cares about her reputation for appearing well-informed, then she may withhold valuable information when

giving advice (Ottaviani and Sorensen, 2006).

If an employee cares about his reputation for being a quick learner, then an “Impetuous Youngsters and Jaded Old-Timers” dynamic can arise (Prendergast and Stole, 1996). In particular, if an employee observes private signals about payoffs of different projects, and smarter employees have more precise information, then smarter employees will put more weight on these private signals. Smarter employees’ outcomes will therefore be more variable, and the market understands this, so there is an incentive for employees to “go out on a limb” by putting excessive weight on their private signals to convince the market they are smart (i.e., “youngsters may be impetuous”). Moreover, reversing a previous decision in the future signals, in part, that a worker’s initial information was wrong, so older workers might inefficiently stick to prior decisions (i.e., “old-timers may be jaded”).

Further Reading Dewatripont, Jewitt, and Tirole (1999b) shows that when there are complementarities between effort and the informativeness of the agent’s output, there may be multiple equilibria. Dewatripont, Jewitt, and Tirole (1999a) explores a more-general two-period model and examines the relationship between the information structure and the incentives the agent faces. It also highlights the difficulties in extending the model beyond two periods with general distributions, since, in general, asymmetric information arises on the equilibrium path. Bonatti and Horner (2016) explores an alternative setting in which effort and the agent’s ability are non-separable, but nevertheless, asymmetric information does not arise on the equilibrium path, in particular because its information structure features all-or-nothing learning. Cisternas (2016) sets up a tractable environment in which asymmetric information in fact arises on the equilibrium path.

The contracting space in the analysis above was very limited—principals could only offer short-term contracts specifying a fixed wage. Gibbons and Murphy (1992) allows for principals to offer (imperfect) short-term performance-based contracts. Such contracts are substitutes for career-concerns incentives and become more important later in a worker’s

career, as the market becomes less impressionable. In principle, we can think of the model above as characterizing the agent’s incentives for a particular long-term contract—the contract implicitly provided by market competition when output is publicly observed. He, Wei, Yu, and Gao (2014) characterizes the agent’s incentives for general long-term contracts in a continuous-time version of this setting and derives optimal long-term contracts. Finally, the firm may have other instruments available for helping shape employees’ career concerns. For example, team design in a setting where individual outputs cannot be individually observed (Bar-Isaac, 2007), and information design in a setting where the agent’s current principal controls the information the market observes about his output (Prat, 2005; Wolitzky, 2012; Horner and Lambert, 2016; Rodina, 2016).

Relational Incentive Contracts (Updated: Sep 26 2017)

If an Agent’s performance is commonly observed only by other members of his organization, or if the market is sure about his intrinsic productivity, then the career concerns motives we previously discussed cannot serve as motivation. However, individuals may form long-term attachments with an organization. In such long-term relationships, “goodwill” can arise as an equilibrium phenomenon, and fear of shattering this goodwill can motivate individuals to perform well and to reward good performance. This intuition is captured in models of relational contracts (informal contracts enforced by relationships). An entire section of the next course in this sequence will be devoted to studying many of the issues that arise in such models, but for now we will look at the workhorse model in the literature to get some of the more general insights.

The workhorse model is an infinitely repeated Principal-Agent game with publicly observed actions. We will characterize the “optimal relational contract” as the equilibrium of the repeated game that either maximizes the Principal’s equilibrium payoffs or the Principal and Agent’s joint equilibrium payoffs. A couple comments are in order at this point. First,

these are applied models of repeated games and therefore tend to focus on situations where the discount factor is not close to 1, asking questions like “how much effort can be sustained in equilibrium?”

Second, such models often have many equilibria, and therefore we will be taking a stance on equilibrium selection in their analysis. The often-made criticism that such models have no predictive power is, as Kandori puts it “... misplaced if we regard the theory of repeated games as a theory of informal contracts. Just as anything can be enforced when the party agrees to sign a binding contract, in repeated games [many outcomes can be] sustained if players agree on an equilibrium. Enforceability of a wide range of outcomes is the essential property of effective contracts, formal or informal.” (Kandori, 2008, p. 7) Put slightly differently, focusing on optimal contracts when discussing formal contract design is analogous to focusing on optimal relational contracts when discussing repeated principal-agent models. Our objective, therefore, will be to derive properties of *optimal* relational contracts.

Description A risk-neutral Principal and risk-neutral Agent interact repeatedly in periods $t = 0, 1, 2, \dots$. In period t , the Agent chooses an effort level $e_t \in E$ at cost $c(e_t) = \frac{c}{2}e_t^2$ that determines output $y_t = e_t \in Y$, which accrues to the Principal. The output can be sold on the product market for price p . At the beginning of date t , the Principal proposes a compensation package to the agent. This compensation consists of a fixed salary s_t and a contingent payment $b_t : E \rightarrow \mathbb{R}$ (with positive values denoting a transfer from the Principal to the Agent and negative values denoting a transfer from the Agent to the Principal), which can depend on the Agent’s effort choice. The Agent can accept the proposal (which we denote by $d_t = 1$) or reject it (which we denote by $d_t = 0$) in favor of an outside option that yields per-period utility \bar{u} for the Agent and $\bar{\pi}$ for the Principal. If the Agent accepts the proposal, the Principal is legally compelled to pay the transfer s_t , but she is not legally compelled to pay the contingent payment b_t .

Timing The stage game has the following five stages

1. P makes A a proposal (b_t, s_t) .
2. A accepts or rejects in favor of outside opportunity yielding \bar{u} to A and $\bar{\pi}$ to P .
3. P pays A an amount s_t .
4. A chooses effort \hat{e}_t at cost $c(\hat{e}_t)$, which is commonly observed.
5. P pays A a transfer \hat{b}_t .

Equilibrium The Principal is not legally required to make the promised payment b_t , so in a one-shot game, she would always choose $\hat{b}_t = 0$ (or analogously, if $b_t < 0$, the Agent is not legally required to pay b_t , so he would choose $\hat{b}_t = 0$). However, since the players are engaged in a long-term relationship and can therefore condition future play on this transfer, nonzero transfers can potentially be sustained as part of an equilibrium.

Whenever we consider repeated games, we will always try to spell out explicitly the variables that players can condition their behavior on. This exercise is tedious but important. Let $h_0^t = \{s_0, d_0, \hat{e}_0, \hat{b}_0, \dots, s_{t-1}, d_{t-1}, \hat{e}_{t-1}, \hat{b}_{t-1}\}$ denote the history up to the beginning of date t . In this game, all variables are commonly observed, so the history up to date t is a public history. We will also adopt the notation $h_s^t = h^t \cup \{s_t\}$, $h_d^t = h_s^t \cup \{d_t\}$, and $h_e^t = h_d^t \cup \{\hat{e}_t\}$, so we can cleanly keep track of within-period histories. (If we analogously defined h_b^t , it would be the same as h_0^{t+1} , so we will refrain from doing so.) Finally, let $\mathcal{H}_0^t, \mathcal{H}_s^t, \mathcal{H}_d^t$, and \mathcal{H}_e^t denote, respectively, the sets of such histories.

Following Levin (2003), we define a **relational contract** to be a complete plan for the relationship. It describes (1) the salary that the Principal should offer the Agent ($h_0^t \mapsto s_t$), (2) whether the Agent should accept the offer ($h_s^t \mapsto d_t$), (3) what effort level the Agent should choose ($h_d^t \mapsto \hat{e}_t$), and (4) what bonus payment the Principal should make ($h_e^t \mapsto \hat{b}_t$). A relational contract is **self-enforcing** if it describes a subgame-perfect equilibrium of the repeated game. An **optimal relational contract** is a self-enforcing relational contract that yields higher equilibrium payoffs for the Principal than any other self-enforcing relational

contract. It is important to note that a relational contract describes behavior on and off the equilibrium path.

Comment. *Early papers in the relational-contracting literature (Bull, 1987; MacLeod and Malcomson, 1989; Baker, Gibbons, and Murphy, 1994) referred to the equilibrium of the game instead as an implicit (as opposed to relational) contract. More recent papers eschew the term implicit, because the term “implicit contracts” has a connotation that seems to emphasize whether agreements are common knowledge, whereas the term “relational contracts” more clearly focuses on whether agreements are enforced formally or must be self-enforcing.*

The Program Though the stage game is relatively simple, and the game has a straightforward repeated structure, solving for the optimal relational contract should in principle seem like a daunting task. There are tons of things that the Principal and Agent can do in this game (the strategy space is quite rich), many of which are consistent with equilibrium play—there are lots of equilibria, some of which may have complicated dynamics. Our objective is to pick out, among all these equilibria, those that maximize the Principal’s equilibrium payoffs.

Thankfully, there are several nice results (many of which are contained in Levin (2003) but have origins in the preceding literature) that make this task achievable. We will proceed in the following steps:

1. We will argue, along the lines of Abreu (1988), that the unique stage game SPNE is an optimal punishment.
2. We will show that optimal reward schedules are “forcing.” That is, they pay the Agent a certain amount if he chooses a particular effort level, and they revert to punishment otherwise. An optimal relational contract will involve an optimal reward scheme.
3. We will then show that distribution and efficiency can be separated out in the stage game. Ex ante transfers have to satisfy participation constraints, but they otherwise do not affect incentives or whether continuation payoffs are self-enforcing.

4. We will show that an optimal relational contract is sequentially optimal on the equilibrium path. Increasing future surplus is good for ex ante surplus, which can be divided in any way, according to (3), and it improves the scope for incentives in the current period. Total future surplus is always maximized in an optimal relational contract, and since the game is a repeated game, this implies that total future surplus is therefore constant in an optimal relational contract.
5. We will then argue that we can restrict attention to stationary relational contracts. By (4), the total future surplus is constant in every period. Contemporaneous payments and the split of continuation payoffs are perfect substitutes for motivating effort provision and bonus payments and for participation. Therefore, we can restrict attention to agreements that “settle up” contemporaneously rather than reward and punish with continuation payoffs.
6. We will then solve for the set of stationary relational contracts, which is not so complicated. This set will contain an optimal relational contract.

In my view, while the restriction to stationary relational contracts is helpful for being able to tractably characterize optimal relational contracts, the important economic insights are actually that the relational contract is sequentially optimal and how this result depends on the separation of distribution and efficiency. The separation of distribution and efficiency in turn depends on several assumptions: risk-neutrality, unrestricted and costless transfers, and a simple information structure. In the next course in this sequence, we will return to these issues and think about settings where one or more of these assumptions is not satisfied.

Step 1 is straightforward. In the unique SPNE of the stage game, the Principal never pays a positive bonus, the Agent exerts zero effort, and he rejects any offer the Principal makes. The associated payoffs are \bar{u} for the Agent and $\bar{\pi}$ for the Principal. It is also straightforward to show that these are also the Agent’s and Principal’s maxmin payoffs, and therefore they constitute an optimal penal code (Abreu, 1988). Define $\bar{s} = \bar{u} + \bar{\pi}$ to be the outside surplus.

Next, consider a relational contract that specifies, in the initial period, payments w and $b(\hat{e})$, an effort level e , and continuation payoffs $u(\hat{e})$ and $\pi(\hat{e})$. The equilibrium payoffs of this relational contract, if accepted are:

$$\begin{aligned} u &= (1 - \delta)(w - c(e) + b(e)) + \delta u(e) \\ \pi &= (1 - \delta)(p \cdot e - w - b(e)) + \delta \pi(e). \end{aligned}$$

Let $s = u + \pi$ be the equilibrium contract surplus. This relational contract is self-enforcing if the following four conditions are satisfied.

1. Participation:

$$u \geq \bar{u}, \pi \geq \bar{\pi}$$

2. Effort-IC:

$$e \in \operatorname{argmax}_{\hat{e}} \{(1 - \delta)(-c(\hat{e}) + b(\hat{e})) + \delta u(\hat{e})\}$$

3. Payment:

$$\begin{aligned} (1 - \delta)(-b(e)) + \delta \pi(e) &\geq \delta \bar{\pi} \\ (1 - \delta)b(e) + \delta u(e) &\geq \delta \bar{u} \end{aligned}$$

4. Self-enforcing continuation contract: $u(e)$ and $\pi(e)$ correspond to a self-enforcing relational contract that will be initiated in the next period.

Step 2: Define the Agent's **reward schedule** under this relational contract by

$$R(\hat{e}) = b(\hat{e}) + \frac{\delta}{1 - \delta} u(\hat{e}).$$

The Agent's no-reneging constraint implies that $R(\hat{e}) \geq \frac{\delta}{1 - \delta} \bar{u}$ for all \hat{e} . Given a proposed effort level e , suppose there is some other effort level \hat{e} such that $R(\hat{e}) > \frac{\delta}{1 - \delta} \bar{u}$. Then

we can define an alternative relational contract in which everything else is the same, but $\tilde{R}(\hat{e}) = R(\hat{e}) - \varepsilon$ for some $\varepsilon > 0$. The payment constraints remain satisfied, and the effort-IC constraint becomes easier to satisfy. Therefore, such a change makes it possible to weakly improve at least one player's equilibrium payoff. Therefore, it is without loss of generality to focus on reward schedules for which $R(\hat{e}) = \frac{\delta}{1-\delta}\bar{u}$ for all $\hat{e} \neq e$.

Step 3: Consider an alternative relational contract in which everything else is the same, but $\tilde{w} = w - \varepsilon$ for some $\varepsilon \neq 0$. This changes the equilibrium payoffs u, π to $\tilde{u}, \tilde{\pi}$ but not the joint surplus s . Further, it does not affect the effort-IC, the payment, or the self-enforcing continuation contract conditions. As long as $\tilde{u} \geq \bar{u}$ and $\tilde{\pi} \geq \bar{\pi}$, then the proposed relational contract is still self-enforcing.

Define the value s^* to be the maximum total surplus generated by any self-enforcing relational contract. The set of possible payoffs under a self-enforcing relational contract is then $\{(u, \pi) : u \geq \bar{u}, \pi \geq \bar{\pi}, u + \pi \leq s^*\}$. For a given relational contract to satisfy the self-enforcing continuation contract condition, it then has to be the case that for any equilibrium effort e ,

$$(u(e), \pi(e)) \in \{(u, \pi) : u \geq \bar{u}, \pi \geq \bar{\pi}, u + \pi \leq s^*\}.$$

Step 4: Suppose the continuation relational contract satisfies $u(e) + \pi(e) < s^*$. Then $\pi(e)$ can be increased in a self-enforcing relational contract, holding everything else the same. Increasing $\pi(e)$ does not affect the effort-IC constraint, it relaxes both the Principal's participation and payment constraints, and it increases equilibrium surplus. The original relational contract is then not optimal. Therefore, any optimal relational contract has to satisfy $s(e) = u(e) + \pi(e) = s^*$.

Step 5: Suppose the proposed relational contract is optimal and generates surplus $s(e)$. By the previous step, it has to be the case that $s(e) = e - c(e) = s^*$. This in turn implies that optimal relational contracts involve the same effort choice, e^* , in each period. Now we want to construct an optimal relational contract that provides the same incentives for the agent to exert effort, for both players to pay promised bonus payments, and also yields continuation

payoffs that are equal to equilibrium payoffs (i.e., not only is the action that is chosen the same in each period, but so are equilibrium payoffs). To do so, suppose an optimal relational contract involves reward scheme $R(\hat{e}) = \frac{\delta}{1-\delta}\bar{u}$ for $\hat{e} \neq e^*$ and

$$R(e^*) = b(e^*) + \frac{\delta}{1-\delta}u(e^*).$$

Now, consider an alternative reward scheme $\tilde{R}(e^*)$ that provides the same incentives to the agent but leaves him with a continuation payoff of u^* :

$$\tilde{R}(e^*) = \tilde{b}(e^*) + \frac{\delta}{1-\delta}u^* = R(e^*).$$

This reward scheme also leaves him with an equilibrium utility of u^*

$$\begin{aligned} u^* &= (1-\delta)(w - c(e^*) + b(e^*)) + \delta u(e^*) = (1-\delta)(w - c(e^*) + R(e^*)) \\ &= (1-\delta)\left(w - c(e^*) + \tilde{R}(e^*)\right) = (1-\delta)\left(w - c(e^*) + \tilde{b}(e^*)\right) + \delta u^*. \end{aligned}$$

Since $\bar{u} \leq u^* \leq s^* - \bar{\pi}$, this alternative relational contract also satisfies the participation constraints.

Further, this alternative relational contract also satisfies all payment constraints, since by construction,

$$\tilde{b}(e^*) + \frac{\delta}{1-\delta}u^* = b(e^*) + \frac{\delta}{1-\delta}u(e^*),$$

and this equality also implies the analogous equality for the Principal (since $s^* = u^* + \pi^*$ and $s^* = u(e^*) + \pi(e^*)$):

$$-\tilde{b}(e^*) + \frac{\delta}{1-\delta}\pi^* = -b(e^*) + \frac{\delta}{1-\delta}\pi(e^*).$$

Finally, the continuation payoffs are (u^*, π^*) , which can themselves be part of this exact same self-enforcing relational contract initiated the following period.

Step 6: The last step allows us to set up a program that we can solve to find an optimal relational contract. A stationary effort level e generates total surplus $s = e - c(e)$. The Agent is willing to choose effort level e if he expects to be paid a bonus b satisfying

$$b + \frac{\delta}{1 - \delta} (u - \bar{u}) \geq c(e).$$

That is, he will choose e as long as his effort costs are less than the bonus b and the change in his continuation payoff that he would experience if he did not choose effort level e . Similarly, the Principal is willing to pay a bonus b if

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

A necessary condition for both of these inequalities to be satisfied is that

$$\frac{\delta}{1 - \delta} (s - \bar{s}) \geq c(e).$$

This condition is also sufficient for an effort level e to be sustainable in a stationary relational contract, since if it is satisfied, there is a b such that the preceding two inequalities are satisfied. This pooled inequality is referred to as the **dynamic-enforcement constraint**.

The Program: Putting all this together, then, an optimal relational contract will involve an effort level that solves

$$\max_e pe - \frac{c}{2}e^2$$

subject to the dynamic-enforcement constraint:

$$\frac{\delta}{1 - \delta} \left(pe - \frac{c}{2}e^2 - \bar{s} \right) \geq \frac{c}{2}e^2.$$

The first-best effort level $e^{FB} = \frac{p}{c}$ solves this problem as long as

$$\frac{\delta}{1-\delta} \left(p e^{FB} - \frac{c}{2} (e^{FB})^2 - \bar{s} \right) \geq \frac{c}{2} (e^{FB})^2,$$

or

$$\delta \geq \frac{p^2}{2p^2 - 2c\bar{s}}.$$

Otherwise, the optimal effort level e^* is the larger solution to the dynamic-enforcement constraint, when it holds with equality:

$$e^* = \frac{p}{c} \left(\delta + \sqrt{\frac{p^2\delta^2 - 2\delta\bar{s}c}{p^2}} \right).$$

For all $\delta < \frac{p^2}{2p^2 - 2c\bar{s}}$, $\delta + \sqrt{\frac{p^2\delta^2 - 2\delta\bar{s}c}{p^2}} < 1$, so $e^* < e^{FB}$.

Comment. *People not familiar or comfortable with these models often try to come up with ways to artificially generate commitment. For example, they might propose something along the lines of, “If the problem is that the Principal doesn’t have the incentives to pay a large bonus when required to, why doesn’t the Principal leave a pot of money with a third-party enforcer that she will lose if she doesn’t pay the bonus?” This proposal seems somewhat compelling, except for the fact that it would only solve the problem if the third-party enforcer could withhold that pot of money from the Principal if and only if the Principal breaks her promise to the Agent. Of course, this would require that the third-party enforcer condition its behavior on whether the Principal and the Agent cooperate. If the third-party enforcer could do this, then the third-party enforcer could presumably also enforce a contract that conditions on these events as well, which would imply that cooperation is contractible. On the other hand, if the third-party enforcer cannot conditionally withhold the money from the Principal, then the Principal’s reneging temptation will consist of the joint temptation to (a) not pay the bonus she promised the agent and (b) recover the pot of money from the third-party enforcer.*

Further Reading The analysis in this section specializes Levin’s (2003) analysis to a setting of perfect public monitoring and no private information about the marginal returns to effort. Levin (2003) shows that in a fairly general class of repeated environments with imperfect public monitoring, if an optimal relational contract exists, there is a stationary relational contract that is optimal. Further, the players’ inability to commit to payments enters the program only through a dynamic enforcement constraint. Using these results, he is able to show how players’ inability to commit to payments shapes optimal incentive contracts in moral-hazard settings and settings in which the agent has private information about his marginal returns to effort.

MacLeod and Malcomson (1998) show that the structure of payments in an optimal relational contract can take the form of contingent bonuses or efficiency wages. Baker, Gibbons, and Murphy (1994) show that formal contracts can complement relational contracts, but they can also crowd out relational contracts. We will explore a number of further issues related to relational-incentive contracts later in the course.

The motivation I gave above begins with the premise that formal contracts are simply not enforceable and asks what *equilibrium* arrangement is best for the parties involved. Another strand of the relational-contracting literature begins with the less-stark premise that formal contracts are costly (but not infinitely so) to write, and informal agreements are less costly (but again, are limited because they must be self-enforcing). Under this view, relational contracts are valuable, because they give parties the ability to adapt to changing circumstances without having to specify in advance just how they will adapt (Macaulay, 1963). Baker, Gibbons, and Murphy (2011) and Barron, Gibbons, Gil, and Murphy (2015) explore implications of relational *adaptation*, and the former paper also considers the question of when adaptation should be governed by a formal contract and when it should be governed through informal agreements.