LEARNING IN RELATIONAL CONTRACTS

Rumen Kostadinov*
New York University

Aditya Kuvalekar†
U of Carlos III Madrid

September, 2017

Abstract

We study relational contracts between a firm and an worker when they face symmetric uncertainty about the match quality. Actions affect learning about the match quality and the firm's payoffs. Because the worker's actions are perfectly observable, the worker cannot bias the firm's beliefs. We show that even when the worker is not protected by limited liability and despite the absence of private information and hidden action, uncertainty about match quality precludes efficiency. The source of inefficiency is the holdup problem arising out of the separation between the entity exerting effort and the entity collecting the output. We characterize the set of all sequential equilibria of the associated game. We show that Pareto Optimal equilibria may involve actions that are dominated in their informational content as well as payoff. Such actions are a modest way for the firm to provide incentives and learn about the match quality, when more efficient ways are not credible. Conditional upon strong performance, we show that the relationships move to a phase where actions that offer better learning and higher payoff are used. In this phase the worker is rewarded with a bonus upon strong performance.

Keywords: Relational contracts, dynamic agency, career concerns

JEL codes: C73,D21,D23,D83,M55

---

*E-mail: `rk1954@nyu.edu` and `avk255@nyu.edu`.

†We are especially grateful to Elliot Lipnowski for countless discussions and comments. We thank Joyee Deb, Debraj Ray, David Pearce and Ennio Stacchetti for valuable insights and constant encouragement. This work has benefited from discussions with Dhruva Bhaskar, Rahul Deb, João Ramos, Nishant Ravi, Evan Sadler, Maher Said and Nikhil Vellodi. We would also like to thank the seminar audience at UCSD, Boston University and CETC, Vancouver.

1 Introduction

We study a dynamic interaction between a firm and a worker with two features: (i) neither the firm nor the worker knows the quality of their match and (ii) there are realistic limits to what sort of contingent contracts the parties can write. The first feature, unknown match quality, is perhaps an aspect of every employment interaction. An employee may deliver splendid results working with one manager and might perform poorly with another. How well a soccer team performs is determined by the chemistry between the coach and his players.

Similarly, limits to contractibility or incompleteness of contracts is a ubiquitous phenomenon. In almost every employment interaction, a firm and a worker observe more than what can be verified in an outside court. In such situations, the parties form "relational contracts" - informal agreements predicated on the non-verifiable dimensions using threats and promises concerning future behavior. When the match quality is unknown, the threats and promises are shaped by the players' beliefs which evolve over time. For example, if the firm and the worker are pessimistic about their match quality then the future rewards they can credibly promise are smaller than when they are optimistic. That is, the nature of the relational contracts itself changes over time as the players learn about the match quality.

We begin by investigating whether relational contracts can achieve the efficient outcome where the firm and worker are merged into a single entity. **?** showed that when the match quality is known, efficiency can be attained in a relational contract if the players are sufficiently patient. Efficiency is also achievable when the match quality is unknown but contracts are complete. However, we show that the interaction of symmetric uncertainty about match quality and limits to contractibility may preclude efficiency.

Our main result is a complete characterization of the optimal relational contract which also highlights two forms of inefficiencies. First, the firm and the worker terminate their relationship at higher beliefs about the match quality relative to the efficient outcome. Of more interest is, perhaps, the second form of inefficiency. The firm learns about the match quality by assigning different tasks to the worker that vary in profitability and informativeness. One would guess that a task which is dominated by another task on both dimensions is useless. Indeed, such tasks are never used in the efficient outcome. However, they are used in the optimal relational contracts at intermediate beliefs when the players are moderately optimistic. This intermediate region resembles "Performance Improvement Plans" (PIPs) that are frequently used in consulting firms. Typically, underperforming employees are put on a PIP where they are given smaller objectives as a chance to redeem themselves. If they do well they are taken out of a PIP but otherwise they are fired.[1]

[1] https://www.thebalance.com/performance-improvement-plan-contents-and-sample-form-1918850

Model: We study an infinite horizon, discrete time game between a firm and a worker. In each period, the firm can fire the worker or offer him a salary. If the worker accepts the salary, he chooses an *observable* action. The output is stochastically determined by the worker's action and a relationship-specific match quality. The players are symmetrically uninformed about their match quality - they form the same beliefs after observing the action and the output. We study the set of all relational contracts i.e. sequential equilibria of this game.

In our setting, an action plays a dual role. As in the standard agency models, an action produces output. But here, it also produces information about the match quality which affects the distribution of output. Hence, we can rank the actions based on their informativeness (in the Blackwell sense) and their profitability - expected output less the effort cost of the action. For most of the paper we analyze an environment with two actions - high and low. Moreover, we assume that the high action is more informative as well as more profitable than the low action. This strengthens our results, as the low action is used in the optimal relational contract, despite being dominated on both dimensions.

Results: Our first important result, Proposition 3, fully characterizes the inefficiency resulting from the interaction of unknown match quality and limits to contractibility. The benchmark for efficiency is the solution to the problem of maximising the joint surplus of the firm and the worker when they are merged into a single entity. Its solution is a simple cutoff policy - quit and take the outside options below a certain cutoff belief about the match quality and choose the high action at higher beliefs. Since there is no private information and no limited liability, prima-facie it is unclear why the players cannot achieve this outcome in the game. We answer this question by providing necessary and sufficient conditions for the efficiency of relational contracts detailed below.

Inefficiency is caused by a holdup problem due to the separation between the firm and the worker when the beliefs are near the efficient cutoff belief. The surplus from experimentation through the high action consists of two parts: the output realised in the current period and the option value of learning, represented by the continuation surplus after accounting for the cost of the action. At the cutoff belief their discounted average equals the sum of the outside options, since both the high action and the outside options are optimal. In a relational contract, the contemporaneous output cannot be used to incentivize the high action. Hence, if the contemporaneous output exceeds the outside options, the continuation surplus is insufficient to compensate the cost of effort and provide continuation payoffs to both parties above their outside options. In this case, the efficient outcome cannot be achieved by any relational contract. Conversely, if the contemporaneous output is low, there is enough continuation surplus to provide incentives and efficiency can be achieved. While this argument applies to our main model where nothing is contractible, Section 5 exhibits extensions such as spot contracts contingent on output and partial commitment on actions to show that the inefficiency remains.

We proceed to characterize the optimal relational contract in our main result, Theorem 1. It exhibits a three-region structure (Figure 1). At low beliefs the players take their outside options, when the belief is high the worker chooses the high action, and at intermediate beliefs the worker chooses the low action. Notably, the low action is used even though it is dominated in terms of output and informativeness and serves no purpose in the efficient benchmark. The reason is that it is a cheaper way to incentivize the worker as it is less costly than the high action. When the firm cannot credibly promise to reimburse the worker for the cost of the high action, the low action provides an alternative which, may be a better alternative than terminating the relationship.

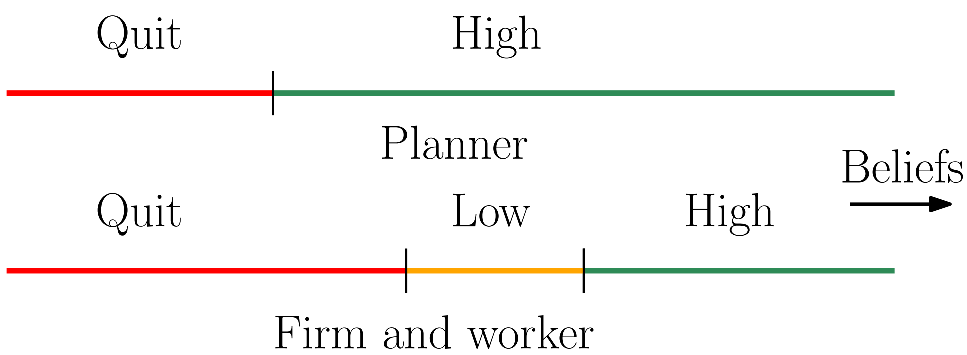

Figure 1: Difference between the planner's problem and optimal relational contract

This reasoning – use the low action when the high action is not incentive compatible and use the high action whenever it is incentive compatible since it is more informative and more profitable – is key to obtaining the three region structure. Indeed, this is the reason why the low action is not used in the efficient benchmark. But, in the game, incentive constraints complicate matters because they make information undesirable at some beliefs. This may make the low action superior *because* it is less informative. In fact, actions that are not ranked in the same way on both profitability and informativeness may lead to irregular optimal policies even in the efficient benchmark.[2] The proof of Theorem 1, which leads to the three-region structure, exhibits an interesting use of the algorithmic procedure due to **?**, APS hereafter. The optimal relational contract is obtained as a fixed point of the APS algorithm. We demonstrate the use of the algorithm by showing that the algorithm preserves a three region structure when operated on an appropriately chosen initial surplus function. Therefore, the three region structure is obtained as a fixed point property.

While a simple model with two actions delivers the insight – actions that are worse for both information and payoff may useful, on path, in a strategic interaction between

[2] A numerical example is available upon request where the information and payoff rankings do not coincide and the associated optimal policy has five regions: Quit-Low-High-Low-High.

a firm and a worker– in Section 5, we look at a setting with continuous action choice to deliver the same.

The three-region structure is consistent with Performance Improvement Plans, as discusesed previously. We interpret the low action as mundane tasks like documentation and testing. These tasks offer less information about the match quality and they are also less profitable. Therefore, it is never optimal to assign well-performing workers to such tasks. But, when the worker starts struggling he is put on a PIP where he is assigned tasks with lower stakes.

The remainder of the paper proceeds as follows. In Section 2, we describe our game and the equilibrium concept. Section 3 sets up the preliminary analysis leading to our main results presented in Section 4. Section 5 details some extensions of the model while Section 6 discusses the connection of our work with some existing literature.

2 Model

We consider an infinitely repeated interaction between a firm (she) and a worker (he), both risk-neutral. The productivity of their relationship is determined by the quality of their match, which can be either good or bad. The parties are symmetrically uninformed: they start with a common prior p_1 that the match quality is good. Time is discrete and denoted by $t = 1, 2, ...$ The firm and the worker have outside options amounting to $\underline{\pi}$ and $\underline{v}$ respectively.

The timing within a period t is as follows. At the beginning of the period the firm and the worker share a common belief p_t that the match quality is good. The firm decides whether to interact with the worker ($d_t^f = 1$) or not ($d_t^f = 0$). If she interacts, she offers a wage $w_t \geq \underline{v}$ to the worker. The worker decides whether to accept ($d_t^w = 1$) or reject ($d_t^w = 0$) the offer. If the firm chooses not to interact with the worker or the worker rejects the firm's offer, the parties terminate the relationship and receive their outside options in each subsequent period. If, instead, the offer is accepted, the wage is paid to the worker and he chooses a level of effort (e_t), which can be High ($\mathbb{H}$) or Low($\mathbb{L}$) and is observed by the firm. The cost of $\mathbb{H}$ to the worker is $c > 0$, while the cost of $\mathbb{L}$ is normalized to 0. The match quality determines how the worker's effort maps into a distribution over output (y_t), as given in Table 1. The outcome of the worker's effort is either Success ($y_t = 1$) or Failure ($y_t = 0$). Since all actions are publicly observable, the firm and the worker form a common posterior p_{t+1} using Bayes rule.

Also, at the end of every period, both the parties observe the realization x_t of a public randomization device. This lets the players publicly randomize based on the realization of the previous period's realization. In period 1, we assume that they players observe a realization x_0 at the beginning. This enables us to convexify the set of equilibrium payoffs at each posterior belief. The timing is summarized in Figure 2.

Effort \ Match quality	Good	Bad
High	$\alpha_{\mathbb{H}}$	$\beta_{\mathbb{H}}$
Low	$\alpha_{\mathbb{L}}$	$\beta_{\mathbb{L}}$

Table 1: Probability of a success given effort and match quality.

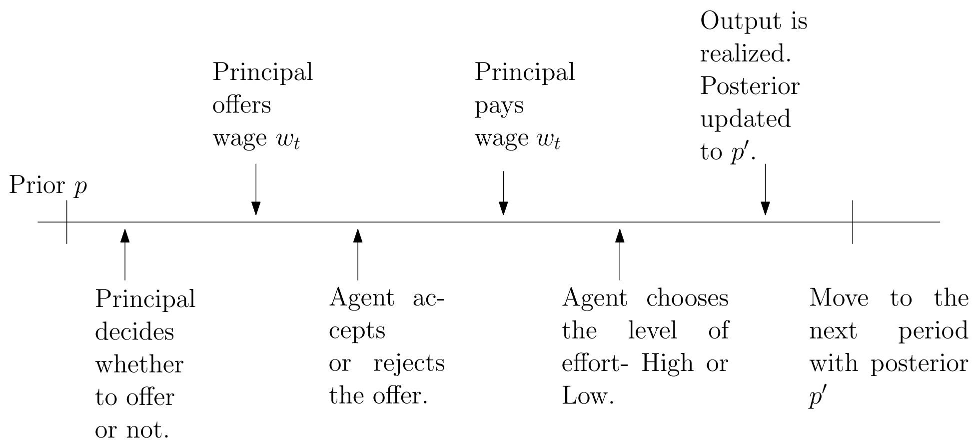

Figure 2: Timeline of one period

2.1 Assumptions

Apart from restricting the probabilities in Table 1 to lie in the interval $(0, 1)$, the following assumptions will be maintained throughout.

Assumption 1. $\alpha_{\mathbb{H}} - c > \underline{\pi} + \underline{v} > \beta_{\mathbb{H}} - c > \beta_{\mathbb{L}}$ *and* $\underline{\pi} + \underline{v} > \alpha_{\mathbb{L}} > \beta_{\mathbb{L}}$.

Assumption 2. $\frac{\alpha_{\mathbb{H}}}{\beta_{\mathbb{H}}} > \frac{\alpha_{\mathbb{L}}}{\beta_{\mathbb{L}}}$.

The above set of assumptions imply the following:

1. $\mathbb{H}$ generates more expected surplus (output net of cost) than $\mathbb{L}$ regardless of the match quality.

2. $\mathbb{H}$ is more informative than $\mathbb{L}$ in the sense of **?**.

3. $\mathbb{L}$ generates less surplus than the outside option regardless of the match quality.

Given a prior p, let $p_e(p) := \alpha_e p + \beta_e(1 - p)$ be the expected output in a period and $p^u_e(p)$ ($p^d_e(p)$) be the posterior belief following effort $e \in \{\mathbb{H}, \mathbb{L}\}$ and success (failure).

According to Bayes' rule,

$$p^u_e(p) = \frac{p\alpha_e}{\alpha_e p + \beta_e(1-p)} \qquad\qquad p^d_e(p) = \frac{p(1-\alpha_e)}{1-(\alpha_e p + \beta_e(1-p))}$$

To say that $\mathbb{H}$ is more informative in the sense of **?** means that

$$p^u_{\mathbb{H}}(p) \geq p^u_{\mathbb{L}}(p) \geq p \geq p^d_{\mathbb{L}}(p) \geq p^d_{\mathbb{H}}(p).$$

In addition, we assume that players are patient enough to sustain $\mathbb{H}$ in equilibrium when the match quality is good. Without this assumption the unique equilibrium would involve the players taking the outside option at every belief.

Assumption 3. $\delta(\alpha_{\mathbb{H}} - c - (\underline{\pi} + \underline{v})) \geq (1-\delta)c$.

2.2 Discussion of the model

Our baseline model assumes that the match quality, realized at the beginning, is static. This need not be the case. The match quality can change when firms undergo technology shocks, changes in management and market conditions. To model these scenarios we could assume that the match quality evolves exogenously according to some Markov process. This would not change the structure of our analysis and the results. This is because, even though the environment is changing, the players observe the same signals about its evolution. Therefore, they share the same beliefs about the match quality at all points in time.

A more realistic depiction of a firm-worker relationship would involve the worker being assigned different projects by the firm. Our model can accommodate this feature fairly easily. Suppose the firm can assign either a Big or a Small project to the worker each period. Big projects carry more responsibility, e.g. managing teams, dealing with clients etc. Small projects can be back-office jobs. The worker chooses his effort on the project. Big projects, due to higher responsibility, can succeed only if the worker exerts high effort. In that case, the probabilities of success are $\alpha_{\mathbb{H}}$ and $\beta_{\mathbb{H}}$ respectively depending on whether the match is good or bad. The probability of success for the Small projects is $\alpha_{\mathbb{L}}$ and $\beta_{\mathbb{L}}$ depending on whether the match is good or bad *regardless of the choice of action*. In this framework, the analysis and the results remain identical to the baseline model.

2.3 Equilibrium

Following the literature, we define a relational contract as a Sequential Equilibrium (SE) with public randomization of the dynamic game between the firm and the worker.

We denote by $h_t := \{x_t, d_t^f, w_t, d_t^w, e_t, y_t\}$ the history in stage t. Whenever one of the players decides not to interact, wage and effort choices will not take place. In order not to complicate the notation we will require players to make these choices even though they opt out[3]. This is inconsequential since these choices are payoff irrelevant when the parties do not interact. Let $h^t := \{h_\tau\}_{\tau=1}^{t-1}$ be the history at the beginning of period t. Let $H^t := \{h^t\}$ be the set of histories up to time t.

A strategy for the firm is a sequence of functions $\{D_t^f, W_t\}_{t=1}^{\infty}$, where $D_t^f : H^t \times [0,1] \to \{0,1\}$ and $W_t : H^t \times [0,1] \to \mathbb{R}$ determine the interaction decision and the wage offer as a function of the history in previous stages. Similarly, a strategy for the worker is a sequence of functions $\{D_t^w, E_t\}_{t=1}^{\infty}$, where $D_t^w : H^t \times [0,1] \times \mathbb{R} \to \{0,1\}$ and $E_t : H^t \times [0,1] \times \mathbb{R} \to \{\mathbb{H}, \mathbb{L}\}$ determine his interaction decision and effort level as a function of the history in the previous stages and the firm's wage offer.

Let $D_t := D_t^f D_t^w$. Following history h^t, the respective continuation payoffs to the firm and the worker from strategies (σ^P, σ^A) when the belief is p_t are given by

$$\pi_t(p_t|h^t, (\sigma^P, \sigma^A)) := \mathbb{E}\left[(1-\delta)\sum_{\tau=t}^{\infty} \delta^{\tau-t}\left(\mathbb{1}_{D_t}(y_t - w_t) + (1-\mathbb{1}_{D_t})\underline{\pi}\right)\right] \tag{1}$$

$$v_t(p_t|h^t, (\sigma^P, \sigma^A)) := \mathbb{E}\left[(1-\delta)\sum_{\tau=t}^{\infty} \delta^{\tau-t}\left(\mathbb{1}_{D_t}(w_t - c\mathbb{1}_{e_t=\mathbb{H}}) + (1-\mathbb{1}_{D_t})\underline{v}\right)\right] \tag{2}$$

The expectation is taken with respect to belief p_t and the continuations of strategies (σ^P, σ^A) following history h^t.

A strategy profile (σ^P, σ^A) is a sequential equilibrium or a relational contract if, following every history (including partial stage game histories), the continuation strategies of the firm and the worker are best responses to each other. Let $\mathcal{E}(p)$ denote the set of SE payoffs of the game starting at prior belief p.

3 Preliminaries

3.1 Efficiency benchmark

The sole conflict in our environment is the separation between the entity exerting effort (worker) and the entity collecting the output (firm). Therefore, our natural notion of efficiency is one where effort is contractible. Alternatively, we could view this as a single worker experimentation problem where the worker chooses a level of effort at each belief to maximiize the total surplus of the firm and the worker. This is a multi-armed bandit problem with two correlated arms: $\mathbb{H}$ and $\mathbb{L}$.

[3] In addition, nature will select an output level, even if the players decide not to interact.

Given a value function f, let $T_e f$ denote the surplus from choosing an effort $e \in \{\mathbb{H}, \mathbb{L}\}$ today and drawing continuations from f according to the probabilty distribution induced by e.

$$T_e f(p) := (1-\delta)(p_e(p) - \mathbb{1}_{e=\mathbb{H}} c) + \delta \mathbb{E}^{e,p}[f(p')]$$
$$\text{where} \quad \mathbb{E}^{e,p}[f(p')] = \left[p_e(p) f(p_e^u) + (1-p_e(p)) f(p_e^d)\right] \quad \forall e \in \{\mathbb{H}, \mathbb{L}\}$$

The optimal policy for this problem specifies an effort level or the outside options at each belief p. Let $G(p)$ denote the payoff associated with the optimal policy. The value function G is the unique fixed point of the following operator $T : B[0,1] \to B[0,1]$ where $B[0,1]$ denotes the set of bounded real-valued functions on $[0,1]$:

$$Tf(p) = \max\{T_{\mathbb{H}} f(p), T_{\mathbb{L}} f(p), \underline{\pi} + \underline{v}\}. \tag{3}$$

The optimal policy has a simple structure. There is a cutoff belief p^{FB} above which the planner uses $\mathbb{H}$ and below which the planner takes the outside options. This comes from the fact that the planner prefers more information, i.e. G is convex. Since $\mathbb{H}$ offers more information as well as more surplus, the planner never chooses $\mathbb{L}$.

Proposition 1. *The efficient outcome is characterised by the unique fixed point G of the operator T. G is increasing and convex. There exists a unique belief $p^{FB} \in (0,1)$ such that the optimal policy is to choose $\mathbb{H}$ above p^{FB} and the outside options below p^{FB}.*

3.2 Recursive Characterization

We proceed to analyse $\mathcal{E}(p)$, the set of equilibrium payoffs of the game starting with prior $p \in [0,1]$. Techniques from **?** (APS) show that sequential equilibria admit a recursive structure with the belief acting as a state variable. Any equilibrium payoff $(v, \pi) \in \mathcal{E}(p)$ can be decomposed into an action profile for the current period and continuation payoffs subject to incentive constraints. Continuation payoffs must belong to $\mathcal{E}(p')$, where p' is the updated belief. There are three types of action profiles which decompose payoffs in the following manner:

- (High effort) The firm offers a wage w, the worker accepts and chooses $\mathbb{H}$. Continuation payoffs are $(v^u, \pi^u) \in \mathcal{E}(p^u_{\mathbb{H}}(p))$ and $(v^d, \pi^d) \in \mathcal{E}(p^d_{\mathbb{H}}(p))$.

$$\begin{aligned}
&\pi = (1-\delta)(p_{\mathbb{H}}(p) - w) + \delta[p_{\mathbb{H}}(p)\pi^u + (1-p_{\mathbb{H}}(p))\pi^d] \\
&v = (1-\delta)(w - c) + \delta[p_{\mathbb{H}}(p) v^u + (1-p_{\mathbb{H}}(p)) v^d] \\
&\pi \geq \underline{\pi}, \qquad v \geq \underline{v} \\
&\delta\left[p_{\mathbb{H}}(p) v^u + (1-p_{\mathbb{H}}(p)) v^d - \underline{v}\right] \geq (1-\delta)c
\end{aligned} \tag{4}$$

- (Low effort) The firm offers a wage w, the worker accepts and chooses $\mathbb{L}$. Continuation payoffs are $(v^u, \pi^u) \in \mathcal{E}(p^u_{\mathbb{L}}(p))$ and $(v^d, \pi^d) \in \mathcal{E}(p^d_{\mathbb{L}}(p))$.

$$\begin{aligned}
&\pi = (1-\delta)(p_{\mathbb{L}}(p) - w) + \delta[p_{\mathbb{L}}(p)\pi^u + (1 - p_{\mathbb{L}}(p))\pi^d] \\
&v = (1-\delta)w + \delta[p_{\mathbb{L}}(p)v^u + (1 - p_{\mathbb{L}}(p))v^d] \\
&\pi \geq \underline{\pi} \qquad v \geq \underline{v}
\end{aligned}$$

- (Outside options) The firm does not interact or the worker rejects the offer.

$$\pi = \underline{\pi} \qquad v = \underline{v}$$

The above characterization relies on the fact that there is an equilibrium where the firm chooses not to interact which gives both parties their outside options (see Lemma 1 in the appendix). This equilibrium is used to threaten deviators.

In light of this punishment, the incentive constraint for the firm is straightforward. Suppose the firm is supposed to offer w and she offers w'. If $w' \geq \underline{v}$, the worst continuation equilibrium for the firm is one where the worker accepts and exerts $\mathbb{L}$ and the outside options are taken in the following period. The resulting payoff for the firm is strictly less than $\underline{\pi}$. If $w' < \underline{v}$, then the worker quits giving the firm a payoff of $\underline{\pi}$. Therefore, the incentives of the firm amount to a participation constraint.

For the worker, an additional constraint (4) is that he needs to compensated through future continuation value if he is exerting $\mathbb{H}$. If he deviates by choosing $\mathbb{L}$, he still collects the wage, saves on the cost c and is punished by receiving $\underline{v}$ tomorrow.

3.3 Optimal Relational Contracts through Surplus Maximization

Notice that the contemporaneous wages can be used to redistribute surplus without affecting incentives, as in **?**. Hence, an optimal relational contract maximizes the sum of the firm and worker's payoffs and is characterized by a surplus function S defined below.

$$S(p) := \max_{(u,v)\in\mathcal{E}(p)} u + v$$

In optimal relational contracts, it is without loss of generality to give the worker the entire future surplus in excess of the firm's outside option. This offers the best incentive to the worker. This observation helps us characterize S as a fixed point of the operator T_* defined in Proposition 2. T_* resembles the operator T in (3) used to obtain the efficient surplus with the addition of an incentive constraint for using $\mathbb{H}$. The incentive constraint implies that T_* is not a contraction. Despite this, we can still obtain S by iteratively applying T_* on an appropriate function. This result exploits the algorithmic version of APS. The APS algorithm obtains the equilibrium payoff set by an iterative application of a set-valued operator on a large enough payoff set. We adapt the algorithm to operate on correspondences for our setting. Then, we exploit the fact that the equilibrium value set for each belief p can be characterized by one number: $S(p)$. The APS operator can be translated to T_* in our setting to operate on functions. Hence, we can use the APS algorithm to obtain S by iteratively applying T_* on a sufficiently large function. As we shall see later, this algorithm lets us establish a number of key properties of S, including the characterization of the optimal relational contract.

We summarize the above in the following proposition.

Proposition 2. *S is a fixed point of the operator T_* defined below.*

$$T_*f := \begin{cases} \max\{T_{\mathbb{H}}f, T_{\mathbb{L}}f, \underline{\pi}+\underline{v}\} & \textit{if } \delta\underbrace{\left[p_{\mathbb{H}}(p)f(p^u_{\mathbb{H}}(p)) + (1-p_{\mathbb{H}}(p))f(p^d_{\mathbb{H}}(p)) - \underline{\pi}\right]}_{\text{Expected future surplus in excess of } \underline{\pi}} \\ & \quad \geq \delta\underline{v} + (1-\delta)c \\ \max\{T_{\mathbb{L}}f, \underline{\pi}+\underline{v}\} & \textit{otherwise} \end{cases}$$

Let $f : [0,1] \to \mathbb{R}$ be a bounded function satisfying $f \geq S$ and $f \geq T_*f$. Then

$$S = \lim_{n\to\infty} T^n_* f$$

where $T^n_*f \geq T^{n+1}_* f$ and the limit is taken pointwise. One particular choice of starting function is $f = G$.

4 Results

4.1 (In)efficiency of relational contracts

We say that relational contracts are inefficient if the optimal relational contract cannot achieve the efficient surplus at some belief i.e. $S(p) < G(p)$ for some $p \in [0,1]$. In the absence of any private information, the only tension is the separation between the entity incuring the cost of effort (worker) and the entity collecting the output (firm). In an environment with known match quality, **?** showed that when the players are sufficiently patient, relational contracts are efficient. The uncertainty about match quality, even though symmetric, may overturn this result as shown in Proposition 3, which characterizes when efficiency is achievable through relational contracts.

Proposition 3. *Relational contracts are inefficient iff $p_{\mathbb{H}}(p^{FB}) > \underline{\pi} + \underline{v}$.*

To illustrate the argument suppose the belief is p^{FB} where the efficient surplus is $\underline{\pi}+\underline{v}$. That is,

$$(1-\delta)\left(p_{\mathbb{H}}(p^{FB}) - c\right) + \delta\mathbb{E}^{\mathbb{H}}G(p') = \underline{\pi} + \underline{v}$$

This means that the cost of $\mathbb{H}$ is exactly compensated by the sum of two factors: the excess of the contemporaneous output over $\underline{\pi}+\underline{v}$ and the excess of the continuation surplus over $\underline{\pi}+\underline{v}$.

$$\underbrace{(1-\delta)\,c}_{\text{Today's cost}} = \underbrace{(1-\delta)[p_{\mathbb{H}}(p^{FB})-(\underline{\pi}+\underline{v})]}_{\text{Today's surplus}} + \underbrace{\delta[\mathbb{E}^{\mathbb{H}}G(p')-(\underline{\pi}+\underline{v})]}_{\text{Tomorrow's surplus}}$$

In any equlibrium, the firm can credibly promise no more than the whole of tomorrow's surplus to the worker. If $p_{\mathbb{H}}(p^{FB}) > \underline{\pi}+\underline{v}$ tomorrow's surplus is less than today's cost. Therefore, at p^{FB} and slightly higher beliefs, the firm will not be able to incentivize the worker to choose $\mathbb{H}$, resulting in inefficiency. When $p_{\mathbb{H}}(p^{FB}) \leq \underline{\pi}+\underline{v}$, a similar argument shows that incentives can be provided to achieve efficiency. It is straightforward to generate parameters that satisfy either condition. Hence, the inefficiency is not a direct consequence of unknown match quality but rather its subtle interaction with incomplete contracts.[4]

More precisely, the source of the inefficiency is a hold-up problem arising out of the separation between the firm and the worker. In equilibrium the worker only derives incentives from future surplus as the firm cannot commit to pay him the output creating a hold-up problem when contemporaneous output must be used to compensate the worker for his effort, i.e. $p_{\mathbb{H}}(p^{FB}) > \underline{\pi}+\underline{v}$. Lastly, one might think that this inefficiency would disappear if we allow for spot contracts contingent on output. However, as we show in Section 5, so long as the worker is protected by limited liability, these inefficiencies persist.

While Proposition 3 places a condition on the cutoff belief in the efficient outcome which is an endogeneous object, there is a simple condition on the primitives of the model which guarantees inefficiency.

Corollary 1. *If $\beta_{\mathbb{H}} \geq \underline{\pi}+\underline{v}$ then relational contracts are inefficient for any $\delta < 1$.*

4.2 Optimal relational contracts

Given the potential inefficiency, the natural next question is the structure of the optimal relational contracts. We answer this question by giving a complete characterization of the optimal relational contracts in our main result below.

Theorem 1. *An optimal relational contract is described by two cutoff beliefs $\underline{p} \geq p^{FB}$ and $\bar{p} \geq \underline{p}$ such that at any belief p*

- *The firm and the worker interact iff $p \geq \underline{p}$.*

[4]While so far we have presented a version with no contracts, in Section 5, we explore various contracting environments to show the robustness of our inefficiency result.

- *The worker exerts* $\mathbb{L}$ *if* $p \in [\underline{p}, \bar{p})$.
- *The worker exerts* $\mathbb{H}$ *if* $p \geq \bar{p}$.

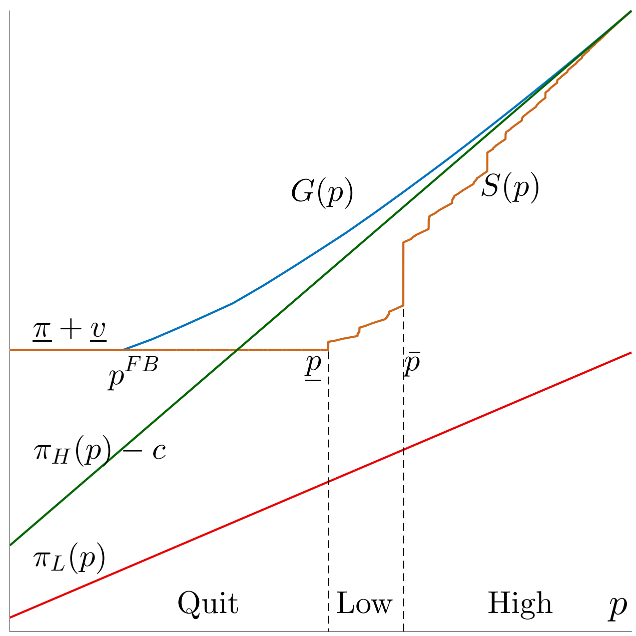

Figure 3: Plot of G and S

The structure of the optimal relational contracts can be interpreted as follows. When the players are sufficiently pessimistic about the match quality, the (expected) value of the relationship is low and learning has a high opportunity cost, so it is optimal to fire the worker. When the players are highly optimistic, the future value of the relationship is sufficient to incentivize the worker to exert $\mathbb{H}$. At intermediate beliefs there is not enough surplus to provide incentives, but the firm would like to retain the worker, as it is reasonably likely that the match quality is good. This is accomplished by the worker choosing $\mathbb{L}$ which is cheaper to incentivize and also offers some learning.

While Theorem 1 does not guarantee that $\mathbb{L}$ is used in a non-degenerate region, Proposition 4 gives conditions under which this is the case.

Proposition 4. *If* $\beta_{\mathbb{L}} \geq \frac{\underline{\pi}+\underline{v}}{1+\delta(\beta_{\mathbb{H}}-(\underline{\pi}+\underline{v}))}$ *then* $\underline{p} < \bar{p}$ *in Theorem 1.*

We proceed to give a sketch of the proof of Theorem 1 that involves an interesting application of the APS algorithm. Intuitively, $\mathbb{H}$ should be preferred to $\mathbb{L}$ at any belief so the worker exerts $\mathbb{L}$ only if the incentive constraint for $\mathbb{H}$ cannot be satisfied. Moreover, incentive provision is easier at higher beliefs because there is more future surplus to promise the worker. This produces the three region structure in Theorem 1. There is a subtlety that complicates this intuition: S, the surplus function in the optimal relational contract, is not convex (Figure 3). Lack of convexity means that more information is not necessarily better. Recall that $\mathbb{L}$ is never used in the efficient benchmark because more information is better and $\mathbb{H}$ dominates $\mathbb{L}$ on both the informativeness and payoff. In a relational contract $\mathbb{L}$ may be preferred to $\mathbb{H}$ at some beliefs where the gains from the inferior information in $\mathbb{L}$ outweigh the payoff losses. Therefore, we may see multiple switches between $\mathbb{L}$ and $\mathbb{H}$ depending on the tradeoff between information and immediate output. However, Theorem 1 rules out such behavior. The key to showing this lies in the APS algorithm as seen in Proposition 2. Since S can be obtained through iterative application of T_* to an appropriate starting function, S inherits properties that are preserved under T_*. While T_* does not preserve convexity (and even continuity) it preserves the weaker property that, barring incentive constraints, $\mathbb{H}$ is preferred to $\mathbb{L}$.

Definition 1. *A function $f : [0, 1] \to \mathbb{R}$ exhibits a monotone policy if $T_{\mathbb{H}} f \geq T_{\mathbb{L}} f$.*

If an increasing function f exhibits a monotone policy, then it is straightforward to see that the optimal policy associated with f will have the three-region structure of Theorem 1. Therefore, the proof boils down to showing that if an increasing function f exhibits a monotone policy, then $T_* f$ is also an increasing function that exhibits a monotone policy. The details of this argument are provided in the Appendix.

The characterisation in Theorem 1 resembles the characterization in the infinitely repeated stochastic partnership game by **?**. There, the state, which evolves exogenously, captures the temptation to cheat. Interestingly, the two models that look outwardly different have a similar intuition behind obtaining the three-region characterization: providing incentives is difficult when future prospects of the relationship look grim. In such situations, parties interact but indulge in inefficient activities, hoping that the relationship could drift towards an optimistic regime. In our model, the state variable that affects the incentives is the belief about the match quality which evolves *endogenously*. Moreover, since actions affect the beliefs, there is a tradeoff between payoff and informativeness which is not present in **?**. This tradeoff can give rise to substantially richer dynamics if we relax the assumption that $\mathbb{H}$ is more informative and more profitable than $\mathbb{L}$ (example available upon request).

4.3 Wage Distribution

Theorem 1 describes the interaction decisions and effort levels in equilibrium. In this section we complete the analysis by characterising equilibrium wages. Risk neutrality implies that the players are indifferent about the timing of the payments as long as the incentive constraints are satisfied. Hence, there is a multiplicity of wage paths. We present one wage path that frontloads the payments as much as possible. Whenever the worker exerts $\mathbb{L}$ the incentive constraints only dictate that his continuation payoffs must be at least $\underline{v}$. Hence, setting them to $\underline{v}$ allows for maximum frontloading of the worker's compensation. When the worker exerts $\mathbb{H}$ there is an additional incentive constraint: he must receive a continuation payoff that compensates him for the cost of effort over and above $\underline{v}$. Hence, maximum frontloading involves a binding incentive constraint where the worker receives an expected continuation payoff of $\underline{v} + \frac{1-\delta}{\delta}c$. This can be decomposed in multiple ways through the continuation payoffs following success and failure. Proposition 5 describes a wage path where the worker receives a constant payment for failure regardless of the belief. The nominal wage following success is decreasing in the belief but this is purely due to the higher probability of success at higher beliefs. Indeed, the expected wage is constant at $\underline{v} + \frac{1}{\delta}c$ as it serves to just compensate the worker for the cost of effort that he has incurred in the previous period[5]. The rest of the worker's payoff is frontloaded to the initial period

Proposition 5. *There exists a wage $\underline{w} \geq \underline{v}$ such that an optimal relational contract with payoffs (π, v) can be supported by the interaction decisions and effort levels in Theorem 1 and the following wages:*

- *If the worker exerts $\mathbb{L}$ the wage in the next period (if the worker exerts effort) is $\underline{v}$.*
- *If the worker exerts $\mathbb{H}$ at belief p, the wages in the next period (if the worker exerts effort) following success and failure respectively are*

$$w^u(p) = \underline{w} + \frac{1}{p_{\mathbb{H}}(p)}\left[\frac{1}{\delta}c - (\underline{w} - \underline{v})\right]$$
$$w^d(p) = \underline{w}$$

- *The wage in the initial period w_0 satisfies*

$$v = (1 - \delta)w_0 + \delta\underline{v}$$

Remark 1. *An important contribution of Proposition 5 is that any equilibrium can be achieved by wages no smaller than $\underline{v}$. Thus, a model with limited liability where the worker must be paid at least his outside option would be no different.*

5 Extensions

Our model could be extended in various directions. We briefly discuss some natural directions to explore what the extensions might entail.

[5]It is possible that following failure the posterior belief is in the region where the outside options are taken so no wage is paid. However, in this case $\hat{v}^d = \underline{v} = w^d(p)$. Hence, instead of being compensated through a wage, he is compensated through the payoff from the outside option.

5.1 Spot contracts contingent on output

As we saw in Section 4.1, the source of the inefficiency is a holdup problem arising from the separation between the firm which collects the output and the worker, who exerts effort. Therefore, one might think that this inefficiency would disappear if the parties could write one period contracts contingent on output. As we show below, limited liability plays a key role in answering this question. If the worker is protected by limited liability output contingent contracts cannot achieve efficiency unless it was already achievable with constant wages. But, when the worker is not protected by limited liability, relational contracts achieve efficiency by spreading the wages upon success and failure to induce high effort.

Proposition 6. *If the firm could commit to one period output contingent contracts, then the players can achieve efficiency if the worker is not protected by limited liability. If the minimum wage the firm must offer the worker for any level of output is $\underline{v}$ then relational contracts are inefficient iff $p_{\mathbb{H}}(p^{FB}) > \underline{\pi} + \underline{v}$.*

5.2 Combination of verifiable and non-verifiable actions

In most organizations there are some contractible aspects of a worker's behavior. Hence, the lack of contractibility in our model may seem extreme. However, it is also common that some aspects of the worker's behavior are not verifiable in an outside court. For example, a teaching assistant's wage can be contingent on the number sections she teaches but not on how well she answers questions. As we show here, the holdup problem from the baseline model remains as long as there is some non-contractibility of the worker's actions. Moreover, in any optimal relational contract, the parties can find the use of dominated actions beneficial for reasons similar to Theorem 1.

We model a mix of contractible and non-contractible actions by considering two-dimensional effort $e \in \{\mathbb{L}, \mathbb{H}\} \times \{\mathbb{L}, \mathbb{H}\}$. We assume that the first component of e is verifiable in court, while the second is not. The cost of effort is separable in the two components and $\mathbb{H}$ is more costly than $\mathbb{L}$ on each of them. Effort is ranked informationally as well as in payoff as follows:

$$(\mathbb{H}, \mathbb{H}) \succ (\mathbb{H}, \mathbb{L}) \succ (\mathbb{L}, \mathbb{H}) \succ (\mathbb{L}, \mathbb{L})$$

Naturally, the efficient outcome involves using $(\mathbb{H}, \mathbb{H})$ above a certain cutoff belief and taking the outside option below. The argument for inefficiency in Proposition 3 is virtually unchanged, as the same holdup problem arises on the non-contractible dimension. Moreover, an analogous result to our Proposition 4 holds showing that an optimal relational contract would involve using $\mathbb{L}$ on the non-contractible dimension.

5.3 Asymmetric information

In our model the firm and the worker are equally uninformed about the match quality. We relax this by assuming the worker has superior information. To simplify matters we will assume the worker holds one of two possible priors: $p_0 - \varepsilon$ and $p_0 + \varepsilon$ which are equally likely and focus on the pure strategy Bayesian equilibria of the game such that the equilibrium set at histories where the firm knows the prior of the worker coincides with the equilibrium set of the baseline model where they share the same prior.

Intuitively, private information should exacerbate the agency problem making it more difficult to achieve efficiency and more likely that $\mathbb{L}$ is used. Indeed, the formal analysis, deferred to the appendix, provides sufficient conditions for the use of $\mathbb{L}$ in optimal relational contracts. Interestingly, this is the exact sufficient condition from Proposition 4 which applies to the baseline model.

5.4 Continuous effort

So far we have analysed a setting with only two efforts one of which is both more informative and more profitable. A natural question is whether a continuous choice of effort would produce similar inefficiencies. In this section we enrich the model by making the choice of effort continuous. Also, a more informative effort may not necessarily be more profitable. As mentioned previously, such correlated bandit problems are very difficult to analyze in discrete time once the information and payoffs ranking do not coincide. However, we show that our central findings: (i)Efficiency being unachievable in a relational contract and (ii)efforts that are worse for both information and payoff are used in an optimal relational contract, are true in this model as well.

Let there be two strictly increasing, concave functions α and β defining the probability of success conditional on the choice of effort (e) and match quality (θ) as follows:

$$\alpha(e) = \text{Prob}(S|\theta = G, \text{effort} = e)$$
$$\beta(e) = \text{Prob}(S|\theta = B, \text{effort} = e).$$

$c(e)$ isthe cost of effort e where c is strictly increasing and convex. As before, the outside options for the firm and worker are denoted by $\underline{\pi}$ and $\underline{v}$. The myopic profit from effort e at prior p is given by

$$\pi(p, e) := p\alpha(e) + (1 - p)\beta(e) - c(e).$$

We adopt the following parametric restrictions.

Assumption 4. *1. $\frac{\alpha(e)}{\beta(e)}$ is increasing in e and $\alpha'(e) > \beta'(e)$*

2. $c(0) = c'(0) = 0$.

3. $\alpha'(1) < c'(1)$

The myopically efficient action at belief p is denoted by $\hat{e}(p)$ given by,

$$\pi_e(p, \hat{e}(p)) = [p\alpha'(\hat{e}(p)) + (1-p)\beta'(\hat{e}(p))] - c'(\hat{e}(p)) = 0.$$

Assumption 4 implies $\hat{e}$ is increasing and $0 < \underline{e} := \hat{e}(0) < \hat{e}(1) =: \bar{e} < 1$. That is, conditional on the match quality being good (bad), the efficient action is $\bar{e}$ ($\underline{e}$) which is strictly less than 1 (larger than 0).

Conditional on prior p and effort e, posteriors are updated to $p_e^u(p)$ and $p_e^d(p)$ respectively. The measure of informativeness, therefore, is the fraction $\frac{\alpha(e)}{\beta(e)}$, which is increasing due to Assumption 4. Hence, higher effort levels are more informative but not necessarily more profitable.

As before, the optimal policy for the efficient outcome is characterized by a cutoff p^{FB} below which the outside option are taken. Also, the value function is convex. Therefore, efforts dominated in both payoff and informational content are not used. In particular, efforts lower than $\underline{e}$ are never used.

In the main model, the holdup problem was created due to the fact that the action chosen at p^{FB}, the first best cutoff, was costly and the cost was borne by the worker. Here too, since the lowest effort chosen is $\underline{e}$, the choice of effort at p^{FB} will involve a cost to be borne by the worker. Therefore, the qualitative nature of the holdup problem is identical to the case of two efforts resulting in efficiency not being achievable in any relational contract. Moreover, a condition analogous to our Prop 4 shows that efforts that are worse for both information and payoff will be used in the optimal relational contract. To be precise, if $\beta(0) \geq \dfrac{\underline{\pi} + \underline{v}}{1 + \delta[\beta(\underline{e}) - (\underline{\pi} + \underline{v})]}$ then, $e = 0$ is used in any optimal relational contract.

While we have highlighted the form of inefficiency resulting in the use of efforts that are inferior for information as well as payoff, in this richer model we could also get the worker choosing an inefficiently higher effort. This can be seen numerically as shown in Figure 4. Notice that for a range of beliefs in the middle, the worker chooses an effort that is higher than the efficient benchmark. Also, as we can see, the optimal effort is non-monotonic in both the efficient benchmark and the optimal relational contract. This non-monotonicity arises because the informational and payoff rankings do not coincide. This is precisely why correlated bandits are difficult to analytically understand in such problems.

6 Related Literature

The early literature on relational contracts established their efficiency in various settings of complete and perfect information when the parties are sufficiently patient (**?**, **?**, **?**). This

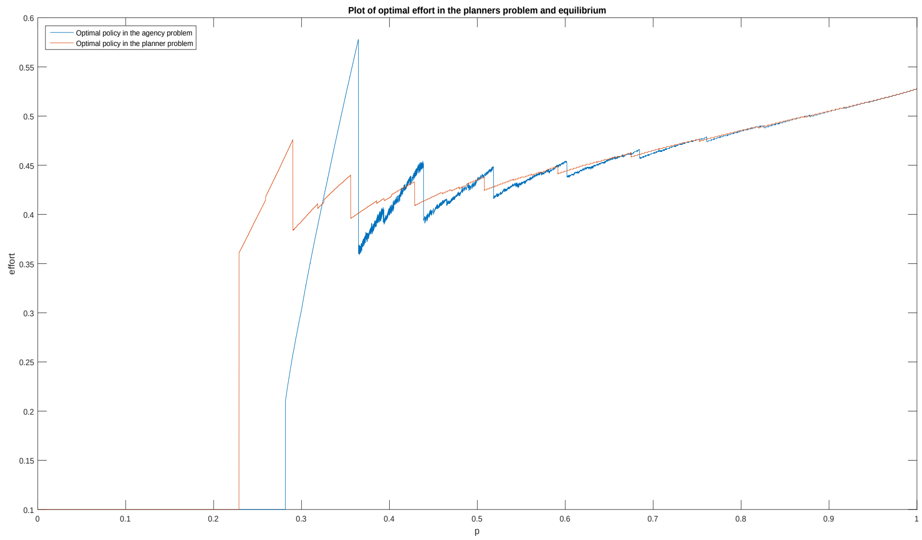

Figure 4: Plot of effort in the planner's and agency problem.

is no longer true in settings with private information such as moral hazard and adverse selection (**?**, **?**). Our paper shows that incomplete but *public* information alone can also create inefficiencies in relational contracts. Persistent uncertainty in relational contracts has also been studied by **?** and **?**. Both papers model learning as an exogenous process, whereas our paper allows the worker to control the flow of information through his effort.

Symmetric uncertainty has traditionally been explored in the career concerns literature starting with **?**. Typically, the agent is paid his marginal product every period, which leaves no room for dynamic incentive provision, e.g. backloaded compensation. Also, the career concerns literature focuses on the *signal jamming* incentives where the worker tries to influence the market's beliefs. Observable action rules out signal jamming.

Symmetric uncertainty has also been given attention in the dynamic contracting literature. With full commitment, the efficient outcome is trivially achieved in our setting. Hence, the dynamic contracting literature is often concerned with moral hazard. **?** analyse a continuous time model with normal learning and find that in order to avoid the ratchet effect (as in **?**,**?**) the contract induces the agent to shirk early on, and when the information about his ability is more precise, the rewards are more tightly linked to the output. These dynamics are similar to our findings that under limited commitment a relationship with perceived low match quality results in low effort and transitions to high effort as prospects improve. Due to the binary state of the match quality in our model, the precision of the signals does not improve over time so we can see switches between high and low effort due to streaks of luck, whereas in **?** the switch to high effort is permanent.

Finally, our conclusions share some similarities with the literature on relationship building which posits that favourable outcomes are achieved by establishing trust through repeated interactions. In **?**, **?**, and **?** these dynamics are driven by asymmetric information. In comparison, our paper contributes an alternative explanation which relies on symmetric information. In **?** information is complete and the parties acquire capital as they interact. They start with low-risk, low-return projects and move to more challenging projects once they have built up capital. The same dynamics can arise in our setting if we interpret an improvement in beliefs as creation of capital. However, the persistence of uncertainty distinguishes our results. Regardless of how many projects have been successful up to a certain point, poor match quality will almost surely result in separation in the long run.

7 Appendix

Proofs from Section 3.1

Proof of Proposition 1. G is the unique solution of the Bellman equation $f = Tf$. Standard contraction mapping arguments establish that G is increasing. To establish the convexity of G, let $f \in B[0,1]$ be convex. It suffices to show that $T_{\mathbb{H}}f$ and $T_{\mathbb{L}}f$ are convex since the maximum of finitely many convex functions is convex. Let $p_1, p_2 \in [0,1]$, $\lambda \in [0,1]$ and $\tilde{p} = \lambda p_1 + (1-\lambda)p_2$. Then

$$\begin{aligned}
T_{\mathbb{H}}f(\tilde{p}) &= (1-\delta)(p_{\mathbb{H}}(\tilde{p}) - c) + \delta \mathbb{E}^{\mathbb{H}} f(\tilde{p}) \\
&\leq (1-\delta)(p_{\mathbb{H}}(\tilde{p} - c) + \delta \mathbb{E}^{\mathbb{H}} \left[\lambda f(p_1) + (1-\lambda) f(p_2)\right] \\
&\leq \lambda \left[(1-\delta)(p_{\mathbb{H}}(p_1) - c) + \mathbb{E}^{\mathbb{H}} f(p_1)\right] + (1-\lambda)\left[(1-\delta)(p_{\mathbb{H}}(p_2) - c) + \delta \mathbb{E}^{\mathbb{H}} f(p_2)\right]
\end{aligned}$$

as required. The argument for the convexity of $T_{\mathbb{L}}f$ is analogous.

The convexity and increasingness of G imply that $T_{\mathbb{H}}G > T_{\mathbb{L}}G$ so $\mathbb{L}$ is never used in the optimal policy. The assumption $\alpha_{\mathbb{H}} - c > \underline{\pi} + \underline{v} > \alpha_{\mathbb{L}}$ implies that $G(1) = \alpha_{\mathbb{H}} - c$ and $\underline{\pi} + \underline{v} > \beta_{\mathbb{H}} - c > \beta_{\mathbb{L}}$ implies $G(0) = \underline{\pi} + \underline{v} > T_{\mathbb{H}}G(0)$. By continuity of G, there exists a maximal belief $p^{FB} \in (0,1)$ such that $G(p^{FB}) = \underline{\pi} + \underline{v}$. Since $T_{\mathbb{H}}G$ is strictly increasing, $G(p) > \underline{\pi} + \underline{v}$ for all $p > p^{FB}$. Hence, the optimal policy is to take the outside options at beliefs $p \leq p^{FB}$ and choose $\mathbb{H}$ at higher beliefs. Clearly, p^{FB} is unique. □

Proofs from Section 3.2

Here, we adapt the results of **?** to our setting. We begin by obtaining the worst equilibrium punishment payoffs. This allows us to obtain the simplified version of the recursive operator of **?** described in the main text.

Lemma 1. *$(\underline{\pi}, \underline{v}) \in \mathcal{E}(p)$ for all p.*

Proof. Consider the following strategies:

- The firm never interacts with the worker.
- The worker accepts wage offer w iff $w \geq \underline{v}$ and always chooses $\mathbb{L}$.

As the worker is never rewarded for choosing $\mathbb{H}$, he chooses $\mathbb{L}$ whenever he accepts the wage offer. Since he expects to receive his outside option in the future he only accepts wage offers above $\underline{v}$. Given that the worker chooses $\mathbb{L}$, the firm does not get enough expected output to justify paying the worker $\underline{v}$ since $p_{\mathbb{L}}(p) \leq \alpha_{\mathbb{L}} < \underline{\pi} + \underline{v}$. Hence, the strategies constitute an SE with payoffs $(\underline{\pi}, \underline{v})$. □

Let $\mathcal{C}$ denote the set of compact-valued payoff correspondences $W : [0,1] \to \mathbb{R}^2$ such that $(\underline{\pi}, \underline{v}) \in W(p)$ and $W(p) \subseteq \{(\pi, v) | (\pi, v) \geq (\underline{\pi}, \underline{v})\}$ for all p. We say payoffs (π, v) are enforceable with respect to $W \in \mathcal{C}$ at belief $p \in [0,1]$ if they can be decomposed by an action profile in the initial period and continuation payoffs as described in the recursive formulation in Section 3.2 except that the continuation payoffs are drawn from W instead of $\mathcal{E}$.

Definition 2. *Let $\mathcal{B}W(p)$ be the convex hull of the set of payoffs (π, v) which are enforceable with respect to $W \in \mathcal{C}$ at belief p.*

It is straightforward to check that $\mathcal{B}W \in \mathcal{C}$ whenever $W \in \mathcal{C}$.

Definition 3. *A correspondence $W \in \mathcal{C}$ is self-generating if $W(p) \subseteq \mathcal{B}W(p)$ for all p.*

Proposition 7 (Self-generation)**.** *If W is self-generating, then $\mathcal{B}W(p) \subseteq \mathcal{E}(p)$ for all p.*

Proof. Let $p \in [0,1]$ and $(\pi, v) \in \mathcal{B}W(p)$. If $(\pi, v) = (\underline{\pi}, \underline{v})$ then Lemma 1 implies $(\pi, v) \in \mathcal{E}(p)$. Otherwise, (π, v) is enforced by effort e, wage w and continuation payoffs $(\pi^u, v^u) \in \mathcal{B}W(p_e^u(p))$ and $(\pi^d, v^d) \in \mathcal{B}W(p_e^d(p))$ since W is self-generating. Construct inductively equilibrium strategies as follows: in the initial period, the firm offers wage w, the worker accepts and exerts effort e. Continuation strategies on the equilibrium path are constructed by decomposing the continuation payoffs inductively in the same manner. Following a deviation by any party, the continuation payoffs in the next period are $(\underline{\pi}, \underline{v})$ regardless of the rest of the history in this stage. If the firm deviates to a different wage offer $w' \geq \underline{v}$ the worker accepts and exerts low effort and if $w' < \underline{v}$ the worker takes his outside option (which is optimal in light of the constant continuation payoffs $(\underline{\pi}, \underline{v})$). Hence, the continuation payoff to any player following a deviation is at most his outside option.An application of the one-shot deviation principle shows that the strategies form an equilibrium. Boundedness of W guarantees that the resulting equilibrium payoffs equal (π, v), as required. $\square$

Proposition 8 (Factorisation)**.** *$\mathcal{B}\mathcal{E}(p) = \mathcal{E}(p)$ for all p.*

Proof. It suffices to show that $\mathcal{E}$ is self-generating. To this end, let $p \in [0,1]$, $(\pi, v) \in \mathcal{E}(p)$ and assume, without loss of generaility, that no public randomisation is needed to support (π, v). We want to show $(\pi, v) \in \mathcal{B}\mathcal{E}(p)$. If $(\pi, v) = (\underline{\pi}, \underline{v})$, this is trivial. Otherwise, (π, v) is generated by some effort e, wage w and continuation payoffs which belong to the equilibrium set at the corresponding updated belief. Hence, we only need to check $\pi \geq \underline{\pi}$, $v \geq \underline{v}$ and the incentive constraint for effort (4) if $e = \mathbb{H}$. The first two conditions are trivial since any equilibrium payoffs must exceed the outside options. The incentive constraint must be satisfied; otherwise the worker can profitably deviate to $\mathbb{L}$. Hence, $(\pi, v) \in \mathcal{B}\mathcal{E}(p)$, as required. $\square$

Proposition 9 (APS algorithm). *Let $W_0 \in \mathcal{C}$ satisfy $W_0(p) \supseteq \mathcal{B}W_0(p)$ and $W_0(p) \supseteq \mathcal{E}(p)$ for all p. Inductively define $W_n = \mathcal{B}W_{n-1}$ for all $n \in \mathbb{N}$. Let $W^*(p) = \cap_{n\in\mathbb{N}} W_n(p)$ for all p. Then $W^*(p) = \mathcal{E}(p)$ for all p.*

Proof. Observe that $\mathcal{B}$ is monotonic: any correspondence $U \in \mathcal{C}$ with $U(p) \subseteq U'(p)$ for all p satisfies $\mathcal{B}U(p) \subseteq \mathcal{B}U'(p)$ for all p. Hence $W_{n-1}(p) \supseteq W_n(p)$ for all n, p and consequently, $W^*(p) \supseteq \mathcal{E}(p)$. To show the converse inclusion it suffices to show W^* is self-generating. To this end, let $(\pi, v) \in W^*(p)$ for some p. Then $(\pi, v) \in \mathcal{B}W_n(p)$ for all n. We want to show $(\pi, v) \in \mathcal{B}W^*(p)$. It is sufficient to consider $(\pi, v) \neq (\underline{\pi}, \underline{v})$. Then, there exists an effort level e such that $(\pi, v) \in \mathcal{B}W_n(p)$ for infinitely many n. Along this subsequence, let w_n denote the wage offered, (π_n^u, v_n^u) and (π_n^d, v_n^d) be the continuation payoffs following success and failure respectively. Without loss of generaility these sequences converge to w, (π^u, v^u) and (π^d, v^d) respectively. Compactness of $W_n(p_e^u)$ and $W_n(p_e^d)$ and decreasingness in n imply $(\pi^u, v^u) \in W^*(p_e^u)$ and $(\pi^d, v^d) \in W^*(p_e^d)$. Hence, (π, v) is enforceable with respect to W^* at belief p through wage w and continuation payoffs (π^u, v^u) and (π^d, v^d), that is $(\pi, v) \in \mathcal{B}W^*(p)$. $\square$

Proofs from Section 3.3

It will be useful to operate on payoff correspondences that, like the equilibrium correspondence, are fully described by the maximum sum of payoffs that can be achieved at every belief.

Definition 4. *Let $\mathcal{D}$ be the set of all correspondences $W \in \mathcal{C}$ such that*

$$W(p) = \{(\pi, v) | (\pi, v) \geq (\underline{\pi}, \underline{v}), \pi + v \leq S^W(p)\} \quad \textit{for all } p$$

where $S^W : [0, 1] \to [\underline{\pi} + \underline{v}, \alpha_{\mathbb{H}} - c]$.

Lemma 2. *If $W \in \mathcal{D}$, then $\mathcal{B}W \in \mathcal{D}$.*

Proof. Let

$$S^{\mathcal{B}W}(p) = \max_{(\pi,v)\in\mathcal{B}W(p)} \pi + v \qquad \text{for all } p.$$

Let $p \in [0, 1]$. Notice that $S^{\mathcal{B}W}(p) \geq \underline{\pi} + \underline{v}$ by definition of the operator $\mathcal{B}$. Moreover, $\alpha_{\mathbb{H}} - c$ is an upper bound on the sum of payoffs in a single period so $S^{\mathcal{B}W}(p) \in [\underline{\pi} + \underline{v}, \alpha_{\mathbb{H}} - c]$. It only remains to show that if $(\pi, v) \geq (\underline{\pi}, \underline{v})$ and $\pi + v = S^{\mathcal{B}W}$ then $(\pi, v) \in \mathcal{B}W(p)$ (the rest follows from taking convex combinations of $(\underline{\pi}, \underline{v})$ and efficient payoffs in $\mathcal{B}W(p)$). The statement clearly holds if $S^{\mathcal{B}W}(p) = \underline{\pi} + \underline{v}$ so suppose $S^{\mathcal{B}W}(p) > \underline{\pi} + \underline{v}$ and consider any such (π, v). There exists a payoff $(\pi', v') \in \mathcal{B}W(p)$ with $\pi + v = S^{\mathcal{B}W}(p)$ which does not involve public randomisation, since $\mathcal{B}W(p)$ is a compact, convex subset of $[\underline{\pi}, \infty) \times [\underline{v}, \infty)$. The payoff (π, v) can be supported by the same effort and continuation payoffs as (π', v') through an adjustment of the wage. $\square$

Lemma 3 states that in optimal relational contracts it is without loss of generality to give the worker the highest equilibrium continuation payoff. This can only improve the incentives for the worker to exert $\mathbb{H}$ and the wage can be used to redistribute the surplus without violating incentives for the firm. The same result holds when $\mathbb{L}$ is enforced but is not needed.

Lemma 3. *Let $W \in \mathcal{D}$, $p \in [0,1]$. Consider a payoff (π, v) such that*

$$\pi + v = \max\{\pi' + v' | (\pi', v') \in \mathcal{B}W(p) \text{ is enforced by } \mathbb{H}\}.$$

Then (π, v) can be enforced by $\mathbb{H}$ and continuation payoffs $\hat{v}^i = S^W(p_e^i) - \underline{\pi}$ for the worker.

Proof. By definition, (π, v) is enforced by some wage w and continuation payoffs $(\pi^i, v^i) \in W(p_{\mathbb{H}}^i)$ for $i = u, d$. Let $\hat{\pi}^i = \underline{\pi}$ for $i = u, d$. Since $W \in \mathcal{D}$ we have $(\hat{\pi}^i, \hat{v}^i) \in W(p_{\mathbb{H}}^i)$ for all i. Consider a wage

$$\hat{w} = w - \frac{\delta}{1-\delta}[p_{\mathbb{H}}(p)(\pi^u - \underline{\pi}) + (1 - p_{\mathbb{H}}(p))(\pi^d - \underline{\pi})].$$

We want to show that (π, v) is also enforced by $\mathbb{H}$, wage $\hat{w}$ and continuation payoffs $(\hat{\pi}^u, \hat{v}^u)$ and $(\hat{\pi}^d, \hat{v}^d)$ following success and failure respectively.

Efficiency of (π, v) implies that $\pi^i + v^i = S(p_{\mathbb{H}}^i)$ for all i so the promise keeping constraints readily hold:

$$\begin{aligned}
\pi &= (1-\delta)[p_{\mathbb{H}}(p) - \hat{w}] + \delta[p_{\mathbb{H}}(p)\hat{\pi}^u + (1 - p_{\mathbb{H}}(p))\hat{\pi}^d] \\
v &= (1-\delta)[\hat{w} - c] + \delta[p_{\mathbb{H}}(p)\hat{v}^u + (1 - p_{\mathbb{H}}(p))\hat{v}^d].
\end{aligned}$$

It remains to check the incentive constraints. Clearly, $(\pi, v) \geq (\underline{\pi}, \underline{v})$ so we only need to check the incentive constraint for effort. But for $i = u, d$

$$S^W(p_{\mathbb{H}}^i) = \pi^i + v^i = \underline{\pi} + \hat{v}^i$$

so $\pi^i \geq \underline{\pi}$ implies $\hat{v}^i \geq v^i$. Since the incentive constraint holds in the original decomposition with continuation payoffs v^u, v^d we have

$$\begin{aligned}
&p_{\mathbb{H}}(p)\hat{v}^u + (1 - p_{\mathbb{H}}(p))\hat{v}^d - \underline{v} \\
\geq &p_{\mathbb{H}}(p)v^u + (1 - p_{\mathbb{H}}(p))v^d - \underline{v} \geq \frac{1-\delta}{\delta}c
\end{aligned}$$

which completes the proof. $\square$

Now we are ready to show the equivalence of operating on the frontier of the equilibrium set through the APS operator $\mathcal{B}$ and operating on the maximum equilibium surplus function through T_*.

Lemma 4. *$S^{\mathcal{B}W}(p) = T_* S^W(p)$ for all p.*

Proof. It suffices to show that for any belief p

- The maximum sum of payoffs in $\mathcal{B}W(p)$ decomposed by $\mathbb{L}$, if it exists, is equal to $T_{\mathbb{L}} S^W(p)$. If no payoffs in $\mathcal{B}W(p)$ are decomposed by $\mathbb{L}$ we have $T_{\mathbb{L}} S^W(p) < \underline{\pi} + \underline{v}$.

- The maximum sum of payoffs in $\mathcal{B}W(p)$ decomposed by $\mathbb{H}$, if it exists, equals $T_{\mathbb{H}} S^W(p)$ and the incentive constraint below holds.

$$p_{\mathbb{H}}(p) S^W(p^u_{\mathbb{H}}(p)) + (1 - p_{\mathbb{H}}(p)) S^W(p^d_{\mathbb{H}}(p)) \geq \underline{\pi} + \underline{v} + \frac{1-\delta}{\delta} c \tag{5}$$

  If no payoffs in $\mathcal{B}W(p)$ are decomposed by $\mathbb{H}$ we have either $T_{\mathbb{H}} S^W(p) < \underline{\pi} + \underline{v}$ or the incentive constraint (5) does not hold.

Let $p \in [0, 1]$. The maximum sum of payoffs in $\mathcal{B}W(p)$ decomposed by e, if it exists, is given by $T_e S^W(p)$; the effort pins down the flow payoff as the wage transfer cancels out and maximality implies that the players share the maximum future surplus, as higher payoffs never hurt incentives. Moreover, in the case of $e = \mathbb{H}$, Lemma 3 states that this maximum sum of payoffs can be achieved by setting continuation payoffs $v^i = S^W(p^i_{\mathbb{H}}) - \underline{\pi}$ for $i = u, d$. Therefore, the incentive constraint

$$p_{\mathbb{H}}(p) v^u + (1 - p_{\mathbb{H}}(p)) v^d \geq \underline{v} + \frac{1-\delta}{\delta} c$$

becomes equivalent to (5).

Hence, it remains to show that when $T_e S^W(p) \geq \underline{\pi} + \underline{v}$, and, in the case of $e = \mathbb{H}$, (5) holds, there exist payoffs in $\mathcal{B}W(p)$ decomposed by effort e. This can be achieved by setting continuation payoffs $\pi^i = \underline{\pi}$ and $v^i = S(p^i_{\mathbb{L}}) - \underline{\pi}$ for $i = u, d$ and adjusting the wage to ensure that each player receives at least his outside option. In the case of $e = \mathbb{H}$, this composition of continuation payoffs implies that (5) is equivalent to the incentive constraint for effort. This completes the proof. $\square$

Now we have all the ingredients to the proof of Proposition 2.

Proof of Proposition 2. Since $\mathcal{E} \in \mathcal{D}$ and $S \equiv S^{\mathcal{E}}$, Lemma 4 and Proposition 8 imply $S(p) = T_* S(p)$ for all p. Thus, S is a fixed point of T_*.

Now take a function f as in the second part of the proposition and let $W_0(p) = \{(\pi, v) \geq (\underline{\pi}, \underline{v}) | \pi + v \leq f(p)\}$ for all p. Hence, $W_0 \in \mathcal{D}$. Lemma 4 and $f \geq T_* f$ imply $W_0(p) \supseteq \mathcal{B}W_0(p)$, whereas $f \geq S$ implies that $W_0(p) \supseteq \mathcal{E}(p)$ for all p. Thus, W_0 satisfies the conditions in Proposition 9 applies. In its notation, we have $W_n \in \mathcal{D}$ for all n by an inductive application of Lemma 2. Thus, $S^{W_n}(p) = T^n_* S^{W_0}(p)$ for all n and p

by Lemma 4. Since $\mathcal{E} \in \mathcal{D}$ and $W_n(p)$ converges to $\mathcal{E}(p)$ for all p, we have that $S^{W_n}(p)$ converges to $S(p)$ for every p. Moreover, $W_n(p) \supseteq W_{n+1}(p)$ implies $T_*^n f \geq T_*^{n+1} f$.

Finally, G satisfies that conditions for a starting function since the equilibrium cannot improve upon the efficient outcome ($G \geq S$) and $G = \max\{T_{\mathbb{H}}G, T_{\mathbb{L}}G, \underline{\pi} + \underline{v}\} \geq T_*G$. $\square$

Proofs from Section 4.1

Proof of Proposition 3. Rearranging

$$G(p^{FB}) = \underline{\pi} + \underline{v} = (1-\delta)(p_{\mathbb{H}}(p^{FB}) - c) + \delta \mathbb{E}^{\mathbb{H},p^{FB}}[G(p')]$$

we obtain

$$(1-\delta)c = (1-\delta)(p_{\mathbb{H}}(p^{FB}) - (\underline{\pi} + \underline{v})) + \delta \mathbb{E}^{\mathbb{H},p^{FB}}[G(p') - (\underline{\pi} + \underline{v})]. \tag{6}$$

Consider the application of T_* to G. If $p_{\mathbb{H}}(p^{FB}) \leq \underline{\pi} + \underline{v}$, (6) implies

$$\delta \mathbb{E}^{\mathbb{H},p^{FB}}[G(p') - (\underline{\pi} + \underline{v})] \geq (1-\delta)c$$

so the incentive constraint in the definition of T_* is satisfied at all beliefs $p \geq p^{FB}$. Hence, $T_*G(p) = \max\{T_{\mathbb{H}}G(p), T_{\mathbb{L}}G(p), \underline{\pi} + \underline{v}\} = G(p)$ for all $p \geq p^{FB}$. Moreover, $\underline{\pi} + \underline{v} \leq T_*G(p) \leq G(p) = \underline{\pi} + \underline{v}$ for all $p \leq p^{FB}$. Hence, $T_*G = G$. The algorithmic part of Proposition 2 implies $S = G$ so relational contracts are not inefficient.

If, instead, $p_{\mathbb{H}}(p^{FB}) > \underline{\pi} + \underline{v}$, (6) implies

$$\delta \mathbb{E}^{\mathbb{H},p^{FB}}[G(p') - (\underline{\pi} + \underline{v})] < (1-\delta)c.$$

By continuity of G and the belief updates $p_{\mathbb{H}}^u$ and $p_{\mathbb{L}}^u$, the above strict inequality holds for some $p > p^{FB}$, hence the incentive constraint for $\mathbb{H}$ is not satisfied at p. Then $T_*G(p) = \max\{T_{\mathbb{L}}G(p), \underline{\pi} + \underline{v}\}$. Since $G(p) = T_{\mathbb{H}}G(p) > T_{\mathbb{L}}G(p)$, Proposition 2 implies $S(p) < G(p)$. Hence, relational contracts are inefficient. $\square$

Proofs from Section 4.2

To prove Theorem 1 we first establish some properties of S similar to, but weaker than, the properties of G in Proposition 1.

Lemma 5. *S is increasing and upper semicontinuous with $S(0) = \underline{\pi} + \underline{v}$ and $S(1) = \alpha_{\mathbb{H}} - c$. There exists a unique belief $p^* \in (0,1]$ such that $S(p) = \underline{\pi} + \underline{v}$ for $p < p^*$ and $S(p) > \underline{\pi} + \underline{v}$ for $p > p^*$.*

Proof of Lemma 5. We first show that S is increasing and upper semicontinuous (usc) using Proposition 2. Since G is a valid starting function which is increasing and usc, it suffices to show that T_* maps increasing and usc functions to increasing and usc functions. Let f be one such function. Notice that the incentive constraint for T_*f given by

$$\delta\left[p_{\mathbb{H}}(p)f(p^u_{\mathbb{H}}(p)) + (1 - p_{\mathbb{H}}(p))f(p^d_{\mathbb{H}}(p)) - \underline{\pi}\right] \geq \delta\underline{v} + (1-\delta)c$$

can be rewritten as $h(p) \geq A$ where A is a constant and h is increasing and usc. Moreover, T_*f is larger if the incentive constraint is satisfied, ceteris paribus. Hence, it is sufficient to show $T_{\mathbb{H}}f$ and $T_{\mathbb{L}}f$ are increasing and usc. This follows directly from the fact that f has the same properties.

The algorithmic part of Proposition 2 with starting function G also shows the values of S at 0 and 1. It suffices to show that $T_*G(0) = \underline{\pi} + \underline{v}$ and $T_*G(1) = \alpha_{\mathbb{H}} - c$. The former follows from $\underline{\pi} + \underline{v} \leq T_*G(0) \leq TG(0) = G(0) = \underline{\pi} + \underline{v}$, while the latter is a consequence of Assumption 3.

Since $\alpha_{\mathbb{H}} - c > \underline{\pi} + \underline{v}$, increasingness of S implies there exists a maximal belief p^* such that $S(p) = \underline{\pi} + \underline{v}$ if $p < p^*$ and $S(p) = \max\{T_{\mathbb{H}}S(p), T_{\mathbb{L}}S(p)\}$ if $p \geq p^*$. That p^* is unique follows from the strict increasingness of $T_{\mathbb{H}}S$, $T_{\mathbb{L}}S$ and h.

□

Next, we establish that the three region structure in Theorem 1 is obtained if the surplus function exhibits a monotone policy.

Lemma 6. *If S exhibits a monotone policy there exist beliefs $\underline{p}, \bar{p}$ such that*

$$S(p) = \begin{cases} \underline{\pi} + \underline{v} & \text{if } p < \underline{p} \\ T_{\mathbb{L}}S(p) & \text{if } \underline{p} \leq p < \bar{p} \\ T_{\mathbb{H}}S(p) & \text{if } p \geq \bar{p} \end{cases}$$

Proof. By Lemma 5 there is a threshold belief p^* such that the players take their outside options at beliefs $p < p^*$ and interact if $p \geq p^*$. Increasingness of S implies that if the incentive constraint is satisfied at some belief it is satisfied for all higher beliefs. Let p^{IC} be the minimum belief at which the incentive constraint for $\mathbb{H}$ is satisfied. Note that such a belief exists due to $S(1) = 1$ and the infimum is attained due to the upper semicontinuity of S. Hence, $T_*f(p) = T_{\mathbb{H}}f(p)$ iff $p \geq \max\{p^*, p^{\text{IC}}\}$. Putting $\underline{p} = p^*$ and $\bar{p} = \max\{p^*, p^{\text{IC}}\}$ completes the proof. □

The final piece of the proof of Theorem 1 is the following lemma which ensures that the monotone policy property is preserved under T_*.

Lemma 7. *Let f be an increasing function that exhibits a monotone policy. Then T_*f is increasing and exhibits a monotone policy.*

Proof. The proof of Lemma 5 showed that $T_* f$ is increasing when f is increasing. Hence, there exists a maximal belief $p^* \leq 1$ such that $T_* f(p) = \underline{\pi} + \underline{v}\ \forall p < p^*$. It remains to show that $T_* f$ exhibits a monotone policy. Let $p \in [0,1]$. We will show that $T_{\mathbb{H}}(T_* f)(p) > T_{\mathbb{L}}(T_* f)(p)$. We examine two cases depending on where the posterior beliefs lie relative to the cutoff p^*.

Firstly, if $p^* > p^d_{\mathbb{L}}(p)$, we have $T_* f(p^d_{\mathbb{L}}) = \underline{\pi} + \underline{v}$. Also, $p^d_{\mathbb{L}} \geq p^d_{\mathbb{H}}$ implies that $T_* f(p^d_{\mathbb{H}}) = \underline{\pi} + \underline{v}$ since $T_* f$ is increasing. Hence,

$$
\begin{aligned}
&T_{\mathbb{H}}(T_* f)(p) - T_{\mathbb{L}}(T_* f)(p) \\
\geq &(1-\delta)(p_{\mathbb{H}}(p) - c) + \delta[p_{\mathbb{H}}(p) T_* f(p^u_{\mathbb{H}}(p)) + (1 - p_{\mathbb{H}}(p))(\underline{\pi} + \underline{v})] \\
&- (1-\delta)p_{\mathbb{L}}(p) - \delta[p_{\mathbb{L}}(p) T_* f(p^u_{\mathbb{L}}(p)) + (1 - p_{\mathbb{L}}(p))(\underline{\pi} + \underline{v})] \\
\geq &(1-\delta)(p_{\mathbb{H}}(p) - c - p_{\mathbb{L}}(p)) > 0.
\end{aligned}
$$

where the second inequality follows from $p_{\mathbb{H}}(p) > p_{\mathbb{L}}(p)$ and $T_* f(p^u_{\mathbb{H}}(p)) \geq T_* f(p^u_{\mathbb{L}}(p)) \geq \underline{\pi} + \underline{v}$.

Secondly, suppose $p^* \leq p^d_{\mathbb{L}}(p)$. Then $T_* f(q') \in \{T_{\mathbb{H}} f(q'), T_{\mathbb{L}} f(q')\}$ for any $q' \in \{p^u_{\mathbb{L}}(p), p^d_{\mathbb{L}}(p)\}$ which implies $T_* f(q') \leq T_{\mathbb{H}} f(q')$ since f exhibits a monotone policy. Thus, $T_{\mathbb{L}}(T_* f)(p) \leq T_{\mathbb{L}}(T_{\mathbb{H}} f)(p)$ and it follows that

$$
\begin{aligned}
&T_{\mathbb{H}}(T_* f)(p) - T_{\mathbb{L}}(T_* f)(p) \\
\geq &T_{\mathbb{H}}(T_{\mathbb{L}} f)(p) - T_{\mathbb{L}}(T_{\mathbb{H}} f)(p) \\
= &(1-\delta)(p_{\mathbb{H}}(p) - c) + \delta(1-\delta)(\mathbb{E}^{\mathbb{H},p}[p_{\mathbb{L}}(p')]) + \delta^2 \mathbb{E}^{\mathbb{H},p} \mathbb{E}^{\mathbb{L},p'}[f(p'')] \\
- &\Big((1-\delta)(p_{\mathbb{L}}(p)) + \delta(1-\delta)(\mathbb{E}^{\mathbb{L},p}[p_{\mathbb{H}}(p') - c]) + \delta^2 \mathbb{E}^{\mathbb{L},p} \mathbb{E}^{\mathbb{H},p'}[f(p'')]\Big) \\
= &(1-\delta)(p_{\mathbb{H}}(p) - c) + \delta(1-\delta) p_{\mathbb{L}}(p) \\
- &\Big((1-\delta)(p_{\mathbb{L}}(p) - c) + \delta(1-\delta)(p_{\mathbb{H}}(p) - c)\Big) \\
= &(1-\delta)^2 (p_{\mathbb{H}}(p) - c - p_{\mathbb{L}}(p)) > 0
\end{aligned}
$$

where $\mathbb{E}^{e,p}[g(p')]$ ($\mathbb{E}^{e,p'}[g(p'')]$) is the expectation of g at the posterior belief p' (p'') given effort e and prior p (p'). This completes the proof. $\square$

Proof of Theorem 1. We know from Proposition 2 that $S = \lim_{n\to\infty} G_n$ where $G_n = T_*^n G$. Since G is increasing and exhibits a monotone policy (Proposition 1), Lemma 7 implies that G_n exhibits a monotone policy for all n. Taking limits yields $T_{\mathbb{H}} S(p) \geq T_{\mathbb{L}} S(p)$ for all p, so S exhibits a monotone policy. Lemma 6 completes the proof. $\square$

Proof of Proposition 4. Suppose, for the sake of contradiction, that $\underline{p} = \bar{p}$. Consider some belief $p = \bar{p} - \varepsilon$ for some small $\varepsilon > 0$. Since the incentive constraint for $\mathbb{H}$ is not satisfied at p, it suffices to show that $\mathbb{L}$ generates more surplus than the outside options.

We know that $S(p^u_{\mathbb{L}}(p)) \geq S(\bar{p})$ whenever ε is small enough. So the joint surplus from $\mathbb{L}$ at belief p is at least

$$\begin{aligned}
&(1-\delta)p_{\mathbb{L}}(p) + \delta\Big[p_{\mathbb{L}}(p)S(\bar{p}) + (1-p_{\mathbb{L}}(p))(\underline{\pi}+\underline{v})\Big] \\
=&(1-\delta)p_{\mathbb{L}}(p) + \delta\Big[p_{\mathbb{L}}(p)\Big((1-\delta)p_{\mathbb{H}}(\bar{p}) + \delta(\underline{\pi}+\underline{v})\Big) + (1-p_{\mathbb{L}}(p))(\underline{\pi}+\underline{v})\Big] \\
=&(1-\delta)p_{\mathbb{L}}(p) + \delta\Big[p_{\mathbb{L}}(p)(1-\delta)\Big(p_{\mathbb{H}}(\bar{p}) - (\underline{\pi}+\underline{v})\Big) + (\underline{\pi}+\underline{v})\Big] \\
=&(1-\delta)p_{\mathbb{L}}(p)\Big(1+\delta(p_{\mathbb{H}}(\bar{p}) - (\underline{\pi}+\underline{v}))\Big) + \delta(\underline{\pi}+\underline{v})
\end{aligned}$$

This exceeds $\underline{\pi}+\underline{v}$ whenever

$$p_{\mathbb{L}}(p)\Big(1+\delta(p_{\mathbb{H}}(\bar{p}) - (\underline{\pi}+\underline{v}))\Big) \geq \underline{\pi}+\underline{v}$$

which is guaranteed if

$$\begin{aligned}
p_{\mathbb{L}}(0)\Big(1+\delta(p_{\mathbb{H}}(0) - (\underline{\pi}+\underline{v}))\Big) &\geq \underline{\pi}+\underline{v} \\
\Leftrightarrow \beta_{\mathbb{L}} &\geq \frac{\underline{\pi}+\underline{v}}{1+\delta(\beta_{\mathbb{H}} - (\underline{\pi}+\underline{v}))}.
\end{aligned}$$

□

Proofs from Section 4.3

Lemma 8. *(a) Without loss of generality whenever the worker chooses $\mathbb{L}$ his continuation payoffs following success and failure are $\underline{v}$.*

(b) There exists $\hat{v}^d$ such that without loss of generality whenever the worker chooses $\mathbb{H}$ his incentive constraint for effort binds and his continuation payoff following failure is $\hat{v}^d$.

Proof. Fix a belief p and consider an equilibrium with payoffs $(\pi, v) \in \mathcal{E}(p)$ where the worker chooses effort e in the initial period, the wage is w and continuation payoffs following success and failure are (π^u, v^u) and (π^d, v^d).

To prove part (a) suppose that $e = \mathbb{L}$. Let $\hat{v}^i = \underline{v}$ and $\hat{\pi}^i = \pi^i + v^i - \hat{v}^i$, $i \in \{u, d\}$. Hence, $(\pi^i, v^i) \in \mathcal{E}(p^i_{\mathbb{L}}(p))$ for all i. Let

$$\hat{w} = w + \frac{\delta}{1-\delta}\Big(p_{\mathbb{L}}(p)(v^u - \hat{v}^u) + (1-p_{\mathbb{L}}(p))(v^d - \hat{v}^d)\Big) \tag{7}$$

and notice that

$$\pi = (1-\delta)(p_{\mathbb{H}}(p) - \hat{w}) + \delta(p_{\mathbb{H}}(p)\hat{\pi}^u + (1-p_{\mathbb{H}}(p))\hat{\pi}^d)$$
$$v = (1-\delta)(\hat{w} - c) + \delta(p_{\mathbb{H}}(p)\hat{v}^u + (1-p_{\mathbb{H}}(p))\hat{v}^d)$$

Thus, we have constructed an equilibrium where in the initial period the wage is $\hat{w}$, the worker exerts $\mathbb{L}$ and his continuation payoffs equal $\underline{v}$.

To show (b) suppose $e = \mathbb{H}$. Since the incentive constraint for effort can be satisfied in equilibrium at belief $\bar{p}$ and the worker's continuation payoffs must satisfy $\bar{v}^i \in [\underline{v}, S(p^i_{\mathbb{H}}(\bar{p})) - \underline{\pi}]$ for all $i \in \{u, d\}$ we must have

$$p_{\mathbb{H}}(\bar{p})\Big(S(p^u_{\mathbb{H}}(\bar{p})) - \underline{\pi}\Big) + (1-p_{\mathbb{H}}(\bar{p}))\Big(S(p^d_{\mathbb{H}}(\bar{p})) - \underline{\pi}\Big) - \underline{v} \geq \frac{1-\delta}{\delta}c$$

Let

$$\hat{v}^d = \min\Big\{\bar{v}^d \in [\underline{v}, S(p^d_{\mathbb{H}}(\bar{p})) - \underline{\pi}]\Big|\exists \bar{v}^u \in [\underline{v}, S(p^u_{\mathbb{H}}(\bar{p})) - \underline{\pi}] \text{ with}$$
$$p_{\mathbb{H}}(\bar{p})\bar{v}^u + (1-p_{\mathbb{H}}(\bar{p}))\bar{v}^d - \underline{v} = \frac{1-\delta}{\delta}c\Big\}$$

Let $\bar{v}^u$ be the associated continuation payoff following success that makes the incentive constraint bind. By the minimality of $\hat{v}^d$ and increasingness of S we must have $\bar{v}^u \geq \hat{v}^d$.

Since $\bar{p}$ is the lowest belief where the incentive constraint for $\mathbb{H}$ is satisfied, we must have $p \geq \bar{p}$. Since $\bar{v}^u \geq \bar{v}^d$ and $p_{\mathbb{H}}$ is increasing, there exist $\hat{v}^u \in [\underline{v}, S(p^u_{\mathbb{H}}(p)) - \underline{\pi}]$ such that $\hat{v}^u \geq \hat{v}^d$ and

$$p_{\mathbb{H}}(p)\hat{v}^u + (1-p_{\mathbb{H}}(p))\hat{v}^d = \frac{1-\delta}{\delta}c$$

Letting $\hat{\pi}^i = \pi^i + v^i - \hat{v}^i$ we have $(\hat{\pi}^i, \hat{v}^i) \in \mathcal{E}(p^i_{\mathbb{H}}(p))$ for all i. So we can define a wage $\hat{w}$ as in (7) to construct an alternative equilibrium with the same payoffs where the incentive constraint for effort binds, the worker exerts $\mathbb{H}$ and receives $\hat{v}^d$ upon failure. □

Proof of Proposition 5. By Lemma 8 there is a payoff-equivalent equilibrium where the worker receives continuation payoff $\underline{v}$ following $\mathbb{L}$, the incentive constraint for effort binds if he exerts $\mathbb{H}$ and his continuation payoff following $\mathbb{H}$ and failure is $\hat{v}^d$.

Consider any period where the worker exerts some effort e. Let the belief at the beginning of the period be p. Suppose the worker exerts effort e' in the next period where the updated belief is $p' \in \{p^u_e(p), p^d_e(p)\}$, and the corresponding continuation payoff for the worker v'. Let the next period wage be w'. If $e' = \mathbb{H}$ the binding incentive constraint for effort implies

$$\begin{aligned} v' &= (1-\delta)(w'-c) + \delta\Big(p_{\mathbb{H}}(p')v^u(p') + (1-p_{\mathbb{H}}(p'))v^d(p')\Big) \\ &= (1-\delta)(w'-c) + \delta\left(\underline{v} + \frac{1-\delta}{\delta}c\right) \\ &= (1-\delta)w' + \delta\underline{v} \end{aligned}$$

and if $e' = \mathbb{L}$ the worker receives a continuation payoff of $\underline{v}$ so

$$v' = (1-\delta)w' + \delta\underline{v}$$

so in both cases

$$w' = \frac{v' - \delta\underline{v}}{1-\delta}$$

If $e = \mathbb{L}$, $v' = \underline{v}$ so the wage the worker receives in the subsequent period is $\underline{v}$. Let $\underline{w} = \frac{\hat{v}^d - \delta\underline{v}}{1-\delta}$. If $e = \mathbb{H}$ the continuation payoff following failure is $\hat{v}^d$ and the associated wage is $\underline{w}$. The continuation payoff following success can be pinned down by the incentive constraint:

$$v^u = \frac{1}{p_{\mathbb{H}}(p)}\left[\frac{1-\delta}{\delta}c - \Big((1-p_{\mathbb{H}}(p))\hat{v}^d - \underline{v}\Big)\right]$$

Substituting into the expression for w' yields the desired expression for the wage.

Finally, if the worker exerts $\mathbb{L}$ in the initial period, his continuation payoff is $\underline{v}$ so the initial wage w_0 satisfies

$$v = (1-\delta)w_0 + \delta\underline{v}.$$

If the worker exerts $\mathbb{H}$ in the initial period then his expected continuation payoff is $\underline{v} + \frac{1-\delta}{\delta}c$ since the incentive constraint binds which leads to the same expression for w_0. □

7.1 Proofs from Section 5.1

Proof of Proposition 6. Let w_s and w_f be the wages upon success and failure and let v be the random continuation value to the worker. The worker's incentive constraint to choose $\mathbb{H}$ over $\mathbb{L}$ is given by

$$\begin{aligned} &(1-\delta)\left[p_{\mathbb{H}}(p)w_s + (1-p_{\mathbb{H}}(p))w_f - c\right] + \delta\mathbb{E}^{\mathbb{H}}v \geq (1-\delta)\left[p_{\mathbb{L}}(p)w_s + (1-p_{\mathbb{L}}(p))w_f\right] + \delta\underline{v} \\ &(1-\delta)(w_s - w_f)(p_{\mathbb{H}}(p) - p_{\mathbb{L}}(p)) + \delta(E^{\mathbb{H}}v - \underline{v}) \geq (1-\delta)c \end{aligned}$$

Without limited liability, the firm can always increase the spread between w_s and w_f to satisfy the above incentive constraint without affecting the expected wage payment $p_{\mathbb{H}}(p)w_s + (1-p_{\mathbb{H}}(p))w_f$. Hence, the incentive constraint for $\mathbb{H}$ can be satisfied whenever the relationship yields more surplus than $\underline{\pi} + \underline{v}$. Therefore, the equilibrium is efficient.

With limited liability, however, the firm can no longer create an arbitrary wage spread since $w_s, w_f \geq \underline{v}$. It is without loss of generality to consider $w_f = \underline{v}$ as this achieves the maximum possible spread given limited liability.

Recall equation (6)

$$(1-\delta)c = (1-\delta)(p_{\mathbb{H}}(p^{FB}) - (\underline{\pi} + \underline{v})) + \delta \mathbb{E}^{\mathbb{H},p^{FB}}[G(p') - (\underline{\pi} + \underline{v})]$$

Suppose $p_{\mathbb{H}}(p^{FB}) \leq \underline{\pi} + \underline{v}$. We want to show $T_* G = G$. Similarly to the proof of Proposition 3 it suffices to show that the incentive constraints for $\mathbb{H}$ can be satisfied at p^{FB} when continuation payoffs are drawn from G. Equation (6) implies

$$\delta \mathbb{E}^{\mathbb{H},p^{FB}}[G(p') - (\underline{\pi} + \underline{v})] \geq (1-\delta)c$$

so we can set the worker's continuation payoffs $v^u, v^d \geq \underline{v}$ following success and failure so that

$$\delta(p_{\mathbb{H}}(p^{FB})(v^u - \underline{v}) + (1 - p_{\mathbb{H}}(p^{FB}))(v^d - \underline{v})) = (1-\delta)c$$

which makes the continuation payoffs for the firm given by $\pi^u = G(p^u(p^{FB})) - v^u$ and $\pi^d = G(p^d(p^{FB})) - v^d$ greater than or equal to $\underline{\pi}$. By setting $w_s = w_f = \underline{v}$ the overall payoff of the worker is $\underline{v}$ and the firm's payoff is $\underline{\pi}$ since $G(p^{FB}) = \underline{\pi} + \underline{v}$ and the path of play is the same as in the first best. Hence, all the incentive constraints for $\mathbb{H}$ are satisfied at p^{FB} and we are done.

We now show that whenever $p_{\mathbb{H}}(p^{FB}) > \underline{\pi} + \underline{v}$ the incentive constraint cannot be satisfied at p^{FB} and it follows similarly that it cannot be satisfied at beliefs slightly above p^{FB}, hence the equilibrium is inefficient. Suppose, for a contradiction that it can be satisfied. Without loss of generality $w_f = \underline{v}$, otherwise the difference between wages for success and failure can be increased without affecting the expected wage while still satisfying the incentive constraint for $\mathbb{H}$. Since $S(p^d_{\mathbb{H}}(p^{FB})) \leq G(p^d_{\mathbb{H}}(p^{FB})) = \underline{\pi} + \underline{v}$ the continuation payoffs following $\mathbb{H}$ and failure must be equal to the outside options. Let (π^u, v^u) denote the continuation payoffs following $\mathbb{H}$ and success. We must have $\pi^u + v^u \leq G(p^u_{\mathbb{H}}(p^{FB}))$.

After plugging in equation (6) the incentive constraint becomes

$$\begin{aligned}
&(1-\delta)[(p_{\mathbb{H}}(p^{FB}) - p_{\mathbb{L}}(p^{FB}))(w_s - \underline{v})] + \delta p_{\mathbb{H}}(p^{FB})(v^u - \underline{v}) \\
\geq &(1-\delta)(p_{\mathbb{H}}(p^{FB}) - (\underline{\pi} + \underline{v})) + \delta \mathbb{E}^{\mathbb{H},p^{FB}}[G(p') - (\underline{\pi} + \underline{v})]
\end{aligned}$$

and we obtain

$$(1-\delta)[(p_{\mathbb{H}}(p^{FB}) - p_{\mathbb{L}}(p^{FB}))(w_s - \underline{v})] \geq (1-\delta)(p_{\mathbb{H}}(p^{FB}) - (\underline{\pi} + \underline{v})) + \delta p_{\mathbb{H}}(p^{FB})(\pi^u - \underline{\pi}). \tag{8}$$

Notice that $p_{\mathbb{H}}(p^{FB}) > \underline{\pi} + \underline{v}$ implies $w_s - \underline{v} > 0$. Now consider the firm's payoff

$$\begin{aligned}\pi &= (1-\delta)[p_{\mathbb{H}}(p^{FB})(1-w_s) + (1-p_{\mathbb{H}}(p^{FB}))(-\underline{v})] + \delta[p_{\mathbb{H}}(p^{FB})\pi^u + (1-p_{\mathbb{H}}(p^{FB}))\underline{\pi}] \\ &= (1-\delta)[-\underline{v} + p_{\mathbb{H}}(p^{FB})(1-(w_s - \underline{v}))] + \delta[\underline{\pi} + p_{\mathbb{H}}(p^{FB})(\pi^u - \underline{\pi})]\end{aligned}$$

Using equation (8) we obtain

$$\begin{aligned}\pi &\leq (1-\delta)[-\underline{v} - p_{\mathbb{L}}(p^{FB})(w_s - \underline{v}) + (\underline{\pi} + \underline{v})] + \delta\underline{\pi} \\ &= \underline{\pi} - (1-\delta)p_{\mathbb{L}}(p^{FB})(w_s - \underline{v}) < \underline{\pi}\end{aligned}$$

which is a contradiction since the firm's payoff must exceed his outside option in any equilibrium.

□

7.2 Proofs from Section 5.3

We begin by developing the asymmetric information model formally. An asymmetric information environment is a set of parameters for the baseline model and a function $\varepsilon : [0,1] \to [0,1]$ specifying how far the beliefs of the optimistic and pessimistic worker are from the firm's belief when the firm considers both types equally likely. This belief spread function must be consistent with the updates in the following sense: for any p, p' for which there exists $n \in \mathbb{N}$, efforts $\{e_1, .., e_n\}$ and outcomes $\{i_1, .., i_n\}$ such that

$$\begin{aligned}p_+ &:= p_{e_n}^{i_n}(p_{e_{n-1}}^{i_{n-1}}..p_{e_1}^{i_1}(p + \varepsilon(p))) \\ p_- &:= p_{e_n}^{i_n}(p_{e_{n-1}}^{i_{n-1}}..p_{e_1}^{i_1}(p - \varepsilon(p)))\end{aligned}$$

satisfy $\frac{1}{2}p_+ + \frac{1}{2}p_- = p'$, we must have $p_+ - p' = p' - p_- = \varepsilon(p')$. This means that if the players start the games from a firm belief p and play a pooling strategy until reaching p', the posteriors of the types of the worker must coincide with their priors when the game is started at p'.

Now we are ready to state the equivalent of Proposition 4 for this model.

Proposition 10. *Fix any asymmetric information environment and suppose*

$$\frac{\beta_{\mathbb{L}}(1 + \delta\beta_{\mathbb{H}})}{1 + \delta\beta_{\mathbb{L}}} \geq \underline{\pi} + \underline{v}.$$

Then there exists a belief p such that any optimal relational contract in the game where the firm's starting belief is p exhibits $\mathbb{L}$ on the equilibrium path.

To prove this result, we develop a recursive representation of the model where the state variables are the firm's posterior belief about the match quality and his belief about the prior of the worker which is either a point mass at one of the possible priors or $(\frac{1}{2}, \frac{1}{2})$ because strategies are pure. The APS procedure operates on correspondences of the type $W : [0,1] \times \{0, \frac{1}{2}, 1\} \to \mathbb{R}^3$. They assign payoffs to the firm and the two types of the worker (according to the prior he holds on the match quality) for each posterior belief of the firm and likelihood of the worker holding the higher of the two possible priors. Notice that the worst equilibrium payoffs to both players are equal to the outside options by the same argument as in Lemma 1. So, as in the baseline model, we can operate on compact-valued correspondences W such that for all $p \in [0,1]$ and $\eta \in \{0, \frac{1}{2}, 1\}$ we have

- $(\underline{\pi}, \underline{v}) \in W(p, \eta)$.
- $(\pi, v) \in W(p, \eta)$ implies $(\pi, v) \geq (\underline{\pi}, \underline{v})$.

Payoffs (π, v_-, v_+) are enforceable with respect to correspondence W at (p, η) if they can be decomposed by actions for the firm and the two types of worker and continuation payoffs drawn from the appropriate posteriors (p', η'). If both types of the worker take the same actions, there are two posterior beliefs p^u and p^d. If the worker is exerting effort, they are determined by Bayes' rule. Continuation payoffs following success (π^u, v_+^u, v_-^u) belong to $W(p^u, \frac{1}{2})$ and continuation payoffs following failure (π^d, v_+^d, v_-^d) are in $W(p^d, \frac{1}{2})$. In case the two types of the worker separate there are updated beliefs (p_+^u, p_+^d) for the optimistic type following success and failure and updated beliefs (p_-^u, p_-^d) for the pessimistic type. If the action reveals that the worker is optimistic, the continuation payoffs satisfy $(\pi_+^i, v_+^i, v_-^i) \in W(p_+^i, 1)$ for $i \in \{u, d\}$. If the action reveals that the worker is the pessimistic type then continuation payoffs satisfy $(\pi_-^i, v_+^i, v_-^i) \in W(p_-^i, 0)$ for $i \in \{u, d\}$. Notice that due to the restriction on equilibrium payoffs once beliefs become degenerate we have $(\pi_+^i, v_+^i) \in \mathcal{E}(p_+^u)$ and $(\pi_-^i, v_-^i) \in \mathcal{E}(p_-^u)$ for all i.

We can review the firm's payoffs π_+, π_- from facing off against each type of worker noting that $\pi = \eta\pi_+ + (1-\eta)\pi_-$. For concreteness we will describe the decomposition of (π_+, v_+). There are three possibilities:[6]

- (High effort) The firm decides to interact and offers wage w. The worker with a high prior accepts and chooses $\mathbb{H}$.

$$\begin{aligned}
&\pi_+ = (1-\delta)(p_{\mathbb{H}}(p) - w) + \delta[p_{\mathbb{H}}(p)\pi_+^u + (1-p_{\mathbb{H}}(p))\pi_+^d] \\
&v_+ = (1-\delta)(w-c) + \delta[p_{\mathbb{H}}(p)v_+^u + (1-p_{\mathbb{H}}(p))v_+^d] \\
&\pi_+ \geq \underline{\pi}, \qquad v_+ \geq \underline{v} \\
&p_{\mathbb{H}}(p)v_+^u + (1-p_{\mathbb{H}}(p))v_+^d - \underline{v} \geq \frac{1-\delta}{\delta}c
\end{aligned}$$

[6]Continuation payoffs in the following decompositions follow the notation set out above. In case of pooling, we set $\pi_+^i = \pi_-^i$ for all $i \in \{u, d\}$.

- (Low effort) The firm offers some wage w, the worker accepts and chooses $\mathbb{L}$. Continuation payoffs are $(\pi^u, v^u) \in W(p^u_{\mathbb{L}}(p))$ and $(\pi^d, v^d) \in W(p^d_{\mathbb{L}}(p))$ such that

$$\pi_+ = (1-\delta)(p_{\mathbb{L}}(p) - w) + \delta[p_{\mathbb{L}}(p)\pi^u_+ + (1 - p_{\mathbb{L}}(p))\pi^d_+]$$
$$v_+ = (1-\delta)w + \delta[p_{\mathbb{L}}(p)v^u_+ + (1 - p_{\mathbb{L}}(p))v^d_+]$$
$$\pi_+ \geq \underline{\pi} \qquad v_+ \geq \underline{v}$$

- (Outside options) $\pi = \underline{\pi}$, $v = \underline{v}$.

Proof of Proposition 10. Towards a contradiction, assume that $\mathbb{L}$ is never played on the path of any equilibrium in state $(p, \frac{1}{2})$ for any p.

The only equilibrium payoffs at firm belief $p = 0$ and $\eta = \frac{1}{2}$ are $(\underline{\pi}, \underline{v}, \underline{v})$ regardless of η. It can be shown that the equilibrium correspondence is upper hemicontinuous in the state variables which implies that there is a lower bound on the firm beliefs at which there exists an equilibrium with at least one of the workers exerting $\mathbb{H}$ (with $\eta = \frac{1}{2}$). Now let $p_1 > 0$ denote the infimum of firm beliefs such that $\mathbb{H}$ is used in an equilibrium starting in state $(p_1, \frac{1}{2})$. By upper hemicontinuity of the equilibrium correspondence, this infimum is attained. Fix such an equilibrium at belief p_1 and use the notation for the continuation payoffs after the initial period established in the review of the APS algorithm. By definition there exists $\varepsilon > 0$ such that the optimistic and pessimistic types of the worker hold beliefs $p_1 + \varepsilon$ and $p_1 - \varepsilon$ about the match quality. Let p be the arithmetic average of $(p^u_{\mathbb{L}})^{-1}(p_1+\varepsilon)$ and $(p^u_{\mathbb{L}})^{-1}(p_1-\varepsilon)$. We will show a contradiction by constructing an equilibrium in state $(p, \frac{1}{2})$ which Pareto improves upon the payoffs $(\underline{\pi}, \underline{v}, \underline{v})$. In what follows, let $\bar{p}$ denote the cutoff for $\mathbb{H}$ in the baseline model.

Since by assumption only $\mathbb{H}$ and the outside option are used on-path, at least of the types of the worker must exert $\mathbb{H}$ in the initial period (otherwise outside options are taken and the state remains the same). This leaves two possibilities for equilibrium play in the initial period: both types exert $\mathbb{H}$ or one exerts $\mathbb{H}$ while the other refuses to interact with the firm.

We begin with the separating case. For concreteness suppose the optimistic type exerts $\mathbb{H}$; the argument for the other case is symmetric. First, we show that the pessimistic type must receive his outside option in *all* subsequent periods. If that were not the case, then $p_1 - \varepsilon \geq \bar{p}$; otherwise following separation no incentives could be provided to the pessimistic type to exert $\mathbb{H}$ since the outside option alone will not move the posterior belief further. Let p^* and ε^* satisfy $p^* + \varepsilon^* = \bar{p}$. Clearly, $p^* < p^1$ so for the contradiction it suffices to show that there is an equilibrium at $(p^*, \frac{1}{2})$ which involves $\mathbb{H}$. Suppose the firm and the optimistic type play the equilibrium strategies from the baseline model which deliver payoffs $(S(\bar{p}) - \underline{v}, \underline{v}) \in \mathcal{E}(\bar{p})$. They can be supported by the wage schedule from Proposition 5, in particular an initial wage of $\underline{v}$. It can be checked that the pessimistic type will not find it profitable to mimic the optimistic type since the wages just compensate the worker for the cost of effort under the optimistic assessment and a higher wage is given for success. Hence, there is a separating equilibrium where the pessimistic type reveals himself by refusing the firm's offer and receives $\underline{v}$ while the optimistic type plays according to an efficient equilibrium of the baseline model at the cutoff belief $\bar{p}$. Both types get $\underline{v}$ but the firm's payoff is strictly greater than his outside option since he gets $\underline{\pi}$ from interacting with the pessimistic type and strictly more from interacting with the optimistic type. This contradiction to the minimality of p^1 establishes that the original separating equilibrium must have $\pi_- = \underline{\pi}$ and $v_- = \underline{v}$.

We proceed to construct an equilibrium in state $(p, \frac{1}{2})$ where the pessimistic type refuses to interact while the optimistic type accepts a wage $w = \underline{v}$ and exerts $\mathbb{L}$. The continuation strategies following $\mathbb{L}$ and success of the firm and the optimistic worker are their strategies from the equilibrium at $(p_1, \frac{1}{2})$[7]. The incentives for the optimistic type are the same since at this point his belief is precisely $p_1 + \varepsilon$. The firm also has no incentive to deviate since in the original equilibrium his expected payoff is at least his outside option and he makes $\underline{\pi}$ from interacting with the pessimistic type, hence he must be making at least $\underline{\pi}$ from the optimistic type. The continuation equilibrium following any other actions and realisations in the initial period is the worst punishment equilibrium giving everyone their outside options. Now we turn to incentives in the initial period. The pessimistic type gets $\underline{v}$ from the prescribed strategies. If he deviates, he gets the contemporaneous wage of $\underline{v}$ and a continuation payoff no greater than $\underline{v}$ since it was not profitable for him to mimic the optimistic type in the original equilibrium. It remains to show the firm and the the optimistic worker have no incentive to deviate. The proof of this is essentially the same as in Proposition 4. Hence, we have constructed an equilibrium at $(p, \frac{1}{2})$ involving $\mathbb{H}$ - a contradiction to the minimality of p_1.

It remains to tackle the pooling case where both types exert $\mathbb{H}$ at $(p_1, \frac{1}{2})$. The incentive constraint for both types implies that their expected continuation payoffs less the contemporaneous cost of effort c exceed $\underline{v}$. Consider altering the strategies by having the firm offer a contemporanoues wage $\underline{v}$ and punishing any deviation with the worst equilibrium yielding $(\underline{\pi}, \underline{v}, \underline{v})$. Both types of the worker have no reason to deviate as it will give them at most $\underline{v}$. The firm receives a contemporaneous payoff of $p_{\mathbb{H}}(p_1) - \underline{v}$ and gets at least $\underline{\pi}$ in the continuation game. Thus, his payoff exceeds $\underline{\pi}$ under the sufficient condition since the latter can only hold when $\beta_{\mathbb{H}} > \underline{\pi} + \underline{v}$. Thus, the alternative strategies also form an equilibrium in state $(p_1, \frac{1}{2})$ which will be used in the construction which follows.

We will construct a pooling equilibrium at $(p, \frac{1}{2})$ where the firm offers $w = \underline{v}$ in the initial period, both types accept, exert $\mathbb{L}$ and upon success play continues as in the above equilibrium. Any other actions and realisations of output lead to the punishment equilibrium with payoffs $(\underline{\pi}, \underline{v}, \underline{v})$. Thus, all we need for incentive compatibility is to

[7] Without loss of generality we can fix some optimal strategy for the pessimistic type. It will not matter for the expected payoff calculations in the initial period since the pessimistic type does not exert $\mathbb{L}$ on path.

show that all types of all players get at least their outside options. This is obvious for the worker since $\mathbb{L}$ has no cost and he gets at least $\underline{v}$ in the continuation game. The firm's payoff exceeds

$$(1-\delta)(p_{\mathbb{L}}(p)-\underline{v})+p_{\mathbb{L}}(p)\delta\Big((1-\delta)(p_{\mathbb{H}}(p_1)-\underline{v})+\delta\underline{\pi}\Big)+(1-p_{\mathbb{L}}(p))\delta\underline{\pi}$$

which exceeds $\underline{\pi}$ if

$$(1-\delta)(p_{\mathbb{L}}(p)-(\underline{\pi}+\underline{v}))+p_{\mathbb{L}}(p)\delta(1-\delta)(p_{\mathbb{H}}(p_1)-(\underline{\pi}+\underline{v}))>0$$

It suffices to show that

$$\beta_{\mathbb{L}}+\beta_{\mathbb{L}}\delta\beta_{\mathbb{H}}\geq(1+\delta)(\underline{\pi}+\underline{v})$$

which is precisely the content of the sufficient condition. Thus we have constructed an equilibrium at $(p,\frac{1}{2})$ exhibiting $\mathbb{H}$ which contradicts the minimality of p_1.

$\square$