

Dynamic contracting with limited commitment and the ratchet effect*

Dino Gerardi[†] Lucas Maestri[‡]

September 2019

Abstract

We study dynamic contracting with adverse selection and limited commitment. A firm (the principal) and a worker (the agent) interact for potentially infinitely many periods. The worker is privately informed about his productivity and the firm can only commit to short-term contracts. The ratchet effect is in place since the firm has the incentive to change the terms of trade and offer more demanding contracts when it learns that the worker is highly productive.

As the parties become arbitrarily patient, the equilibrium outcome takes one of two forms. If the prior probability of the worker being productive is low, the firm offers a pooling contract and no information is ever revealed. In contrast, if this prior probability is high, the firm fires the unproductive worker at the beginning of the relationship.

Keywords: Dynamic contracting; limited commitment; ratchet effect.

JEL: D80; D82; D86.

*We are grateful to the editor, three anonymous referees, Luca Anderlini, Paolo Ghirardato, Daniel Gottlieb, Felipe Iachan, Ignacio Monzón, Juan Morales, César Santos, and various seminar audiences for useful comments. The financial support of the European Research Council (Consolidator Grant INFOASYM-MARKETS) is gratefully acknowledged. Lucas Maestri acknowledges that this study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001. Lucas Maestri is also grateful for the financial support of CNPq - National Council for Scientific and Technological Development.

[†]Collegio Carlo Alberto, Università di Torino; dino.gerardi@carloalberto.org.

[‡]FGV EPGE - Escola Brasileira de Economia e Finanças; lucas.maestri@fgv.br.

1 Introduction

This paper contributes to the literature on the ratchet effect by analyzing an infinite horizon contracting problem with short-term contracts. We frame the analysis in the context of a labor relationship between a worker and a firm. In each period, the worker can produce a good of quality $q \in [0, 1]$ at a cost that is linear in q . The worker is privately informed about his (persistent) marginal cost, which is low with prior probability p_0 and high with probability $1 - p_0$. The firm can only commit to short-term contracts, which indicate the payment that the worker receives, in the current period, if he produces a good of a specified quality. In each period in which the worker is employed, the firm offers a menu with finitely many contracts. The worker can either accept one contract in the menu or reject all contracts and end the relationship.

We show that when the discount factor is not too high, the firm is able to extract the worker's private information independently of the value of the prior. In particular, if the prior p_0 is high, the firm offers a *firing* menu in every period. A firing menu contains only one contract, which specifies the efficient quality when the cost is low and yields a payoff equal to zero to the worker. The contract is only accepted by the low-cost worker. Thus, the firm learns the worker's cost in the first period by firing the high-cost worker.

When the prior p_0 is not too high, the firm employs a *sequentially screening* procedure. The firm offers two contracts until it discovers the worker's cost, which occurs in finitely many periods. During the screening procedure, the high-cost worker accepts the first contract while the low-cost worker randomizes between the two contracts. Once the screening process is complete and the firm discovers the cost, the worker delivers the efficient quality and obtains a zero payoff.

When the parties are sufficiently patient, the sequentially screening procedure is infeasible. Consider the last period of the screening procedure in which the firm offers a menu that fully separates the two types of worker (in the sense that each type accepts a different contract). The low-cost worker can guarantee a large *future* payoff by mimicking the high-cost worker. Because of this, it is impossible to design two contracts that simultaneously satisfy the truth-telling constraints of the two types of worker. Only a sufficiently generous contract can prevent the low-cost worker from imitating the high-cost worker. But in this case the high-cost worker has an incentive to adopt the "take the money and run" strategy

(i.e., accept the contract designed for the low-cost worker and then quit the relationship).¹

The firm could, in principle, adopt more complex dynamic screening strategies. For instance, in the initial phase of the relationship both types of worker could accept different contracts in the menu with positive, but different, probabilities. Then the firm could use the available information to induce partial or even complete separation (through a firing menu). To investigate the feasibility and optimality of such strategies, we analyze the limiting outcome, as the parties become arbitrarily patient, of all perfect Bayesian equilibria.

We show that the limiting equilibrium outcome is unique and takes a very simple form. If the prior is below a certain threshold \hat{p} , then, in every period, the firm offers the most profitable contract that the high-cost worker is willing to accept. Both types of worker accept the contract (i.e., they pool) and there is no learning. In contrast, if the prior is above \hat{p} , the firm offers the firing menu and the high-cost worker quits the relationship without delay. In both cases, the limiting equilibrium allocation is inefficient.

Our results illustrate that when the parties are sufficiently patient, the firm can only screen the worker's type by firing the high-cost worker. The driving forces behind our findings are similar to those which prevent full separation. When the discount factor is large, it is very costly for the firm to separate the two types of worker and continue the relationship with both of them. A lasting relationship with the high-cost worker provides strong incentives to the low-cost worker to misrepresent his information. Using this fact, we show that the firm would not benefit from engaging in partial screening strategies, even if such policies were feasible.

Our benchmark model assumes that the relationship ends when the worker rejects all the contracts in the firm's menu. This modeling assumption captures situations in which the parties are committed to terminate the relationship upon disagreement. Of course, one can also imagine situations in which the relationship continues even when the parties do not reach an agreement. Therefore, we analyze an extension of the model that allows for rehiring. We study the infinitely repeated game in which, in each period, the firm proposes a menu of contracts from which the worker has to select at most one. We first show that the complete-information version of this game admits a folk theorem. Although the firm has the bargaining power to make offers, the worker can obtain large payoffs by rejecting unfavorable contracts. This is possible because the acceptance of unfavorable contracts by

¹A similar result appears in Laffont and Tirole (1990).

the worker triggers a continuation equilibrium in which the firm implements an efficient allocation that yields a zero payoff to the worker. We use these findings from the complete-information game to show that a version of the folk theorem holds for our model with rehiring.² In particular, when the parties are sufficiently patient, the firm can obtain a payoff arbitrarily close to the payoff of the optimal mechanism with commitment.

This paper contributes to the literature on repeated adverse-selection with limited commitment pioneered by Freixas, Guesnerie, and Tirole (1985), Gibbons (1987), and Laffont and Tirole (1987, 1988). In these seminal papers, the parties interact for two periods. One of the main findings is that there is partial separation of the agent's types in the first period (i.e., the equilibrium is semi-pooling) and full separation in the second and final period. Therefore the outcome of two-period environments presents gradual information revelation. In contrast, our paper shows that when the relationship is infinitely repeated and the prior is low, the equilibrium allocation is close to a pooling allocation when the parties are patient.

Hart and Tirole (1988) analyze a dynamic model in which the seller makes a rental offer to the buyer in every period. The buyer's valuation for the good is private information and can take on two values, both of which are larger than the seller's cost of producing the good. As the parties become sufficiently patient, the equilibrium allocation converges to the efficient allocation, in which both types of buyers consume the good in every period. Note that for high values of the prior, this pooling allocation coincides with the seller's optimal mechanism under full commitment (i.e., lack of commitment is not detrimental to the seller's payoff). In a recent paper, Beccuti and Möller (2018) extend Hart and Tirole's analysis to the case in which the seller is more patient than the buyer. Halac (2012) studies a relational contract model in which the principal is privately informed about his outside option. When the uninformed party has the bargaining power, Coasian forces lead to a pooling outcome when the parties are sufficiently patient. Our work differs from these papers in two respects. First, in our model, the agent's private information is necessary to determine the best course of action and, therefore, pooling allocations are never optimal for the firm under full commitment. Second, we analyze environments in which the ratchet effect leads to inefficiencies.³

²This finding is reminiscent of earlier contributions to repeated games with incomplete information and simultaneous moves (see Pęski (2008) and the references therein).

³Our work analyzes the relationship between two infinitely lived players. In the context of political economy, several papers study the effects of limited commitment in repeated interactions between one principal

Our work is also related to the literature on renegotiation. The seminal paper by Laffont and Tirole (1990) analyzes a two-period model. Recently, Strulovici (2017) and Maestri (2017) study renegotiation in infinite horizon models. These studies find that equilibrium allocations become efficient as the parties become arbitrarily patient. In contrast, in our model the limit allocation is inefficient whenever the firing allocation is not a commitment solution.

Bhaskar (2014) studies learning in a dynamic model in which the principal and the agent are ex-ante symmetrically informed about the job’s difficulty. When the agent’s effort is unobservable, it is impossible for the principal to design a contract that induces an interior effort level in the first period. Bhaskar and Mailath (2019) consider a related dynamic model and show that inducing high effort becomes prohibitively costly for the principal as the parties become arbitrarily patient. Therefore, the ratchet effect imposes stringent constraints on the learning process of the relationship. In contrast, our paper assumes adverse-selection and no exogenous learning and concludes that the ratchet effect imposes constraints on the amount of private information that is revealed in a dynamic relationship.

There is also a connection between our paper and the literature on durable goods monopoly under limited commitment. Ausubel and Deneckere (1989) study a model in which the seller posts prices and obtain a folk theorem for the “no gap” case. In our context, a folk theorem holds when rehiring is possible. Skreta (2006, 2015) analyzes more general selling mechanisms and shows that posting a price is the seller’s optimal strategy. In these studies, the relationship between the buyer and the seller ends as soon as the durable good is traded, while in our model the parties can make a new transaction every period.

Finally, a number of authors have identified situations in which the ratchet effect is mitigated. Kanemoto and MacLeod (1992) argue that competition for secondhand workers guarantees the existence of efficient piece-rate contracts in long-term relationships. Carmichael and MacLeod (2000) show that if entry in a market is difficult, then it is possible to sustain cooperation between an infinitely lived firm and a stream of short lived workers. Our findings suggest that rehiring is another possible remedy to the ratchet effect.

The rest of the paper is organized as follows. We present the model in Section 2. In Section 3, we briefly discuss the mechanism design problem with commitment. In Section 4,

and a continuum of privately informed agents (see, among others, Acemoglu, Golosov, and Tsyvinski (2010), Farhi, Sleet, Werning, and Yeltekin (2012), and Scheuer and Wolitzky (2016)).

we show existence of equilibria and provide conditions under which all private information is revealed. Section 5 contains the main result which completely characterizes the unique equilibrium outcome when the parties are arbitrarily patient. In Section 6, we analyze the extension of the model in which rehiring is possible. Section 7 concludes. Most proofs are relegated to a number of appendices.

2 The Model

We study a dynamic principal-agent model with adverse selection and short-term contracts, framed in the context of a labor relationship between a firm and a worker.

The worker is privately informed about his (persistent) type, which is equal to L with prior probability $p_0 \in (0, 1)$, and equal to H with probability $1 - p_0$. The firm and the worker interact for potentially infinitely many periods. In each period, the worker of type $i \in \{H, L\}$ can produce a good of quality $q \in [0, 1]$ at a cost of $\theta_i q$, where $0 < \theta_L < \theta_H$. We refer to the low type L (high type H) as the low (high) cost worker. We write $\Delta\theta := \theta_H - \theta_L$. The worker bears an additional cost $\alpha > 0$ in every period in which he interacts with the firm. The cost α can be interpreted as the per-period payoff of an outside option available to the worker if he ends the relationship (in Section 7, we discuss the case $\alpha = 0$).

The firm's valuation of a good of quality q is $v(q)$. The function $v : [0, 1] \rightarrow \mathbb{R}_+$ is twice continuously differentiable, increasing, strictly concave, and satisfies $v(0) = 0$.⁴

Both parties' preferences are linear in money. When the worker produces a good of quality q and the firm makes a transfer equal to x , the payoff of type $i \in \{H, L\}$ is $x - \theta_i q - \alpha$, while the firm's payoff is $v(q) - x$.

We let q_i^* , $i \in \{H, L\}$, denote the efficient quality produced by type i :

$$q_i^* = \arg \max_{q \in [0, 1]} (v(q) - \theta_i q)$$

To make the problem interesting, we assume

$$v(q_H^*) - \theta_H q_H^* - \alpha > 0$$

⁴The concavity of $v(\cdot)$ guarantees that the firm's screening problem in the proof of Proposition 1 is well behaved. The concavity also allows us to derive a number of useful bounds in the proof of Proposition 2. Finally, the assumption $v'(0) < \infty$ implies that for large values of the prior, the solution to the mechanism design problem with commitment is to fire the high type (see Section 3). This is used in the proof of the main result.

This assumption guarantees that the firm prefers hiring the high-cost worker over collecting its outside option, which yields a payoff equal to zero. Moreover, we assume that $q_H^* \in (0, 1)$ and, therefore $q_L^* > q_H^*$.⁵ In this case, the efficient allocation varies with worker's type.

The firm and the worker play the following game. At the beginning of period $t = 0, 1, \dots$, the firm offers a menu m_t of contracts to the worker. Each contract is of the form (x_t, q_t) and specifies the transfer x_t paid by the firm and the quality $q_t \in [0, 1]$ that the worker must produce. We assume that the quality is verifiable and, thus, each contract is enforceable. After receiving the menu m_t , the worker has two options: (i) selecting a contract from the menu; or (ii) rejecting all the contracts and quitting the relationship. In the first case, the game moves to the next period $t + 1$. In the second case, the game ends and both parties obtain a continuation payoff equal to zero. The parties discount future payoffs at the common discount factor $\delta \in (0, 1)$.

We let $\mathcal{M} = \bigcup_{j=1}^M (\mathbb{R} \times [0, 1])^j$ denote the set of available menus, where $M \in \{2, 3, \dots\}$ is an exogenous upper bound to the number of contracts that a menu can contain. The restriction $M \geq 2$ guarantees that the menus can contain two contracts (so that it is possible for the firm to separate the two worker types). When the firm offers the menu m_t , the set of actions available to the worker is $m_t \cup \{\emptyset\}$, where \emptyset denotes the choice of rejecting all the contracts in m_t and quitting. We let a_t denote the agent's decision in period t .

For every $t = 1, 2, \dots$, a period- t (non-final) public history $h^t = (m_0, a_0, \dots, m_{t-1}, a_{t-1})$ consists of all the menus offered by the firm in the previous periods $\tau = 0, \dots, t - 1$, as well as all the worker's decisions, provided that he never chose to quit (i.e., $a_\tau \neq \emptyset$ for every $\tau = 0, \dots, t - 1$). We let $H^0 = \{h^0\}$ denote the set containing the empty history h^0 . We write H^t for the set of all period- t public histories. Finally, $\mathcal{H} = \bigcup_{t=0,1,\dots} H^t$ is the set of all (non-final) public histories.

A behavior strategy σ^F for the firm is a sequence $\{\sigma_t^F\}$, where σ_t^F is a function from H^t into $\Delta(\mathcal{M})$, mapping the history h^t into a (possibly random) menu. A behavior strategy (σ^H, σ^L) for the worker is a sequence $\{(\sigma_t^H, \sigma_t^L)\}$, where σ_t^i , $i \in \{H, L\}$, associates to every pair $(h^t, m_t) \in H^t \times \mathcal{M}$ a probability distribution over the set $m_t \cup \{\emptyset\}$. We write $\sigma = (\sigma^F, \sigma^H, \sigma^L)$ for a strategy profile. Finally, we let $\mu = \{\mu(h^t), \mu(h^t, m_t)\}_{h^t \in \mathcal{H}, m_t \in \mathcal{M}}$ denote the firm's system of beliefs, where $\mu(h^t)$ and $\mu(h^t, m_t)$ represent the probability

⁵In particular, we use this assumption in the proof of Proposition 1 to construct a sequence of separating contracts.

that the firm assigns, at the history h^t and (h^t, m_t) , respectively, to the event that the worker's type is equal to L .

Our solution concept is perfect Bayesian equilibrium (PBE, or equilibrium henceforth), formally defined below.

Definition 1 *A PBE of our game is a strategy profile σ and a system of beliefs μ such that:*

- i) σ is sequentially rational given μ ;*
- ii) for every history $(h^t, m_t) \in H^t \times \mathcal{M}$, $\mu(h^t, m_t) = \mu(h^t)$;*
- iii) for every history $(h^t, m_t) \in H^t \times \mathcal{M}$ and for every action $a_t \in m_t \cup \{\emptyset\}$, if*

$$(1 - \mu(h^t)) \sigma_t^H(a_t | h^t, m_t) + \mu(h^t) \sigma_t^L(a_t | h^t, m_t) > 0$$

then the belief $\mu(h^t, m_t, a_t)$ is derived from $\mu(h^t)$ according to Bayes' rule:

$$\mu(h^t, m_t, a_t) = \frac{\mu(h^t) \sigma_t^L(a_t | h^t, m_t)}{(1 - \mu(h^t)) \sigma_t^H(a_t | h^t, m_t) + \mu(h^t) \sigma_t^L(a_t | h^t, m_t)}$$

In addition to sequential rationality and Bayesian updating whenever possible (i.e., including off-path histories (h^t, m_t, a_t) that are reached with positive probability given (h^t, m_t)), the concept of PBE imposes the “no signaling what you don't know” condition (Fudenberg and Tirole (1991)) in the sense that the firm does not revise its belief after proposing a menu.

Given a strategy profile σ and a system of beliefs μ , for each history h^t we let $V_F(h^t; (\sigma, \mu))$ denote the firm's continuation payoff at h^t . We also let $\mathbb{T} \in \mathbf{N} \cup \{\infty\}$ denote the random period in which the relationship terminates (we set $\mathbb{T} = \infty$ if the worker remains employed forever).⁶ Then we have:

$$V_F(h^t; (\sigma, \mu)) := \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t}^{\mathbb{T}-1} \delta^{\tau-t} (v(q_\tau) - x_\tau) \mid h^t \right]$$

where $\mathbb{E}_{(\sigma, \mu)}[Y | h^t]$ represents the conditional expected value (given h^t) of the random variable Y given the strategy profile σ and the system of beliefs μ . Analogously, for every

⁶Here, and in what follows, we use $\mathbf{N} = \{0, 1, \dots\}$ to denote the set of nonnegative integers and adopt the convention that $\sum_{\tau=t}^{t-1} \delta^{\tau-t} = 0$.

history h^t we let $W_i(h^t; (\sigma, \mu))$ denote the expected continuation payoff at h^t of the worker of type $i \in \{H, L\}$. We have:

$$W_i(h^t; (\sigma, \mu)) := \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t}^{\mathbb{T}-1} \delta^{\tau-t} (x_\tau - \theta_i q_\tau - \alpha) \mid i, h^t \right]$$

To simplify the notation, we omit the argument (σ, μ) and write $V_F(h^t)$ and $W_i(h^t)$ when there is no ambiguity. We also use $V_F(h^t, m_t)$ and $W_i(h^t, m_t)$, $i \in \{H, L\}$, to denote the firm and worker's payoff at the history (h_t, m_t) .

For $i \in \{H, L\}$, and $q \in [0, 1]$, we let

$$\pi_i(q) := v(q) - \theta_i q - \alpha$$

denote the firm's profits when the quality is q , the worker is of type i and the firm pays the reservation wage $\theta_i q + \alpha$. Therefore, $\pi_i(q_i^*)$ represents the highest level of profits that the firm can achieve from the interaction with type i . As $\pi_L(q_L^*) > \pi_H(q_H^*)$, let $\hat{p} \in (0, 1)$ be defined by $\pi_H(q_H^*) = \hat{p} \pi_L(q_L^*)$.

We conclude this section with a simple result that provides a lower bound to the firm's payoff under any PBE.

Lemma 1 *Fix a PBE (σ, μ) . For every history $h^t \in \mathcal{H}$, we have:*

$$V_F(h^t; (\sigma, \mu)) \geq \max \{ \pi_H(q_H^*), \mu(h^t) \pi_L(q_L^*) \}$$

Proof of Lemma 1.

By contradiction, suppose that there exist a PBE (σ, μ) , a history h^t , and $\varepsilon > 0$ such that

$$V_F(h^t; (\sigma, \mu)) < \max \{ \pi_H(q_H^*), \mu(h^t) \pi_L(q_L^*) \} - \varepsilon$$

Suppose that $\pi_H(q_H^*) > \mu(h^t) \pi_L(q_L^*)$. If the firm offers the menu $\{(\theta_H q_H^* + \alpha + \frac{\varepsilon}{2}, q_H^*)\}$ in every period $t, t+1, \dots$ (notice that both types strictly prefer to accept the contract in the menu rather than quit the relationship), then its continuation payoff will be equal to

$$\pi_H(q_H^*) - \frac{\varepsilon}{2} > V_F(h^t; (\sigma, \mu))$$

which is a contradiction.

Next, consider the case $\pi_H(q_H^*) \leq \mu(h^t) \pi_L(q_L^*)$. The firm can guarantee a continuation payoff at least equal to

$$\mu(h^t) \left(\pi_L(q_L^*) - \frac{\varepsilon}{2} \right) > V_F(h^t; (\sigma, \mu))$$

by offering the menu $\{(\theta_L q_L^* + \alpha + \frac{\varepsilon}{2}, q_L^*)\}$ in every period $t, t+1, \dots$ (in equilibrium, the low type must accept the contract in the menu). ■

Intuitively, the following two options are always available to the firm. The first option is to stop learning and offer $(\theta_H q_H^* + \alpha, q_H^*)$, the most profitable contract in the class of contracts that are accepted by both types of worker. The second option is to fire the high-cost worker and interact only with the low-cost worker. In this case, the most profitable contract is $(\theta_L q_L^* + \alpha, q_L^*)$.

3 The Commitment Allocation

It is useful to review the benchmark model in which the firm can fully commit to a sequence of menus (m_0, m_1, \dots) . This benchmark provides an upper bound to the firm's profits in the game with limited commitment. It is well known that the solution to the firm's commitment problem is to replicate the optimal static mechanism (see, for example, Chapter 1 in Laffont and Tirole, 1993).

The optimal static mechanism is as follows. There exists a critical value $p^C \in (\hat{p}, 1)$ such that if $p_0 > p^C$, the optimal menu (with commitment) is unique and equal to $\{(\theta_L q_L^* + \alpha, q_L^*)\}$.⁷ The low-cost worker accepts the contract in the menu while the high-cost worker rejects it. Thus, the firm's profits are equal to $p_0 \pi_L(q_L^*)$.

On the other hand, if $p_0 < p^C$, then the unique optimal menu is

$$\{(x_H^C, q_H^C), (x_L^C, q_L^C)\} = \{(\theta_H q_H^* + \alpha, q_H^*), (\theta_L q_L^* + \Delta\theta q_H^C + \alpha, q_L^*)\} \quad (1)$$

for some $q_H^C \in (0, q_H^*)$. The high-cost worker accepts the first contract and obtains a payoff equal to zero. The low-cost worker is indifferent between the two contracts (therefore, he obtains a payoff equal to $\Delta\theta q_H^C$) and accepts the second contract. In this case, the firm's commitment profits are equal to:

$$p_0 [v(q_L^*) - \theta_L q_L^* - \Delta\theta q_H^C - \alpha] + (1 - p_0) [v(q_H^C) - \theta_H q_H^C - \alpha]$$

⁷To see why $p^C > \hat{p}$, let $V_F^C(p)$ be the commitment payoff of the firm when the prior is p . The function $V_F^C(\cdot)$ is strictly increasing. If $p^C \leq \hat{p}$, then we obtain the following contradiction: $V_F^C(p^C) = p^C \pi_L(q_L^*) \leq \hat{p} \pi_L(q_L^*) = \pi_H(q_H^*) = V_F^C(0)$.

Finally, if $p_0 = p^C$, then there are two optimal deterministic mechanisms: $\{(\theta_L q_L^* + \alpha, q_L^*)\}$ and $\{(x_H^C, q_H^C), (x_L^C, q_L^C)\}$ as in equation (1). In addition, there is a continuum of optimal random mechanisms, since any randomization between the two optimal deterministic mechanisms is also an optimal mechanism.

Suppose that $p_0 \leq p^C$. In the dynamic game with limited commitment it is impossible to implement the optimal mechanism of the form $\{(x_H^C, q_H^C), (x_L^C, q_L^C)\}$ in every period. This is because, as shown in Lemma 1, the firm's continuation payoff must be equal to $\pi_i(q_i^*)$ as soon as the firm discovers that the worker is of type i .⁸

Finally, the firm's payoff in any PBE must be equal to $p_0 \pi_L(q_L^*)$ whenever $p_0 \geq p^C$. It cannot be smaller, because of Lemma 1, and it cannot be larger, because $\{(\theta_L q_L^* + \alpha, q_L^*)\}$ is an optimal mechanism with commitment. Therefore, if $p_0 \geq p^C$, all PBE share the following feature: the high type quits the relationship in the first period, while the low type accepts the contract $(\theta_L q_L^* + \alpha, q_L^*)$ in every period.

4 Existence and Learning

In this section, we show the existence of PBE for generic values of the parameters. We also identify the conditions under which the firm is able to learn the worker's type (with and without firing). In particular, if the parties are impatient, then learning is possible for any prior. In contrast, if the parties are sufficiently patient, then learning only takes place when the firm is willing to fire the high-cost worker.

We start with a general result that holds in every PBE: the low-cost worker's relationship with the firm lasts forever.

Lemma 2 *Fix a PBE (σ, μ) and an arbitrary history h^t . For every menu m_t offered by the firm at h^t , we have*

$$\sum_{(x_t, q_t) \in m_t} \sigma_t^L((x_t, q_t) | h^t, m_t) = 1$$

We omit the proof of this simple result and provide only an intuitive discussion. Suppose, by contradiction, that there are a PBE (σ, μ) and a history h^t at which the low type rejects all the contracts in the firm's menu with a positive probability. This implies that the interaction with the high type must yield a strictly positive continuation payoff

⁸Notice that $\pi_H(q_H^C) < \pi_H(q_H^*)$ since $q_H^C < q_H^*$. Also $v(q_L^*) - \theta_L q_L^* - \alpha - \Delta \theta q_H^C < \pi_L(q_L^*)$ since $q_H^C > 0$.

to the firm, otherwise its continuation payoff would be strictly smaller than $\mu(h^t) \pi_L(q_L^*)$, contradicting Lemma 1. Clearly, a strictly positive continuation payoff is possible only if the high type is expected to deliver a strictly positive (discounted) quality in the future. This and the individual rationality of the high type's behavior imply that the low type's decision to quit is not optimal, as he can guarantee a strictly positive payoff by imitating the high-cost worker at h^t and in every future period.

Suppose that the firm is interested in separating the two types and learning the worker's cost. This requires the existence of two decisions, one of which is taken only by the high type, while the other is taken only by the low type. After observing the first (second) decision, the firm becomes convinced that the worker's cost is high (low).

In light of Lemma 2, there are two ways in which separation can take place in equilibrium. One possibility is *separation with firing*: the high type quits the relationship and one of the contracts in the firm's menu is only accepted by the low type. The other possibility is *separation with employment*: one contract in the firm's menu is only accepted by the high type, while another contract is only accepted by the low type.

Our next result shows that separation with employment cannot occur for values of the discount factor larger than the critical threshold $\hat{\delta} := \frac{1}{1+q_H^*}$.

Lemma 3 *Suppose that $\delta > \hat{\delta}$ and let (σ, μ) be an arbitrary PBE of the game. It is impossible to find a history (h^t, m_t) (on or off-path) satisfying the following properties:*

- i) $\mu(h^t) \in (0, 1)$;*
- ii) there exists a contract (x_H, q_H) in m_t for which $\sigma_t^H((x_H, q_H) | h^t, m_t) > 0$ and $\sigma_t^L((x_H, q_H) | h^t, m_t) = 0$;*
- iii) there exists a contract (x_L, q_L) in m_t for which $\sigma_t^L((x_L, q_L) | h^t, m_t) > 0$ and $\sigma_t^H((x_L, q_L) | h^t, m_t) = 0$.*

We provide the intuition for this result (see the Supplemental Material for the proof). By contradiction, suppose that at h^t , the belief is nondegenerate and the firm's menu contains a contract (x_i, q_i) that is accepted with positive probability only by the type $i \in \{H, L\}$. Following the acceptance of this contract, the firm's belief will assign probability one to type i .⁹ Furthermore, in equilibrium, type i will select the efficient contract $(\theta_i q_i^* + \alpha, q_i^*)$ in every period after t . If the discount factor is above $\hat{\delta}$, it is impossible to find two contracts,

⁹Notice that in a PBE, the beliefs must satisfy this condition at all histories, including those that are off-path.

(x_H, q_H) and (x_L, q_L) , to satisfy the two incentive compatibility constraints. In fact, either the low type prefers to imitate the high type (at h^t and in every future period), or the high type has an incentive to adopt the “take the money and run” strategy (i.e., the strategy of accepting the generous contract (x_L, q_L) and then quitting).

We are now ready to state the main result of this section, which establishes the (generic) existence of PBE.

Proposition 1 *For generic values of the parameters, there exists a PBE.*

The proof of Proposition 1 (in Appendix A) shows how to construct a PBE for all values of δ outside a set of discount factors which can contain at most two elements (the values of these two elements depend on the primitives $\theta_H, \theta_L, \alpha, v(\cdot)$).¹⁰ For the remainder of the paper, we assume that the discount factor δ does not belong to this (possibly empty) set.

The equilibrium that we construct satisfies a number of properties. First, the equilibrium is “almost Markovian,” in the sense that the parties’ behavior in period t depends on the firm’s belief and their actions in period $t - 1$ (the history up to period $t - 2$ affects the behavior in period t only through the belief). Second, the high type plays a pure strategy and his equilibrium payoff is equal to zero. Third, the menu proposed by the firm (at any history) contains at most two contracts. Finally, the firm adopts a deterministic behavior at on-path histories.

We now introduce some definitions to illustrate our equilibrium. First, we say that there is a *pooling* allocation if the firm offers the menu $\{(\theta_H q_H^* + \alpha, q_H^*)\}$ in every period and both types accept the contract $(\theta_H q_H^* + \alpha, q_H^*)$ (with probability one). We also say that there is a *firing* allocation if the firm offers the menu $\{(\theta_L q_L^* + \alpha, q_L^*)\}$ in every period, the high type quits in the first period, and the low type accepts the contract $(\theta_L q_L^* + \alpha, q_L^*)$ in every period. Finally, we say that there is a *sequentially screening* allocation if the firm offers a menu with two contracts in every period in which its belief is nondegenerate. Furthermore, the high type accepts the first contract with probability one, while the low type accepts the second contract with a strictly positive probability (if this probability is less than one, the low type randomizes between the two contracts). Therefore, in a sequentially screening allocation either the firm learns that the worker’s type is low or it becomes more confident that the worker’s type is high.

¹⁰Our formal argument does not cover the values of δ at which the mapping $V^1 : [0, 1] \rightarrow \mathbb{R}$ defined in equation (33) (see the Supplemental Material) satisfies simultaneously $V^1(0) = \pi_H(q_H^*)$ and $\partial_+ V^1(0) = 0$. We show that there can be at most two such values of δ .

To illustrate our construction, it is convenient to distinguish between the case $\delta \leq \hat{\delta}$ and the case $\delta > \hat{\delta}$. We start with the first case. We assume (without loss) that the firm offers the menu $\{(\theta_H q_H^* + \alpha, q_H^*)\}$ when the belief p is equal to zero. Also, the firm offers the menu $\{(\theta_L q_L^* + \alpha, q_L^*)\}$ when $p \geq p^C$. For any belief $p \in (0, p^C)$, we first provide the following three options to the firm: i) offering a pooling menu, i.e., a menu with one contract that is accepted by both types; ii) offering a firing menu, i.e., a menu with one contract that induces separation with firing (where the low type accepts the contract while the high type quits); iii) offering a menu with two contracts to induce separation with employment (i.e., with probability one, the two types choose different contracts).

Clearly, the optimal pooling menu is $\{(\theta_H q_H^* + \alpha, q_H^*)\}$, while the optimal firing menu is $\{(\theta_L q_L^* + \alpha, q_L^*)\}$. In case iii), we choose the two contracts to maximize the firm's payoff subject to the incentive compatibility (IC) and the individual rationality (IR) constraints. Notice that after separation with employment, the firm's belief is either zero or one. In both cases, the firm's behavior is known and we can compute the two types' continuation payoffs. As in the standard mechanism design problem with commitment, the optimal menu in case iii) is such that both the low type's IC constraint and the high type's IR constraint are binding.

We construct the firm's value function $V(\cdot; 1)$ and the low type's payoff correspondence $\Phi(\cdot; 1)$ when the firm is forced to choose one of the three options above.¹¹ We take $V(\cdot; 1)$ and $\Phi(\cdot; 1)$ as given and offer the firm the possibility of *probabilistic* separation with employment. This means that the firm offers two contracts. The high type accepts the first contract with probability one, while the low type randomizes between the contracts. After this round of probabilistic separation, the firm is again forced to use the three options above and, therefore, the parties continuation payoffs are given by $V(\cdot; 1)$ and $\Phi(\cdot; 1)$. In the probabilistic separation phase, we select the two contracts and the low type's behavior (i.e., the probability of accepting each contract) to maximize the firm's payoff subject, of course, to the IC and IR constraints. As usual, the solution to the optimization problem satisfies the low type's IC constraint with equality and, therefore, randomizing between the two contracts is indeed optimal for the low-cost worker.

The possibility of probabilistic separation defines a new value function $V(\cdot; 2)$ and a new payoff correspondence $\Phi(\cdot; 2)$. If $V(p; 2) = V(p; 1)$ for every $p \in [0, 1]$, then we stop

¹¹For some values of the beliefs p , the solution to the firm's problem is not unique and different solutions generally yield different payoffs to the low type (hence, we use the correspondence $\Phi(\cdot; 1)$).

the process, as the firm does not benefit from probabilistic separation. On the other hand, if $V(p; 2) > V(p; 1)$ (it is possible to construct examples for which this is the case), then we allow for an additional round of probabilistic separation with employment.

We continue the process (allowing, at each iteration, for a new round of probabilistic separation) until we find a fixed point $(V(\cdot), \Phi(\cdot))$. We show that for generic values of the parameters a fixed point exists and is achieved after finitely many iterations. Moreover, our proof shows that if we fix the parameters $(\theta_H, \theta_L, \alpha, v(\cdot))$, then there is T such that for generic discount factors smaller than $1 - q_H^*/q_L^*$, the number of iterations is smaller than T (see Corollary 1 below for the implications of this result).

The pair $(V(\cdot), \Phi(\cdot))$ allows us to construct a simple equilibrium. For each belief p , the parties behave according to the solution of the firm's optimization problem (which yields the payoff $V(p)$ to the firm). The solution consists of the optimal menu and the decision of the worker. In particular, if the optimal menu is $\{(\theta_H q_H^* + \alpha, q_H^*)\}$, then both types accept the contract. If the optimal menu is $\{(\theta_L q_L^* + \alpha, q_L^*)\}$, then only the low type accepts the contract. Finally, if the optimal menu contains two contracts, then the high type accepts the first contract (with probability one) while the low type accepts the second contract with probability in $(0, 1]$ (this probability is part of the solution to the optimization problem).

The proof of Proposition 1 also specifies the parties' off-path behavior and shows that unilateral deviations are not profitable.

We conclude the discussion of the case of a low discount factor pointing out a property of our equilibria. Fix a PBE (σ, μ) . We say that there is *full learning* by period t if for any $t' \geq t$ and for any *on-path* public history $h^{t'}$, the belief $\mu(h^{t'})$ is either zero or one. This means that all the uncertainty about the worker's ability is resolved by period t .

Our construction shows that when the parties are not too patient, the firm never chooses the pooling allocation. Depending on the value of the prior, the firm prefers either the sequentially screening allocation (if the prior is low) or the firing allocation (if the prior is high). In both cases, there is full learning. Moreover, our proof shows that the number of periods until the worker completely reveals his private information is uniformly bounded.¹² Formally, we obtain the following result.

Corollary 1 *Fix the parameters $(\theta_H, \theta_L, \alpha, v(\cdot))$. There exists $T \in \{1, 2, \dots\}$ such that for any prior p_0 and for generic values of the discount factor smaller than $1 - q_H^*/q_L^*$, there*

¹²This follows from the argument provided at the end of the proof of Lemma 8.

exists a PBE with full learning by period T .

Finally, as δ shrinks to zero, the firm's equilibrium payoff converges to the payoff of the optimal mechanism with commitment. This is a general property that holds for all PBE.¹³

We now turn to the case the case $\delta > \hat{\delta}$. Recall that in this case separation with employment is not feasible. Therefore, the firm is unable to implement a sequentially screening allocation. As a result, the firm chooses between the pooling and the firing allocation. The equilibrium that we construct takes a very simple form. If the prior is weakly higher than \hat{p} , the firm offers the optimal firing menu $\{(\theta_L q_L^* + \alpha, q_L^*)\}$ in every period and the high-cost worker quits in the first period. Thus, the equilibrium is with full learning by period one. In contrast, if the prior is lower than \hat{p} , the firm offers the optimal pooling menu $\{(\theta_H q_H^* + \alpha, q_H^*)\}$ in every period and never updates (along the equilibrium path) its belief.

The analysis in this section shows that when the parties are sufficiently patient, the ratchet effect has a strong impact on equilibrium behavior. In particular, it suggests that, for sufficiently large values of δ , the firm can learn the worker's cost only by firing the high type. As we will see in the next section, this is not a special feature of our equilibrium, but a much more general result.

5 Limit Uniqueness

In the last section, we constructed a simple equilibrium for the case with sufficiently patient parties ($\delta > \hat{\delta}$). This equilibrium implements the pooling allocation when the prior is lower than \hat{p} and the firing allocation when the prior is higher than \hat{p} . Our construction relies on the fact that a sequentially screening allocation is not feasible when the discount factor is above $\hat{\delta}$. However, the firm could, in principle, employ more complex dynamic screening strategies. For instance, one could imagine an equilibrium in which two or more contracts in the firm's menu are accepted with positive, but different, probabilities by the two types. In this case, the firm's belief could increase without jumping to one, as occurs in a sequentially screening allocation. This raises the question of whether there are other equilibrium outcomes, in addition to the one identified in Section 4. Moreover, can the firm do better than just offering the optimal pooling menu or the optimal firing menu?

¹³For brevity, we omit the proof of this simple finding.

We show that in the limit, as the parties become arbitrarily patient, there exists a unique equilibrium outcome. This outcome coincides with the equilibrium outcome described in Section 4 (for the case $\delta > \hat{\delta}$). First, if the prior is above \hat{p} , the limiting equilibrium allocation is firing and the high type quits the relationship without delay.¹⁴ In contrast, if the prior is below \hat{p} , the limiting equilibrium allocation is pooling and there is no learning.¹⁵

5.1 High Belief Case: $p > \hat{p}$

We first study the limiting equilibrium outcome when the prior is above \hat{p} . To do so, we define the notion of *firing region*. Recall that \mathbb{T} denotes the random time at which the worker quits the relationship.

Definition 2 *The interval $[p, 1]$ is a firing region if there exist $\bar{K} > 0$ and $\bar{\delta} < 1$ such that the following holds: Fix $\delta > \bar{\delta}$, an arbitrary PBE (σ, μ) , and a history h^t at which $\mu(h^t) \geq p$. Then we have:*

- i) $\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t}^{\mathbb{T}-1} \delta^{\tau-t} |h^t, H \right] < \bar{K} (1 - \delta);$
- ii) $V_F(h^t; (\sigma, \mu), H) < \bar{K} (1 - \delta);$
- iii) $W_L(h^t; (\sigma, \mu)) < \bar{K} (1 - \delta).$

When the firm's belief falls in a firing region, the expected discounted length of the firm's relationship with the high type vanishes as the parties become arbitrarily patient (property i). Furthermore, both the firm's payoff from the interaction with the high type and the low type's payoff shrink to zero as δ goes to one (properties ii and iii, respectively). Although not part of its definition, a firing region satisfies the additional property that the equilibrium allocation must converge to the firing allocation.¹⁶

We are now ready to state the main result of this section.

¹⁴Recall that with commitment the high type is fired if the prior is above p^C and that p^C is larger than \hat{p} . Therefore, limited commitment enlarges (in the limit, as δ goes to one) the region of beliefs for which the equilibrium allocation is firing.

¹⁵When the prior is equal to \hat{p} , the limiting equilibrium outcome is not uniquely pinned down, as there are equilibria implementing the pooling allocation, the firing allocation and convex combinations of such allocations.

¹⁶Formally, we can take \bar{K} such that for every $\delta > \bar{\delta}$, every PBE (σ, μ) , and every history h^t at which $\mu(h^t) \geq p$, we also have $\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t |q_t - q_L^*| |h^t, L \right] < \bar{K} (1 - \delta)$, and $\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t (x_t - \theta_L q_L^* - \alpha) |h^t, L \right] < \bar{K} (1 - \delta)$.

Proposition 2 For every $p > \hat{p}$ the interval $[p, 1]$ is a firing region.

The proof of Proposition 2 is by induction. The interval $[p^C, 1]$ is a firing region since, for any δ , the unique equilibrium allocation is firing. This is the first step of the proof. In the inductive step, we take an arbitrary firing region $[p, 1]$ and extend it to include beliefs lower than p .

The following preliminary result plays a key role in the proof of the inductive step.

Lemma 4 Suppose that $[p, 1]$ is a firing region. There exist $K > 0$ and $\tilde{\delta} < 1$ such that for every $\delta > \tilde{\delta}$, for every PBE (σ, μ) , and for every history h^t , with $\mu(h^t) < p$, the following holds: Suppose that at h^t the firm offers a menu m_t containing a contract (x_t^L, q_t^L) accepted with positive probability and for which

$$\mu(h^t, m_t, (x_t^L, q_t^L)) \geq p$$

Then we have:

- i) $\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t}^{\mathbb{T}-1} \delta^{\tau-t} |h^t, m_t, H \right] < K(1 - \delta)$;
- ii) $V_F(h^t, m_t; (\sigma, \mu), H) < K(1 - \delta)$;
- iii) $W_L(h^t, m_t; (\sigma, \mu)) < K(1 - \delta)$.

Consider a history h^t at which the firm offers a menu m_t containing a contract that leads to a firing region. Lemma 4 provides bounds for the length of the high type's relationship and for the continuation payoffs conditional on the menu m_t being offered.

The proof of Lemma 4 is in Appendix B. The driving force behind Lemma 4 is similar to that behind Lemma 3, which establishes that separation with employment cannot occur for large values of δ . The intuition is as follows. Following the acceptance of a contract (x_t^L, q_t^L) that leads to the firing region, the low type's continuation payoff is close to zero. Suppose the firm's relationship with the high type is long lasting. In this case, only a large transfer x_t^L can prevent the low type from mimicking the high type. But then it becomes profitable for the high type to accept the contract (x_t^L, q_t^L) and then quit.

We now turn to the inductive step. For every $p \geq \hat{p}$ let $f(p) \in [0, p - \hat{p}]$ be defined by

$$\frac{f(p)}{p} \pi'_H(0) + \left(1 - \frac{f(p)}{p}\right) \pi_H(q_H^*) = (p - f(p)) \pi_L(q_L^*) \quad (2)$$

The function $f : [\hat{p}, 1] \rightarrow [0, 1 - \hat{p}]$ is strictly increasing and satisfies $f(\hat{p}) = 0$ and $f(p) < p - \hat{p}$ for every $p > \hat{p}$.¹⁷

¹⁷Recall that the function $\pi_H(\cdot)$ is concave and, therefore, $\pi'_H(0) > \pi'_H(0) q_H^* \geq \pi_H(q_H^*)$.

Lemma 5 *Suppose that the interval $[p, 1]$, $p \in (\hat{p}, 1)$, is a firing region. Then $\left[p - \frac{f(p)}{2}, 1\right]$ is also a firing region.*

The rest of the section provides the proof of Lemma 5 which consists of several steps. We outline in detail each step and relegate some technical arguments to Appendix B.

The continuation play starting at some history h^t is an equilibrium of the original game (when the prior is equal to the firm's belief at h^t). It is therefore without loss of generality (and convenient in terms of notation) to establish the three properties of a firing region for the initial history h^0 .

We first show that the expected discounted length of the high type's relationship shrinks to zero (property i of Definition 2). This is the main part of the proof of Lemma 5 because, as we show later, the remaining two properties follow from the first one.

Any equilibrium (σ, μ) must satisfy the following two conditions. First, the firm's equilibrium payoff $V_F(h^0)$ must be at least equal to $p_0\pi_L(q_L^*)$. Second, the low type must prefer his strategy to mimicking the high type. Thus, $W_L(h^0) \geq W_{LH}(h^0)$, where $W_{LH}(h^0)$ denotes the low type's continuation payoff at h^0 if he mimics the high type (at every history). We show that if $[p, 1]$ is a firing region, the prior p_0 is above $p - \frac{f(p)}{2}$ and δ is close to one, the two conditions just mentioned can be simultaneously satisfied only if the expected discounted length of the high type's relationship is close to zero.

We proceed in four steps. The payoffs $V_F(h^0)$, $W_L(h^0)$, $W_{LH}(h^0)$, and the expected length of the relationship are complicated objects since they depend on the entire history of the game. Instead of working directly with these variables we replace them with some bounds which are easy to express and compare. We do this in the first three steps. In Step 1, we show that it is without loss of generality to restrict attention to a class of simple equilibria. In Step 2, we provide bounds to the continuation payoffs and to the length of the relationship with the high type. These bounds are derived by changing the timing of the transfers in the equilibrium of the game. In Step 3 we introduce an auxiliary game which allows for a simpler expression of these bounds. Finally, in Step 4 we use these bounds to establish property i of Definition 2.

Step 1: Restriction to a class of simple equilibria

We now show that to establish property i of Definition 2 is without loss of generality to restrict attention to equilibria in which (i) the firm's strategy in the first period is pure (i.e., the firm does not randomize among different menus at $t = 0$); and (ii) the high type's

equilibrium payoff is equal to zero. To see why restriction (i) is without loss, suppose that $((\sigma^F, \sigma^H, \sigma^L), \mu)$ is a PBE and m_0 is a menu offered with positive probability by the firm at $t = 0$ ($\sigma_0^F(m_0|h_0) > 0$). Let $\tilde{\sigma}^F$ be the strategy which is identical to σ^F in every period except the first one, in which instead the firm offers the menu m_0 with certainty ($\tilde{\sigma}_0^F(m_0|h_0) = 1$). The assessment $((\tilde{\sigma}^F, \sigma^H, \sigma^L), \mu)$ is also a PBE. Furthermore, the outcome of the equilibrium $((\tilde{\sigma}^F, \sigma^H, \sigma^L), \mu)$ coincides with the continuation outcome of $((\sigma^F, \sigma^H, \sigma^L), \mu)$ after the firm proposes the contract m_0 . Therefore, if it is impossible to establish the first property for arbitrary PBE, then it is also impossible to establish the property for the class of equilibria which satisfy restriction (i).

We show next that restriction (ii) is also without loss of generality. Suppose that (σ, μ) is a PBE in which the firm offers the menu m_0 (with probability one) in the first period and which yields a strictly positive payoff $W_H(h^0; (\sigma, \mu))$ to the high type. Then it is possible to construct a new PBE $(\tilde{\sigma}, \tilde{\mu})$ which is outcome equivalent to (σ, μ) , except for the fact that the first-period transfers are uniformly decreased by $(1 - \delta)^{-1} W_H(h^0; (\sigma, \mu))$. In other words, in the first period the firm replaces every contract (x_0, q_0) in the menu m_0 with the contract $(x_0 - (1 - \delta)^{-1} W_H(h^0; (\sigma, \mu)), q_0)$.¹⁸ Finally, notice that the first property of a firing region does not depend on equilibrium transfers.

Step 2: Change in the timing of transfers

This step provides bounds to the payoffs and to the length of the relationship that only depend on the quality of the goods delivered by the worker. Let $\tilde{\mathbb{T}} \in \mathbf{N} \cup \{\infty\}$ denote the random time that stops the play at the first history $(h^{\tilde{\mathbb{T}}}, m_{\tilde{\mathbb{T}}})$ at which the menu $m_{\tilde{\mathbb{T}}}$ contains a contract $(x_{\tilde{\mathbb{T}}}, q_{\tilde{\mathbb{T}}})$ accepted with positive probability and for which $\mu(h^{\tilde{\mathbb{T}}}, m_{\tilde{\mathbb{T}}}, (x_{\tilde{\mathbb{T}}}, q_{\tilde{\mathbb{T}}})) \geq p$ (we set $\tilde{\mathbb{T}} = \infty$ if the event does not occur in finite time). We claim that there exist $K > 0$ and $\bar{\delta} < 1$ such that for $\delta > \bar{\delta}$ we have:

$$\begin{aligned} V_F(h^0) &\leq \bar{V}_F(h^0) := \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{\mathbb{T}}-1} \delta^t \pi_H(q_t) + \mathbb{I}_{\{\tilde{\mathbb{T}} < \infty\}} \delta^{\tilde{\mathbb{T}}} \mu(h^{\tilde{\mathbb{T}}}) \pi_L(q_L^*) \right] + K(1 - \delta) \\ W_L(h^0) &\leq \bar{W}_L(h^0) := \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{\mathbb{T}}-1} \delta^t \Delta \theta_{q_t} |L \right] + K(1 - \delta) \end{aligned} \tag{3}$$

¹⁸In the new equilibrium $(\tilde{\sigma}, \tilde{\mu})$, each type of worker accepts the contract $(x_0 - (1 - \delta)^{-1} W_H(h^0; (\sigma, \mu)), q_0)$ with the same probability with which he accepts the contract (x_0, q_0) in the original equilibrium (σ, μ) .

$$W_{LH}(h^0) \geq \underline{W}_{LH}(h^0) := \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \Delta \theta q_t | H \right] - K(1 - \delta)$$

and

$$\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t | h^0, H \right] \leq \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t | h^0, H \right] + K(1 - \delta) \quad (4)$$

The bounds on the payoffs provided in (3) are derived as follows (the formal derivation is provided in the appendix). Recall that the high type's equilibrium payoff $W_H(h^0)$ is equal to zero. Let m_0 be the menu offered by the firm at h^0 and consider any contract $(x_0, q_0) \in m_0$ accepted by the high type with positive probability. Let h^1 denote the history $(h^0, m_0, (x_0, q_0))$. If $W_H(h^1) = 0$, then we have $x_0 = \theta_H q_0 + \alpha$. If instead $W_H(h^1) > 0$, then we increase the transfer x_0 by the amount $\frac{\delta}{(1-\delta)} W_H(h^1)$. The new transfer is equal to $\theta_H q_0 + \alpha$. At the same time, for every menu m_1 offered at h^1 , we decrease all the transfers of the contracts in m_1 by the amount $\frac{1}{(1-\delta)} W_H(h^1, m_1)$. These changes leave all parties' continuation payoffs unchanged. We repeat this procedure in periods $1, \dots, \tilde{T} - 1$. Thus, for every $t = 0, \dots, \tilde{T} - 1$, the new transfer in period t is equal to $\theta_H q_t + \alpha$. Furthermore, it follows from Lemma 4 that at the history $(h^{\tilde{T}}, m_{\tilde{T}})$, if δ is close to one, the firm's continuation payoff is close to $\mu(h^{\tilde{T}}) \pi_L q_L^*$, while the worker's payoff is close to zero.¹⁹ Combining the expressions of the new transfers with the findings on the continuation payoffs we obtain the bounds in (3).

The bound on the expected length of the high type's relationship (inequality 4) follows directly from Lemma 4 (property i).

Step 3: The auxiliary game

We introduce an auxiliary game (a direct mechanism) which replicates the equilibrium outcome from period 0 through period $\tilde{T} - 1$ and implements the payoffs $\bar{V}_F(h^0)$, $\bar{W}_L(h^0)$, and $\underline{W}_{LH}(h^0)$. In particular, we consider a mechanism in which the worker reveals his type to a designer. The designer selects a history and one of two messages, \mathbf{m}_0 and \mathbf{m}_p (see below for more details about the mechanism). Let Υ_z , for $z \in \{0, p\}$, denote the expected discounted length of the relationship conditional on the message \mathbf{m}_z , and let \tilde{q}_z denote the

¹⁹In fact, the high type's payoff is bounded above by the low type's payoff which, in turn, is close to zero (see property iii in Lemma 4).

expected discounted quality of the good conditional on \mathbf{m}_z . Using the auxiliary game, we show that

$$\begin{aligned}\bar{V}_F(h^0) &\leq \check{V}_F(h^0) := \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H(\tilde{q}_0) + \frac{p_0}{p} [\Upsilon_p \pi_H(\tilde{q}_p) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K(1 - \delta) \\ \bar{W}_L(h^0) &= \Upsilon_p \tilde{q}_p \Delta\theta + K(1 - \delta) \\ \underline{W}_{LH}(h^0) &= \left(\frac{1}{1-p_0}\right) \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \tilde{q}_0 \Delta\theta + \left(\frac{1-p}{1-p_0}\right) \left(\frac{p_0}{p}\right) \Upsilon_p \tilde{q}_p \Delta\theta - K(1 - \delta)\end{aligned}\tag{5}$$

and

$$\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{T-1} \delta^t |h^0, H \right] \leq \left(\frac{1}{1-p_0}\right) \left(1 - \frac{p_0}{p}\right) \Upsilon_0 + \left(\frac{1-p}{1-p_0}\right) \left(\frac{p_0}{p}\right) \Upsilon_p + K(1 - \delta)\tag{6}$$

This greatly simplifies our problem since it reduces to four the number of endogenous variables (Υ_z and \tilde{q}_z , $z \in \{0, p\}$) in the expressions of the bounds. We then use the expressions in (5) and (6) to show that the inequalities $\check{V}_F(h^0) \geq p_0 \pi_L(q_L^*)$ (recall that $\check{V}_F(h^0) \geq \bar{V}_F(h^0) \geq V_F(h^0)$) and $\bar{W}_L(h^0) \geq \underline{W}_{LH}(h^0)$ can be satisfied only if the right hand side of inequality (6) shrinks to zero (as δ goes to one), establishing therefore the first property of a firing region.

We now describe the direct mechanism. The worker reveals his private information to a designer who, in turn, chooses an outcome and reports it to the firm. The outcome consists of a history $h^{\tilde{\Pi}}$ of the game and a message in $\{\mathbf{m}_0, \mathbf{m}_p\}$. The designer chooses the outcomes in such a way that upon receiving message \mathbf{m}_z the firm's belief is equal to z . Moreover, the likelihoods of the histories reflect those from the equilibrium of the game. Specifically, if the worker announces the low type, then the designer chooses the outcome $(h^{\tilde{\Pi}}, \mathbf{m}_p)$ with probability $\Pr(h^{\tilde{\Pi}}|L)$, the probability of the history $h^{\tilde{\Pi}}$ when the worker's type is low and the parties play the equilibrium (σ, μ) . On the other hand, if the worker announces the high type, then the designer chooses the outcome $(h^{\tilde{\Pi}}, \mathbf{m}_p)$ with probability $\Pr(h^{\tilde{\Pi}}|H) \frac{\mu(h^{\tilde{\Pi}})^{(1-p)}}{(1-\mu(h^{\tilde{\Pi}}))^p}$ and the outcome $(h^{\tilde{\Pi}}, \mathbf{m}_0)$ with probability $\Pr(h^{\tilde{\Pi}}|H) \left[1 - \frac{\mu(h^{\tilde{\Pi}})^{(1-p)}}{(1-\mu(h^{\tilde{\Pi}}))^p}\right]$.

We now turn to the payoffs of the firm and the low type. The firm's payoff depends only on the outcome and not on the worker's report. Consider an arbitrary history $h^{\tilde{\Pi}} = (m_0, (x_0, q_0), \dots, m_{\tilde{\Pi}}, (x_{\tilde{\Pi}-1}, q_{\tilde{\Pi}-1}))$. If the outcome is $(h^{\tilde{\Pi}}, \mathbf{m}_p)$, the firm's payoff is equal

to:

$$(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \pi_H(q_t) + \mathbb{I}_{\{\tilde{T} < \infty\}} \delta^{\tilde{T}} p \pi_L(q_L^*) + K(1 - \delta)$$

If the outcome is $(h^{\tilde{T}}, \mathbf{m}_0)$, the firm's payoff is equal to:

$$(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \pi_H(q_t) + K(1 - \delta)$$

We conclude that if every type reveals his type truthfully, then the firm's expected payoff is equal to $\bar{V}_F(h^0)$.

Consider now the low type. His payoff depends both on the outcome and on his report. First, if the outcome is either $(h^{\tilde{T}}, \mathbf{m}_p)$ or $(h^{\tilde{T}}, \mathbf{m}_0)$, with $h^{\tilde{T}} = (m_0, (x_0, q_0), \dots, m_{\tilde{T}}, (x_{\tilde{T}-1}, q_{\tilde{T}-1}))$, then the low type obtains a payoff equal to

$$(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \Delta \theta q_t$$

In addition, the low type obtains an extra payoff equal to $K(1 - \delta)$ if he is honest, and equal to $-K(1 - \delta)$ if he lies to the designer. It follows that the low type's expected payoff is equal to $\bar{W}_L(h^0)$ if he reveals his type truthfully, and equal to $\underline{W}_{LH}(h^0)$ if he lies to the designer. Finally, we assume that the high type is committed to reporting truthfully his type to the designer.

It is natural to ask why we introduced the messages \mathbf{m}_0 and \mathbf{m}_p in the mechanism, given that they do not affect the worker's payoffs and the firm is a passive player in the auxiliary game. The reason for this is that the additional messages allow us to classify all the histories $h^{\tilde{T}}$ into two large classes, depending on whether they are associated with the message \mathbf{m}_0 or the message \mathbf{m}_p . Recall that when the firm observes the message \mathbf{m}_z , $z \in \{0, p\}$, its belief is equal to z . Thus, by the martingale property of the beliefs (see Aumann and Maschler (1995) and Kamenica and Gentzkow (2011)), we conclude that the probability of observing the message \mathbf{m}_0 is equal to $(1 - \frac{p_0}{p})$, while the probability of observing the message \mathbf{m}_p is equal $\frac{p_0}{p}$. Below, we use this fact to express the bounds on the continuation payoffs and the length of the relationship in terms of the endogenous variables Υ_z and \tilde{q}_z , $z \in \{0, p\}$.

First, observe that we can rewrite the firm's payoff $\bar{V}_F(h^0)$ as follows:

$$\begin{aligned} \bar{V}_F(h^0) &= \left(1 - \frac{p_0}{p}\right) \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\hat{T}-1} \delta^t \pi_H(q_t) | \mathbf{m}_0 \right] + \\ &\frac{p_0}{p} \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\hat{T}-1} \delta^t \pi_H(q_t) + \mathbb{I}_{\{\hat{T} < \infty\}} \delta^{\hat{T}} p \pi_L(q_L^*) | \mathbf{m}_p \right] + K(1 - \delta) \end{aligned} \quad (7)$$

We now turn to the low type's payoffs. Fix an arbitrary outcome $(h^{\hat{T}}, \mathbf{m}_p)$ and let $\Pr(h^{\hat{T}}, \mathbf{m}_p)$ denote the ex-ante probability of the outcome. Recall that the firm's belief upon observing the outcome $(h^{\hat{T}}, \mathbf{m}_p)$ is equal to p . This immediately implies:

$$\Pr(h^{\hat{T}}, \mathbf{m}_p) = \frac{p_0}{p} \Pr(h^{\hat{T}}, \mathbf{m}_p | L) = \frac{1 - p_0}{1 - p} \Pr(h^{\hat{T}}, \mathbf{m}_p | H)$$

We conclude that the outcome $(h^{\hat{T}}, \mathbf{m}_p)$ is reached with probability $\frac{p}{p_0} \Pr(h^{\hat{T}}, \mathbf{m}_p)$ when the worker announces that his type is low, and with probability $\frac{1-p}{1-p_0} \Pr(h^{\hat{T}}, \mathbf{m}_p)$ when the worker announces that his type is high.

Similarly, an outcome $(h^{\hat{T}}, \mathbf{m}_0)$ is reached with probability $\frac{1}{1-p_0} \Pr(h^{\hat{T}}, \mathbf{m}_0)$ if the worker announces the high type and with probability zero if the worker announces the low type ($\Pr(h^{\hat{T}}, \mathbf{m}_0)$ denotes the ex-ante probability of the outcome).

Combining these observations, we can rewrite the low type's payoffs as follows:

$$\begin{aligned} \bar{W}_L(h^0) &= \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\hat{T}-1} \delta^t \Delta \theta q_t | \mathbf{m}_p \right] + K(1 - \delta) \\ \bar{W}_{LH}(h^0) &= \left(\frac{1}{1-p_0}\right) \left(1 - \frac{p_0}{p}\right) \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\hat{T}-1} \delta^t \Delta \theta q_t | \mathbf{m}_0 \right] + \\ &\left(\frac{1-p}{1-p_0}\right) \left(\frac{p_0}{p}\right) \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\hat{T}-1} \delta^t \Delta \theta q_t | \mathbf{m}_p \right] - K(1 - \delta) \end{aligned}$$

We next provide the formal definitions of Υ_z and \tilde{q}_z for $z \in \{0, p\}$:

$$\Upsilon_z = \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\hat{T}-1} \delta^t | \mathbf{m}_z \right]$$

and

$$\tilde{q}_z = \frac{1}{\Upsilon_z} \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t q_t | \mathbf{m}_z \right]$$

if $\Upsilon_z \neq 0$, and $\tilde{q}_z = 0$ otherwise.

The definitions of Υ_z and \tilde{q}_z , $z \in \{0, p\}$, allow us to express $\bar{W}_L(h^0)$ and $\underline{W}_{LH}(h^0)$ as in (5). The inequality $\bar{V}_F(h^0) \leq \check{V}_F(h^0)$ in (5) follows from equation (7), the concavity of the function $\pi_H(\cdot)$, and Jensen's inequality.

Finally, using the definitions of Υ_0 and Υ_p and inequality (4), we are able to bound the length of the high type's relationship as in (6).

Step 4: Bounding the expected length of the high type's relationship

The following claim establishes the first property of a firing region.

Claim 1 *Fix $K > 0$ and $p > \hat{p}$. There exists $K' > 0$ such that, for every $p_0 \in \left[p - \frac{f(p)}{2}, p \right]$, for every δ , and for every $(\Upsilon_z, \tilde{q}_z) \in [0, 1]^2$, $z \in \{0, p\}$, the inequalities*

$$\left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H(\tilde{q}_0) + \frac{p_0}{p} [\Upsilon_p \pi_H(\tilde{q}_p) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K(1 - \delta) \geq p_0 \pi_L(q_L^*) \quad (8)$$

$$\Upsilon_p \tilde{q}_p \Delta\theta + K(1 - \delta) \geq \left(\frac{1}{1 - p_0}\right) \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \tilde{q}_0 \Delta\theta + \left(\frac{1 - p}{1 - p_0}\right) \left(\frac{p_0}{p}\right) \Upsilon_p \tilde{q}_p \Delta\theta - K(1 - \delta) \quad (9)$$

are simultaneously satisfied only if

$$\left(\frac{1}{1 - p_0}\right) \left(1 - \frac{p_0}{p}\right) \Upsilon_0 + \left(\frac{1 - p}{1 - p_0}\right) \left(\frac{p_0}{p}\right) \Upsilon_p + K(1 - \delta) \leq K'(1 - \delta) \quad (10)$$

The inequalities (8) and (9) capture the constraints $\check{V}_F(h^0) \geq p_0 \pi_L(q_L^*)$ and $\bar{W}_L(h^0) \geq \underline{W}_{LH}(h^0)$, respectively, while the left hand side of inequality (10) represents the upper bound to the expected discounted length of the high type's relationship.

The proof of Claim 1 is tedious and relegated to the appendix. The logic behind this claim is better understood when one considers the problem of maximizing $\check{V}_F(h^0)$ (with respect to Υ_z and \tilde{q}_z , $z \in \{0, p\}$) subject to the low type's incentive compatibility constraint: $\bar{W}_L(h^0) \geq \underline{W}_{LH}(h^0)$. Clearly, $\check{V}_F(h^0)$ is maximized by setting Υ_0 equal to one, \tilde{q}_0 equal to q_H^* , and Υ_p equal to zero (recall that $p > \hat{p}$ and, therefore, $p \pi_L(q_L^*) > \pi_H(q_H^*) \geq \pi_H(\tilde{q}_p)$ for any \tilde{q}_p). However, this would violate the low type's incentive compatibility constraint.

Hence the following trade-off emerges. To increase the firm's payoff by increasing Υ_0 while also satisfying the incentive compatibility constraint (9), it is also necessary to increase Υ_p , which decreases the firm's payoff. Notice that when the prior p_0 is close to p , $\left(1 - \frac{p_0}{p}\right) \Upsilon_0$ and \tilde{q}_0 have a small impact both on the firm's payoff and the constraint. In contrast, Υ_p has a small impact on the constraint, and a large (negative) impact on the firm's payoff. We conclude that for δ and p_0 sufficiently large, the optimal values of $\left(1 - \frac{p_0}{p}\right) \Upsilon_0$ and Υ_p must be close to zero. Therefore, if we could maximize the firm's payoff, subject to $\bar{W}_L(h^0) \geq \underline{W}_{LH}(h^0)$, the solution would be close to a firing allocation, yielding a payoff close to $p_0 \pi_L(q_L^*)$. For the same reason, any allocation that satisfies $\bar{W}_L(h^0) \geq \underline{W}_{LH}(h^0)$ and that is not close to a firing allocation leads to a payoff for the firm smaller than $p_0 \pi_L(q_L^*)$, hence violating (8).

We have shown that the first property of a firing region holds. In particular, there exist $\bar{K} > 0$ and $\bar{\delta} < 1$ such that for any $\delta > \bar{\delta}$, any prior p_0 above $p - \frac{f(p)}{2}$, and any PBE (σ, μ) , we have $E_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t |h^0, H \right] \leq (1 - \delta) \bar{K}$. We now turn to the remaining two properties. To verify the second property (the firm's payoff conditional on type H shrinks to zero weakly faster than $1 - \delta$), notice that

$$V_F(h^0; H) \leq v(1) \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t |h^0, H \right] \leq v(1) \bar{K} (1 - \delta)$$

Finally, we use the result above to bound the low type's continuation payoff $W_L(h^0)$ (third property). We have

$$p_0 \pi_L(q_L^*) \leq V_F(h^0) \leq (1 - p_0) V_F(h^0; H) + p_0 (\pi_L(q_L^*) - W_L(h^0))$$

which implies

$$W_L(h^0) \leq \frac{1 - p_0}{p_0} V_F(h^0; H) < \frac{1 - \hat{p}}{\hat{p}} V_F(h^0; H) \leq \frac{1 - \hat{p}}{\hat{p}} v(1) \bar{K} (1 - \delta)$$

This concludes the proof of Lemma 5.

5.2 Low Belief Case: $p < \hat{p}$

In this section, we characterize the limiting equilibrium outcome when the prior is lower than \hat{p} . To do so, we first define the notion of *pooling region*.

Definition 3 *The interval $[0, p]$ is a pooling region if for every $\varepsilon > 0$ there exists $\bar{\delta} < 1$ such that the following holds: Fix $\delta > \bar{\delta}$, an arbitrary PBE (σ, μ) , and a history h^t at which $\mu(h^t) \leq p$. Then we have:*

$$i) \mathbb{E}_{(\sigma, \mu)} [\delta^{\mathbb{T}}] < \varepsilon;$$

$$ii) \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t |q_t - q_H^*| \right] < \varepsilon;$$

$$iii) \text{ For } i \in \{H, L\}, \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t (x_t - \theta_H q_H^* - \alpha) |i \right] < \varepsilon.$$

When the belief is in a pooling region, the equilibrium allocation converges to the pooling allocation as δ goes to one. Our next result shows that all beliefs lower than \hat{p} belong to a pooling region.

Proposition 3 *For every $p < \hat{p}$ the interval $[0, p]$ is a pooling region.*

Before turning to the proof of Proposition 3, we establish a preliminary result. We show that the equilibrium belief cannot grow too quickly around \hat{p} when the parties are sufficiently patient.

Lemma 6 *For every $\varepsilon > 0$ there exists $\bar{\delta} < 1$ such that the following holds: Fix $\delta > \bar{\delta}$, and a PBE (σ, μ) . There does not exist a history h^t with $\mu(h^t) < \hat{p} - \varepsilon$ at which the firm offers a menu m_t containing a contract (x_t, q_t) accepted with positive probability and such that $\mu(h^t, m_t, (x_t, q_t)) > \hat{p} + \varepsilon$.*

The proof of Lemma 6 is in Appendix B. Recall that for every $\varepsilon > 0$, the interval $[\hat{p} + \varepsilon, 1]$ is a firing region. Therefore, it follows from Lemma 4 that if δ is close to one and the belief jumps from $\mu(h^t) < \hat{p} - \varepsilon$ to $\mu(h^t, m_t, (x_t, q_t)) > \hat{p} + \varepsilon$, the firm's continuation payoff at h^t must be close to $\mu(h^t) \pi_L(q_L^*)$. But then the firm's payoff would be smaller than $\pi_H(q_H^*)$ (since $\mu(h^t) < \hat{p} - \varepsilon$ and $\hat{p} \pi_L(q_L^*) = \pi_H(q_H^*)$), contradicting Lemma 1.

We now outline the proof of Proposition 3 (see Appendix B for the formal proof). As in the previous section, we simplify the notation and establish the three properties of a pooling region for the initial history h^0 .

The first two properties do not depend on equilibrium transfers. Therefore, it is without loss of generality to establish these properties for the class of equilibria in which (i) the

firm's strategy in the first period is pure; and (ii) the high type's equilibrium payoff is equal to zero.²⁰

We start with the first property and fix $p < \hat{p}$. By contradiction, let us assume that there exists a sequence $\{\delta_n, p_{0,n}, (\sigma_n, \mu_n)\}_{n=1}^\infty$ such that δ_n converges to one, $p_{0,n} \in [0, p]$, (σ_n, μ_n) is a PBE of the game with discount factor equal to δ_n and prior equal to $p_{0,n}$, and $\lim_{n \rightarrow \infty} \mathbb{E}_{(\sigma_n, \mu_n)} [\delta_n^{\tilde{T}}] = \xi > 0$. We show that for n sufficiently large it is strictly profitable for the low type to deviate from the equilibrium strategy and mimic the high type. To ease the notation, in what follows we suppress the index n and write δ , p_0 , and (σ, μ) to denote an arbitrary element of the sequence.

Below, we proceed as follows. First, we replace the low type's payoffs (from the equilibrium strategy and the deviation) with bounds which depend on the expected discounted time at which the firm's belief falls for the first time in the firing region. We then compute this expected discounted time conditional on the worker's type. With this we show that when the parties are sufficiently patient it is strictly profitable for the low type to deviate and mimic the high type.

Fix a small ε and now let $\tilde{T} \in \mathbf{N} \cup \{\infty\}$ denote the random time that stops the play at the first history $(h^{\tilde{T}}, m_{\tilde{T}})$ at which the menu $m_{\tilde{T}}$ contains a contract $(x_{\tilde{T}}, q_{\tilde{T}})$ that is accepted with a positive probability and for which $\mu(h^{\tilde{T}}, m_{\tilde{T}}, (x_{\tilde{T}}, q_{\tilde{T}})) \geq \hat{p} + \varepsilon$. Recall that $[\hat{p} + \varepsilon, 1]$ is a firing region and $W_H(h^0) = 0$. We proceed similarly as in Section 5.1 by making a change in the timing of the transfers and using Lemma 4 to bound the continuation payoffs at $(h^{\tilde{T}}, m_{\tilde{T}})$ (see the discussion after inequality (4)). For δ close to one the low type's payoffs $W_L(h^0)$ and $W_{LH}(h^0)$ are bounded as follows:

$$\begin{aligned} W_L(h^0) &\leq \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \Delta \theta_{q_t} | L \right] + \varepsilon \\ W_{LH}(h^0) &\geq \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \Delta \theta_{q_t} | H \right] - \varepsilon \end{aligned} \tag{11}$$

Also, for δ close to one the firm's payoff is bounded by:

$$V_F(h^0) \leq \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \pi_H(q_t) + \mathbb{I}_{\{\tilde{T} < \infty\}} \delta^{\tilde{T}} \mu(h^{\tilde{T}}) \pi_L(q_L^*) \right] + \varepsilon \leq$$

²⁰The proof of this claim is identical to the proof provided in Section 5.1 and, therefore, omitted.

$$\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \pi_H(q_t) + \mathbb{I}_{\{\tilde{T} < \infty\}} \delta^{\tilde{T}} (\hat{p} + \varepsilon) \pi_L(q_L^*) \right] + \varepsilon$$

where the second inequality holds because the belief at the history $h^{\tilde{T}}$ is bounded above (by definition) by $\hat{p} + \varepsilon$.

Notice that ε can be arbitrarily small. Therefore, since $\pi_H(q_H^*) = \hat{p} \pi_L(q_L^*)$, and q_H^* is the unique maximizer of $\pi_H(\cdot)$, the inequality above implies that

$$\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t |q_H^* - q_t| \right] \approx 0 \quad (12)$$

for δ sufficiently large. If this were not the case, then the firm's payoff would be strictly smaller than $\pi_H(q_H^*)$ (again, for δ sufficiently large).

Combining (11) with (12), we obtain that for δ close to one the upper bound of $W_L(h^0)$ is close to

$$\Delta \theta q_H^* \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t |L| \right] = \Delta \theta q_H^* \left(1 - \mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} |L| \right] \right) \quad (13)$$

while the lower bound of $W_{LH}(h^0)$ is close to

$$\Delta \theta q_H^* \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t |H| \right] = \Delta \theta q_H^* \left(1 - \mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} |H| \right] \right) \quad (14)$$

We now approximate $\mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} |L| \right]$ and $\mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} |H| \right]$. We do this in two steps. First, we express $\mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} |i| \right]$ as function of $\mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} \right]$. Then we approximate $\mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} \right]$.

Let $\Pr(h^{\tilde{T}})$ denote the (ex-ante) probability of reaching the history $h^{\tilde{T}}$. Then $h^{\tilde{T}}$ is reached with a probability $\frac{\mu(h^{\tilde{T}})}{p_0} \Pr(h^{\tilde{T}})$ if the low type follows his strategy σ^L , and with a probability $\frac{1-\mu(h^{\tilde{T}})}{1-p_0} \Pr(h^{\tilde{T}})$ if he mimics the high type and plays the strategy σ^H . Also, it follows from Lemma 6 that for δ close to one $\mu(h^{\tilde{T}})$ must be close to \hat{p} . Therefore, for δ sufficiently large, we have

$$\begin{aligned} \mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} |L| \right] &\approx \frac{\hat{p}}{p_0} \mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} \right] \\ \mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} |H| \right] &\approx \frac{1-\hat{p}}{1-p_0} \mathbb{E}_{(\sigma, \mu)} \left[\delta^{\tilde{T}} \right] \end{aligned} \quad (15)$$

Next, we examine the relationship between $\mathbb{E}_{(\sigma,\mu)}[\delta^{\mathbb{T}}]$ and $\mathbb{E}_{(\sigma,\mu)}[\delta^{\hat{\mathbb{T}}}]$ when δ is close to one. First, at the history $(h^{\hat{\mathbb{T}}}, m_{\hat{\mathbb{T}}})$ the expected discounted length of the relationship with the high type is close to zero (see Lemma 4 and recall that the interval $[\hat{p} + \varepsilon, 1]$ is a firing region). Second, the firm's belief at $h^{\hat{\mathbb{T}}}$ must be close to \hat{p} (a value of $\mu(h^{\hat{\mathbb{T}}})$ far away from \hat{p} would contradict Lemma 6). Finally, recall that for δ large, $\mathbb{E}_{(\sigma,\mu)}[\delta^{\mathbb{T}}]$ is close (by assumption) to ξ . Putting these observations together and using Bayes' rule, we conclude that $\mathbb{E}_{(\sigma,\mu)}[\delta^{\hat{\mathbb{T}}}]$ is close to $\frac{\xi}{1-\hat{p}}$ for δ close to one.

Combining (13), (14), (15), and the last observation, we conclude that for δ close to one the upper bound of $W_L(h^0)$ is close to

$$\Delta\theta q_H^* \left(1 - \frac{\hat{p}}{p_0} \frac{\xi}{1-\hat{p}}\right) \leq \Delta\theta q_H^* \left(1 - \frac{\hat{p}}{p} \frac{\xi}{1-\hat{p}}\right)$$

(where the inequality follows from $p_0 \leq p$), while the lower bound of $W_{LH}(h^0)$ is close to

$$\Delta\theta q_H^* \left(1 - \frac{1-\hat{p}}{1-p_0} \frac{\xi}{1-\hat{p}}\right) \geq \Delta\theta q_H^* \left(1 - \frac{1-\hat{p}}{1-p} \frac{\xi}{1-\hat{p}}\right)$$

Hence, since $p < \hat{p}$ and $\xi > 0$, we have

$$\Delta\theta q_H^* \left(1 - \frac{\hat{p}}{p} \frac{\xi}{1-\hat{p}}\right) < \Delta\theta q_H^* \left(1 - \frac{1-\hat{p}}{1-p} \frac{\xi}{1-\hat{p}}\right)$$

which implies the existence of a profitable deviation for values of δ close to one.

Finally, the second and third property of a pooling region are direct consequences of the first property. Intuitively, if the high type never quits the relationship, the best option for the firm is to implement the best pooling allocation.

6 Rehiring

In the model analyzed so far, the worker's decision to reject all the contracts in the menu is an irreversible action that ends the relationship. In other words, the firm cannot rehire the worker after a period of unemployment. As we argued in the previous two sections, this impairs the firm's ability to screen the worker. Once the worker reveals his type, his continuation payoff must be equal to zero. The firm can afford to pay the reservation wage, because the worker has no alternative than to end the relationship.

This logic does not apply when rehiring is possible. In this case, the worker can credibly threaten the firm to reject offers that pay slightly above the reservation wage, because he expects to obtain a large payoff in the rest of the relationship. As we will see below, with rehiring, it is possible to sustain equilibrium outcomes in which the worker's payoff remains strictly positive even when his type is known to the other party. This, in turn, makes it easier for the firm to screen the worker.

There are different ways to break the link between the decision to reject all the contracts and the decision to end the relationship. One possibility is to assume that the relationship lasts forever and quitting is not allowed. Another possibility is to add to the benchmark model the option for the worker to reject all the contracts and remain in the relationship. In the rest of the section, we analyze these extensions of the model.

We start with the infinitely repeated game in which, in each period, the firm proposes a menu of contracts. The worker either accepts a contract in the menu or rejects all of them. Both parties obtain a payoff equal to zero in each period in which the worker rejects all the contracts in the menu.

Consider the standard mechanism design problem with commitment. We say that the payoffs $(V_{F,H}, V_{F,L}, W_H, W_L)$ are *incentive-compatible* and *ex-post strictly individually rational* if there exists an incentive compatible direct mechanism $\{(x_H, q_H), (x_L, q_L)\}$, $(x_i, q_i) \in \mathbb{R}_{++} \times [0, 1]$ for $i \in \{H, L\}$, satisfying:

- i) For $i \in \{H, L\}$, the firm's payoff $V_{F,i}$ when the worker is of type i is strictly positive: $V_{F,i} := v(q_i) - x_i > 0$;
- ii) For $i \in \{H, L\}$, type i 's payoff is strictly positive: $W_i := x_i - \theta_i q_i - \alpha > 0$.

The main result of this section is a folk theorem. We show that any profile of incentive-compatible and ex-post strictly individual rational payoffs can be achieved in the infinitely repeated game when the parties are sufficiently patient.

Proposition 4 *For every tuple $(V_{F,H}, V_{F,L}, W_H, W_L) \in \mathbb{R}_{++}^4$ of incentive-compatible and ex-post strictly individually rational payoffs there exists $\delta^\dagger \in (0, 1)$ such that for every $\delta \geq \delta^\dagger$ there exists a PBE (of the infinitely repeated game) that leads to such payoffs.*

The proof of Proposition 4 is in the Supplemental Material. Fix a tuple $(V_{F,H}, V_{F,L}, W_H, W_L)$ of incentive-compatible and ex-post strictly individually rational payoffs and let $\{(x_H, q_H), (x_L, q_L)\}$ denote the corresponding direct mechanism. We construct an equilibrium which consists of two phases. The screening phase takes place in the first period when the firm offers

the menu $\{(x_H, q_H), (x_L, q_L)\}$. Each type $i \in \{H, L\}$ selects the menu (x_i, q_i) and the firm learns the worker's type. The post-screening phase with type $i \in \{H, L\}$ begins in the second period and implements the contract (x_i, q_i) in every period.

In equilibrium, the firm never updates its belief in the post-screening phase. It is therefore necessary to show that in the game with *complete* information with type i , there exists an equilibrium that implements (x_i, q_i) in every period (when the parties are sufficiently patient). The following lemma establishes this important result.

Lemma 7 *Consider the infinitely repeated game with complete information in which the firm interacts with type $i \in \{H, L\}$. Let (x_i, q_i) be a contract yielding the payoff $V_{F,i} = v(q_i) - x_i > 0$ to the firm and the payoff $W_i = x_i - \theta_i q_i - \alpha > 0$ to the worker. There exists $\delta^\dagger \in (0, 1)$ such that for every $\delta \geq \delta^\dagger$ there exists a subgame perfect equilibrium that leads to the payoffs $(V_{F,i}, W_i)$.*

Proof of Lemma 7.

Fix $\varepsilon \in (0, \min\{V_{F,i}, W_i\})$. Let (\underline{x}_i, q_i^*) , $\underline{x}_i = \alpha + \theta_i q_i^* + \frac{\varepsilon}{2}$, denote the efficient contract that yields the payoff $\frac{\varepsilon}{2}$ to the worker. Also, let (\bar{x}_i, q_i^*) , $\bar{x}_i = v(q_i^*) - \frac{\varepsilon}{2}$, denote the efficient contract that yields the payoff $\frac{\varepsilon}{2}$ to the firm.

Consider the following strategy profile, generated by a simple three states automaton.

State $(i, 0)$: This is the initial state. The automaton prescribes that the firm offers the menu $\{(x_i, q_i)\}$ and the worker accepts the contract (x_i, q_i) . The state remains $(i, 0)$ unless there is a deviation by the firm, in which case the state changes to $(i, 1)$, irrespective of the worker's decision. When the firm deviates and offers a menu different from $\{(x_i, q_i)\}$, the worker accepts the contract which maximizes his current payoff, provided that this is positive (here and in what follows, we require the worker to select the contract with the smallest index if there are multiple contracts yielding the highest current payoff). Finally, the worker rejects all the contracts if they all yield a negative payoff.

State $(i, 1)$: The automaton prescribes that the firm offers the menu $\{(\bar{x}_i, q_i^*)\}$ and the worker accepts (\bar{x}_i, q_i^*) . If the firm offers $\{(\bar{x}_i, q_i^*)\}$, the state remains $(i, 1)$ irrespective of the worker's decision. Suppose instead that the firm deviates and offers the menu $m \neq \{(\bar{x}_i, q_i^*)\}$. The worker rejects every contract (x, q) with $x < v(1) + \alpha$ and selects, among the remaining ones, the contract that yields the highest current payoff, provided that this is positive. If the worker accepts a contract (x, q) with $x < v(1) + \alpha$, the state changes to $(i, 2)$. In all other cases, the state remains $(i, 1)$.

State $(i, 2)$: The automaton prescribes that the firm offers the menu $\{(\underline{x}_i, q_i^*)\}$ and the worker accepts (\underline{x}_i, q_i^*) . The state remains $(i, 2)$ unless there is a deviation by the firm. In this case, the state changes to $(i, 1)$, irrespective of the worker's decision. When the firm deviates and offers a menu different from $\{(\underline{x}_i, q_i^*)\}$, the worker accepts the contract which maximizes his current payoff, provided that this is positive.

This concludes the description of the automaton. We now verify that this strategy profile is a subgame perfect equilibrium for values of δ sufficiently large. The firm's equilibrium payoffs are $V_{F,i}$ in state $(i, 0)$, $\frac{\varepsilon}{2}$ in state $(i, 1)$, and $v(q_i^*) - \underline{x}_i$ in state $(i, 2)$. The worker's equilibrium payoffs are W_i in state $(i, 0)$, $\bar{x}_i - \theta_i q_i^* - \alpha$ in state $(i, 1)$, and $\frac{\varepsilon}{2}$ in state $(i, 2)$. Thus, the firm obtains the largest payoff in state $(i, 2)$ and the lowest payoff in state $(i, 1)$, while the worker obtains the largest payoff in state $(i, 1)$ and the lowest payoff in state $(i, 2)$. In fact, notice that $V_{F,i} + W_i = \pi_i(q_i)$ and $\pi_i(q_i^*) \geq \pi_i(q_i)$ for every q_i . This, and the fact that $\varepsilon \in (0, \min\{V_{F,i}, W_i\})$, imply:

$$\begin{aligned} v(q_i^*) - \underline{x}_i &> V_{F,i} > v(q_i^*) - \bar{x}_i = \frac{\varepsilon}{2} \\ \bar{x}_i - \theta_i q_i^* - \alpha &> W_i > \underline{x}_i - \theta_i q_i^* - \alpha = \frac{\varepsilon}{2} \end{aligned}$$

Thus, for δ sufficiently large, one-shot deviations from the automaton described above are not profitable. ■

Recall that the mechanism $\{(x_H, q_H), (x_L, q_L)\}$ is incentive compatible. Thus, if the post-screening phase with type $i \in \{H, L\}$ implements the contract (x_i, q_i) in every period, then in the screening phase both types of the worker have an incentive to reveal their types. The rest of the proof of Proposition 4 provides a complete description of the equilibrium strategies and beliefs and verifies that sequential rationality is satisfied at all histories.

Consider now a different extension of the benchmark model in which the worker has the option to reject all the contracts and remain in the relationship. The result of Proposition 4 extends to this setup. To see this, notice that the strategy of rejecting all the contracts and quitting is weakly dominated by the strategy of rejecting all the contracts and remaining in the relationship. Once we remove the strategy of quitting, we are back to the infinitely repeated game analyzed above.

In summary, our analysis shows that when the relationship can continue after the worker rejects all the contracts (either because the game is infinitely repeated, or because the worker has the option to reject all the contracts and stay in the relationship), many screening opportunities are available to the firm, including separation with employment. This is

in contrast with the benchmark model in which the relationship ends automatically after the worker rejects all the contracts. In this case, only the pooling and the firing allocations can be implemented when the parties are sufficiently patient.

7 Concluding Remarks

We study a dynamic-contracting model with adverse selection and limited commitment. In our benchmark model, the relationship ends when the worker rejects all the contracts from the firm's menu. We characterize the limiting equilibrium outcome as the parties become arbitrarily patient. If the prior probability that the worker has a low cost is low, the firm offers a pooling contract in every period. In contrast, if the prior is high, the firm fires the high-cost worker at the beginning of the relationship.

We conduct our analysis under the assumption that the worker bears a cost $\alpha > 0$ when he is employed. The assumption simplifies the analysis (it is easier to link the expected length of the relationship with the expected quality of the good provided by the worker) but is not crucial. In fact, all our results extend to the case in which α is equal to zero.²¹

In this paper, the worker's action is verifiable. In some situations, the agent's effort leads to stochastic outcomes and monitoring is thus imperfect. This adds a moral hazard component to the screening problem. We leave this interesting extension for future research.

Appendix A: Proof of Proposition 1

In this appendix we prove the existence of PBE for generic values of the parameters. In particular, we construct an equilibrium in which the high type's payoff is equal to zero. Although not Markovian, in our PBE the firm and the low type's equilibrium continuation payoffs depend on the firm's belief. We use the function $V : [0, 1] \rightarrow \mathbb{R}_+$ to denote the firm's payoff (as function of its belief) and the correspondence $\Phi : [0, 1] \rightrightarrows \mathbb{R}_+$ to denote the set of payoffs of the low type. To simplify the notation, we write $\Phi(p) = z$ for $\Phi(p) = \{z\}$. We also write $\min \Phi(p)$ ($\max \Phi(p)$) to denote the smallest (largest) element of $\Phi(p)$.

For $\delta > \hat{\delta}$ we define V and Φ as follows:

$$V(p) = \max \{ \pi_H(q_H^*), p\pi_L(q_L^*) \} \tag{16}$$

²¹See the working paper at <https://www.carloalberto.org/wp-content/uploads/2015/02/no.401.pdf>

$$\Phi(p) = \begin{cases} \Delta\theta q_H^* & \text{for } p < \hat{p} \\ [0, \Delta\theta q_H^*] & \text{for } p = \hat{p} \\ 0 & \text{for } p > \hat{p} \end{cases} \quad (17)$$

where \hat{p} satisfies $\pi_H(q_H^*) = \hat{p}\pi_L(q_L^*)$.

Next, we show that for generic values of δ smaller than $\hat{\delta}$ there exists a pair (V, Φ) satisfying a number of properties.

Lemma 8 *Fix the parameters $(\theta_H, \theta_L, \alpha, v(\cdot))$. For all but at most two values of δ in $(0, \hat{\delta}]$, there exists a pair (V, Φ) satisfying the following conditions:*

- i) V is continuous and Φ is upper hemicontinuous;*
- ii) there exists $\bar{p} \in (0, 1)$ such that $V(p) = p\pi_L^*$ and $\Phi(p) = 0$ for $p > \bar{p}$;*
- iii) there exists $\underline{p} \in [0, \bar{p}]$ such that $V(p) = \pi_H^*$ for $p \leq \underline{p}$, $\Phi(p) = \Delta\theta q_H^*$ for $p < \underline{p}$, and $\Delta\theta q_H^* = \max \Phi(\underline{p})$;*
- iv) $V(p) = \tilde{V}(p)$ for $p \in [\underline{p}, \bar{p}]$, and $V(p) > \tilde{V}(p)$ for $p \in (0, \underline{p}) \cup (\bar{p}, 1)$, where $\tilde{V}(p)$ is defined by*

$$\begin{aligned} \tilde{V}(p) = & \max_{(q_H, q_L) \in [0, 1]^2, x \in \mathbb{R}, \tilde{p} \leq p} \frac{1-p}{1-\tilde{p}} [(1-\delta)\pi_H(q_H) + \delta V(\tilde{p})] + \\ & \frac{p-\tilde{p}}{1-\tilde{p}} [(1-\delta)(v(q_L) - x) + \delta\pi_L(q_L^*)] \\ \text{s.t. } & x - \theta_H q_L - \alpha \leq 0 \\ & (1-\delta)(x - \theta_L q_L - \alpha) = (1-\delta)\Delta\theta q_H + \delta \min \Phi(\tilde{p}) \end{aligned} \quad (18)$$

- v) If $\min \Phi(p) < \Delta\theta q_H^*$ and $v \in [\min \Phi(p), \Delta\theta q_H^*]$, there exists $p' \in [0, p]$ such that $v \in \Phi(p')$.*

The proof of Lemma 8 is in the Supplemental Material.

We are now ready to conclude the proof of Proposition 1. First, we consider generic values of $\delta \leq \hat{\delta}$ and use the pair $(V(\cdot), \Phi(\cdot))$ defined in Lemma 8 to construct the equilibrium strategies. We assume that $\underline{p} < \bar{p}$ (the case $\underline{p} = \bar{p} = \hat{p}$ will be discussed below together with the case $\delta > \hat{\delta}$).

For every $p \in [0, 1]$, we construct a set of menus $\mathbf{m}(p)$.

If $p < \underline{p}$, the set $\mathbf{m}(p)$ contains only the menu $m(p) = \{(\theta_H q_H^* + \alpha, q_H^*)\}$. If $p > \bar{p}$, $\mathbf{m}(p)$ contains only the menu $m(p) = \{(\theta_L q_L^* + \alpha, q_L^*)\}$.

Consider now $p \in [\underline{p}, \bar{p}]$, and let $(q_H(p), q_L(p), \tilde{p}(p))$ denote the solution to the optimization problem (18) such that

$$(1-\delta)\Delta\theta q_H(p) + \delta \min \Phi(\tilde{p}(p)) = \min \Phi(p)$$

We let $m(p)$ denote the menus containing the contracts $(x_H(p), q_H(p))$ and $(x_L(p), q_L(p))$, where the payments $x_H(p)$ and $x_L(p)$ are given by:

$$\begin{aligned} x_H(p) &= \theta_H q_H(p) + \alpha \\ x_L(p) &= \theta_L q_L(p) + \alpha + \Delta\theta q_H(p) + \frac{\delta}{1-\delta} \min \Phi(\tilde{p}(p)) \end{aligned}$$

If $\max \Phi(p) = \min \Phi(p)$, then $\mathbf{m}(p) = \{m(p)\}$. If $\max \Phi(p) > \min \Phi(p)$ and $p \in (\underline{p}, \bar{p})$, then we let $(q'_H(p), q'_L(p), \tilde{p}'(p))$ denote the solution to the optimization problem (18) such that

$$(1 - \delta) \Delta\theta q'_H(p) + \delta \min \Phi(\tilde{p}'(p)) = \max \Phi(p) \quad (19)$$

We also let $m'(p)$ denote the menus containing the contracts $(x'_H(p), q'_H(p))$ and $(x'_L(p), q'_L(p))$, where the payments $x'_H(p)$ and $x'_L(p)$ are given by:

$$\begin{aligned} x'_H(p) &= \theta_H q'_H(p) + \alpha \\ x'_L(p) &= \theta_L q'_L(p) + \alpha + \Delta\theta q'_H(p) + \frac{\delta}{1-\delta} \min \Phi(\tilde{p}'(p)) \end{aligned} \quad (20)$$

In this case, we set $\mathbf{m}(p) = \{m(p), m'(p)\}$.

If $\max \Phi(\underline{p}) > \min \Phi(\underline{p})$, then we set $m'(\underline{p}) = \{(\theta_H q_H^* + \alpha, q_H^*)\}$ and $\mathbf{m}(\underline{p}) = \{m(\underline{p}), m'(\underline{p})\}$.

Finally, we consider \bar{p} and set $m(\bar{p}) = \{(\theta_L q_L^* + \alpha, q_L^*)\}$. If $\max \Phi(\bar{p}) = \min \Phi(\bar{p})$, then $\mathbf{m}(\bar{p}) = \{m(\bar{p})\}$. Otherwise, we set $m'(\bar{p}) = \{(x'_H(\bar{p}), q'_H(\bar{p})), (x'_L(\bar{p}), q'_L(\bar{p}))\}$ (see equations (19) and (20)) and $\mathbf{m}(\bar{p}) = \{m(\bar{p}), m'(\bar{p})\}$.

The equilibrium strategies are described in terms of the state which consists of a belief $p \in [0, 1]$ and a continuation payoff $v \in \Phi(p)$. The initial state is $(p_0, \min \Phi(p_0))$, where p_0 is the prior.

Consider an arbitrary public history h^t and suppose the state is (p, v) . In equilibrium, if $v = \min \Phi(p)$, then the firm offers the menu $m(p)$. On the other hand, if $v > \min \Phi(p)$, then the firm randomizes between the two menus in $\mathbf{m}(p)$ and proposes $m(p)$ with probability β defined by

$$\beta \min \Phi(p) + (1 - \beta) \max \Phi(p) = v$$

We now turn to the worker's strategy. Consider a public history h^t in which the firm's belief $\mu(h^t)$ is equal to p . Let $m = ((x_1, q_1), \dots, (x_k, q_k))$ denote the menu offered by the firm, and for $i \in \{H, L\}$ define

$$(\bar{x}_i, \bar{q}_i) := \arg \max_{j=1, \dots, k} x_j - \theta_i q_j - \alpha$$

If $\bar{x}_L - \theta_L \bar{q}_L - \alpha < 0$, then both types reject all the contracts in the menu and quit the relationship. Furthermore, if the worker accepts a contract, then the firm's belief will be equal to one (in other words, the new state will be $(1, 0)$).

Suppose that $\bar{x}_L - \theta_L \bar{q}_L - \alpha \geq 0$ and $\bar{x}_H - \theta_H \bar{q}_H - \alpha < 0$. In this case, the low type picks the contract (\bar{x}_L, \bar{q}_L) , while the high type quits the relationship. Again, if the worker accepts a contract, the firm's belief will be equal to one.

We now turn to the case $\bar{x}_H - \theta_H \bar{q}_H - \alpha \geq 0$, and distinguish among three different possibilities. First, assume that the contracts (\bar{x}_L, \bar{q}_L) and (\bar{x}_H, \bar{q}_H) are such that

$$(1 - \delta) (\bar{x}_L - \theta_L \bar{q}_L - \alpha) \leq (1 - \delta) (\bar{x}_H - \theta_L \bar{q}_H - \alpha) + \delta \min \Phi(p)$$

In this case, both types accept the contract (\bar{x}_H, \bar{q}_H) , and the firm's belief remains unchanged. If the worker accepts any other contract, the firm's belief will jump to one.

Second, if

$$(1 - \delta) (\bar{x}_L - \theta_L \bar{q}_L - \alpha) \geq (1 - \delta) (\bar{x}_H - \theta_L \bar{q}_H - \alpha) + \delta \Delta \theta q_H^*$$

then type $i \in \{H, L\}$ accepts the contract (\bar{x}_i, \bar{q}_i) . The firm's belief will become zero if the worker accepts the contract (\bar{x}_H, \bar{q}_H) , and one if the worker accepts any other contract.

Finally, assume that

$$\frac{(1 - \delta) [(\bar{x}_L - \theta_L \bar{q}_L) - (\bar{x}_H - \theta_L \bar{q}_H)]}{\delta} \in (\min \Phi(p), \Delta \theta q_H^*)$$

and thus

$$(1 - \delta) (\bar{x}_L - \theta_L \bar{q}_L - \alpha) = (1 - \delta) (\bar{x}_H - \theta_L \bar{q}_H - \alpha) + \delta \left[\tilde{\beta} \min \Phi(p') + (1 - \tilde{\beta}) \max \Phi(p') \right]$$

for some $p' \leq p$ and some $\tilde{\beta} \in [0, 1]$. In this case, the high type accepts the contract (\bar{x}_H, \bar{q}_H) , while the low type chooses the contract (\bar{x}_H, \bar{q}_H) with probability $\frac{p'}{1-p'} \frac{1-p}{p}$, and the contract (\bar{x}_L, \bar{q}_L) with probability $1 - \frac{p'}{1-p'} \frac{1-p}{p}$. Following the acceptance of the contract (\bar{x}_H, \bar{q}_H) the new state will be $(p', \tilde{\beta} \min \Phi(p') + (1 - \tilde{\beta}) \max \Phi(p'))$. If the worker accepts the contract (\bar{x}_L, \bar{q}_L) or any other contract, the firm's belief will be equal to one.

The above strategy profile, together with the firm's belief, constitute a PBE. The sequential rationality of the firm's strategy follows from the construction of the pair (V, Φ) . In equilibrium, the high type behaves myopically and maximizes his period- t payoff at any

history h^t . This behavior is indeed optimal since the high type's continuation payoff (computed at the beginning of period $t + 1$) is equal to zero after any public history. Finally, notice that when the low type randomizes, all the contracts in the strategy's support yield the same expected payoff (and this is greater than the payoff of any other contract).

We now briefly turn to the case $\delta > \hat{\delta}$ and the case $\underline{p} = \bar{p} = \hat{p}$ (when $\delta \leq \hat{\delta}$). Recall the definitions of V and Φ in equations (16) and (17), respectively. For every belief p , we define the set of menus $\mathbf{m}(p)$ as follows. If $p < \hat{p}$, the set $\mathbf{m}(p)$ contains only the menu $m(p) = \{(\theta_H q_H^* + \alpha, q_H^*)\}$. If $p > \hat{p}$, $\mathbf{m}(p)$ contains only the menu $m(p) = \{(\theta_L q_L^* + \alpha, q_L^*)\}$. Finally, the set $\mathbf{m}(\hat{p})$ contains both the menu $m(\hat{p}) = \{(\theta_L q_L^* + \alpha, q_L^*)\}$ and the menu $m'(\hat{p}) = \{(\theta_H q_H^* + \alpha, q_H^*)\}$. The equilibrium strategies and beliefs are defined similarly to the case $\underline{p} < \bar{p}$ above and we omit the details. ■

Appendix B

Proof of Lemma 4.

Fix a PBE (σ, μ) and a history h^t ($\mu(h^t) < p$) at which the firm offers a menu m_t with the properties described in the statement of the lemma. First, notice that if the high type rejects all the contracts in m_t with probability one (i.e., the high type quits), then the high type's length of the relationship, the firm's payoff (conditional on the high type), and the low type's payoff are all equal to zero (if the low type's payoff is strictly positive, the firm's payoff would fall below $\mu(h^t) \pi_L(q_L^*)$).

Consider now the case in which the high type accepts a contract in m_t , say (x_t^H, q_t^H) , with positive probability. We let h_H^{t+1} denote the history $(h^t, m_t, (x_t^H, q_t^H))$. We also let h_L^{t+1} denote the history $(h^t, m_t, (x_t^L, q_t^L))$.

The fact that type $i \in \{H, L\}$ accepts with positive probability the contract (x_t^i, q_t^i) implies that

$$(1 - \delta) (x_t^H - \theta_H q_t^H - \alpha) + \delta W_H(h_H^{t+1}) \geq (1 - \delta) (x_t^L - \theta_H q_t^L - \alpha) + \delta W_H(h_L^{t+1})$$

and

$$(1 - \delta) (x_t^L - \theta_L q_t^L - \alpha) + \delta W_L(h_L^{t+1}) \geq (1 - \delta) (x_t^H - \theta_L q_t^H - \alpha) + \delta W_L(h_H^{t+1})$$

We add the two incentive compatibility constraints and obtain

$$(1 - \delta) \Delta \theta (q_t^L - q_t^H) + \delta (W_L(h_L^{t+1}) - W_H(h_L^{t+1})) \geq \delta (W_L(h_H^{t+1}) - W_H(h_H^{t+1}))$$

Recall that $\mu(h_L^{t+1}) \geq p$ and that $[p, 1]$ is a firing region. Therefore, there exist \bar{K} and $\bar{\delta} < 1$ such that $W_L(h_L^{t+1}) \leq \bar{K}(1 - \delta)$ for $\delta > \bar{\delta}$. Of course, $W_H(h_L^{t+1}) \geq 0$. This, together with the above inequality, implies:

$$(1 - \delta) \Delta\theta (q_t^L - q_t^H) + \bar{K}(1 - \delta) \geq \delta (W_L(h_H^{t+1}) - W_H(h_H^{t+1})) \quad (21)$$

We now let $\mathcal{D}_H(h_H^{t+1}) := \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t+1}^{\mathbb{T}-1} \delta^{\tau-t-1} |h_H^{t+1}, H \right]$ denote the expected discounted time, computed at h_H^{t+1} , until the high type quits. Our next goal is to provide an upper bound to $\mathcal{D}_H(h_H^{t+1})$. Thus, without loss, assume that $\mathcal{D}_H(h_H^{t+1})$ is strictly positive. We let

$$\mathcal{Q}_{t+1} = \frac{\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t+1}^{\mathbb{T}-1} \delta^{\tau-t-1} q_\tau |h_H^{t+1}, H \right]}{\mathcal{D}_H(h_H^{t+1})}$$

the expected discounted total quality provided by the high type at the history h_H^{t+1} .

Using Jensen's inequality (recall that the function $\pi(\cdot)$ is concave), we can bound the firm's continuation payoff (conditional on type H) as follows:

$$\begin{aligned} V_F(h_H^{t+1}; H) &\leq \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t+1}^{\mathbb{T}-1} \delta^{\tau-t-1} \pi(q_\tau) |h_H^{t+1}, H \right] \leq \\ &\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t+1}^{\mathbb{T}-1} \delta^{\tau-t-1} \pi(\mathcal{Q}_{t+1}) |h_H^{t+1}, H \right] = \mathcal{D}_H(h_H^{t+1}) \pi(\mathcal{Q}_{t+1}) \end{aligned}$$

Let $\check{q}_H \in (0, q_H^*)$ be such that $\pi_H(\check{q}_H) = 0$ and notice that $\pi_H(q) < 0$ for every $q < \check{q}_H$. This implies that $\mathcal{Q}_{t+1} \geq \check{q}_H$. In fact, if the last inequality is violated, then $V_F(h_H^{t+1}; H)$ is strictly negative, and $V_F(h_H^{t+1})$ is strictly less than $\mu(h_H^{t+1}) \pi_L(q_L^*)$, contradicting Lemma 1. Notice that one strategy available to the low type is to imitate the high type's behavior (in every period). Therefore, we conclude that

$$W_L(h_H^{t+1}) - W_H(h_H^{t+1}) \geq \Delta\theta \mathcal{D}_H(h_H^{t+1}) \mathcal{Q}_{t+1} \geq \Delta\theta \mathcal{D}_H(h_H^{t+1}) \check{q}_H$$

Combining the inequality above with inequality (21) we obtain

$$\mathcal{D}_H(h_H^{t+1}) \leq \frac{(1 - \delta)(q_t^L - q_t^H)}{\delta \check{q}_H} + \frac{\bar{K}(1 - \delta)}{\delta \Delta\theta \check{q}_H} \leq \frac{(1 - \delta)}{\delta \check{q}_H} \left(1 + \frac{\bar{K}}{\Delta\theta} \right)$$

Hence, for $\delta > \frac{1}{2}$ we have

$$\mathcal{D}_H(h_H^{t+1}) \leq \frac{2(1 - \delta)}{\check{q}_H} \left(1 + \frac{\bar{K}}{\Delta\theta} \right)$$

This, in turn, implies that (for $\delta > \frac{1}{2}$)

$$\mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{\tau=t}^{T-1} \delta^{\tau-t} |h^t, m_t, H \right] \leq (1 - \delta) + \mathcal{D}_H(h_H^{t+1}) \leq \left(1 + \frac{2}{\bar{q}_H} + \frac{2\tilde{K}}{\Delta\theta\bar{q}_H}\right) (1 - \delta) := \tilde{K} (1 - \delta)$$

and establishes part i).

To verify property ii), notice that the inequality above implies that the firm's continuation payoff $V_F(h^t, m_t; (\sigma, \mu), H)$ is bounded above by $v(1) \tilde{K} (1 - \delta)$.

Finally, we turn to property iii). The analysis above implies that

$$V_F(h^t, m_t) \leq \mu(h^t) [\pi_L(q_L^*) - W_L(h^t, m_t)] + v(1) \tilde{K} (1 - \delta) \quad (22)$$

Let δ' be such that

$$v(1) \tilde{K} (1 - \delta') = \frac{\pi_H(q_H^*)}{4}$$

and notice that for $\delta \geq \tilde{\delta} = \max\{\delta', \bar{\delta}, \frac{1}{2}\}$ and $\mu(h^t) \leq \frac{\pi_H(q_H^*)}{2\pi_L(q_L^*)}$

$$\mu(h^t) \pi_L(q_L^*) + v(1) \tilde{K} (1 - \delta) \leq \frac{3}{4} \pi_H(q_H^*)$$

It follows that if the firm offers the menu m_t at the history h^t and $\delta \geq \tilde{\delta}$, then $\mu(h^t) > \frac{\pi_H(q_H^*)}{2\pi_L(q_L^*)}$. Finally, recall that $V_F(h^t, m_t)$ is bounded below by $\mu(h^t) \pi_L(q_L^*)$. This and inequality (22) imply

$$W_L(h^t, m_t) \leq \frac{v(1)}{\mu(h^t)} \tilde{K} (1 - \delta) \leq \frac{2\pi_L(q_L^*) v(1) \tilde{K}}{\pi_H(q_H^*)} (1 - \delta)$$

This shows that there exists $K > 0$ satisfying the three properties in Lemma 4. ■

Derivation of the inequalities in (3).

First, notice that $W_H(h^0) = 0$ implies that the menu m_0 offered at h^0 yields a continuation payoff of zero to the high type. Let $m_0 = ((x_0^1, q_0^1), \dots, (x_0^k, q_0^k))$ be the menu offered at h^0 . For $i \in \{H, L\}$, let m_0^i denote the set of contracts in m_0 accepted with positive probability by the type i . If there exists a contract (x_0^j, q_0^j) in $m_0^L \setminus m_0^H$, then $\tilde{\mathbb{T}}$ coincides with 0 (in fact, the firm's belief jumps to one if the worker accepts the contract (x_0^j, q_0^j)) and there is nothing to prove. Therefore, assume that $m_0^L \subseteq m_0^H$ and recall that m_0^L is non-empty.

For every contract (x_0^j, q_0^j) in m_0^H , we let $h_j^1 = (h^0, m_0, (x_0^j, q_0^j))$ denote the history in which the worker accepts the contract (x_0^j, q_0^j) in period 0, and recall that $W_H(h_j^1) \geq 0$ represents the high-worker's payoff at h_j^1 . For every contract $(x_0^j, q_0^j) \in m_0^H$ we replace the payment x_0^j with the payment

$$\tilde{x}_0^j = x_0^j + \frac{\delta}{1-\delta} W_H(h_j^1)$$

which, clearly, implies $\tilde{x}_0^j = \theta_H q_0^j + \alpha$.

To keep the parties' payoffs unchanged, we also modify the payments in period 1. In particular, for every $(x_0^j, q_0^j) \in m_0^H$ consider the history h_j^1 . Let m_1 denote a menu offered at h_j^1 with positive probability. We subtract $\frac{1}{1-\delta} W_H(h_j^1, m_1)$, the high type's continuation payoff at the moment that the menu m_1 is offered, from the payment of every contract in m_1 . Notice that this yields to the high type a continuation payoff (evaluated at the beginning of period 1) equal to zero.

We recursively apply the procedure outlined above to periods $t = 1, \dots, \tilde{T}$ (i.e., we increase the payments in period t and, at the same time, decrease the payments in period $t + 1$). By construction, every contract (x_t, q_t) accepted with positive probability by the high type in period $t = 0, \dots, \tilde{T} - 1$ is replaced with the contract $(\theta_H q_t + \alpha, q_t)$, while the payments in every menu offered in period \tilde{T} are uniformly decreased by the high type's continuation payoff. Finally, in every period $t = 0, \dots, \tilde{T} - 1$, the set of contracts accepted by the low type is contained in the set of contracts accepted by the high type (this follows from the definition of \tilde{T}).

The change in the timing of the transfers yields:

$$\begin{aligned} V_F(h^0) &= \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \pi_H(q_t) + \mathbb{I}_{\{\tilde{T} < \infty\}} \delta^{\tilde{T}} \left(V_F(h^{\tilde{T}}, m_{\tilde{T}}) + W_H(h^{\tilde{T}}, m_{\tilde{T}}) \right) \mid h^0 \right] \\ W_L(h^0) &= \mathbb{E}_{(\sigma, \mu)} \left[(1 - \delta) \sum_{t=0}^{\tilde{T}-1} \delta^t \Delta \theta q_t + \mathbb{I}_{\{\tilde{T} < \infty\}} \delta^{\tilde{T}} \left(W_L(h^{\tilde{T}}, m_{\tilde{T}}) - W_H(h^{\tilde{T}}, m_{\tilde{T}}) \right) \mid h^0, L \right] \\ W_{LH}(h^0) &\geq \mathbb{E}_{(\sigma, \mu)} \left[\sum_{t=0}^{\tilde{T}-1} \delta^t \Delta \theta q_t - \mathbb{I}_{\{\tilde{T} < \infty\}} \delta^{\tilde{T}} W_H(h^{\tilde{T}}, m_{\tilde{T}}) \mid h^0, H \right] \end{aligned}$$

The payoff $W_H(h^{\tilde{T}}, m_{\tilde{T}})$ is bounded above by $W_L(h^{\tilde{T}}, m_{\tilde{T}})$. Therefore, it follows from Lemma 4 and from the fact that $[p, 1]$ is a firing region that there exist K and $\bar{\delta} < 1$ such

that for any $\delta > \bar{\delta}$ and for any PBE (σ, μ) the payoffs $V_F(h^0)$, $W_L(h^0)$, and $W_{LH}(h^0)$ satisfy the inequalities in (3).

Proof of Claim 1.

First, assume that $\tilde{q}_0 \leq \frac{\check{q}_H}{2}$ and notice that $\pi_H(\tilde{q}_0) < \pi_H(\frac{\check{q}_H}{2}) < 0$ (recall that $\check{q}_H \in (0, q_H^*)$ satisfies $\pi_H(\check{q}_H) = 0$). Also, notice that $\hat{p} < p_0 < p$. If inequality (8) is satisfied, then we have

$$\begin{aligned} 0 \leq & \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H(\tilde{q}_0) + \frac{p_0}{p} [\Upsilon_p \pi_H(\tilde{q}_p) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K(1 - \delta) - p_0 \pi_L(q_L^*) \leq \\ & \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H\left(\frac{\check{q}_H}{2}\right) + \frac{p_0}{p} \Upsilon_p [\pi_H(q_H^*) - p \pi_L(q_L^*)] + K(1 - \delta) \leq \\ & \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H\left(\frac{\check{q}_H}{2}\right) + K(1 - \delta) \end{aligned}$$

Putting together this and $p_0 < p^C$ we obtain

$$\left(\frac{1}{1 - p_0}\right) \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \leq -\frac{K(1 - \delta)}{(1 - p^C) \pi_H\left(\frac{\check{q}_H}{2}\right)} \quad (23)$$

Similarly, we obtain

$$\begin{aligned} 0 \leq & \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H(\tilde{q}_0) + \frac{p_0}{p} [\Upsilon_p \pi_H(\tilde{q}_p) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K(1 - \delta) - p_0 \pi_L(q_L^*) \leq \\ & \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H(\tilde{q}_0) + \frac{p_0}{p} \Upsilon_p [\pi_H(q_H^*) - p \pi_L(q_L^*)] + K(1 - \delta) \leq \\ & \frac{p_0}{p} \Upsilon_p [\pi_H(q_H^*) - p \pi_L(q_L^*)] + K(1 - \delta) \leq \\ & \left(1 - \frac{f(p)}{2p}\right) \Upsilon_p [\pi_H(q_H^*) - p \pi_L(q_L^*)] + K(1 - \delta) \end{aligned}$$

where the last inequality follows from $p_0 \in \left[p - \frac{f(p)}{2}, p\right]$ and $\pi_H(q_H^*) - p \pi_L(q_L^*) < 0$.

Hence, we have:

$$\left(\frac{1 - p}{1 - p_0}\right) \left(\frac{p_0}{p}\right) \Upsilon_p \leq \Upsilon_p \leq \frac{K(1 - \delta)}{\left(1 - \frac{f(p)}{2p}\right) [p \pi_L(q_L^*) - \pi_H(q_H^*)]} \quad (24)$$

For the case $\tilde{q}_0 \leq \frac{\check{q}_H}{2}$, inequalities (23) and (24) imply the result.

We now move to the case $\tilde{q}_0 > \frac{\check{q}_H}{2}$. It follows from the concavity of $\pi_H(\cdot)$ that $\pi_H(\tilde{q}_0) \leq \pi_H(0) + \pi_H'(0) \tilde{q}_0 \leq \pi_H'(0) \tilde{q}_0$. Also, $\pi_H(\tilde{q}_p) < \pi_H(q_H^*)$ for any $\tilde{q}_p \neq q_H^*$. Thus, we have

$$\begin{aligned} & \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H(\tilde{q}_0) + \frac{p_0}{p} [\Upsilon_p \pi_H(\tilde{q}_p) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K(1 - \delta) \leq \\ & \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H'(0) \tilde{q}_0 + \frac{p_0}{p} [\Upsilon_p \pi_H(q_H^*) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K(1 - \delta) \end{aligned} \quad (25)$$

Suppose that inequality (9) holds. Clearly, the inequality continues to hold if we replace \tilde{q}_p with one. This allows us to conclude that \tilde{q}_0 , Υ_0 , and Υ_p must satisfy

$$\Upsilon_0 \tilde{q}_0 \leq \Upsilon_p + \frac{2K(1-\delta)}{\Delta\theta \left(\frac{1}{1-p_0}\right) \left(1 - \frac{p_0}{p}\right)} \quad (26)$$

Combining inequalities (25) and (26) we obtain (recall that $p_0 > \hat{p}$):

$$\begin{aligned} & \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \pi_H(\tilde{q}_0) + \frac{p_0}{p} [\Upsilon_p \pi_H(\tilde{q}_p) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K(1 - \delta) \leq \\ & \left(1 - \frac{p_0}{p}\right) \pi'_H(0) \Upsilon_p + \frac{p_0}{p} [\Upsilon_p \pi_H(q_H^*) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K \left(1 + \frac{2\pi'_H(0)(1-\hat{p})}{\Delta\theta}\right) (1 - \delta) \end{aligned}$$

We define $K_1 := \left(1 + \frac{2\pi'_H(0)(1-\hat{p})}{\Delta\theta}\right)$. It follows from inequality (8) and the inequality above that

$$p_0 \pi_L(q_L^*) \leq \left(1 - \frac{p_0}{p}\right) \pi'_H(0) \Upsilon_p + \frac{p_0}{p} [\Upsilon_p \pi_H(q_H^*) + (1 - \Upsilon_p) p \pi_L(q_L^*)] + K_1(1 - \delta)$$

which leads to:

$$\begin{aligned} 0 & \leq \Upsilon_p \left[\left(1 - \frac{p_0}{p}\right) \pi'_H(0) + \frac{p_0}{p} [\pi_H(q_H^*) - p \pi_L(q_L^*)] \right] + K_1(1 - \delta) \leq \\ & \Upsilon_p \left[\left(1 - \frac{p - \frac{f(p)}{2}}{p}\right) \pi'_H(0) + \frac{p - \frac{f(p)}{2}}{p} [\pi_H(q_H^*) - p \pi_L(q_L^*)] \right] + K_1(1 - \delta) \end{aligned} \quad (27)$$

The second inequality holds because the expression $\left(1 - \frac{p_0}{p}\right) \pi'_H(0) + \frac{p_0}{p} [\pi_H(q_H^*) - p \pi_L(q_L^*)]$ is affine in p_0 and is negative for $p_0 = p$ and equal to zero for $p_0 = p - f(p)$ (see the definition of the function $f(\cdot)$ in equation (2)). Also, recall that $p_0 \in \left[p - \frac{f(p)}{2}, p\right]$.

From inequality (27) we obtain:

$$\Upsilon_p \leq \frac{K_1}{\left(p - \frac{f(p)}{2}\right) \pi_L(q_L^*) - \left(\frac{f(p)}{2p}\right) \pi'_H(0) - \left(1 - \frac{f(p)}{2p}\right) \pi_H(q_H^*)} (1 - \delta) := K_2(1 - \delta)$$

and, thus,

$$\left(\frac{1-p}{1-p_0}\right) \left(\frac{p_0}{p}\right) \Upsilon_p \leq \Upsilon_p \leq K_2(1 - \delta) \quad (28)$$

Finally, using (26) and (28) and $p_0 < p^C$ we have:

$$\begin{aligned} \left(\frac{1}{1-p_0}\right) \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \tilde{q}_0 & \leq \left(\frac{1}{1-p^C}\right) \left(1 - \frac{p_0}{p}\right) \Upsilon_p + \frac{2K(1-\delta)}{\Delta\theta} \leq \\ & \left(\frac{1}{1-p^C}\right) \Upsilon_p + \frac{2K(1-\delta)}{\Delta\theta} \leq \left(\frac{1}{1-p^C}\right) K_2(1 - \delta) + \frac{2K(1-\delta)}{\Delta\theta} \end{aligned}$$

Recall that $\tilde{q}_0 > \frac{\check{q}_H}{2}$. It follows from the last inequality that

$$\left(\frac{1}{1-p_0}\right) \left(1 - \frac{p_0}{p}\right) \Upsilon_0 \leq \left(\frac{2}{\check{q}_H}\right) \left[\frac{K_2}{1-p^C} + \frac{2K}{\Delta\theta}\right] (1-\delta)$$

which coupled with (28) implies the result. ■

Proof of Lemma 6.

Fix a PBE (σ, μ) and a history h^t as described in the statement of the lemma. The firm's continuation payoff after offering the menu m_t is equal to

$$V_F(h^t, m_t) = (1 - \mu(h^t)) V_F(h^t, m_t; H) + \mu(h^t) V_F(h^t, m_t; L)$$

Recall from Proposition 2 that $[\hat{p} + \varepsilon, 1]$ is a firing region. Therefore, it follows from Lemma 4 that there exist \bar{K} and $\bar{\delta} > 1 - \frac{\varepsilon\pi_L(q_L^*)}{2\bar{K}}$ such that $\delta > \bar{\delta}$ implies $V_F(h^t, m_t; H) \leq \bar{K}(1-\delta)$. This, in turn, implies

$$V_F(h^t, m_t) \leq (\hat{p} - \varepsilon) \pi_L(q_L^*) + \bar{K}(1-\delta) < \left(\hat{p} - \frac{\varepsilon}{2}\right) \pi_L(q_L^*) < \pi_H(q_H^*)$$

contradicting Lemma 1. ■

Proof of Proposition 3.

We start by proving the first property of a pooling region. Without loss of generality, we establish the property at the initial history h^0 and restrict attention to equilibria in which (i) the firm's strategy in the first period is pure; and (ii) the high type's equilibrium payoff is equal to zero.

Fix $p < \hat{p}$. By contradiction, suppose that there exists a sequence $\{\delta_n, p_{0,n}, (\sigma_n, \mu_n)\}_{n=1}^\infty$ such that δ_n converges to one, $p_{0,n} \in [0, p]$, (σ_n, μ_n) is a PBE of the game with discount factor equal to δ_n and prior equal to $p_{0,n}$ and

$$\lim_{n \rightarrow \infty} \mathbb{E}_{(\sigma_n, \mu_n)} [\delta_n^{\mathbb{T}}] = \xi > 0 \tag{29}$$

Let $\underline{k} := \left\lfloor \frac{2}{1-\hat{p}} \right\rfloor$, and for $k = \underline{k}, \underline{k} + 1, \dots$, let $\tilde{\mathbb{T}}_k \leq \infty$ be the random time that stops the play at the first history $(h^{\tilde{\mathbb{T}}}, m_{\tilde{\mathbb{T}}})$ at which the menu $m_{\tilde{\mathbb{T}}}$ contains a contract $(x_{\tilde{\mathbb{T}}}, q_{\tilde{\mathbb{T}}})$ accepted with positive probability and for which $\mu(h^{\tilde{\mathbb{T}}}, m_{\tilde{\mathbb{T}}}, (x_{\tilde{\mathbb{T}}}, q_{\tilde{\mathbb{T}}})) \geq \hat{p} + \frac{1}{k}$. As usual, we set $\tilde{\mathbb{T}}_k = \infty$ if the event does not occur in finite time.

It follows from Lemma 4 that for every $k \geq \underline{k}$ there exist $n_k^1 \in \mathbf{N}$ and K_k^1 such that for every $n \geq n_k^1$ the PBE (σ_n, μ_n) satisfies the following property. If the firm offers a menu with a contract that is accepted with positive probability and leads to a belief weakly higher than $\hat{p} + \frac{1}{k}$, then the expected discounted time until the high type quits the relationship is bounded above by $K_k^1 (1 - \delta_n)$. Thus, for $n \geq n_k^1$ we have:

$$\left| \mathbb{E}_{(\sigma_n, \mu_n)} [\delta_n^{\tilde{\mathbb{T}}}] - \mathbb{E}_{(\sigma_n, \mu_n)} \left[\mathbb{I}_{\{\tilde{\mathbb{T}}_k < \infty\}} \left(1 - \mu_n \left(h^{\tilde{\mathbb{T}}_k} \right) \right) \delta_n^{\tilde{\mathbb{T}}_k} \right] \right| \leq K_k^1 (1 - \delta_n) \quad (30)$$

Next, recall that $[\hat{p} + \frac{1}{k}]$ is a firing region and Lemma 4 (property ii) provides an upper bound to the firm's continuation payoff when it offers a menu with a contract that leads to a firing region. Therefore, for every $k \geq \underline{k}$ there exist $n_k^2 \in \mathbf{N}$ and K_k^2 such that for every $n \geq n_k^2$ the firm's equilibrium payoff is bounded as follows:²²

$$V_F(h^0; (\sigma_n, \mu_n)) \leq \mathbb{E}_{(\sigma_n, \mu_n)} \left[\mathbb{I}_{\{\tilde{\mathbb{T}}_k < \infty\}} \left[(1 - \delta_n) \sum_{t=0}^{\tilde{\mathbb{T}}_k - 1} \delta_n^t \pi_H(q_t) + \delta_n^{\tilde{\mathbb{T}}_k} \mu_n \left(h^{\tilde{\mathbb{T}}_k} \right) \pi_L(q_L^*) \right] + \mathbb{I}_{\{\tilde{\mathbb{T}}_k = \infty\}} (1 - \delta_n) \sum_{t=0}^{\tilde{\mathbb{T}}_k - 1} \delta_n^t \pi_H(q_t) \right] + K_k^2 (1 - \delta_n)$$

Notice that when $\tilde{\mathbb{T}}_k < \infty$, the belief $\mu_n \left(h^{\tilde{\mathbb{T}}_k} \right)$ is, by definition, lower than $\hat{p} + \frac{1}{k}$, and, therefore, we have:

$$\mu_n \left(h^{\tilde{\mathbb{T}}_k} \right) \pi_L(q_L^*) < \left(\hat{p} + \frac{1}{k} \right) \pi_L(q_L^*) = \pi_H(q_H^*) + \frac{1}{k} \pi_L(q_L^*)$$

Combining the last two inequalities, for every $n \geq n_k^2$ we obtain:

$$\begin{aligned} V_F(h^0; (\sigma_n, \mu_n)) &\leq \mathbb{E}_{(\sigma_n, \mu_n)} \left[\mathbb{I}_{\{\tilde{\mathbb{T}}_k < \infty\}} \left[(1 - \delta_n) \sum_{t=0}^{\tilde{\mathbb{T}}_k - 1} \delta_n^t \pi_H(q_t) + \delta_n^{\tilde{\mathbb{T}}_k} \pi_H(q_H^*) \right] + \right. \\ &\quad \left. \mathbb{I}_{\{\tilde{\mathbb{T}}_k = \infty\}} (1 - \delta_n) \sum_{t=0}^{\tilde{\mathbb{T}}_k - 1} \delta_n^t \pi_H(q_t) \right] + K_k^2 (1 - \delta_n) + \frac{1}{k} \pi_L(q_L^*) = \\ &\pi_H(q_H^*) - \mathbb{E}_{(\sigma_n, \mu_n)} \left[(1 - \delta_n) \sum_{t=0}^{\tilde{\mathbb{T}}_k - 1} \delta_n^t [\pi_H(q_H^*) - \pi_H(q_t)] \right] + K_k^2 (1 - \delta_n) + \frac{1}{k} \pi_L(q_L^*) \end{aligned}$$

²²This bound is derived using the same procedure as the one outlined above in the derivation of the inequalities in (3).

This and the fact that the firm's payoff is bounded below by $\pi_H(q_H^*)$ lead to the following result. For every $k \geq \underline{k}$,

$$\limsup_{n \rightarrow \infty} \mathbb{E}_{(\sigma_n, \mu_n)} \left[(1 - \delta_n) \sum_{t=0}^{\hat{\mathbb{T}}_k - 1} \delta_n^t [\pi_H(q_H^*) - \pi_H(q_t)] \right] \leq \frac{1}{k} \pi_L(q_L^*) \quad (31)$$

Inequality (31) implies that for every $\eta > 0$ there exists $k_\eta \in \mathbf{N}$ such that for every $k \geq k_\eta$ there exists $\hat{n}_k \in \mathbf{N}$ such that for $n \geq \hat{n}_k$ we have

$$\mathbb{E}_{(\sigma_n, \mu_n)} \left[(1 - \delta_n) \sum_{t=0}^{\hat{\mathbb{T}}_k - 1} \delta_n^t |q_H^* - q_t| \right] \leq \eta \quad (32)$$

Furthermore, k_η and \hat{n}_k are such that for every $k \geq k_\eta$ and every $n \geq \hat{n}_k$

$$\frac{\xi}{1 - \hat{p}} - \eta \leq \mathbb{E}_{(\sigma_n, \mu_n)} \left[\mathbb{I}_{\{\hat{\mathbb{T}}_k < \infty\}} \delta_n^{\hat{\mathbb{T}}_k} \right] \leq \frac{\xi}{1 - \hat{p}} + \eta$$

The above result is a consequence of equality (29), inequality (30), and Lemma 6.

Fix $\varepsilon \in \left(0, \frac{\Delta\theta q_H^* \xi}{4(1-\hat{p})(1+\Delta\theta q_H^*)} \left(\frac{\hat{p}}{p} - \frac{1-\hat{p}}{1-p}\right)\right)$. Recall that for every k , the interval $[\hat{p} + \frac{1}{k}, 1]$ is a firing region and Lemma 4 (property iii) provides an upper bound to the low type's continuation payoff when the firm's menu contains a contract that leads to a firing region. Finally, recall that if a history h^t is reached with probability $\Pr(h^t)$ under (σ_n, μ_n) , then that history is reached with probability $\frac{\mu_n(h^t)}{p_{0,n}} \Pr(h^t)$ if the worker behaves according to σ_n^L , and with probability $\frac{(1-\mu_n(h^t))}{(1-p_{0,n})} \Pr(h^t)$ if the worker behaves according to σ_n^H .

Putting together these observations and the last three inequalities we conclude that there exist \tilde{k} and \tilde{n} such that for every $n \geq \tilde{n}$ the low type obtains a payoff of at most

$$\Delta\theta q_H^* \left(1 - \mathbb{E}_{(\sigma_n, \mu_n)} \left[\delta_n^{\hat{\mathbb{T}}_k} | \sigma_n^L \right] \right) + \varepsilon \leq \Delta\theta q_H^* \left(1 - \frac{\hat{p}}{p_{0,n}} \frac{\xi}{1 - \hat{p}} + \varepsilon\right) + \varepsilon \leq \Delta\theta q_H^* \left(1 - \frac{\hat{p}}{p} \frac{\xi}{1 - \hat{p}} + \varepsilon\right) + \varepsilon$$

when he behaves according to σ_n^L , and a payoff of at least

$$\begin{aligned} \Delta\theta q_H^* \left(1 - \mathbb{E}_{(\sigma_n, \mu_n)} \left[\delta_n^{\hat{\mathbb{T}}_k} | \sigma_n^H \right] \right) - \varepsilon &\geq \Delta\theta q_H^* \left(1 - \frac{1-\hat{p}}{1-p_{0,n}} \frac{\xi}{1-\hat{p}} - \varepsilon\right) - \varepsilon \geq \\ &\Delta\theta q_H^* \left(1 - \frac{1-\hat{p}}{1-p} \frac{\xi}{1-\hat{p}} - \varepsilon\right) - \varepsilon \end{aligned}$$

when he behaves according to σ_n^H . Notice that

$$\begin{aligned} \Delta\theta q_H^* \left(1 - \frac{1-\hat{p}}{1-p} \frac{\xi}{1-\hat{p}} - \varepsilon\right) - \varepsilon - \Delta\theta q_H^* \left(1 - \frac{\hat{p}}{p} \frac{\xi}{1-\hat{p}} + \varepsilon\right) - \varepsilon &= \\ \Delta\theta q_H^* \frac{\xi}{1-\hat{p}} \left(\frac{\hat{p}}{p} - \frac{1-\hat{p}}{1-p}\right) - 2\varepsilon(1 + \Delta\theta q_H^*) &> 0 \end{aligned}$$

which implies that for n sufficiently large the low type has an incentive to deviate and follow σ_n^H instead of the equilibrium strategy σ_n^L .

We now turn to the second property of a pooling region. It is again without loss of generality to restrict attention to equilibria in which (i) the firm's strategy in the first period is pure; and (ii) the high type's equilibrium payoff is equal to zero. The second property follows directly from the first property and inequality (32).

Finally, we establish the third property. Assume, towards a contradiction, that there exists a sequence $\{\delta_n, p_{0,n}, (\sigma_n, \mu_n)\}_{n=1}^\infty$ such that δ_n converges to one, $p_{0,n} \in [0, p]$, (σ_n, μ_n) is an arbitrary PBE of the game with discount factor equal to δ_n and prior equal to $p_{0,n}$

$$\lim_{n \rightarrow \infty} \mathbb{E}_{(\sigma_n, \mu_n)} \left[(1 - \delta_n) \sum_{t=0}^{T-1} \delta_n^t (x_t - \theta_H q_H^* - \alpha) | i \right] = \tilde{\xi} > 0$$

for some $i \in \{H, L\}$. Using the first two properties of a pooling region we conclude that

$$\limsup_{n \rightarrow \infty} V_F(h^0; (\sigma_n, \mu_n)) \leq \pi_H (q_H^*) - \min\{p_0, 1 - p_0\} \tilde{\xi} < \pi_H (q_H^*)$$

which leads to a contradiction and concludes the proof. ■

References

- [1] Acemoglu, D., M. Golosov, and A. Tsyvinski (2010): "Dynamic Mirrlees Taxation under Political Economy Constraints," *Review of Economic Studies*, 77, 841-881.
- [2] Aumann, R. J., and M. B. Maschler (1995): *Repeated Games with Incomplete Information*, The MIT Press.
- [3] Ausubel, L. M., and R. J. Deneckere (1989): "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica*, 57, 511-531.
- [4] Beccuti, J., and M. Möller (2018): "Dynamic Adverse Selection with a Patient Seller," *Journal of Economic Theory*, 173, 95-117.
- [5] Bhaskar, V. (2014): "The Ratchet Effect Re-examined: A Learning Perspective," mimeo.
- [6] Bhaskar, V. and G. Mailath (2019): "The Curse of Long Horizons," *Journal of Mathematical Economics*, 82, 74-89.

- [7] Carmichael, L., and B. MacLeod (2000): “Worker Cooperation and the Ratchet Effect,” *Journal of Labor Economics*, 18, 1-19.
- [8] Farhi, E., C. Sleeter, I. Werning, and S. Yeltekin (2012): “Non-linear Capital Taxation Without Commitment,” *Review of Economic Studies*, 79, 1469-1493.
- [9] Freixas, X., R. Guesnerie, and J. Tirole (1985): “Planning under Incomplete Information and the Ratchet Effect,” *The Review of Economic Studies*, 52, 173-191.
- [10] Fudenberg, D., and J. Tirole (1991): “Perfect Bayesian Equilibrium and Sequential Equilibrium,” *Journal of Economic Theory*, 53, 236-260.
- [11] Gibbons, R. (1987): “Piece-Rate Incentive Schemes,” *Journal of Labor Economics*, 5, 413-429.
- [12] Halac, M. (2012): “Relational Contracts and the Value of Relationships,” *American Economic Review*, 102, 750-779.
- [13] Hart, O., and J. Tirole (1988): “Contract Renegotiation and Coasian Dynamics,” *Review of Economic Studies*, 55, 509-540.
- [14] Kamenica, E., and M. Gentzkow (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590-2615.
- [15] Kanemoto, Y., and B. MacLeod (1992): “The Ratchet Effect and the Market for Secondhand Workers,” *Journal of Labor Economics*, 10, 85-98.
- [16] Laffont, J.J., and J. Tirole (1987): “Comparative Statics of the Optimal Dynamic Incentive Contract,” *European Economic Review*, 31, 901-926.
- [17] Laffont, J.J., and J. Tirole (1988): “The Dynamics of Incentive Contracts,” *Econometrica*, 56, 1153-1175.
- [18] Laffont, J.J., and J. Tirole (1990): “Adverse Selection and Renegotiation in Procurement,” *Review of Economic Studies*, 57, 597-625.
- [19] Laffont, J.J., and J. Tirole (1993): *A Theory of Incentives in Procurement and Regulation*, The MIT Press.

- [20] Maestri, L. (2017): “Dynamic Contracting under Adverse Selection and Renegotiation,” *Journal of Economic Theory*, 171, 136-173.
- [21] Peški, M. (2008): “Repeated Games with Incomplete Information on One Side,” *Theoretical Economics*, 3, 29-84.
- [22] Scheuer, F., and A. Wolitzky (2016): “Capital Taxation under Political Constraints,” *American Economic Review*, 106, 2304-2328.
- [23] Skreta, V. (2006): “Sequentially Optimal Mechanisms,” *Review of Economic Studies*, 73, 1085-1111.
- [24] Skreta, V. (2015): “Optimal Auction Design under Non-Commitment,” *Journal of Economic Theory*, 159, 854-890.
- [25] Strulovici, B. (2017): “Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts,” *Econometrica*, 85, 585-616.