Hiding Lemons among Peaches: Optimal Retention and Promotion Policy Design

Andrew B. Choi and Takuma Habu*

6 November 2025

– Click here for latest version –

Abstract

How should an employer design its retention and promotion policy when personnel decisions reveal private information about its workers to the labour market? We show it is always optimal to over-retain workers—inefficiently retaining workers who would be more productive elsewhere—to dampen the signalling effect of retention. In contrast, the optimal policy may over- or under-promote workers depending on the production technologies of competing firms. Our results shed light on the role of job titles and the Peter Principle. To solve for the optimal policy, we develop an approach to characterise a sender's commitment payoff in signalling games.

---

*Andrew B. Choi: Department of Economics, University of Michigan, andybchoi@gmail.com. Takuma Habu: Cornell SC Johnson College of Business, Cornell University, takumahabu@cornell.edu. We thank Marco Battaglini, Raphael Boleslavsky, Bart de Koning, Yi Chen, Thomas Jungbauer, Emir Kamenica, Michael Gibbs, Robert Gibbons, Justin Johnson, Rohit Lamba, Xiao Lin, Elliot Lipnowski, Doron Ravid, Collin Raymond, Lars Stole, Michael Waldman, and Kai Hao Yang for helpful comments.

1 Introduction

How should a firm retain or promote workers who can leave for another employer? When an employer has private information about its workers, retention or promotion decisions can drive up the wage because they serve as positive signals about workers' abilities to other potential employers.[1] In this paper, we explore whether a firm would dampen the positive signalling effect by designing a policy that retains or promotes some "lemon" workers along with some "peach" workers despite the costs of doing so.

We show that the firm always finds it optimal to put in place a policy that over-retains workers; that is, the firm retains workers who would be more productive elsewhere. However, whether the firm's optimal policy over- or under-promotes workers depends on how other potential employers value workers. The results differ for retention and promotion policies because promoting a worker entails an opportunity cost—the firm loses output from having the worker in the previous position—that is not present with retention. Our results provide a novel rationale for why workplaces abound with stories of incompetent workers and managers:[2] whenever the labour market cannot distinguish incompetent workers from competent ones, it can be profitable for the firm to tolerate the presence of incompetent workers as a means to minimise wage costs.[3]

To our knowledge, our paper is the first to study the signalling effects of retention and promotion jointly.[4] Moreover, our main result on promotion stands in contrast to the existing literature beginning with Waldman (1984) that finds that the signalling effect of promotion always leads to under-promotions (Waldman, 1984; Milgrom and Oster, 1987; Gibbons and Katz, 1991; Bernhardt, 1995; Ishida, 2004; Owan, 2004; Gibbons and Waldman, 2006; DeVaro and Waldman, 2012). The difference arises because we consider the *ex*

[1]Empirically, the literature has found evidence for the presence of signalling effect of promotions (DeVaro and Waldman, 2012; Bognanno and Melero, 2016; Cassidy, DeVaro and Kauhanen, 2016; Dato et al., 2016). Others also find evidence for asymmetric learning in labour markets (e.g., Gibbons and Katz, 1991; Schonberg, 2007; Pinkston, 2009; Kahn, 2013).

[2]For example, the well-known *Peter Principle* (Peter and Hull, 1969) from the management literature postulates that people in a hierarchy tend to rise to "a level of respective incompetence." We describe how our results relate to the Peter Principle precisely in section 6.1.

[3]In a matching model with a secondary market, Vohra (2024) demonstrates that a firm may hire less qualified workers to deter poaching.

[4]The labour literature has long recognised the inefficiency arising from the signalling effect of promotions, starting with Waldman (1984)—indeed, the term "promotion as signals" has become widely used (Gibbons and Roberts, 2013). With respect to retention, beginning with Greenwald (1986), the literature has predominantly studied the adverse selection effect arising from employers' strategic responses to outside offers. Gibbons and Katz (1991) is an exception. While their model captures both the signalling and adverse selection effects of retention, they do not consider promotions.

ante incentives of a firm operating in a labour market among *heterogeneous* firms.

Our model introduces multiple jobs (à la Costa, 1988) and heterogeneous firms to the setup of Waldman (1984). An incumbent firm employs a unit mass of junior workers with differing abilities. The firm must decide who to retain and, among those retained, who to promote to a senior position. The output from a worker in a given position increases linearly with the worker's ability. Relative to the junior position, the senior position's output increases more steeply with the worker's ability. A worker produces more in the senior position than in the junior position if and only if the worker's ability is above a threshold. The workers only care about the wage.[5] The labour market consists of heterogeneous outside firms with affine production technologies. To retain a worker (either in a junior or a senior position), the incumbent firm pays a wage equal to the worker's outside option—the highest wage the worker can command from outside firms. The outside firms only observe whether each worker was retained and, if so, whether the worker was promoted.[6]

The incumbent firm's problem is to maximise profits by designing an *HR policy* that specifies retention and promotion probabilities as functions of a worker's ability. We are interested in how the firm's optimal policy differs from the efficient benchmark that would be obtained under complete information, in which workers are employed by firms where they would be most productive.

Our first main finding (Theorem 3) is that, relative to the efficient benchmark, any optimal HR policy over-retains workers, and no optimal HR policy under-retains workers. That is, the firm always retains some workers who would be more productive at another firm and never lets go of workers who are most productive at the incumbent firm. Retaining an extra worker is beneficial for the firm because it lowers the average ability—and thus the wage cost—of all retained workers; however, doing so is also costly because the extra worker's output would be lower than the wage. We show that, on the margin, the benefit strictly dominates the cost, implying that the firm would over-retain workers. The result is driven by heterogeneous production technologies across outside firms, implying that workers' outside options are not only increasing but also convex in ability.

Our second main finding (Theorem 4) is that an optimal HR policy may exhibit either over-promotion or under-promotion (or even both if the optimal policy is non-deterministic).

[5] Workers may care about retention and promotions for reasons other than wage (Baker, Jensen and Murphy, 1988). Like Waldman (1984), we assume away this aspect.

[6] Following Waldman (1984), we do not introduce adverse selection issues related to how the incumbent firm responds to a wage offer from other firms (Greenwald, 1986; Gibbons and Katz, 1991) or allow the workers to possess private information (Costa, 1988).

The same rationale for over-retention means that over-promoting always increases the firm's profits from the promoted workers. However, unlike over-retention, over-promotion involves an opportunity cost—the additionally promoted worker, who had the highest ability among all junior workers, no longer works in the junior position. We demonstrate how the net effect of these two opposing forces depends on the nature of heterogeneity in the production technologies of outside firms. In particular, our results suggest that overpromotion is more likely among more educated workers, which is consistent with DeVaro and Waldman's (2012) empirical finding that higher-educated workers are more likely to be promoted.

As an alternative way for the firm to attenuate the signalling effect of retention (with and without promotions), we consider the firm's incentive to design jobs by merging the junior and senior positions. We show that an employer cannot strictly benefit from merging positions, meaning the firm prefers its workers to specialise. Finally, we also discuss the firm's incentive to create vacuous job titles to provide more information to the labour market.

Our model is a signalling game in which the sender (i.e., the firm) commits to a signalling action. Thus, the firm's problem of finding an optimal HR policy corresponds to finding the Sender's signalling strategy that maximises the sender's ex ante payoff. We make a technical contribution to the literature by developing a new approach to solving for the Sender's commitment payoff and strategy in signalling games.[7] Specifically, we construct an auxiliary information design game that provides an upper bound to the commitment payoff, and show that the bound is tight whenever the sender has access to a "rich" set of signalling actions. We show that our model satisfies the relevant richness condition, which allows us to recover the optimal HR policy from the solution to the auxiliary information design game.

Moreover, the equivalence between the two games also means that the firm can indirectly implement optimal signalling strategies by committing to Blackwell experiments (Kamenica and Gentzkow, 2011; Mathevet, Pearce and Stacchetti, 2024; Best and Quigley, 2024; Deb, Pai and Said, 2023) instead of committing directly to HR policies (Milgrom and Roberts, 1988; Waldman, 2003). The result also allows us to relax the commitment assumption using ideas from the information design literature (Lin and Liu, 2024) and from the repeated games literature (Fudenberg and Levine, 1992; Wolitzky and Luo, 2025).

[7] We explore some other applications of our approach in Appendix B.

2 The HR policy design problem

A firm employs a unit mass of junior workers.[8] Each worker has a *type*, interpreted as the worker's ability, distributed according to a CDF F that has full support on $\Theta := [0,1]$ with density f. The firm can *let go*, *retain* (in the junior position), or *promote* (to a senior position) its workers. We assume that each worker's output is affine and strictly increasing in the worker's type. Letting $j=1$ and 2 denote the junior and senior positions, respectively, the output of a type-θ worker in position $j \in \{1,2\}$ is given by

$$v_j(\theta) := s_j + k_j\theta,$$

where $s_j \in \mathbb{R}$ and $k_j \in \mathbb{R}_{++}$. The intercept s_j can be interpreted as firm-and-position-specific human capital, and the slope k_j as the marginal product of the worker's ability in position j. Following Costa (1988), we assume that the more senior position has a higher marginal product but not a uniformly higher total product; i.e., $k_2 > k_1$ and $v_2(\theta) \geq v_1(\theta)$ if and only if $\theta \geq \overline{\theta}_1$ for some $\overline{\theta}_1 \in (0,1)$.[9] We say that the firm assigns position $j=0$ to mean that the firm lets go of the worker; in this case, there is zero output, $v_0(\theta) := 0$.

There is a set I of *outside firms* with heterogeneous production technologies. A worker of type θ working at outside firm $i \in I$ produces an output of $c^i(\theta)$. For each i, we assume that c^i is affine and strictly increasing in θ. Define $c(\theta) := \sup\{c^i(\theta) : i \in I\}$ to be a type-θ worker's highest output across all outside firms.[10] As a supremum of strictly increasing affine functions, $c(\cdot)$ is strictly increasing and convex. To ease the interpretation and presentation of results, we assume that $c(\cdot)$ is, in fact, strictly convex and differentiable. We refer to the outside firm that maximises the output of a type-θ worker as the *θ-outside firm*. The convexity of $c(\cdot)$ reflects that the θ-outside firm's output is more sensitive to workers' abilities when θ is higher—this property will be the key driver of our results.

For the main body of the paper, we further assume $c(\cdot)$ to be such that workers whose type is below $\overline{\theta}_0 \in (0,\overline{\theta}_1)$ are valued by some outside firm more than by the incumbent, and workers whose type is above $\overline{\theta}_1$ are valued the most by the incumbent firm. This oft-made

[8]Similar to Waldman (1984), we could introduce a "pre-stage" game in which firms that are symmetrically informed about a pool of workers initially compete to hire workers. Firms subsequently privately learn about the abilities of the workers they hire, and workers gain firm-specific skills as they work. The game we introduce is then a subgame that follows after this pre-stage game for a particular firm.

[9]More precisely, we assume that $\overline{\theta}_1 \equiv \frac{s_1-s_2}{k_2-k_1} \in (0,1)$.

[10]In case an outside firm has multiple positions, we treat the production technology associated with each position as that of distinct firms.

Figure 1: Value of workers.

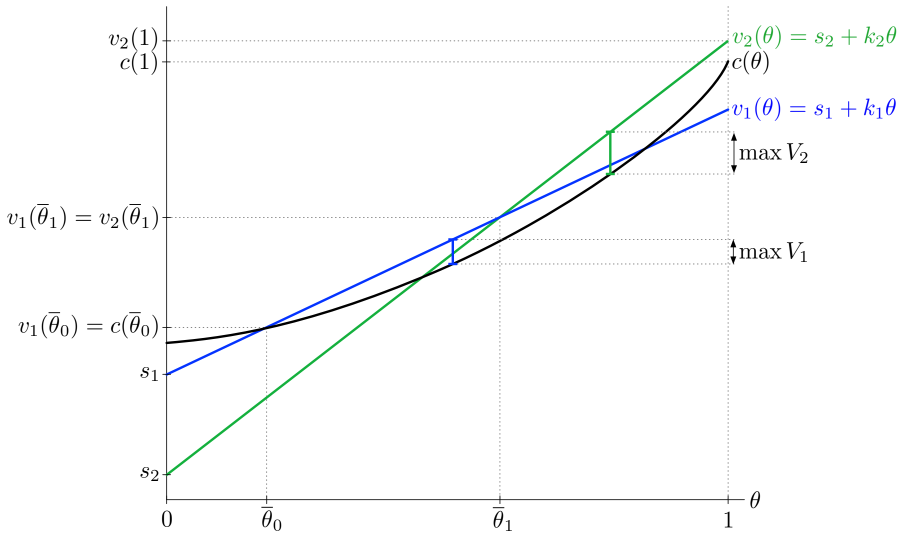

assumption (e.g., Waldman, 1984; Gibbons and Katz, 1991) ensures that it is possible for the incumbent firm to retain some of its workers and may be justified by factors such as workers' incumbent-firm-specific human capital, switching costs for workers, hiring and training costs for outside firms, or simply by the incumbent firm's superior production technology.[11]

For an allocation of workers to firms to be efficient, each worker must be employed at a firm (and in a position) that maximises the worker's output. Thus, our assumptions mean that it is efficient for the firm to let go of workers whose type is below $\overline{\theta}_0$, retain workers whose type lies in $[\overline{\theta}_0, \overline{\theta}_1]$ in the junior level, and (to retain) and promote workers whose type is above $\overline{\theta}_1$.

A worker's payoff is simply the wage. The firm's payoff from paying a wage w to a type-θ worker in position j is $v_j(\theta) - w$. As we explain in more detail below, we follow Waldman (1984) and view the firm as paying its workers just enough to deter other firms from poaching them. Thus, $c(\theta)$ can be interpreted as the wage cost for the firm of a type-θ worker if the labour market could observe θ. We define $V_j := v_j - c$ as the firm's profit in such a case and assume that $\max V_2 > \max V_1$ to capture the idea that seniors can potentially be more profitable than junior workers. Figure 1 gives a typical example of the firm's production technology that satisfies our assumptions.

We study the firm's problem of designing an *HR policy*, $\varphi : \Theta \to \Delta(\{0, 1, 2\})$, that

[11] We discuss the case in which outside firms value sufficiently high types of workers more than the incumbent firm in Section 6.4.

(possibly stochastically) lets go, retains at the junior position, or promotes each type of worker. Let $\Psi := (\Delta(\{0,1,2\}))^{\Theta}$ denote the set of all HR policies. We assume that the market can distinguish between workers based only on their assigned positions. Thus, to retain a worker in position $j \in \{1,2\}$, it suffices for the incumbent employer to just outbid the market and pay a wage equal to the market's expectation of the types of worker in position j.[12] Given a policy φ, the expected type of a worker in position j is given by $\mathbb{E}_{\varphi}[\theta|j] = \frac{\int_{\Theta} \theta \varphi(j|\theta) \mathrm{d}F(\theta)}{\int_{\Theta} \varphi(j|\theta) \mathrm{d}F(\theta)}$. Then, the firm's problem can be written as:

$$W^* := \max_{\varphi \in \Psi} \int_{\Theta} \sum_{j=1}^{2} \left[v_j(\theta) - c\left(\mathbb{E}_{\varphi}\left[\tilde{\theta}|j\right]\right) \right] \varphi(j|\theta) \,\mathrm{d}F(\theta). \tag{1}$$

We do not impose incentive compatibility constraints on the firm, so the firm is allowed to adopt HR policies that, ex post, result in losses for some types of workers. The lack of such constraints makes the firm's problem distinct from existing promotion-as-signals models and, as we explain in Section 5, allows for the possibility of over-retention and over-promotion by the firm.

3 Solving for the firm's optimal HR policy

The HR policy design problem is an example of a more general problem in signalling games of finding the sender's highest expected payoff when the sender can commit to signalling actions. In this section, we develop an approach to solving the general problem and then apply it to the HR policy design problem.[13] As such, readers who are interested in the solution to the HR policy design problem can safely skip the following subsection without affecting the flow or readability of the rest of the paper.

3.1 A general solution approach

Our approach is to appeal to an auxiliary information design game that effectively decouples the informational and payoff-relevant contents of signalling actions. While it is

[12] We implicitly assume a Bertrand-type competition for the workers so that $c(\cdot)$, strictly speaking, represents the second highest output of workers among firms in the labour market (see, for example, Blume, 2003). We assume that workers break ties in favour of the incumbent employer.

[13] We are not aware of existing methods that can be directly applied to solve the firm's problem, (1). The closest we are aware of is Boleslavsky and Shadmehr (2024) who provide a method for solving a similar problem when our richness condition fails.

possible to directly solve the original signalling game in some simple cases, our approach applies in more general environments.[14]

Consider a general signalling game ("base game") between a Sender (S) and a Receiver (R) in which the sender can commit to signalling actions. Let $\theta \in \Theta$ denote the sender's type, $a_i \in A_i$ denote each player $i \in \{S, R\}$'s action, and $u_i(\theta, a_S, a_R)$ denote player i's ex post payoff. The Sender commits to an action strategy $\sigma : \Theta \to \Delta A_S$. Then, after observing a_S, the Receiver updates his belief using Bayes' rule and takes an action a_R sequentially rationally. Let W^* denote the maximum ex ante payoff that the Sender can obtain from an optimal choice of σ.

Next, consider an auxiliary information design game ("auxiliary game") where, instead of committing to an action strategy, the Sender commits to a public experiment $\pi : \Theta \to \Delta M$, where M is a rich message space (e.g., $M = [0,1]$). Then, the realisation $m \in M$ is publicly observed, beliefs are updated using Bayes' rule, and the Sender chooses an action a_S sequentially rationally. Finally, having observed the Sender's action a_S as well as the realisation m from the experiment, the Receiver chooses an action a_R sequentially rationally. Let $\widehat{W}^*$ denote the maximum ex ante payoff that the Sender can obtain from an optimal choice of π.

The public nature of the experiment in the auxiliary game means that the information design game can be formulated as a standard Bayesian persuasion problem (Kamenica and Gentzkow, 2011). We take a belief-based approach and let $A_R^*(a_S, \mu)$ denote the set of actions that are best responses for the Receiver with belief $\mu \in \Delta\Theta$ after observing the Sender choose action $a_S \in A_S$. The highest payoffs for the Sender from choosing action a_S and inducing belief μ is then

$$\omega(a_S, \mu) := \max_{a_R \in A_R^*(a_S, \mu)} \int_\Theta u_S(\theta, a_S, a_R) \, \mathrm{d}\mu(\theta).$$

Finally, since the Sender best responds to any posterior belief μ in the auxiliary game, the Sender's value function is an upper envelope of the family of functions $\{\omega(a_S, \cdot)\}_{a_S \in A_S}$, given by $w(\mu) := \max_{a_S \in A_S} \omega(a_S, \mu)$. Equipped with $w(\cdot)$, the Sender's problem is a standard Bayesian persuasion problem:

$$\widehat{W}^* := \max_{\tau \in \Delta\Delta\Theta} w(\mu) \text{ s.t. } \int_{\Delta\Theta} \mu \, \mathrm{d}\tau(\mu) = \mu_0.$$

[14] In Appendix B, we apply our approach to some well-known signalling games.

It follows that standard methods from information design can be used to solve for τ^* that attains $\widehat{W}$ and derive experiments π^* that induce τ^*.[15]

Our first technical result is that the Sender's payoff is always higher in the auxiliary game.

Theorem 1. *Sender prefers to commit to experiments over signalling actions; i.e.,* $\widehat{W}^* \geq W^*$.

We briefly sketch the proof.[16] Note that the Sender can always induce the same distribution of posterior beliefs in the auxiliary game as in the base game by treating the signalling actions as messages. However, unlike in the base game, the Sender cannot commit to her actions in the auxiliary game, which has two implications. First, it implies that the auxiliary game is not a relaxation of the base game, so the result should not be obvious. Second, it implies that the sender's strategy in the base game is not necessarily sequentially rational in the auxiliary game. Nevertheless, in the auxiliary game, one can let the sender take a sequentially rational action that maximises her payoff, given her beliefs, and fixing the receiver's strategy to be the same as in the base game.[17] The result follows because the sender's payoff must be weakly higher from such a strategy.

Our next result is that the sender's payoffs in the two games are equal if her action space is sufficiently rich. Towards that goal, let us say that two actions $a_S, a'_S \in A_S$ are *copies* if they are payoff equivalent; i.e., $u_i(\cdot, a_S \cdot,) = u_i(\cdot, a'_S \cdot,)$ for each $i \in \{S, R\}$. We refer to a tuple of sequentially rational strategy profiles and a belief map that obeys Bayes' rule that attains W and $\widehat{W}$ as *solutions* to the base and auxiliary games, respectively. We say that the Sender's signalling action space is *sufficiently rich* if there exists a solution to the auxiliary game in which there are as many copies of each on-path signalling action as there are posterior beliefs in which the Sender is willing to take the action. Intuitively, when the Sender has access to a rich signalling action space, she can use copies of actions to convey information without payoff consequences, allowing her to replicate her payoff in the auxiliary game. This intuition leads to the following result.

Theorem 2. *If the Sender's signalling action space is sufficiently rich, then* $\widehat{W}^* = W^*$, *and there exists a solution to the auxiliary game which can be used to construct a solution to*

[15]A standard proof also means that a solution τ^* always exists that induces at most $\min\{|\Theta|, |A_S||A_R|\}$ number of posterior beliefs.

[16]All proofs are in the appendix.

[17]This also implies that the Sender does not benefit from the additional ability to commit to action strategies in the auxiliary game.

the base game.

Theorem 2 is particularly useful in applications in which the Sender has control over the set of actions she can use to signal. For example, employers are often free to define job titles (as relevant to our application), and sellers can set the names of the products they sell. In such cases, one can recover the solution to the base game by solving the auxiliary game.

We note that it is easy to check whether the sender's signalling action space is sufficiently rich: One first solves the Bayesian persuasion problem associated with the auxiliary game, and then simply checks whether there are as many copies of an action as there are the number of posterior beliefs in which the Sender is willing to take the action. Moreover, if the Sender takes a different action for each posterior belief, then there only needs to be one copy of each action for the richness condition to be satisfied. Below, we give two conditions on the Sender's value functions $\{\omega(a_S,\cdot)\}_{a_S\in A_S}$ that ensure that the Sender's action space is sufficiently rich.

Proposition 1. *The Sender's action space, A_S, is sufficiently rich in each of the following cases: (i) $\omega(a_S,\cdot)$ is concave for each $a_S\in A_S$; (ii) $\omega(a_S,\cdot)$ is convex for each $a_S\in A_S$ and there exists $a_S^\theta\in\arg\max_{a_S\in A_S}\omega(a_S,\delta_\theta)$ that is distinct for each $\theta\in\Theta$.*

The first case implies that, fixing the Sender's action, she never benefits from providing additional information. This is sufficient to ensure that the action space is sufficiently rich. In the second case, fixing the Sender's action, she wants to fully reveal the state and so she must have access to enough signalling actions to be able to do so. In applying the general solution approach to our application of designing HR policies, we will appeal to case (i) from Proposition 1.

3.2 Solving the HR policy design problem

We adopt the general approach described above to solve the HR policy design problem by first solving the relevant auxiliary information design game that effectively decouples informational and payoff-relevant contents of job assignment. In the auxiliary game, the firm initially commits to a public Blackwell experiment about the workers' abilities. We think of such an experiment as a test of whether a worker has reached a particular milestone or passed an exam, and the result as a job title or a qualification. Because the result of the experiment is publicly observed, the firm has no private information about workers' types in this auxiliary game. Having observed the outcome of the experiment, the firm assigns a position to each worker in a sequentially rational manner. The outside firms, having observed the experiment and its result as well as the job assignment for each worker, compete over the workers.

Finding the incumbent-firm-preferred equilibrium of the auxiliary game above reduces to a Bayesian persuasion problem that we describe now. Because production technologies are affine in θ, both the firm's and the market's decisions depend only on the posterior mean induced by the experiment. Moreover, because the firm and the labour market are symmetrically informed about workers' types, the assignment of positions provides no more information than that provided by the experiment. Thus, if the posterior mean of a group of workers is E, the firm must pay $c(E)$ to retain them, regardless of promotion. Consequently, for each realised posterior, the firm assigns workers in that group to a position that maximizes the expected payoff. Thus, the firm's payoff from inducing posterior mean E is

$$V(E) := \left[\max_{j\in\{1,2\}} v_j(E) - c(E)\right]_+,$$

where $[\cdot]_+ \equiv \max\{\cdot,0\}$. We refer to $V(\cdot)$ as the firm's value function. Finally, because of the well-known fact that the firm can induce a distribution of posterior means if and only if it is a mean-preserving contraction of the prior distribution, we can write the firm's information design problem as

$$\widehat{W}^* = \max_{G\in\mathrm{MPC}(F)} \int_\Theta V(E)\,\mathrm{d}G(E), \tag{2}$$

where $\mathrm{MPC}(F)$ denotes the set of all mean-preserving contractions of F.

We now establish an equivalence between the firm's original problem and the auxiliary game—a key result that we rely upon to solve for the firm's optimal HR policy.

Corollary 1. *The value of the firm's problem,* (1)*, in the original game equals the value of the firm's problem,* (2)*, in the related game; i.e.,* $W^* = \widehat{W}^*$*. Moreover, if* G^* *solves* (2)*, then there exists an HR policy* $\varphi^* \in \Psi$ *that solves the firm's original problem,* (1)*, and induces* G^* *as the distribution of posterior means.*

The result is a direct application of the general solution approach. Intuitively, the results hold because the firm can always induce the same distribution of posterior beliefs in the auxiliary game as in the base game by treating the signalling actions as messages, and the cost function $c(\cdot)$ is convex in the market's belief, so the firm has no reason to provide more information to the market than is revealed by job assignment.

To proceed, we further simplify the firm's information design problem by showing that it is without loss to only consider experiments with three messages. Let $\mathscr{F}$ denote the set of all tuples $(\theta_0, \theta_1, E_2) \in \Theta^3$ which satisfy $\theta_0 \in [0, \overline{\theta}_0]$, $\theta_1 \in [\theta_0, 1]$, and $\mathbb{E}[\theta|\theta \geq \theta_0] \leq E_2 \leq \mathbb{E}[\theta|\theta \geq \theta_1]$. Given $(\theta_0, \theta_1, E_2) \in \mathscr{F}$, define $G_{\theta_0,\theta_1,E_2}$ to be the distribution satisfying (i) $G \in \text{MPC}(F)$; (ii) the mass on $[0, \theta_0)$ from F is put on $\mathbb{E}[\theta|\theta < \theta_0]$; and (iii) the mass on $[\theta_0, 1]$ from F is put on $\mathbb{E}[\theta|\theta_0 \leq \theta < \theta_1]$ and on E_2. We call such a distribution *simple*. The lemma below shows that it suffices to restrict attention to simple distributions.

Lemma 1. *There exists a simple G^* that solves the information design problem,* (2)*; i.e.,*

$$\widehat{W}^* = \max_{(\theta_0,\theta_1,E_2)\in\mathscr{F}} \int_\Theta V(E)\, \mathrm{d}G_{\theta_0,\theta_1,E_2}(\theta). \tag{3}$$

We use results from Arieli et al. (2023) to prove Lemma 1. In particular, they show that information design problems, such as the firm's problem (2), always have a solution that belongs in a certain family of CDFs, which contains the set of simple CDFs. We then use the structure of V to show that a simple CDF always solves the firm's problem. Importantly, because (3) is a maximisation problem with respect to three real-valued parameters, Lemma 1 tells us that we can solve the firm's information design problem, (2), and characterise optimal distributions of posterior beliefs using a first-order approach.[18]

4 Inefficiencies of optimal HR policies

Our main results concern the nature of inefficiencies associated with the firm's optimal retention and promotion decisions.[19] To be precise, we say that an HR policy φ exhibits *over-retention* (resp. *under-retention*) if φ retains workers whose types are below (resp. above) the efficient retention cutoff, $\overline{\theta}_0$, with positive probability. Under these definitions, an HR policy exhibits efficient retention if and only if it does not exhibit over-retention or under-retention, and it can exhibit both over- and under-retention if it over-retains some types of workers and under-retains other types of workers. Analogously, we say that φ exhibits *over-promotion* (resp. *under-promotion*) if φ promotes workers whose types are below (resp. above) the efficient promotion cutoff, $\overline{\theta}_1$, with positive probability.

[18]The objective function need not be concave nor differentiable everywhere, so the first-order conditions may not be sufficient for optimality. Nevertheless, we are able to solve the problem using a first-order approach by using the special properties of our environment and results by Dworczak and Martini (2019) and Arieli et al. (2023).

[19]In the appendix, we provide a characterisation of the firm's optimal HR policies.

4.1 Over-retention is optimal

We first establish that we should generally expect firms to over-retain workers except in a knife-edge case where the optimal policy retains workers efficiently.

Theorem 3. *An optimal HR policy exists with efficient retention if and only if conditions*

$$V_2'(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0]) \geq V_1'(\overline{\theta}_0) \tag{4}$$

and

$$V_2(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0]) = V_2'(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0])(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0] - \overline{\theta}_0) \tag{5}$$

hold. Moreover, under-retention is never optimal. Therefore, if conditions (4) *or* (5) *do not hold, then any optimal HR policy exhibits over-retention.*

The underlying intuition for Theorem 3 is most clearly demonstrated by considering a cutoff HR policy in which the firm retains and promotes a worker if and only if the worker's type is above a retention cutoff θ_0 and a promotion cutoff $\theta_1 > \theta_0$, respectively. Such a cutoff policy can be implemented with an information structure that separately pools intervals $[0, \theta_0)$, $[\theta_0, \theta_1)$, and $[\theta_1, 1]$. The firm's profits from the cutoff policy is given by

$$\begin{aligned} \widehat{V}(\theta_0, \theta_1) = & \int_{\theta_0}^{\theta_1} \left(v_1(\theta) - c\left(\mathbb{E}\left[\tilde{\theta}|\theta_0 \leq \tilde{\theta} < \theta_1\right]\right)\right) f(\theta)\,\mathrm{d}\theta \\ & + \int_{\theta_1}^{1} \left(v_2(\theta) - c\left(\mathbb{E}\left[\tilde{\theta}|\tilde{\theta} \geq \theta_1\right]\right)\right) f(\theta)\,\mathrm{d}\theta, \end{aligned} \tag{6}$$

where the first and second lines are profits associated with workers in the junior and senior positions, respectively.

Let us first understand why under-retention is never optimal. Figure 2 below shows the case in which the firm adopts a retention cutoff that is strictly above the efficient retention cutoff, $\theta_0 > \overline{\theta}_0$. The two grey-shaded triangles in the figure represent the profits from each type of worker in the junior position under such a policy.

Figure 2: Under-retention is not optimal.

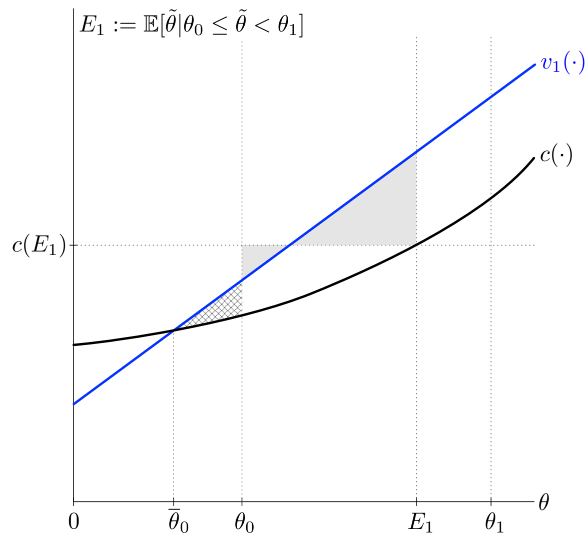

Now consider modifying the information structure associated with the cutoff HR policy so that it fully reveals the types of workers in the interval $(\overline{\theta}_0, \theta_0)$. With this modification, the firm obtains strictly positive profits, $v_1(\theta) - c(\theta) > 0$, for all types of workers in the interval—the hatched area in the figure above—while profits associated with other types of workers remain unchanged. Since the firm can always make such a strictly profitable modification to any information structure associated with an HR policy that exhibits under-retention, it follows that under-retention is never optimal for the firm.[20] In other words, under-retention is suboptimal because it implies the firm "gives up" workers even though it values the workers more than any outside firm.

To see why over-retention is generally optimal, consider the firm's incentive to marginally over-retain by choosing a retention cutoff θ_0 that is just below the efficient cutoff $\overline{\theta}_0$. By doing so, the firm loses from paying the marginal worker more than the output but gains from lowering the average wage of all junior workers. However, the average junior worker, having a higher ability than the marginally retained worker, is at risk of being poached by an outside firm that places a higher *marginal value* on worker ability. This raises the importance of lowering the market's belief about the average junior worker and ensures that the gain dominates the loss from over-retention.

To see this more clearly, consider the derivative of $\widehat{V}_1(\theta_0, \theta_1)$ with respect to θ_0 evaluated at the efficient retention cutoff $\overline{\theta}_0$ while assuming that $E_1 \equiv \mathbb{E}[\theta|\overline{\theta}_0 \leq \theta < \theta_1] \in$

[20]We note that the modified information structure is itself not optimal for the firm.

$(\overline{\theta}_0, \overline{\theta}_1)$ (which can be shown to hold in the relevant cases). Leibniz rule gives that

$$\left.\frac{\partial \widehat{V}(\cdot)}{\partial \theta_0}\right|_{\theta=\overline{\theta}_0} = -\left[f\left(\overline{\theta}_0\right)\left(v_1\left(\overline{\theta}_0\right) - c\left(E_1\right)\right) + \int_{\overline{\theta}_0}^{\theta_1} c'\left(E_1\right)\frac{\partial E_1}{\partial \theta_0} f(\theta)\mathrm{d}\theta\right], \tag{7}$$

The first term (inside the square brackets) in (7) is the direct effect on the firm's profit from letting go of a type-$\overline{\theta}_0$ worker: the firm pays $c(E_1)$ to retain a type-$\overline{\theta}_0$ worker that produces $v_1(\overline{\theta}_0)$ in output. The second term reflects the signalling effect of a change in the retention threshold on the firm's wage cost,[21] and the integral sums the change in the wage rate associated with all junior workers. Since the first term is strictly negative and the second term is strictly positive, the sign of (7) is a priori ambiguous. However, rewrite (7) as follows:

$$\begin{aligned}\left.\frac{\partial \widehat{V}(\cdot)}{\partial \theta_0}\right|_{\theta=\overline{\theta}_0} &= f\left(\overline{\theta}_0\right)\left[\left(c\left(E_1\right) - v_1\left(\overline{\theta}_0\right)\right) - c'\left(E_1\right)\left(E_1 - \overline{\theta}_0\right)\right] \\ &= f\left(\overline{\theta}_0\right)\left[\left(c\left(E_1\right) - c\left(\overline{\theta}_0\right)\right) - \left(c_{E_1}\left(E_1\right) - c_{E_1}\left(\overline{\theta}_0\right)\right)\right],\end{aligned} \tag{8}$$

where we substituted the expression for $\frac{\partial E_1}{\partial \theta_0}$ into (7), and the second line uses the fact that $c(\overline{\theta}_0) = v_1(\overline{\theta}_0)$ by definition of $\overline{\theta}_0$ and $c'(E_1)$ is the slope of the production technology of the E_1-outside firm, c_{E_1}. As evident from Figure 3, the convexity of c (which arises from the fact that θ-outside firm's output is more sensitive to workers' abilities when θ is higher) implies that the second difference in (8) is strictly greater than the first. Consequently, the derivative (7) is strictly negative meaning that the firm has a strict incentive to over-retain.

[21] $\frac{\partial E_1}{\partial \theta_0}$ captures the changes in the market's expectation about types of junior workers, and multiplying by $c'(E_1)$ gives the marginal change in market wage.

Figure 3: Over-retention.

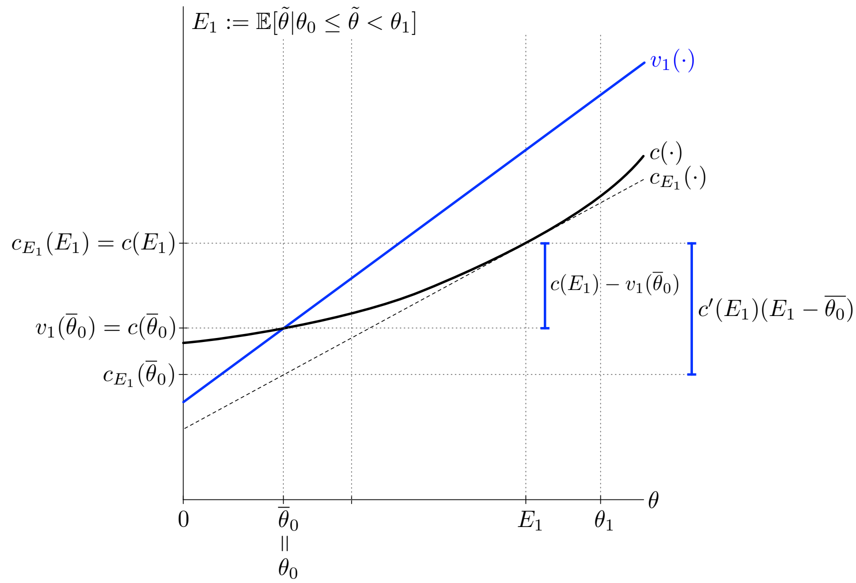

In proving Theorem 3, we also show that there is over-retention even when the optimal G is associated with $(\theta_0, \theta_1, E_2)$ for which $E_2 < \mathbb{E}[\theta|\theta \geq \theta_1]$ (which includes cases in which optimal HR policy is not deterministic). In this case, the effect of changing the retention threshold is more complex, because changing the retention threshold also affects the mass of workers in position $j = 2$.[22] Nevertheless, the same forces described above lead the firm to over-retain workers.

4.2 Over-promotion is sometimes optimal

We now show that the nature of the inefficiency associated with promotion decisions is more nuanced compared to retention decisions.

In general, there is a multiplicity of optimal HR policies because many information structures can induce a particular distribution of posterior means.[23] To ensure that over-promotion is not driven by unreasonable HR policies,[24] we focus attention on optimal

[22]To see why, in general, the firm's expected payoff from a simple distribution $G_{(\theta_0,\theta_1,E_2)}$ is

$$(1-F(\theta_0))\left[V(\mathbb{E}[\theta|\theta_0 \leq \theta < \theta_1])\,\alpha(\theta_0,\theta_1,E_2) + V(\mathbb{E}[\theta|\theta \geq \theta_1])\,(1-\alpha(\theta_0,\theta_1,E_2))\right],$$

where $\alpha(\theta_0, \theta_1, E_2)$ represents the proportion of retained workers who are assigned to position $j = 1$. Hence, changing the retention threshold θ_0 now also affects the mass of workers who are retained in position $j = 2$.

[23]See Example 3 by Arieli et al. (2023).

[24]In contrast, the over-retention result (Theorem 3) holds for any optimal HR policy so that, in particular, it holds for the "reasonable" optimal HR policies.

policies with the intuitive property that higher-type workers are more likely to be assigned to a "higher" position. We refer to such HR policies as being *monotone*.[25] Monotone HR policies have the desirable property that, in an enriched environment in which worker types are endogenous, higher-type workers would not have the incentive to pretend to be lower-type workers (for example, by shirking).[26] As we demonstrate later, it will also turn out to be related to the types of HR policies that can be implemented even when firms have less than full commitment power. The following lemma establishes that it is without loss of optimality to focus on monotone HR policies.

Lemma 2. *Any simple distribution of posterior means can be induced by a monotone HR policy.*

The proof uses that simple distributions are special cases of a class of distributions of posterior means which Arieli et al. (2023) show can be induced by a monotone information structure. For example, given a tuple $(\theta_0, \theta_1, E_2) \in \mathscr{F}$, the following HR policy φ induces $G_{\theta_0,\theta_1,E_2}$: φ lets go of the workers whose types lie in $[0, \theta_0]$ with probability one, promotes workers whose types lie in $[\theta_1, 1]$ with probability one, and randomise between promoting and not promoting (but retaining) workers whose types lie in $[\theta_0, \theta_1]$. Thus, we interpret θ_0 as a *retention threshold* and θ_1 as a *always-promote threshold*, where we make the distinction to account for the possible randomisation.

The following theorem characterises the nature of inefficiencies associated with promotion.

Theorem 4. *An optimal monotone HR policy that is deterministic exhibits either over-promotion or under-promotion but not both. An optimal monotone HR policy that is not deterministic may exhibit only over-promotion, only under-promotion, or both.*

---

[25] To define this property formally, note that an HR policy φ, when viewed as an information structure (i.e., with $M = \{0,1,2\}$ and $\varphi = \pi$), induces three possible posterior means, $\{E_j\}_{j=0}^2$, where E_j is the posterior mean after observing $j \in M$. Let $\widehat{\varphi}(\theta) \in \Delta(\{E_j\}_{j=0}^2)$ denote the distribution of posterior means such that

$$\widehat{\varphi}(E|\theta) = \begin{cases} \varphi(j|\theta) & \text{if } \exists j \in \{0,1,2\},\ E = E_j, \\ 0 & \text{otherwise.} \end{cases}$$

With this notation, an HR policy φ is *monotone* if, for any $\theta' \geq \theta$, $\widehat{\varphi}(\theta')$ first-order stochastically dominates $\widehat{\varphi}(\theta)$. Observe that if φ is deterministic, then the support of $\widehat{\varphi}(\cdot)$ is a singleton set, and thus, monotonicity reduces to requiring posterior means to be higher for higher types of workers.

[26] Goldstein and Leitner (2018) and Mensch (2021) provide discussions on why and when the monotonicity property might be desirable.

Once again, the intuition for why the results differ between retention and promotion decisions is most clear when considering a deterministic monotone HR policy. Thus, consider again the firm's expected payoff from such a policy given by equation (6). Observe that a change in the always-promote threshold θ_1 affects not only senior workers (i.e., $j = 2$) but also junior workers (i.e., $j = 1$). In contrast, changing the retention threshold only affects the junior workers because the firm's average profit-per-worker from letting go of workers (i.e., $j = 0$) is unaffected by its retention decisions—it always equals zero. Hence, compared to the case of over-retention, there is an additional force to consider in the case of over-promotion.

To be more precise, let us consider the derivative of $\widehat{V}$ with respect to θ_1 evaluated at the efficient promotion threshold $\overline{\theta}_1$ assuming that $E_1 \equiv \mathbb{E}[\theta|\theta_0 \leq \theta < \overline{\theta}_1] \in (\overline{\theta}_0, \overline{\theta}_1)$ and $E_2 \equiv \mathbb{E}[\theta|\theta \geq \overline{\theta}_1] \in (\overline{\theta}_1, 1)$ (which can be shown to hold in the relevant cases):

$$\begin{aligned} \left.\frac{\partial \widehat{V}(\cdot)}{\partial \theta_1}\right|_{\theta=\overline{\theta}_1} = & f\left(\overline{\theta}_1\right)\left[\left(v_1\left(\overline{\theta}_1\right) - c\left(E_1\right)\right) - \left(c_{E_1}\left(\overline{\theta}_1\right) - c_{E_1}\left(E_1\right)\right)\right], \\ & + f\left(\overline{\theta}_1\right)\left[\left(c\left(E_2\right) - v_2\left(\overline{\theta}_1\right)\right) - \left(c_{E_2}\left(E_2\right) - c_{E_2}\left(\overline{\theta}_1\right)\right)\right], \end{aligned} \tag{9}$$

where the first and second lines are changes to the firm's profits from junior and senior workers, respectively.

Now, suppose the firm slightly over-promotes by promoting a worker whose type is just below $\overline{\theta}_1$. Although such over-promotion increases the firm's total output from senior workers, the increase is exactly cancelled out by the reduction in total output from junior workers because the firm's net gain, $v_2(\overline{\theta}_1) - v_1(\overline{\theta}_1)$, equals zero by the definition of the efficient promotion cutoff $\overline{\theta}_1$. However, the promoted marginal worker, who was previously paid $c(E_1)$, must now be paid $c(E_2)$. Hence, over-promotion is profitable for the firm only if the signalling effects are greater than the increased wage cost $c(E_2) - c(E_1) > 0$.

Over-promotion reduces wage rates for workers in both junior and senior positions,[27] and as in the case of retention, the magnitude of the reduction depends on the production technologies of E_1- and E_2-outside firms. The convexity of c is sufficient to ensure that the second line in (9) is strictly negative.[28] However, convexity is not sufficient to determine the sign of the first term, which depends on the relative marginal products of workers at E_1- and E_2-outside firms.

[27]Recall $E_1 \equiv \mathbb{E}[\theta|\theta_0 \leq \theta < \overline{\theta}_1]$ and $E_2 \equiv \mathbb{E}[\theta|\theta \geq \overline{\theta}_1]]$ so that they are both decreasing in $\overline{\theta}_1$.

[28]We note that strict convexity of c is not necessary to sign the second line.

Figure 4: Under-promotion is optimal.

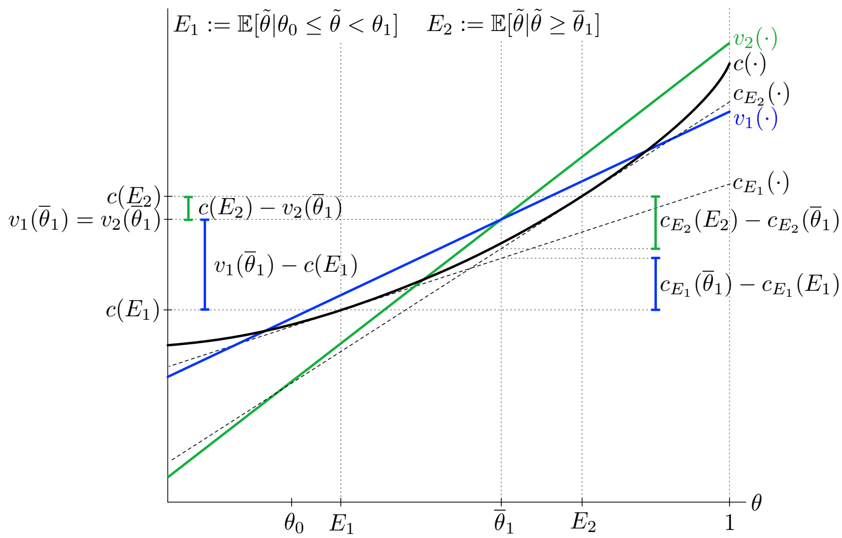

In particular, if the marginal product of the E_1-outside firm is low compared to that of the E_2-outside firm, the increased wage cost associated with the marginal worker can dominate the reduction in wage rates from signalling effects. This case is shown in Figure 4 above in which the increase in wage cost from over-promotion (the sum of the vertical bars to the left) is greater than the benefit from the signalling effect (the sum of the vertical bars to the right). In this case, the firm prefers to under-promote instead of over-promote workers. Alternatively, it can be the case that the marginal product of workers is relatively high at the E_1-outside firm so that the signalling effects dominate the increased wage cost from over-promotion so that the firm would find it profitable to over-promote workers.

In the case where no deterministic HR policy is optimal, an optimal monotone policy exists with a retention threshold of $\theta_0^* \le \overline{\theta}_0$ and an always-promotion threshold of $\theta_1^* > \theta_0^*$. If $\theta_1^* \le \overline{\theta}_1$, such an HR policy over-promotes workers because workers of types in the interval $[\theta_0^*, \theta_1^*)$ are promoted with positive probability and workers in the interval $[\theta_1^*, \overline{\theta}_1]$ are promoted with probability one. If, instead, $\theta_1^* > \overline{\theta}_1$, then the policy both over- and under-promotes workers because it promotes workers of types in the interval $[\theta_0^*, \theta_1^*)$ with positive probability but doesn't promote workers of types in the interval $[\overline{\theta}_1, \theta_1^*)$ with probability one.

The proposition below fully characterises the cases in which a deterministic HR policy is optimal.

Proposition 2. *A deterministic HR policy is optimal if and only if there does not exist a* *tuple* $(\theta_0^b, E_1^b, E_2^b)$ *with* $\theta_0^b \leq E_1^b \leq E_2^b$ *such that*

$$c\left(E_j^b\right) - v_j\left(\theta_0^b\right) = c'\left(E_j^b\right)\left(E_j^b - \theta_0^b\right)\ \forall j = \{1,2\}, \tag{10}$$

$$k_2 - k_1 = c'\left(E_2^b\right) - c'\left(E_1^b\right), \tag{11}$$

$$E_1^b \leq \mathbb{E}\left[\theta|\theta \geq \theta_0^b\right] < E_2^b < \mathbb{E}\left[\theta|\theta \geq \theta_1^b\right], \tag{12}$$

where θ_1^b *satisfies* $\mathbb{E}[\theta|\theta_0^b \leq \theta < \theta_1^b] = E_1^b$.

To understand the proposition, observe that existence of a tuple $(\theta_0^b, E_1^b, E_2^b)$ satisfying condition (12) implies existence of a simple CDF associated with a tuple $(\theta_0^b, \theta_1^b, E_2^b)$ where $\mathbb{E}[\theta|\theta \geq \theta_0^b] < E_2^b < \mathbb{E}[\theta|\theta \geq \theta_1^b]$, and it can be shown that no deterministic HR policy can induce such a simple CDF. Roughly, condition (12) requires a sufficient mass of workers who are "good enough" so that the firm can reduce the wage cost of senior workers by sometimes promoting workers who are not "good enough." Conditions (10) and (11), which are restrictions on v_1, v_2 and c (but not on F), ensure that such a simple CDF is optimal in the firm's information design problem. Specifically, condition (10) ensures that setting the retention threshold at θ_0^b is optimal for the firm that chooses an always-promote threshold of θ_1^b.[29] Condition (11), which requires the firm and the outside firms to value the additional marginal product from senior over junior workers equally, ensures that setting the always-promote threshold at θ_1^b is optimal when inducing posterior means for junior and senior workers of E_1^b and E_2^b, respectively.

The discussion above implies that we can guarantee existence of an optimal deterministic HR policy when: either no tuple exists that satisfies conditions (10) and (11); or, even if such a tuple exists, the prior distribution F is such that condition (12) does not hold. In the appendix, we show that for a tuple $(\theta_0^b, E_1^b, E_2^b)$ satisfying conditions (10) and (11) to exist, $c(\cdot)$ must be sufficiently convex but not too convex. Thus, to benefit from randomisation, the production technologies of outside firms cannot vary too widely.

4.3 When does the firm over-promote?

We focus on the case in which an optimal monotone HR policy is deterministic so that the optimal monotone HR policy either exhibits over-promotion or under-promotion but not

[29] For each $j \in \{1,2\}$, the condition (10) ensures that setting the retention threshold at θ_0^b is optimal for a firm that chooses between letting go of workers and retaining workers in position j while inducing a posterior mean of E_j^b.

both. To state when the firm would over- or under-promote, let $\overline{E}_1^e$ denote the expected type of junior workers that makes it optimal for the firm to choose the efficient promotion threshold of $\overline{\theta}_1$, and let $\overline{\theta}_0^e$ denote the optimal retention threshold in case junior worker's expected type is $\overline{E}_1^e$.[30] We note that $\overline{E}_1^e$ and $\overline{\theta}_0^e$ need not be consistent with F so that, in general, we would expect $\overline{E}_1^e \neq \mathbb{E}[\theta|\overline{\theta}_0^e \leq \theta < \overline{\theta}_1]$.

Proposition 3. *Suppose a deterministic optimal HR policy φ^* exists.*

(i) Suppose there exists a tuple $(\theta_0^b, E_1^b, E_2^b)$ with $\theta_0^b \leq E_1^b \leq E_2^b$ satisfying conditions (10) *and* (11)*. Let θ_1^b be such that $\mathbb{E}[\theta|\theta_0^b \leq \theta < \theta_1^b] = E_1^b$. Then, φ^* exhibits over-promotion if $E_2^b < \mathbb{E}[\theta|\theta \geq \theta_1^b]$ and either (a) $\mathbb{E}[\theta|\theta \geq \overline{\theta}_1] > E_2^b$ or (b) $\mathbb{E}[\theta|\theta \geq \overline{\theta}_1] \leq E_2^b$ and $\mathbb{E}[\theta|\overline{\theta}_0^e \leq \theta < \overline{\theta}_1] > \overline{E}_1^e$; otherwise, φ^* exhibits under-promotion.*

(ii) Suppose no such tuple exists. Then, φ^* exhibits over-promotion if $\mathbb{E}[\theta|\overline{\theta}_0^e \leq \theta < \overline{\theta}_1] > \overline{E}_1^e$; otherwise, φ^* exhibits under-promotion.

Recall that conditions (10) and (11) are restrictions on v_1, v_2 and c but not on F, and that a tuple satisfying these conditions exists if c is convex but not too convex; in other words, when the production technologies of outside firms are not too varied. Thus, case (ii) in Proposition 3 above says that the firm would over-promote if production technologies of outside firms are sufficiently varied and the expected type of junior workers if the firm were to follow a threshold HR policy $(\overline{\theta}_0^e, \overline{\theta}_1)$ is sufficiently high. The latter condition ensures that the effect of lowering the promotion threshold on the profits from workers in the junior position is sufficiently small. Case (i) in Proposition 3 concerns the environment in which the production technologies of outside firms are not too varied. Here, over-promotion occurs if there is a sufficiently large mass of high-type workers (so that $\mathbb{E}[\theta|\theta \geq \theta_1^b], \mathbb{E}[\theta|\theta \geq \overline{\theta}_1] > E_2^b$) which, in particular, implies promoting efficiently results in wage costs that are "too high" for the firm. Alternatively, over-promotion can occur if, as in case (i), the expected type of junior workers if the firm were to follow a threshold HR policy $(\overline{\theta}_0^e, \overline{\theta}_1)$ is sufficiently high.

To better understand when a firm over-promotes, we also study the firm's incentive to over-promote as we change the environment. To that end, we say that tuples of primitives (v_2, c, F) and $(\tilde{v}_2, \tilde{c}, \tilde{F})$ (satisfying our assumptions) are *comparable* if they both imply a common efficient retention and promotion thresholds, and if they both satisfy conditions

[30]Concretely, $\overline{E}_1^e$ sets the derivative $\partial\widehat{V}/\partial\theta_1$ evaluated at $\theta = \overline{\theta}_1$ to zero when expected type of junior and senior workers are $\overline{E}_1^e$ and $\mathbb{E}[\theta|\theta \geq \overline{\theta}_1]$ respectively, and $\overline{\theta}_0^e$ sets the derivative $\partial\widehat{V}/\partial\theta_0$ evaluated at $\theta = \overline{\theta}_0^e$ to zero when expected type of junior workers is $\overline{E}_1^e$.

for either case (i)(b) or (ii) of Proposition 3. Given such comparable tuples of primitives, we say that over-promotion is *more likely* under (v_2, c, F) than under $(\tilde{v}_2, \tilde{c}, \tilde{F})$ if the change can result in a switch from the optimal monotone HR policy exhibiting under-promotion to over-promotion but not the other way around.[31]

We first contemplate a change in the production technologies of outside firms. Specifically, we consider making comparable changes that alter $c'(\overline{E}_2)$, where $\overline{E}_2 := \mathbb{E}[\theta|\theta \geq \overline{\theta}_1]$. To isolate the signalling effect, we ensure $c(\overline{E}_2)$ and $c(1)$ remain unchanged so that the change does not directly alter the wage cost for the firm from promoting efficiently, and we also leave $c(\cdot)$ unchanged on $[0, \overline{\theta}_1]$ to avoid directly affecting the firm's choice of retention threshold.

Proposition 4. *Suppose no deterministic monotone HR policy is optimal. Consider comparable (v_2, c, F) and $(v_2, \tilde{c}, F)$ such that $c(\theta) = \tilde{c}(\theta)$ for $\theta \in [0, \overline{\theta}_1] \cup \{\overline{E}_2, 1\}$. Then, over-promotion is more likely under $(v_2, \tilde{c}, F)$ than under (v_2, c, F) if $\tilde{c}'(\overline{E}_2) > c'(\overline{E}_2)$.*

In words, the above proposition says over-promotion is more likely with a greater marginal product of $\overline{E}_2$-outside firm (i.e., the firm that values efficiently promoted workers the highest). The result follows from the fact that the greater marginal product of such a firm leads to a greater magnification of the signalling effect associated with promotions.

To demonstrate that the incentive to over- or under-promote is largely driven by the technologies of outside firms, we also consider a change in the marginal product of the incumbent firm's production technology, k_2 (with a corresponding change in s_2 to ensure that the efficient promotion threshold remains unchanged). The following proposition shows that such a comparable change does not affect the firm's incentive to over- or under-promote.

Proposition 5. *Suppose no deterministic monotone HR policy is optimal. Consider comparable (v_2, c, F) and $(\tilde{v}_2, c, F)$. Then, over-promotion is no more likely under $(\tilde{v}_2, c, F)$ than under (v_2, c, F).*

Propositions 4 and 5 reinforce the insight that whether there is over-promotion or under-promotion is driven by the technologies of outside firms, rather than the technology of the incumbent employer.

[31]More concretely, we say that over-promotion is more likely under (v_2, c, F) than under $(\tilde{v}_2, \tilde{c}, \tilde{F})$ if the derivative $\frac{\partial \widehat{V}}{\partial \theta_1}$ evaluated at $(\overline{\theta}_0^e, \overline{\theta}_1)$ is smaller under (v, c, F) than under $(\tilde{v}_2, \tilde{c}, \tilde{F})$, where $\overline{\theta}_0^e$ is such that the first-order condition with respect to θ_0 holds. That is, $\frac{\partial \widehat{V}}{\partial \theta_0}$ evaluated at $(\overline{\theta}_0^e, \overline{\theta}_1)$ is zero if $\overline{\theta}_0^e > 0$ or strictly negative if $\overline{\theta}_0^e = 0$. Observe that there cannot be a switch from under- to over-promotion between two tuples of primitives that satisfy condition (i) of Proposition 3.

We can also use our model to study how the nature of inefficiency associated with optimal HR policies changes when the workers have different characteristics that are publicly observable; e.g., workers may have different education levels or professional qualifications. To that end, suppose there are two sets of workers and their prior distribution of types is ordered according to their likelihood ratios.[32] We show that over-promotion is more likely for the set of workers whose distribution dominates another in the likelihood-ratio order.

Proposition 6. *Suppose no deterministic monotone HR policy is optimal. Consider comparable* (v_2, c, F) *and* $(v_2, c, \tilde{F})$*. Then, over-promotion is more likely under* $(v_2, c, \tilde{F})$ *than under* (v_2, c, F) *if* $\tilde{F}$ *dominates* F *in the likelihood ratio order.*

For example, if workers with a higher education level are "better" than those with a lower education level in the sense of likelihood ratio order, over-promotion is more likely among workers who are more educated. This observation is consistent with DeVaro and Waldman (2012) who empirically find that, fixing ability, workers with higher education levels are more likely to be promoted.

5 Relaxing the firm's commitment power

The firm can directly commit to HR policies in a number of ways. For example, the firm can delegate its retention and promotion decisions to an HR department (Milgrom and Roberts, 1988) or establish a reputation for following certain policies (Waldman, 2003). Indeed, firms such as AIG and Amazon let go (resp. promote) workers below (resp. above) a certain percentile, and others promote workers based on tenure.[33]

Even when direct commitment is not possible, the equivalence between the original signalling game and the information design game implies that the firm can also indirectly implement optimal HR policies by committing to public Blackwell experiments about its workers.

For example, the firm could establish a clear career path for workers in which progress depends on the workers reaching certain milestones (or passing some tests) as a way to implement optimal HR policies. Indeed, in many industries, firms disclose to potential

[32] Recall that, given two absolutely continuous CDFs F and H, with support $[0,1]$ and respective densities f and h, the CDF H dominates F in the likelihood ratio order if and only if $\frac{h(\cdot)}{f(\cdot)}$ is increasing.

[33] These policies may also serve to incentivise workers to exert more effort.

workers how they can progress through the ranks.[34] Alternatively, the firm could also implement optimal experiments by committing to contracts à la Deb, Pai and Said (2023).[35] In this light, the resistance from HR departments that line managers often face in trying to get their managees promoted "faster" or firing managees can be viewed as an optimal response for the firm to manipulate the signalling effects of promotion. Nevertheless one may wonder: To what extent do our results depend on the firm's ability to commit?

We consider two ways that further demonstrate how else the firm can partially implement optimal HR policies. First, we consider the case in which the employer lacks the power to commit to HR policies, but the labour market can observe the distribution of positions held by its workers after it has assigned them. That is, in addition to observing the firm's choice of HR policy φ and the position of the workers $j \sim \varphi(\cdot|\theta)$, the labour market also observes the distribution of positions after the firm has made its retention and promotion decisions. Firms sometimes disclose such information voluntarily,[36] or disclose them as mandated by accounting rules or laws,[37] alternatively, they are available publicly through websites such as LinkedIn or through services that provide intelligence about HR structures of companies.[38]

Importantly, public knowledge of the distribution of assigned positions effectively allows the market to detect (and respond to) deviations by the firm from its stated HR policy, φ, if deviations result in a distribution of positions that is inconsistent with φ. Following Lin and Liu (2024)'s arguments, the possibility for the market to respond to deviations allows the firm to effectively commit to HR policies that result in the same marginal distribution of messages as that of the announced one.

To state the firm's problem when the firm has this partial commitment power, let $D(\varphi) \subseteq$

[34] Career fairs at colleges and universities are ripe with employers explaining how graduates can expect to progress through the rank in the company.

[35] The Principal, the Agent, and the Receiver, in their setting, correspond to the firm, a line manager, and the labour market in our setting, where the line manager is an agent of the firm who can observe the workers' types. We would add to the setup of Deb, Pai and Said (2023) that the firm is able to assign positions once it has gathered information about workers from the line manager.

[36] For example, certain types of companies and institutions, such as economic consultancies and academic departments, provide a public list of their staff in various positions.

[37] For example, in the UK, the Companies Act 2006 (section 411) specifies that firms above certain size "must also disclose the average number of persons within each category of persons so employed." Some diversity, equity and inclusion reporting requirements stipulate companies to disclose demographic information about their employees (e.g., The European Union's Directive 2014/95/EU).

[38] For example, companies such as Crayon collect "Competitive Intelligence" to "provide the most complete picture of what your competitors are up to."

Ψ denote the set of HR policies that have the same marginal distribution as $\varphi \in \Psi$; i.e.,

$$D(\varphi) := \left\{ \varphi' \in \Psi : \int_{\Theta} \varphi'(\cdot|\theta)\, \mathrm{d}F(\theta) = \int_{\Theta} \varphi(\cdot|\theta)\, \mathrm{d}F(\theta) \right\}.$$

The firm's problem is to choose an HR policy that maximises its ex ante expected payoff:

$$\max_{\varphi \in \Psi} \int_{\Theta} \sum_{j=1}^{2} \left[v_j(\theta) - c\left(\mathbb{E}_{\varphi}\left[\tilde{\theta} | j \right] \right) \right] \varphi(j|\theta)\, \mathrm{d}F(\theta) \tag{13}$$

$$\text{s.t. } \varphi \in \arg\max_{\varphi' \in D(\varphi)} \int_{\Theta} \sum_{j=1}^{2} \left[v_j(\theta) - c\left(\mathbb{E}_{\varphi}\left[\tilde{\theta} | j \right] \right) \right] \varphi'(j|\theta)\, \mathrm{d}F(\theta). \tag{14}$$

The constraint (14), which we refer to as the firm's *credibility constraint*, ensures that the firm, after announcing that it will follow an HR policy φ, would not deviate to some other HR policy that has the same marginal distribution of positions when the labour market uses φ to update beliefs about workers. We say that an HR policy φ is *credible* if φ satisfies the credibility constraint.

In solving (13), we first show that we can restrict attention to deterministic and monotone HR policies.

Lemma 3. *An HR policy is credible if and only if it is deterministic and monotone.*

An immediate corollary is that the ability to partially commit is sufficient for the firm to attain the full-commitment payoff whenever there exists a deterministic optimal HR policy.

Corollary 2. *Optimal HR policy can be implemented under partial commitment if and only if a deterministic optimal HR policy exists.*

Note that an HR policy φ that is both deterministic and monotone is a simple threshold policy. Thus, Lemma 3 simplifies the problem (1) to one of finding optimal retention and promotion thresholds. Using this simplification, we prove that the nature of inefficiency remains unchanged.

Proposition 7. *Any HR policy φ^* that solves* (13) *is a threshold policy with retention and promotion thresholds, θ_0^* and $\theta_1^* \geq \theta_0^*$ respectively. In addition, φ^* exhibits over-retention, and it may exhibit over- or under-promotion but not both.*

We also consider a reputational foundation of the firm's ability to commit, as alluded to by Spence (1973) and Waldman (2003). To this end, results by Fudenberg and Levine

(1992) imply that, when a signalling game is repeated frequently, and states and sender's actions are observed by the receiver at the end of each repetition, the sender can attain her pure strategy Stackelberg payoff. More recently, Wolitzky and Luo (2025) show that, under commonly assumed payoff structures, the sender can attain her pure strategy commitment payoff if her ex ante optimal signalling strategy is (pure and) monotone, even when only actions are observed. While these results do not directly apply to our settings,[39] they suggest that the firm can implement an optimal HR policy via reputation whenever a deterministic HR policy exists (recall Proposition 9).

We note that some direct or indirect ability to commit to HR policies is necessary to ensure that firm can implement HR policy that exhibits over-retention and/or over-promotion. For example, suppose we require HR policy to not make a loss with respect to any type of worker by adding the following set of constraints to the firm's original problem:

$$\operatorname{supp}(\varphi(\cdot|\theta)) \in \underset{j\in\{0,1,2\}}{\arg\max}\, v_j(\theta) - c\left(\mathbb{E}_{\varphi}\left[\tilde{\theta}|j\right]\right) \;\forall\theta\in\Theta.$$

These incentive compatibility constraints that are standard in signalling games remove the possibility of over-retention and over-promotion because doing so necessarily results in the firm making losses on the types of workers it over-retains or over-promotes. In fact, without the ability to commit, we obtain the classic result as Waldman (1984) does that the firm always under-promotes (and also that the firm under-retains).[40]

6 Discussions

6.1 Peter Principle

In a best-selling book, Peter and Hull (1969) introduce the idea of the *Peter Principle* which they define as the 'fact' that: "In a hierarchy, every employee tends to rise to his level of incompetence." The book provides ample (mostly fictional) examples of competent workers being promoted to positions for which they lack competence. In the economics

[39]Fudenberg and Levine (1992) assume action spaces to be finite, while Wolitzky and Luo (2025) further assume the state space to be finite.

[40]To see why, recall Figure 2 which shows the case in which the firm slightly over-retains by choosing $\theta_0 > \overline{\theta}_0$ and makes on some workers represented by the lower shaded triangle. In fact, the firm can eliminate losses by choosing a higher retention threshold (i.e., under-retaining further) so that the outside option of the expected type of a junior worker, $c(\mathbb{E}[\theta|\theta_0 \leq \theta < \theta_1])$, exactly equals the value of the marginally retained junior worker, $v_1(\theta_0)$.

literature, Benson, Li and Shue (2019) view the Principle as a consequence of firms facing a "costly trade-off between promoting the best potential managers and incentivising workers" that results in "firms promot[ing] workers who excel in their current roles, at the expense of promoting those who would make the best managers." Alternatively, Lazear (2004) demonstrates that, even without such trade-offs, the Peter Principle can hold due to mean reversion when promotions are based on stochastic performance. In contrast, our results—that it is (resp. can be) profitable, as a means to minimise wage costs, for the firm to retain (resp. promote) workers to a position in which they would be less productive—provide a novel rationale for the Peter Principle that operates even in environments without the presence of worker incentive tradeoffs and stochastic performance. In particular, in our model, the workers' types represent their abilities that are positively and deterministically associated with performance in both junior (e.g., sales) and senior (e.g., manager) positions. Our results therefore suggest that the Peter Principle would hold in practice beyond the setting of sales workers that previous empirical studies have focused on such as in low-powered backoffice or government jobs.

6.2 Job design

An alternative way to attenuate the positive signalling effect of promotion is to combine the jobs associated with positions 1 and 2 into one role, say to position $j = m$ (for *merged* position) as such merging of positions prevents the labour market from distinguishing workers based on whether they are retained as a junior or promoted to be a senior. However, designing jobs in this manner would presumably be inefficient for the firm,[41] and we capture this inefficiency by assuming that a type-θ worker's output in the merged position is a weighted average of v_1 and v_2; i.e.,

$$v_m(\theta) = \alpha v_1(\theta) + (1-\alpha) v_2(\theta),$$

where $\alpha \in (0,1)$.[42] However, the firm cannot benefit from merging jobs in this manner. To see this, define $V_m(E) := [v_m(E) - c(E)]_+$ as the firm's value function G^* denote a solution

[41] We note that the firm benefits from the ability to covertly assign positions $j \in \{1,2\}$ to attenuate the signalling effect without suffering a cost in terms of worker output.

[42] The results also hold when α depends continuously on θ.

to the firm's problem when merging the two positions. But because $V_m \leq V$, we have

$$\begin{aligned}\int_\Theta V_m(E)\,\mathrm{d}G^*(E) &\leq \int_\Theta V(E)\,\mathrm{d}G^*(E) \\ &\leq \max_{G\in\mathrm{MPC}(F)} \int_\Theta V(E)\,\mathrm{d}G(E) = \widehat{V}^*.\end{aligned}$$

Hence, the firm cannot do better by merging positions. Thus, while the firm benefits from obfuscating workers' abilities to the labour market by misallocating across positions, it prefers to specialise workers into positions.

6.3 Vacuous job titles

In the National Longitudinal Survey of Youth in 1990, Pergamit and Veum (1999) find that approximately a third of those who are promoted "continued to perform basically the same duties as before."[43] While we would interpret such a finding as suggestive of the fact that firms do indeed over-promote workers, it is also possible that the finding represents evidence of a disconnect in the relationship between job titles and job responsibilities.

In the information design game that we use to solve the firm's original problem, the firm can use messages ($m \in M$) to provide more information than is possible via the assignment of positions alone. However, we show that the firm has no incentive to do so; in our model, the firm sends at most one message for each position.

In other settings, it is possible that the firm strictly benefits from providing more information by assigning multiple job titles to the same job responsibility; i.e., by assigning vacuous job titles.[44] Such cases arise if there is convexity in the worker's output from a given position j in such a way that the firm's value function $\omega(j,\cdot)$ is not necessarily concave.

6.4 Labour market that values the workers more

In our main setup, we assumed that it is efficient for the current employer of the firm to retain the most productive workers; i.e., $v_2(1) \geq c(1)$. However, it is also possible that high-θ workers should be employed by other firms that value them more; i.e., $v_2(1) < c(1)$. In this case, define $\overline{\theta}_2 := \max\{\theta \in \Theta : v_2(\theta) = c(\theta)\}$ as the efficient threshold for letting

[43] Somewhat more recently, Kosteas (2011) reports a similar finding based on the 1996–2006 waves of the National Longitudinal Surveys of Youth 1979 cohort.

[44] In the terminology from section 3.1, firms use copies of signalling actions.

go of workers rather than retaining them in position $j = 2$. Then, when the prior mean about workers' types is sufficiently high, it can now be optimal for the firm to adopt a deterministic HR policy characterised by a retention threshold $\theta_2^* \geq \overline{\theta}_2$ in which all workers of type $\theta < \theta_2^*$ are retained and assigned to position $j = 2$ and all workers of type $\theta \geq \theta_2^*$ are let go.

6.5 Beyond two jobs

Our approach to solving for the optimal HR policies extends easily to cases in which the firm can assign workers to more than two positions. Suppose jobs are ordered so that: $v_k(\cdot)$ has both a lower intercept and a greater slope than $v_{k-1}(\cdot)$, and their intersection $\overline{\theta}_{k-1} \in (0,1)$ increases with k. Then, the tradeoffs involved between any two consecutive positions would resemble that between $j = 1$ and 2 that we described above. Additional care must be taken to determine whether the firm would find it optimal to assign at least some workers to each position.

6.6 Counteroffers

To simplify exposition and to be consistent with Waldman (1984), we do not explicitly model the wage-setting game and assume that the firm commits to offering a wage to each worker so as to ensure that the worker cannot obtain a strictly higher wage from other firms in the labour market. An alternative interpretation of this assumption is that the firm is always willing to make a counteroffer up to the average output of workers in position j, $v_j(\mathbb{E}_\varphi[\theta|j])$ if an outside firm were to make a poaching offer of say w to a worker in position $j \in \{1,2\}$. This assumption is reasonable in the information design game as the firm learns no more information about the worker than the market. However, to the extent that the firm learns the workers' types before assigning positions to them, the firm may be tempted to let go of workers who are worth less than w to the firm; i.e., let go of workers in position j whose types θ is such that $v_j(\theta) \leq w$. This type of counteroffer behaviour by the firm introduces an adverse selection problem for those attempting to poach workers from the firm (Greenwald, 1986; Gibbons and Katz, 1991). While a detailed analysis of the model with adverse selection is beyond the scope of this paper, we note that the firm benefits from letting go of workers strategically because doing so lowers the poaching offers the labour market is willing to make for the workers.

6.7 Relation with Waldman (1984) and Gibbons and Katz (1991)

Waldman (1984) studies how the firm's decision to promote or not promote a worker affects the wage it must pay in order to retain workers. In Waldman (1984), the firm assigns job $j \in \{1,2\}$ to the worker, where the firm's direct payoff from assigning job j to a type-θ worker is given by $(s_1, k_1) = (x(1+s), 0)$ and $(s_2, k_2) = (0, 1+s)$ for some $s > 0$ and $x \in (\mathbb{E}[\theta], 1]$. The labour market's value for a type-θ worker is simply $c(\theta) := \theta$. Thus, the firm's value function in Waldman (1984) is given by

$$V(E) = (1+s)\max\{E, x\} - E = \begin{cases} (1+s)x - \theta & \text{if } \theta < \overline{\theta}_1 \\ s\theta & \text{if } \theta \geq \overline{\theta}_1 \end{cases}, \tag{15}$$

where $\overline{\theta}_1 = x$. Observe that $V(\cdot) > 0$ so that the efficient outcome is for the incumbent firm to retain all workers. Because V is convex, a full revelation is optimal for the firm. However, because V is piecewise linear, the firm can attain the optimal payoff by simply setting the promotion threshold to be $\overline{\theta}_1$. Thus, the optimal promotion rule in Waldman (1984) is efficient and does not exhibit under- or over-promotion.

Gibbons and Katz (1991) study a model that can capture both the signalling effect of retention à la Waldman (1984) and the adverse-selection effect of retention à la Greenwald (1986). The firm decides whether to retain or let go of the worker. The firm's payoffs from assigning each position are characterised, respectively, by $(s_0, k_0) = (0,0)$ and $(s_1, k_1) = (s, 1)$, where $0 < s < \mathbb{E}[\theta]$. The labour market's value for a type-θ worker is again $c(\theta) := \theta$. Hence, the firm's value function is simply $V(E) = s$. Observe that it is always efficient for the firm to retain workers of all types.

Observe that both Waldman (1984) and Gibbons and Katz (1991) assume a common production technology for outside firms. The analysis here, therefore, highlights that when the firm can commit to HR policies, inefficiencies in HR policies arise from the fact that heterogeneity in the production technologies of outside firms, which, in turn, leads to convexity in the workers' outside options.

7 Conclusion

In his seminal study that began the literature on signalling games, Spence (1973) focuses on the job market as a "paradigm case" of "markets in which signaling takes place and in which the primary signalers are relatively numerous and in the market sufficiently infrequently that they are not expected to (and therefore do not) invest in acquiring signaling reputations." In other words, Spence viewed the job market as a prototypical environment in which the sender cannot commit to signalling actions. Yet, the same no-commitment assumption has been maintained in many important economic environments that have since been studied in which there are only a *few* senders that operate *frequently*: e.g., government agencies (Vickers, 1986), politicians and committees (Banks, 1990; Cameron and Rosendorff, 1993), long-lived firms (Myers and Majluf, 1984; Milgrom and Roberts, 1986; Waldman, 1984). In this paper, we develop a new approach that allows one to answer when, why and how the sender in these environments can benefit from committing to signalling actions.

We apply our approach to study an employer's decision to retain and promote workers. We show that employers benefit from committing to over-retain and over-promote workers to reduce wage costs by dampening the positive signalling effects of retention and promotion, despite the obvious costs of retaining and promoting incompetent workers. While we explore other uses of our approach in the appendix, we hope that our approach, which connects the literature on signalling and information design, is fruitful in studying other applications.

References

Arieli, Itai, Yakov Babichenko, Rann Smorodinsky, and Takuro Yamashita. 2023. "Optimal persuasion via bi-pooling." *Theoretical Economics*, 18(1): 15–36.

Baker, George P., Michael C. Jensen, and Kevin J. Murphy. 1988. "Compensation and Incentives: Practice vs. Theory." *The Journal of Finance*, 43(3): 593–616.

Banks, Jeffrey S. 1990. "Monopoly Agenda Control and Asymmetric Information." *The Quarterly Journal of Economics*, 105(2): 445.

Benson, Alan, Danielle Li, and Kelly Shue. 2019. "Promotions and the Peter Principle." *The Quarterly Journal of Economics*, 134(4): 2085–2134.

Bernhardt, Dan. 1995. "Strategic Promotion and Compensation." *The Review of Economic Studies*, 62(2): 315–339.

Best, James, and Daniel Quigley. 2024. "Persuasion for the Long Run." *Journal of Political Economy*, 132(5): 1740–1791.

Blume, Andreas. 2003. "Bertrand without fudge." *Economics Letters*, 78(2): 167–168.

Bognanno, Michael, and Eduardo Melero. 2016. "Promotion Signals, Experience, and Education." *Journal of Economics & Management Strategy*, 25(1): 111–132.

Boleslavsky, Raphael, and Mehdi Shadmehr. 2024. "Signaling with Commitment." *Working Paper*.

Cameron, Charles M., and B.Peter Rosendorff. 1993. "A Signaling Theory of Congressional Oversight." *Games and Economic Behavior*, 5(1): 44–70.

Cassidy, Hugh, Jed DeVaro, and Antti Kauhanen. 2016. "Promotion signaling, gender, and turnover: New theory and evidence." *Journal of Economic Behavior & Organization*, 126: 140–166.

Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *The Quarterly Journal of Economics*, 102(2): 179–222.

Costa, Joan E. Ricart I. 1988. "Managerial Task Assignment and Promotions." *Econometrica*, 56(2): 449–466.

Dato, Simon, Andreas Grunewald, Matthias KrÃ€kel, and Daniel MÃŒller. 2016. "Asymmetric employer information, promotions, and the wage policy of firms." *Games and Economic Behavior*, 100: 273–300.

Deb, Rahul, Mallesh Pai, and Maher Said. 2023. "Indirect Persuasion." *Working Paper*.

DeVaro, Jed, and Michael Waldman. 2012. "The Signaling Role of Promotion: Further Theory and Empirical Evidence." *Journal of Labor Economics*, 30(1): 91–147.

Dworczak, Piotr, and Giorgio Martini. 2019. "The Simple Economics of Optimal Persuasion." *Journal of Political Economy*, 127(5): 1993–2048.

Fudenberg, Drew, and David K. Levine. 1992. "Maintaining a Reputation when Strategies are Imperfectly Observed." *The Review of Economic Studies*, 59(3): 561–579.

Fudenberg, Drew, and Jean Tirole. 1991. *Game Theory.* MIT Press.

Gentzkow, Matthew, and Emir Kamenica. 2016. "A Rothschild-Stiglitz Approach to Bayesian Persuasion." *American Economic Review*, 106(5): 597–601.

Gibbons, Robert. 1992. *A Primer in Game Theory.* Pearson.

Gibbons, Robert, and John Roberts, ed. 2013. *The Handbook of Organizational Economics.* . 1 ed., Princeton University Press.

Gibbons, Robert, and Lawrence F. Katz. 1991. "Layoffs and Lemons." *Journal of Labor Economics*, 9(4): 351–380.

Gibbons, Robert, and Michael Waldman. 2006. "Enriching a Theory of Wage and Promotion Dynamics inside Firms." *Journal of Labor Economics*, 24(1): 59–107.

Goldstein, Itay, and Yaron Leitner. 2018. "Stress tests and information disclosure." *Journal of Economic Theory*, 177: 34–69.

Greenwald, Bruce C. 1986. "Adverse Selection in the Labour Market." *The Review of Economic Studies*, 53(3): 325–347.

Ishida, Junichiro. 2004. "Signaling and strategically delayed promotion." *Labour Economics*, 11(6): 687–700.

Kahn, Lisa B. 2013. "Asymmetric Information between Employers." *American Economic Journal: Applied Economics*, 5(4): 165–205.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review*, 101(6): 2590–2615.

Kolotilin, Anton. 2018. "Optimal information disclosure: A linear programming approach." *Theoretical Economics*, 13(2): 607–635.

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. 2017. "Persuasion of a Privately Informed Receiver." *Econometrica*, 85(6): 1949–1964.

Kosteas, Vasilios D. 2011. "Job Satisfaction and Promotions." *Industrial Relations: A Journal of Economy and Society*, 50(1): 174–194.

Lazear, Edward P. 2004. "The Peter Principle: A Theory of Decline." *Journal of Political Economy*, 112(S1): S141–S163.

Lin, Xiao, and Ce Liu. 2024. "Credible Persuasion." *Journal of Political Economy*, 132(7): 2228–2273.

Mathevet, Laurent, David Pearce, and Ennio Stacchetti. 2024. "Reputation and Information Design." *Working Paper*.

Mensch, Jeffrey. 2021. "Monotone persuasion." *Games and Economic Behavior*, 130: 521–542.

Milgrom, Paul, and John Roberts. 1986. "Price and Advertising Signals of Product Quality." *Journal of Political Economy*, 94(4): 796–821.

Milgrom, Paul, and John Roberts. 1988. "An Economic Approach to Influence Activities in Organizations." *American Journal of Sociology*, 94: S154–S179.

Milgrom, Paul, and Sharon Oster. 1987. "Job Discrimination, Market Forces , and the Invisibility Hypothesis." *The Quarterly Journal of Economics*, 102(3): 453.

Myers, Stewart C., and Nicholas S. Majluf. 1984. "Corporate financing and investment decisions when firms have information that investors do not have." *Journal of Financial Economics*, 13(2): 187–221.

Owan, Hideo. 2004. "Promotion, Turnover, Earnings, and Firm-Sponsored Training." *Journal of Labor Economics*, 22(4): 955–978.

Pergamit, Michael R., and Jonathan R. Veum. 1999. "What is a Promotion?" *Industrial and Labor Relations Review*, 52(4): 581.

Peter, Laurence J., and Raymond Hull. 1969. *The Peter Principle.* New York:William Morrow.

Pinkston, Joshua C. 2009. "A Model of Asymmetric Employer Learning with Testable Implications." *Review of Economic Studies*, 76(1): 367–394.

Rothschild, Michael, and Joseph Stiglitz. n.d.. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *The Quarterly Journal of Economics*, 90(4): 629.

Schonberg, Uta. 2007. "Testing for Asymmetric Employer Learning." *Journal of Labor Economics*, 25: 651–691.

Shaked, Moshe, and George J. Shanthikumar, ed. 2007. "Stochastic Orders." , ed. Moshe Shaked and George J. Shanthikumar . 1 ed. Springer New York, NY.

Spence, Michael. 1973. "Job Market Signaling." *The Quarterly Journal of Economics*, 87(3): 355–374.

Vickers, John. 1986. "Signalling in a model of monetary policy with incomplete information." *Oxford Economic Papers*, 38(3): 443–455.

Vohra, Akhil. 2024. "Unraveling and Inefficient Matching." *American Economic Journal: Microeconomics (forthcoming).*

Waldman, Michael. 1984. "Job Assignments, Signalling, and Efficiency." *The RAND Journal of Economics*, 15(2): 255.

Waldman, Michael. 2003. "Ex Ante versus Ex Post Optimal Promotion Rules: The Case of Internal Promotion." *Economic Inquiry*, 41(1): 27–41.

Wolitzky, Alexander, and Daniel Luo. 2025. "Marginal Reputation." *Econometrica (forthcoming).*

A Appendix: Proofs

A.1 A general solution approach

We first formally define the equilibrium concepts for the base and auxiliary games.

Base game. Let $\sigma : \Theta \to \Delta A_S$ denote the Sender's signalling strategy, $\rho : A_S \to \Delta A_R$ denote the Receiver's action strategy, and let $\mu : A_S \to \Delta\Theta$ be a belief map. We say that a tuple (σ, ρ, μ) is: *Bayes consistent* (BC) if μ is obtained by updating μ_0 using σ via Bayes' rule whenever possible; and *Receiver-IC* (RIC) if ρ maximises the Receiver's expected payoff given belief μ. Letting $W(\sigma, \rho)$ denote the Sender's expected payoff from a profile of strategies (σ, ρ), a solution to the base game solves the following problem:

$$W^* := \sup_{(\sigma', \rho', \mu')} W\left(\sigma', \rho'\right) \text{ s.t. } \left(\sigma', \rho', \mu'\right) \text{ satisfying BC and RIC.} \tag{16}$$

Auxiliary game. Let $\widehat{\sigma} : M \to \Delta A_S$ denote the Sender's action strategy, $\widehat{\rho} : A_S \times M \to \Delta A_R$ denote the Receiver's action strategy, and let $\widehat{\mu} : M \to \Delta\Theta$ be a belief map.[45] Given an experiment π, we say that a tuple $(\widehat{\sigma}, \widehat{\rho}, \widehat{\mu})$ is: *Bayes consistent given* π (π-BC) if $\widehat{\mu}$ is obtained by updating μ_0 with respect to π; *Receiver-IC given* π (π-RIC) if $\widehat{\rho}$ maximises Receiver's expected payoff given belief $\widehat{\mu}$; *Sender-IC given* π (π-SIC) if $\widehat{\sigma} : M \to \Delta A_S$ maximises Sender's expected payoff given $\widehat{\mu}$ and $\widehat{\rho}$; i.e., for each $m \in M$. We call a tuple satisfying π-BC, π-RIC and π-SIC a π-*PBE*. Letting $\widehat{W}_i(\pi, \widehat{\sigma}, \widehat{\rho})$ the Sender's payoff from experiment π and a profile of strategies $(\widehat{\sigma}, \widehat{\rho})$, a solution to the auxiliary game solves the following problem:

$$\widehat{W}^* := \sup_{(\pi', \widehat{\sigma}', \widehat{\rho}', \widehat{\mu}')} \widehat{W}\left(\pi', \widehat{\sigma}', \widehat{\rho}'\right) \text{ s.t. } \left(\widehat{\sigma}', \widehat{\rho}', \widehat{\mu}'\right) \text{ is a } \pi'\text{-PBE.} \tag{17}$$

Proof of Theorem 1. Fix a solution to the base game, $(\sigma^*, \rho^*, \mu^*)$. We will construct a tuple $(\pi, \widehat{\sigma}, \widehat{\rho}, \widehat{\mu})$ such that $(\widehat{\sigma}, \widehat{\rho}, \widehat{\mu})$ is a π-PBE and $\widehat{W}(\pi, \widehat{\sigma}, \widehat{\rho}) \geq W(\sigma^*, \rho^*)$. Optimality of $\widehat{W}^*$ then implies $\widehat{W}^* \geq W^*$. Without loss, let $M = A_S$. To distinguish between elements of M and A_S, we write $\mathtt{a}_S(m)$ to denote the unique $a_S \in A_S$ associated with $m \in M$ and $\mathtt{m}(a_S)$ to denote the unique $m \in M$ associated with $a_S \in A_S$.

[45] Because the Sender's action strategy does not depend on θ, the Receiver's belief about the state after observing (a_S, m) is the same as after observing just m.

We first claim that, for all $a_S \in A_S$,

$$V\left(a_S, \mu^*\left(\cdot|a_S\right)\right) = \sum_{\theta\in\Theta}\sum_{a_R\in A_R} u_S\left(\theta, a_S, a_R\right)\rho^*\left(a_R|a_S\right)\mu^*\left(\theta|a_S\right).$$

To prove this, suppose by way of contradiction that, for some $a_S \in A_S$, the left-hand side is strictly greater than the right-hand side (by construction, the left-hand side is always weakly greater than the right-hand side). This means that there exists ρ' that attains the left-hand side and is optimal for the Receiver given belief $\mu^*(\cdot|a_S)$. Hence, (σ^*, ρ', μ^*) is BC and RIC. Moreover,

$$\begin{aligned} W\left(\sigma^*, \rho'\right) &= \sum_{a_S\in A_S}\left[\sum_{\theta\in\Theta}\sum_{a_R\in A_R} u_i\left(\theta, a_S, a_R\right)\rho'\left(a_R|a_S\right)\mu^*\left(\theta|a_S\right)\right]\sigma^*\left(a_S\right) \\ &> \sum_{a_S\in A_S}\left[\sum_{\theta\in\Theta}\sum_{a_R\in A_R} u_i\left(\theta, a_S, a_R\right)\rho^*\left(a_R|a_S\right)\mu^*\left(\theta|a_S\right)\right]\sigma^*\left(a_S\right) = W\left(\sigma^*, \rho^*\right). \end{aligned}$$

But above contradicts that $(\sigma^*, \rho^*, \mu^*)$ solves the base game.

Define, for all $a_S \in A_S$, $\pi(\mathtt{m}(a_S)|\cdot) := \sigma^*(a_S|\cdot)$, $\widehat{\mu}(\cdot|\mathtt{m}(a_S)) := \mu^*(\cdot|a_S)$, $\widehat{\sigma}(\cdot|\mathtt{m}(a_S)) := \delta_{\{a_S^*\}}(\cdot)$, and $\widehat{\rho}(\cdot|a_S, m) := \rho^*(\cdot|a_S)$ for all $(a_S, m) \in A_S \times M$, where, for all $a_S \in A_S$, $a_S^* \in \arg\max_{a_S' \in A_S} \omega(a_S', \mu^*(\cdot|a_S))$. In words, we construct the Blackwell experiment π so as to induce the same distribution of posterior beliefs as σ^* in the signalling game. The Receiver's strategy in the information design game $\widehat{\rho}$ is the Sender-preferred strategy that is optimal for the Receiver after observing a_S and having belief $\widehat{\mu}(\cdot|\mathtt{m}(a_S)) = \mu^*(\cdot|a_S)$. Finally, in the auxiliary game, the Sender takes action a_S^* that maximises her expected payoff when she has belief $\widehat{\mu}(\cdot|\mathtt{m}(a_S))$ given the Receiver follows $\widehat{\rho}$.

Now, observe that $\widehat{\sigma}$ is π-SIC and $\widehat{\rho}$ is $\widehat{\pi}$-RIC by construction. To check for π-BC, consider first $m \in M$ such that $\sum_{\theta\in\Theta}\pi\left(m|\theta\right)\mu_0\left(\theta\right) > 0$. By definition of π, we must have $\sum_{\theta\in\Theta}\sigma^*(\mathtt{a}_S(m)|\theta)\mu_0(\theta) > 0$ so that

$$\widehat{\mu}\left(\cdot|m\right) \equiv \mu^*\left(\cdot|\mathtt{a}_S\left(m\right)\right) = \frac{\sigma^*\left(\mathtt{a}_S\left(m\right)|\cdot\right)\mu_0\left(\cdot\right)}{\sum_{\theta\in\Theta}\sigma^*\left(\mathtt{a}_S\left(m\right)|\theta\right)\mu_0\left(\theta\right)} = \frac{\pi\left(m|\cdot\right)\mu_0\left(\cdot\right)}{\sum_{\theta\in\Theta}\pi\left(m|\theta\right)\mu_0\left(\theta\right)},$$

where the first equality follows from the fact that $(\sigma^*, \rho^*, \mu^*)$ is BIC and the second equality follow from the definition. Hence, $\widehat{\mu}(\cdot|m)$ is obtained via Bayes' rule. If, instead, $\sum_{\theta\in\Theta}\pi\left(m|\theta\right)\mu_0\left(\theta\right) = 0$, then $\sum_{\theta\in\Theta}\sigma^*(a_S|\theta)\mu_0(\theta) = 0$ so that $\widehat{\mu}(\cdot|m)$ can be set arbitrarily—in particular, it can be set to equal $\mu^*(\cdot|\mathtt{a}_S(m))$.

It remains to show that $\widehat{W}(\pi,\widehat{\sigma},\widehat{\rho}) \geq W^*$. Let $\sigma^*(\cdot) := \sum_{\theta\in\Theta}\sigma^*(\cdot|\theta)\mu_0(\theta)$ and observe

$$\begin{aligned}
\widehat{W}(\pi,\widehat{\sigma},\widehat{\rho}) &= \sum_{\theta\in\Theta}\sum_{m\in A_S}\sum_{a_R\in A_R} u_S(\theta,a_S^*,a_R)\,\widehat{\rho}\left(a_R|a_S^*,\mu^*(\cdot|\mathtt{a}_S(m))\right)\sigma^*(\mathtt{a}_S(m)|\theta)\,\mu_0(\theta) \\
&= \sum_{m\in A_S}\left[\sum_{\theta\in\Theta}\sum_{a_R\in A_R} u_S(\theta,a_S^*,a_R)\,\widehat{\rho}\left(a_R|a_S^*,\mu^*(\cdot|\mathtt{a}_S(m))\right)\mu^*(\theta|\mathtt{a}_S(m))\right]\sigma^*(\mathtt{a}_S(m)) \\
&\geq \sum_{m\in A_S}\left[\sum_{\theta\in\Theta}\sum_{a_R\in A_R} u_S(\theta,a_S,a_R)\,\rho^*(a_R|a_S)\,\mu^*(\theta|\mathtt{a}_S(m))\right]\sigma^*(\mathtt{a}_S(m)) \\
&= \sum_{\theta\in\Theta}\sum_{a_S\in A_S}\sum_{a_R\in A_R} u_S(\theta,a_S,a_R)\,\rho^*(a_R|a_S)\,\sigma^*(a_S|\theta)\,\mu_0(\theta) = W^*,
\end{aligned}$$

where the inequality follows from the fact that $\widehat{\rho}$ is the Sender-preferred selection among all Receiver-optimal action strategies which, in particular, includes $\rho^*(\cdot|a_S)$ because

$$\underset{a_R\in A_R}{\arg\max}\sum_{\theta\in\Theta} u_R(\theta,a_S,a_R)\,\widehat{\mu}(\theta|\mathtt{m}(a_S)) = \underset{a_R\in A_R}{\arg\max}\sum_{\theta\in\Theta} u_R(\theta,a_S,a_R)\,\mu^*(\theta|a_S),$$

which, in turn, follows from the fact that posterior beliefs are the same between the signalling game and the information design game (i.e., $\widehat{\mu}(\cdot|\mathtt{m}(a_S)) = \mu^*(\cdot|a_S)$). ■

Toward proving Theorem 2, we introduce a canonical class of π-PBE in which the Sender does not randomise over actions after observing a message and every message induces different posterior beliefs. The following establishes that there always exists a payoff-equivalent π-PBE with these properties. We refer to any π-PBE from this canonical class as an *S-pure* π-PBE.

Lemma 4. *Fix π and suppose $(\widehat{\sigma},\widehat{\rho},\widehat{\mu})$ is a π-equilibrium. Then, there exists a payoff equivalent π'-equilibrium $(\widehat{\sigma}',\widehat{\rho}',\widehat{\mu}')$ such that (i) $\widehat{\sigma}'(\cdot|m)$ is degenerate for all $m \in \text{supp}(\pi)$, and (ii) $\widehat{\mu}(\cdot|m) \neq \widehat{\mu}(\cdot|m')$ for all distinct $m,m' \in \text{supp}(\pi)$.*

Proof. (i) Suppose that for some $m \in \text{supp}(\pi)$ and some distinct $a_S^1, a_S^2 \in A_S$, $\widehat{\sigma}(a_S^1|m) > 0$ and $\widehat{\sigma}(a_S^2|m) > 0$; i.e., the Sender takes actions a_S^1 and a_S^2 with positive probability after observing message m. Note that π-SIC implies that the Sender must be indifferent between actions a_S^1 and a_S^2 after observing m under belief $\widehat{\mu}(\cdot|m)$ while taking into account the fact that the Receiver's response may differ; i.e.,

$$\sum_{\theta\in\Theta}\sum_{a_R\in A_R} u_S\left(\theta,a_S^1,a_R\right)\widehat{\rho}\left(a_R|a_S^1,m\right)\widehat{\mu}(\theta|m) = \sum_{\theta\in\Theta}\sum_{a_R\in A_R} u_S\left(\theta,a_S^2,a_R\right)\widehat{\rho}\left(a_R|a_S^2,m\right)\widehat{\mu}(\theta|m).$$

Consider an alternative Sender action strategy $\widehat{\sigma}'$ obtained by pooling actions a_S^1 and a_S^2:

$$\widehat{\sigma}'\left(a_S'|m'\right)=\begin{cases}\widehat{\sigma}\left(a_S^1|m\right)+\widehat{\sigma}\left(a_S^2|m\right) & \text{if } \left(a_S',m'\right)=\left(a_S^1,m\right),\\ 0 & \text{if } \left(a_S',m'\right)=\left(a_S^2,m\right),\\ \widehat{\sigma}\left(a_S'|m'\right) & \text{otherwise.}\end{cases}$$

Since $\mu(\cdot|m)$ is unaffected by the change, it remains the case that $\widehat{\rho}(a_R|a_S^1,m)$ is optimal for the Receiver. It follows that Sender's payoff from a_S^1—which was the same as her payoff from a_S^2 by π-SIC—is unchanged. Therefore, $(\widehat{\sigma}',\widehat{\rho},\widehat{\mu})$ is a π-equilibrium. Applying the procedure to $(\widehat{\sigma},\widehat{\rho},\widehat{\mu})$ as needed ensures that $\widehat{\sigma}'(\cdot|m)$ is degenerate for all $m\in\text{supp}(\pi)$.

(ii) Suppose there exist distinct $m_1,m_2\in\text{supp}(\pi)$ that induces the same posterior belief and suppose we consider π' such that

$$\pi'\left(m|\cdot\right)=\begin{cases}\pi\left(m_1|\cdot\right)+\pi\left(m_2|\cdot\right) & \text{if } m=m_1,\\ 0 & \text{if } m=m_2,\\ \pi\left(m|\cdot\right) & \text{otherwise.}\end{cases}$$

and let $\widehat{\mu}'$ be given by Bayes rule for all $m\in\text{supp}(\pi')$. Note that $\widehat{\mu}'(\cdot|m_1)=\widehat{\mu}(\cdot|m_1)=\widehat{\mu}(\cdot|m_2)$. By (i), we may assume that $\widehat{\sigma}(a_S^1|m_1)=\widehat{\sigma}(a_S^2|m_2)=1$ for some $a_S^1,a_S^2\in A_S$. By π-SIC, we may pool actions a_S^1 and a_S^2 as in part (i) of the proof. The same argument then implies that $(\widehat{\sigma}',\widehat{\rho}',\widehat{\mu}')$ is a π'-equilibrium where $\widehat{\rho}'(\cdot|a_S,m)=\widehat{\rho}(\cdot|a_S,m)$ for all $m\in\text{supp}(\pi')$ and a_S such that $\widehat{\sigma}'(a_S|m)=1$. ■

In proving Theorem 2, we explcitly show how a solution to the base game can be constructed from a solution to the auxiliary game.

Proof of Theorem 2. When the action space is sufficiently rich, we can take the action space of the Sender in the base game to be $A_S^M:=A_S\times M$ (instead of A_S) so that, for any distinct $m,m'\in M$, a_S^m and $a_S^{m'}$ are payoff equivalent but they are distinguishable from one another.

Fix $(\pi^*,\widehat{\sigma}^*,\widehat{\rho}^*,\widehat{\mu}^*)$ that is a solution to the auxiliary game. We will construct a tuple (σ^*,ρ^*,μ^*) satisfying BC, RIC, and $\widehat{W}^*=W(\sigma^*,\rho^*)$ as follows: for all $a_S^m\in A_S^M$, $\sigma^*\left(a_S^m|\cdot\right):=\widehat{\sigma}^*\left(a_S|m\right)\pi^*\left(m|\cdot\right)$, $\rho^*\left(\cdot|a_S^m\right):=\widehat{\rho}^*\left(\cdot|a_S,m\right)$, and $\mu^*\left(\cdot|a_S^m\right):=\widehat{\mu}^*\left(\cdot|m\right)$. By construction,

$\widehat{W}^* = W(\sigma^*, \rho^*)$ because

$$
\begin{aligned}
W(\sigma^*, \rho^*) &= \sum_{\theta\in\Theta} \sum_{a_S^m \in A_S^M} \sum_{a_R\in A_R} u_S(\theta, a_S, a_R)\, \widehat{\rho}^*(a_R|a_S, m)\, \widehat{\sigma}^*(a_S|m)\, \pi^*(m|\theta)\, \mu_0(\theta) \\
&= \sum_{\theta\in\Theta} \sum_{m\in M} \sum_{a_S\in A_S} \sum_{a_R\in A_R} u_S(\theta, a_S, a_R)\, \widehat{\rho}^*(a_R|a_S, m)\, \widehat{\sigma}^*(a_S|m)\, \pi^*(m|\theta)\, \mu_0(\theta) \\
&= \widehat{W}(\pi^*, \widehat{\sigma}^*, \widehat{\rho}^*) = \widehat{W}^*.
\end{aligned}
$$

That μ^* is BC follows from the fact that $\widehat{\mu}$ is π^*-BC: for any $a_S^m \in \operatorname{supp}(\sigma^*)$ so that $\sigma^*(a_S|m) > 0$, we have

$$
\frac{\sigma^*\left(a_S^m|\cdot\right)\mu_0(\cdot)}{\sum_{\theta\in\Theta}\sigma^*\left(a_S^m|\theta\right)\mu_0(\theta)} = \frac{\widehat{\sigma}^*(a_S|m)}{\widehat{\sigma}^*(a_S|m)} \frac{\pi^*(m|\cdot)\,\mu_0(\cdot)}{\sum_{\theta\in\Theta}\pi^*(m|\theta)\,\mu_0(\theta)} = \widehat{\mu}^*(\cdot|m),
$$

where the last equality follows from the fact that $\widehat{\mu}$ is π^*-BC. That $\widehat{\rho}^*$ is π^*-RIC ensures ρ is RIC because $\mu(\cdot|a_S^m) = \widehat{\mu}^*(\cdot|m)$. This proves that $\widehat{W}^* \leq W^*$. To see that $\widehat{W}^* \geq W^*$, observe that we can adopt the proof of Theorem 1 by extending the message space in the auxiliary game to be A_S^M.

Let $(\pi^*, \widehat{\sigma}^*, \widehat{\rho}^*, \widehat{\mu}^*)$ be a solution the auxiliary information design problem. By Lemma 4, we can take π^*-equilibrium to be S-pure. Construct a tuple (σ, ρ, μ) as follows: for all $(a_S, m) \in A_S \times M$ such that $\sum_{\theta\in\Theta}\widehat{\sigma}^*(a_S|m)\,\pi^*(m|\theta)\mu_0(\theta) > 0$,

$$
\begin{aligned}
\sigma^*(a_S^m|\cdot) &:= \widehat{\sigma}^*(a_S|m)\,\pi^*(m|\cdot) \\
\rho^*(\cdot|a_S^m) &:= \widehat{\rho}^*(\cdot|a_S, m), \\
\mu^*(\cdot|a_S^m) &:= \widehat{\mu}^*(\cdot|m).
\end{aligned}
$$

Observe that this tuple satisfies BC and RIC, and moreover, $W(\sigma, \rho) = \widehat{W}^*$. Hence, $\widehat{W}^* = W^*$. ■

Proof of Proposition 1. (i) Let τ^* be the solution to the information design problem. We want to verify that the Sender's optimal action to each $\mu \in \operatorname{supp}(\tau^*)$ is distinct. Toward a contradiction, suppose there exist distinct $\mu, \mu' \in \operatorname{supp}(\tau^*)$ in which the Sender takes action $a_S^* \in A_S$. Since $\max \omega(a_S, \cdot)$ is concave,

$$
\max \omega\left(a_S^*, \lambda\mu + (1-\lambda)\mu'\right) \geq \lambda \max \omega(a_S^*, \mu) + (1-\lambda) \max \omega\left(a_S^*, \mu'\right)
$$

for any $\lambda \in (0,1)$; i.e., inducing $\lambda\mu + (1-\lambda)\mu'$ leads to (weakly) higher payoff for the Sender.

(ii) Recall that the maximum of two convex functions is convex, so that $w(\mu)$ is convex in this case. Standard arguments in information design mean fully revealing the state is optimal. The last condition ensures that the Sender has enough actions to be able to fully reveal the state. ■

A.2 Firm's HR design problem

A.2.1 Theorem 3

Formally, we say that an HR policy $\varphi : \Theta \to \Delta(\{0,1,2\})$ exhibits:

- *over-retention* if $\int_0^{\overline{\theta}_0} \varphi(0|\theta)\mathrm{d}F(\theta) < F(\overline{\theta}_0)$;
- *under-retention* if $\int_{\overline{\theta}_0}^1 \varphi(0|\theta)\mathrm{d}F(\theta) > 0$;
- *efficient retention* if $\int_0^{\overline{\theta}_0} \varphi(0|\theta)\mathrm{d}F(\theta) = F(\overline{\theta}_0)$ and $\int_{\overline{\theta}_0}^1 \varphi(0|\theta)\mathrm{d}F(\theta) = 0$.

These definitions allow for an HR policy to exhibit both over-retention and under-retention.

Toward proving Theorem 3, we first characterise the case in which efficient retention is possible.

Lemma 5. *An optimal HR policy exists with efficient retention if and only if the efficient retention threshold $\overline{\theta}_0$ satisfies $V_2'(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0]) \geq V_1'(\overline{\theta}_0)$ and $V_2(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0]) = V_2'(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0])(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0] - \overline{\theta}_0)$.*

Proof. Let G be a solution to the firm's information design problem and suppose it can be induced by an HR policy that exhibits efficient retention. By Fact 2, the support function p of G equals zero on $[0,\overline{\theta}_0]$. Since $p \geq V$, p must have a kink with a right-derivative greater than $V_1'(\overline{\theta}_0)$. Since $\mathrm{supp}(G) \subseteq [\overline{\theta}_0, 1]$, by Fact 3(ii), p must be tangent to V at some $E \in (\overline{\theta}_0, 1]$. Because V_1 is concave, it must be that $E \in (\overline{\theta}_1, 1]$. The concavity of V_2 means that such a point E, if it exists, must be unique. Finally, observe that condition $V_2(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0]) = V_2'(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0])(\mathbb{E}[\theta|\theta \geq \overline{\theta}_0] - \overline{\theta}_0)$ ensures $E = \mathbb{E}[\theta|\theta \geq \overline{\theta}_0]$. That the converse also holds follows from Fact 2. ■

Let us now prove the key lemma that delivers Theorem 3.

Lemma 6. *No optimal HR policy exhibits under-retention.*

Proof. Let φ be an optimal HR policy that exhibits under-retention. Then, there exists a subset $S \subseteq [\overline{\theta}_0, 1]$ who are let go with strictly positive probability. In the associated information design game, consider the G that is induced by φ. Suppose we modify G by fully revealing the types of workers in S whenever they are to be let go. Since $V(\theta) > 0$ for all $\theta \in S$, the firm can assign them to positions that maximise their output to obtain a strictly positive payoff, rather than a payoff of zero from letting go of them. Moreover, the firm's profits from retained workers are unaffected by this modification of G. Thus, the firm's payoff must be strictly higher with this modification. By Lemma 1, the firm can attain at least this payoff from an HR policy, which contradicts the optimality of φ. ■

We now prove Theorem 3, which says that any optimal HR policy (monotone or otherwise) exhibits over-retention.

Proof of Theorem 3. The first two statements are simply Lemma 5 and 6, respectively. If conditions (4) and (5) do not hold, then no HR policy with efficient retention is optimal. That is, any optimal HR policy exhibits inefficient retention. It must also exhibit no under-retention by Lemma 6. By definitions of efficient retention and under-retention, it follows that any optimal HR policy must exhibit over-retention. ■

A.2.2 Lemma 2

Proof of Lemma 2. Let G denote the simple CDF associated with a tuple $(\theta_0, \theta_1, E_2) \in \mathscr{F}$. Because G is a bi-pooling distribution, by lemma 2 by Arieli et al. (2023), G can be induced by a monotone information structure π. By construction, G induces at most three posterior means, E_0, E_1, E_2, such that $E_0 < E_1 \leq E_2$. Hence, it can be implemented by an HR policy φ by setting $\varphi(j|\theta) = \pi(m_j|\theta)$ for all $\theta \in \Theta$ and $j \in \{0,1,2\}$, where each m_j induces posterior mean E_j. ■

A.2.3 Theorem 4

We say that an HR policy $\varphi : \Theta \to \Delta(\{0,1,2\})$ exhibits:

- *over-promotion* if $\int_0^{\overline{\theta}_1} \varphi(2|\theta)\mathrm{d}F(\theta) > 0$;
- *under-promotion* if $\int_{\overline{\theta}_1}^1 \varphi(2|\theta)\mathrm{d}F(\theta) < 1 - F(\overline{\theta}_1)$;
- *efficient promotion* if $\int_0^{\overline{\theta}_1} \varphi(2|\theta)\mathrm{d}F(\theta) = 0$ and $\int_{\overline{\theta}_1}^1 \varphi(2|\theta)\mathrm{d}F(\theta) = 1 - F(\overline{\theta}_1)$

Note again that these definitions allow for an HR policy to exhibit both over-retention and under-retention. Recall that any simple G can be induced by a monotone HR policy and that, if G is pooling (vs bi-pooling), then it can be induced by a deterministic monotone HR policy.

We first note that if the degenerate distribution of posterior means (i.e., F) is optimal, then there is over-promotion (resp. under-promotion) if the firm assigns all workers to $j = 2$ (resp. $j < 2$) under the prior belief. Toward proving Theorem 4, let us consider the case in which $p_2^o > p_1^o$.

Proposition 8. *Suppose $p_2^o > p_1^o$. Suppose no optimal monotone HR policy is deterministic. Then, optimal monotone HR policies can exhibit only over-promotion, only under-promotion, or both over- and under-promotion.*

Proof. If no optimal monotone HR policy is deterministic, then optimal G must bi-pool the interval $[\theta_0^b, 1]$ to $\mathbb{E}[\theta|\theta_0^b < \theta < \theta_1^b]$ and $E_2 \in (\mathbb{E}[\theta|\theta > \theta_0^b], \mathbb{E}[\theta|\theta \geq \theta_1^b])$. A monotone HR policy that induces G is one that: (i) assigns $j = 2$ with probability one to workers whose types are in $[\theta_1^b, 1]$, and (ii) assigns $j = 1$ with some probability β to workers whose types are in $[\theta_0^b, \theta_1^b]$ and assigns $j = 2$ with complementary probability. If $\overline{\theta}_1 \geq \theta_1^b$, then this monotone HR policy exhibits only over-promotion. However, if $\overline{\theta}_1 < \theta_1^b$, this monotone HR policy exhibits both over-promotion (because it sometimes assigns $j = 2$ to workers whose types are in $[\theta_0^b, \overline{\theta}_1]$) and under-promotion (because it sometimes assigns $j = 1$ to workers whose types are in $[\overline{\theta}_1, \theta_1^b]$). To show that monotone HR policy can exhibit only under-promotion, let $\tilde{t}_1^b$ be such that $\mathbb{E}[\theta|\theta \geq \tilde{t}_1^b] = E_2^b$. If $\mathbb{E}[\theta|\theta_0^b < \theta < \tilde{\theta}_1^b] < E_1^b$, following the same argument as in the proof of Lemma 2 by Arieli et al. (2023), we can induce a bi-pooling distribution by a monotone HR policy that: (i) assigns $j = 1$ with probability one to workers whose types are in $[\theta_0^b, \tilde{\theta}_1^b]$, and (ii) assigns $j = 1$ with some probability β to workers whose types are in $[\tilde{\theta}_1^b, 1]$ and assigns $j = 2$ with complementary probability. In this case, if $\overline{\theta}_1 \leq \tilde{t}_0^b$, there is only under-promotion. ∎

We now focus on the case in which an optimal HR policy is deterministic. To state the result formally, let $\overline{E}_2 := \mathbb{E}[\theta|\theta \geq \overline{\theta}_1]$. When $\overline{E}_2 < E_2^b$, define $\overline{E}_1^e$ be such that

$$V\left(\overline{E}_1^e\right) + V'\left(\overline{E}_1^e\right)\left(\overline{\theta}_1 - \overline{E}_1^e\right) = V\left(\overline{E}_2\right) + V'\left(\overline{E}_2\right)\left(\overline{\theta}_1 - \overline{E}_2\right). \qquad (18)$$

and let

$$\overline{\theta}_0^e = \begin{cases} 0 & \text{if } V\left(\overline{E}_1^e\right) \geq V'\left(\overline{E}_1^e\right)\overline{E}_1^e, \\ \overline{E}_1^e - \frac{V\left(\overline{E}_1^e\right)}{V'\left(\overline{E}_1^e\right)} & \text{otherwise.} \end{cases} \tag{19}$$

As Figure 5 shows, given $\overline{E}_2 < E_2^b$, $\overline{E}_1^e$ is the point such that the tangent to V at $\overline{E}_1^e$ intersects the tangent of V at $\overline{E}_2$ at $(\overline{\theta}_1, V(\overline{E}_2) + V'(\overline{E}_2)(\overline{\theta}_1 - \overline{E}_2))$. Figures 5a and 5b show the cases in which $\overline{\theta}_0^e = 0$ and $\overline{\theta}_0^e > 0$, respectively.

Figure 5: Over- vs under-promotion.

(a) Case: $\overline{\theta}_0^e = 0$.

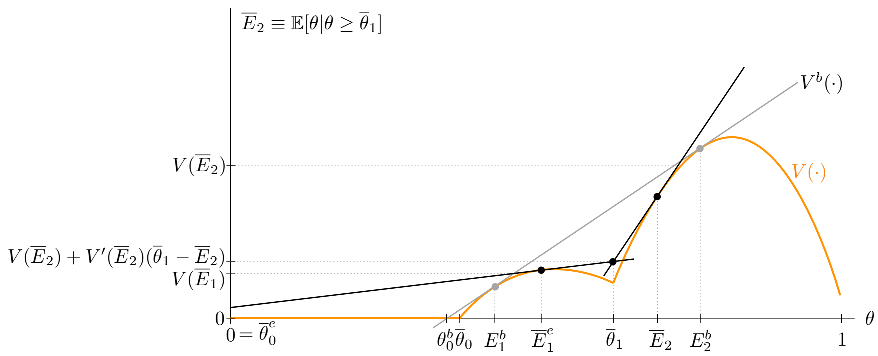

(b) Case: $\overline{\theta}_0^e = \overline{E}_1^e - \frac{V(\overline{E}_1^e)}{V'(\overline{E}_1^e)}$.

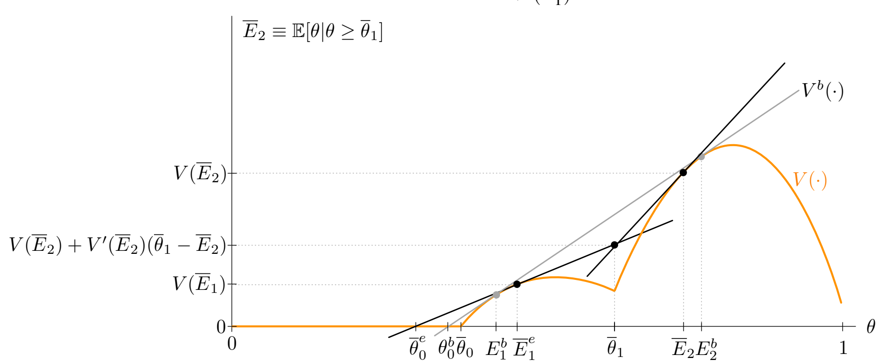

Recall that we say that a deterministic monotone HR policy exhibits over-promotion (resp. under-promotion) if the associated promotion cutoff is below (resp. above) the efficient promotion threshold.

Proposition 9. *Suppose optimal monotone HR policy, φ^*, is deterministic.*

(i) Suppose V has a bi-tangent. Then, φ^* exhibits over-promotion if $E_2^b < \mathbb{E}[\theta|\theta \geq \theta_1^b]$ and either (a) $\overline{E}_2 > E_2^b$ or (b) $\overline{E}_2 < E_2^b$ and $\mathbb{E}[\theta|\overline{\theta}_0^e < \theta < \overline{\theta}_1] > \overline{E}_1^e$; otherwise, φ^*

exhibits under-promotion.

(ii) Suppose no bi-tangent of V exists. Then, φ^* exhibits over-promotion if $\mathbb{E}[\theta|\overline{\theta}_0^e < \theta < \overline{\theta}_1] > \overline{E}_1^e$; otherwise φ^* exhibits under-promotion.

Proof. If V has a bi-tangent and a deterministic monotone HR policy is optimal, it must be the case that (30) does not hold and $E_2^b > \mathbb{E}[\theta|\theta \geq \theta_1^b]$.

Suppose $\overline{E}_2 > E_2^b$. Toward a contradiction, suppose $\theta_1^* \geq \overline{\theta}_1$. By definition of E_2^b as (one of) the bi-tangency points of V and the fact that $V|_{[\overline{\theta}_0,\overline{\theta}_1]}$ and $V|_{[\overline{\theta}_1,1]}$ are concave, the tangent of V at $\overline{E}_2 \in (\overline{\theta}_1, 1)$ must lie strictly above V on $[\overline{\theta}_0, \overline{\theta}_1]$ and equal zero at some $\theta_0 < \overline{\theta}_0$ as can be seen from the figure below.

Figure 6: $\overline{E}_2 > E_2^b$: Over-promotion.

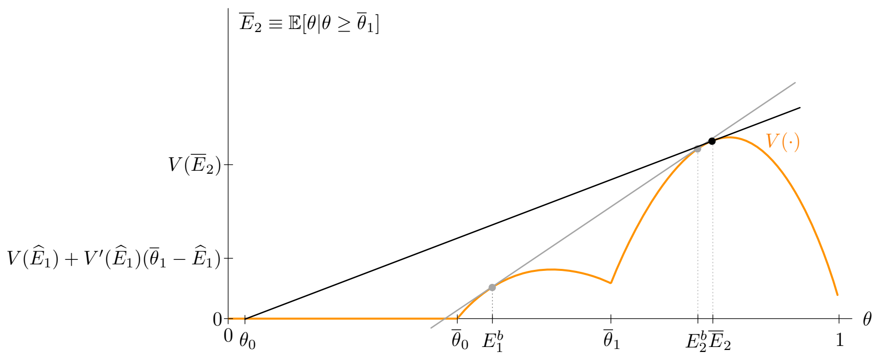

Thus, any G supported by p such that $p|_{[\overline{\theta}_1,1]}$ is given by the tangent on V at $\overline{E}_2$ must pool the mass on $[\theta_0, 1]$ at $\mathbb{E}[\theta|\theta \geq \theta_0]$. But because $\mathbb{E}[\theta|\theta \geq \theta_0] < \overline{E}_2 \equiv \mathbb{E}[\theta|\theta \geq \overline{\theta}_1]$, such a support function is not a tangent to V at $\mathbb{E}[\theta|\theta \geq \theta_0]$; i.e., no such G is optimal. Hence, we must have $\theta_1^* < \overline{\theta}_1$.

Now suppose $\overline{E}_2 < E_2^b$. Consider first the case $\overline{E}_1' := \mathbb{E}[\theta|\overline{\theta}_0^e < \theta < \overline{\theta}_1] > \overline{E}_1^e$. Then, concavity of $V|_{[\overline{\theta}_0,\overline{\theta}_1]}$ implies that the tangent to V at $\overline{E}_1'$ evaluated at $\overline{\theta}_1$ must lie strictly below the intersection of the tangents to V at $\overline{E}_1^e$ and $\overline{E}_2$as can be seen in the figure below.

Figure 7: $\overline{E}_2 < E_2^b$ and $\overline{E}_1' > \overline{E}_1^e$: Over-promotion.

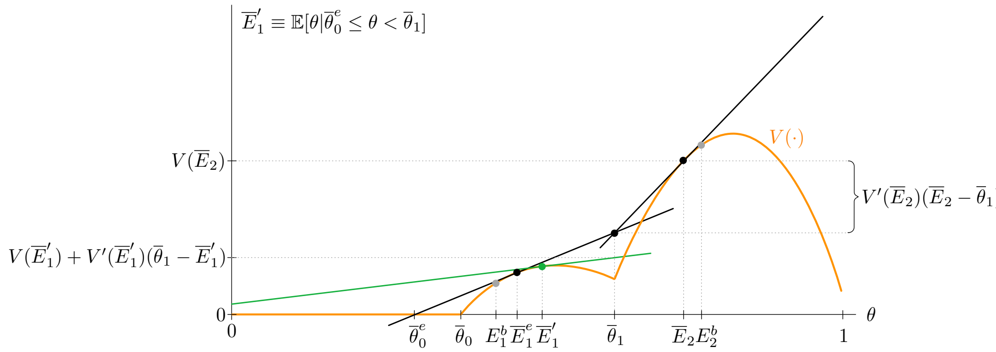

But this means that

$$\begin{aligned}
&\frac{\mathrm{d}\widehat{V}\left(\overline{\theta}_0^e, \overline{\theta}_1, \overline{E}_2\right)}{\mathrm{d}\theta_1} \\
&= f\left(\overline{\theta}_1\right)\left[\left(V\left(\overline{E}_1'\right) + V'\left(\overline{E}_1'\right)\left(\overline{\theta}_1 - \overline{E}_1'\right)\right) - \left(V\left(\overline{E}_2\right) - V'\left(\overline{E}_2\right)\left(\overline{E}_2 - \overline{\theta}_1\right)\right)\right] \qquad (20)
\end{aligned}$$

is strictly negative and so $\theta_1^* < \overline{\theta}_1$; i.e., there is over-promotion. If $\overline{E}_1' < \overline{E}_1^e$, analogous arguments yields that $\theta_1^* > \overline{\theta}_1$ (see figure below) so that there is under-promotion.

Figure 8: $\overline{E}_2 < E_2^b$ and $\overline{E}_1' < \overline{E}_1^e$: Under-promotion.

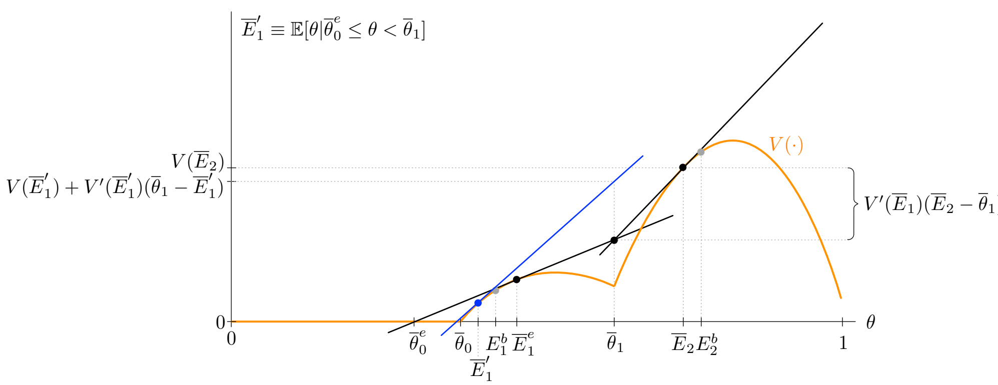

The case for when $E_2^b < \mathbb{E}[\theta|\theta \geq \theta_1^b]$ follows immediately from Theorem 5.

For the case in which V does not have a bi-tangent, observe that the argument for why there is over- or under-promotion when $\overline{E}_2 < E_2^b$ did not rely on bi-tangents. Hence, the same argument still goes through even when V does not have a bi-tangent. ■

We note that the condition $\overline{E}_2 > E_2^b$ holds if there is a sufficient mass of workers whose types are above E_2^b. The condition $\mathbb{E}[\theta|\overline{\theta}_0^e \leq \theta < \overline{\theta}_1] > \overline{E}_1^e$ holds if there is a sufficient mass of workers whose types are in $[\overline{E}_1^e, \overline{\theta}_1]$.

Proof of Theorem 4. Case: $p_2^o > p_1^o$. Proposition 8 deals with the case when no optimal HR policy is optimal, and Proposition 9 deals with the case when optimal HR policy is deterministic.

Case: $p_2^o \leq p_1^o$. From Theorem 6 in Appendix C, the only interesting case when $p_1^o \geq p_2^o$ is if $\mathbb{E}[\theta] \in [E_1^b, E_2^b]$. In this case, the analysis of whether there is over- or under-promotion is the same as the case in which $p_2^o > p_1^o$ when the firm does not let go of any workers. ■

A.2.4 Proposition 2

Proof of Proposition 2. We first note that any simple G associated with $(\theta_0, \theta_1, E_2)$ that is a pooling distribution can be induced by a deterministic monotone HR policy that lets go of workers of types in $[0, \theta_0)$, retains in the junior position workers of types in $[\theta_0, \theta_1)$, and promotes to the senior position workers of types in $[\theta_1, 1]$. However, it follows from Lemma 3 by Arieli et al. (2023) (along with an argument analogous to the proof of Lemma 2 above) that G cannot be induced by any deterministic monotone HR policy if and only if G has a bi-pooling interval on $[\theta_0, 1]$. By Fact 3(ii), G has a bi-pooling interval on $[\theta_0, 1]$ if there exists two bi-tangecy points $E_1^b := \mathbb{E}[\theta|\theta_0 \leq \theta < \theta_1]$ and $E_2^b := E_2$ to V that satisfy

$$\theta \leq E_1^b \leq \mathbb{E}[\theta|\theta > \theta_0] \leq E_2^b \leq \mathbb{E}[\theta|\theta \geq \theta_1] \leq 1.$$

Letting $\theta_0^b := \theta_0$ gives condition (iii). For E_1^b and E_2^b to be bi-tangency points of V, it must be that $V'(E_1^b) = V'(E_2^b)$, which corresponds to condition (i), and that the bi-tangent to V equals zero at θ_0^b, which corresponds to condition (ii). ■

A.2.5 Comparative statics

Proof of Proposition 4. Suppose $\overline{E}_2 < E_2^b$. Fix comparable primitives (v_2, c, F) and $(v_2, \tilde{c}, F)$ that satisfy the hypothesis of the proposition and $\tilde{c}'(\overline{E}_2) > c'(\overline{E}_2)$. Let $\tilde{V}(\cdot) := [\max_{j\in\{1,2\}} v_j(\cdot) - \tilde{c}(\cdot)]_+$. Observe that V and $\tilde{V}$ coincide on $[0, \overline{\theta}_1] \cup \{\overline{E}_2, 1\}$. Moreover, $\tilde{V}'(\overline{E}_2) = k_2 - \tilde{c}'(\overline{E}_2) < k_2 - c'(\overline{E}_2) = V'(\overline{E}_2)$ and $\tilde{V}(\overline{E}_2) = V(\overline{E}_2)$ mean that the tangent to $\tilde{V}$ at $\overline{E}_2$ cuts $\overline{\theta}_1$ above the tangent to V at $\overline{E}_2$ (see figure below which shows the case when $\mathbb{E}[\theta|\overline{\theta}_0^e < \theta < \overline{\theta}_1] > \overline{\theta}_1^e$).

Figure 9: Comparative statics on c.

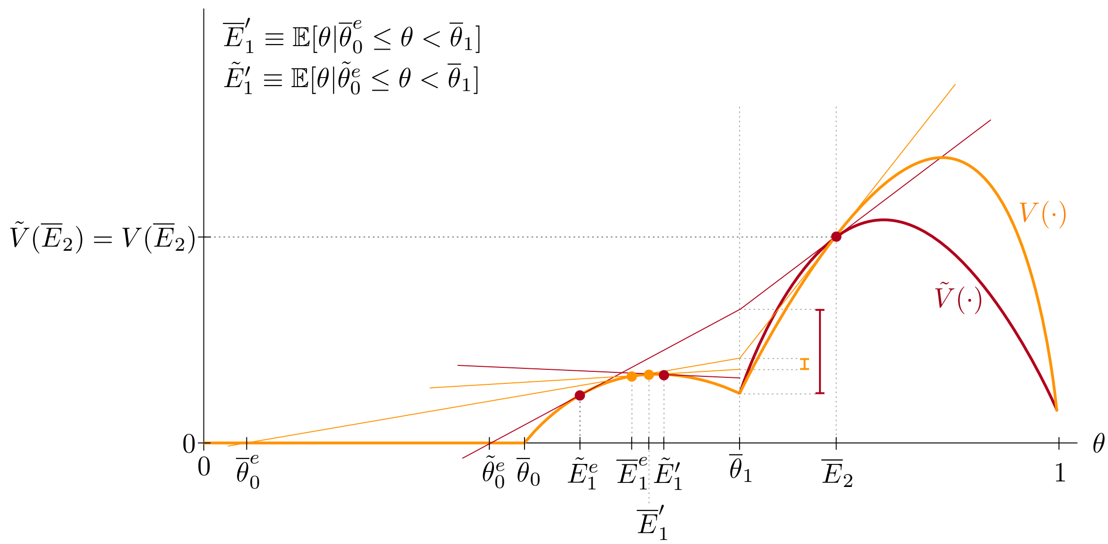

Given concavity of V on $[\overline{\theta}_0, \overline{\theta}_1]$, we must have $\tilde{E}_1^e < \overline{E}_1^e$, where $\tilde{E}_1^e$ solves (18) (with V replaced with $\tilde{V}$). In turn, $\tilde{\theta}_0^e$ defined via (19) (with $\overline{E}_1^e$ replaced with $\tilde{E}_1^e$) must be such that $\tilde{\theta}_0^e \geq \overline{\theta}_0^e$. Therefore, $\tilde{E}_1' := \mathbb{E}[\theta|\tilde{\theta}_0^e < \theta < \overline{\theta}_1] \geq \overline{E}_1' := \mathbb{E}[\theta|\overline{\theta}_0^e < \theta < \overline{\theta}_1]$. Hence, the derivative (20) must decrease. That is, over-promotion becomes more likely. ■

Proof of Proposition 5. Recall that $\overline{\theta}_1 \equiv \frac{s_1 - s_2}{k_2 - k_1}$ and consider an increase in k_2 to $\tilde{k}_2 = k_2 + \varepsilon$ for some $\varepsilon > 0$. Then, to ensure that the efficient promotion threshold remains unchanged, we need

$$\overline{\theta}_1 = \frac{s_1 - \tilde{s}_2}{\tilde{k}_2 - k_1} \Leftrightarrow \tilde{s}_2 = s_1 - \overline{\theta}_1 (k_2 - k_1) - \varepsilon \overline{\theta}_1 = s_2 - \varepsilon \overline{\theta}_1.$$

Let $\tilde{V}$ denote the firm's value function under $(\tilde{k}_2, \tilde{s}_2)$:

$$\tilde{V}(\theta) = \begin{cases} V(\theta) & \text{if } \theta \leq \overline{\theta}_1, \\ \tilde{s}_2 + \tilde{k}_2 \theta - c(\theta) = v_2(\theta) - c(\theta) + \varepsilon \left(\theta - \overline{\theta}_1\right) & \text{if } \theta > \overline{\theta}_1. \end{cases}$$

Note that the tangent to $\tilde{V}$ at $\overline{E}_2 \equiv \mathbb{E}[\theta|\theta \geq \overline{\theta}_1]$ is given by

$$\begin{aligned} \tilde{V}'\left(\overline{E}_2\right)\left(\theta - \overline{E}_2\right) + \tilde{V}\left(\overline{E}_2\right) &= \left(\tilde{k}_2 - c'\left(\overline{E}_2\right)\right)\left(\theta - \overline{E}_2\right) + \tilde{v}_2\left(\overline{E}_2\right) - c\left(\overline{E}_2\right) \\ &= \left(k_2 - c'\left(\overline{E}_2\right)\right)\left(\theta - \overline{E}_2\right) + \varepsilon\left(\theta - \overline{E}_2\right) \\ &\quad + v_2\left(\overline{E}_2\right) - c\left(\overline{E}_2\right) + \varepsilon\left(\overline{E}_2 - \overline{\theta}_1\right) \\ &= V'\left(\overline{E}_2\right)\left(\theta - \overline{E}_2\right) + V\left(\overline{E}_2\right) + \varepsilon\left(\theta - \overline{\theta}_1\right). \end{aligned}$$

Hence, the tangent to $\tilde{V}$ at $\overline{E}_2$ and the tangent to V at $\overline{E}_2$ intersect exactly at $\overline{\theta}_1$. But then the derivative (20) does not change. This is shown in the figure below.

Figure 10: Comparative statics on v_2.

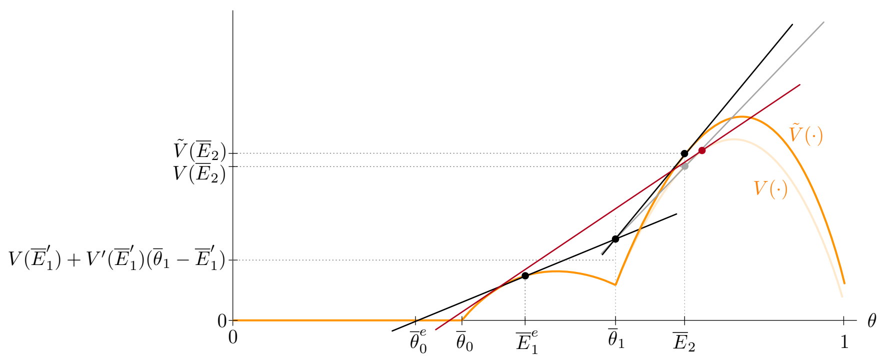

Note further that since $\tilde{V}'|_{[\overline{\theta}_1,1]} > \tilde{V}|_{[\overline{\theta}_1,1]}$, the bi-tangency points of $\tilde{V}$ (if they exist), denoted $\tilde{E}_1^b$ and $\tilde{E}_2^b$, satisfy $\tilde{E}_1^b < E_1^b$ and $\tilde{E}_2^b > E_2^b$. Because $\overline{E}_2 < E_2^b$, it must be $\overline{E}_2 < \tilde{E}_2^b$ so that the nature of inefficiency remains unchanged. ■

The following is a well-known result that tells us that if two CDFs F and H are ordered according to their likelihood ratios, then their conditional distributions are first-order stochastically ordered.[46]

Fact 1. *Let F and H be two continuous CDF with support $[0,1]$. Then,*

$$F(\cdot|a < X < b) \geq H(\cdot|a < X < b) \;\forall a,b \in [0,1] : b > a$$

if and only if $\frac{h(\cdot)}{f(\cdot)}$ is increasing.

It follows that if F and H are two prior distributions of workers' types and that $\frac{h(\cdot)}{f(\cdot)}$ is increasing (i.e., H dominates F in the likelihood ratio order), then $\mathbb{E}_H[\theta|\theta_0 < \theta < \theta_1] \geq \mathbb{E}_F[\theta|\theta_0 < \theta < \theta_1]$ for any $\theta_0, \theta_1 \in [0,1]$ with $\theta_1 > \theta_0$.

Proof of Proposition 6. By definition, that $\tilde{F}$ dominates F in the likelihood ratio order means $\frac{\tilde{f}(\cdot)}{f(\cdot)}$ is increasing. Hence, by Fact 1, $\overline{E}_2(\tilde{F}) := \mathbb{E}_{\tilde{F}}[\theta|\theta \geq \overline{\theta}_1] \geq \overline{E}_2(F) := E_F[\theta|\theta \geq \overline{\theta}_1]$. Moreover, because V is concave on $[\overline{\theta}_1, 1]$,

$$V\left(\overline{E}_2\left(\tilde{F}\right)\right) - V'\left(\overline{E}_2\left(\tilde{F}\right)\right)\left(\overline{E}_2\left(\tilde{F}\right) - \overline{\theta}_1\right) > V\left(\overline{E}_2(F)\right) - V'\left(\overline{E}_2(F)\right)\left(\overline{E}_2(F) - \overline{\theta}_1\right).$$

[46] See, for example, Shaked and Shanthikumar (2007), Theorem 1.C.5.

Define $\overline{\theta}_0^e(\tilde{F})$ and $\overline{\theta}_0^e(F)$ via (19) respectively for $\overline{E}_2(\tilde{F})$ and $\overline{E}_2(F)$. By concavity of V on $[\overline{\theta}_0, \overline{\theta}_1]$, we must have $\overline{\theta}_0^e(\tilde{F}) > \overline{\theta}_0^e(F)$. Hence,

$$\mathbb{E}_F[\theta|\overline{\theta}_0^e(F) < \theta < \overline{\theta}_1] \leq \mathbb{E}_F[\theta|\overline{\theta}_0^e(\tilde{F}) < \theta < \overline{\theta}_1] \leq \mathbb{E}_{\tilde{F}}[\theta|\overline{\theta}_0^e(\tilde{F}) < \theta < \overline{\theta}_1],$$

where the last inequality follows from Fact 1. Hence, the derivative (20) must be smaller under $\tilde{F}$ than under F. That is, over-promotion is more likely under $\tilde{F}$. ■

A.3 Partial commitment

Proof of Lemma 3. Fix an HR policy, φ. The result follows if we can prove the claim: there exists measurable subsets $S, S' \subseteq \theta$ such that (i) $F(S) = F(S') = \ell > 0$; (ii) $\sup S' \leq \widehat{\theta} := \inf S$; and (iii) $\varphi(\cdot|\theta) = \delta_{j'}$ for all $\theta \in S$ and $\varphi(\cdot|\theta') = \delta_j$ for all $\theta' \in S'$ with $j > j'$; then the firm's direct payoff is weakly greater from $\widehat{\varphi}$ that equals φ except that $\varphi(\cdot|\theta) = \delta_j$ for all $\theta \in S$ and $\varphi(\cdot|\theta') = \delta_{j'}$ for all $\theta' \in S'$. Condition (i) means that S and S' have the same size, (ii) means that S' consists of workers whose types are below that of any worker in S, and (iii) means that φ assigns a higher position to lower types of workers who are in S'. The claim says that the firm is better off by assigning lower types of workers to a lower position. To prove the claim, observe that $v_j(\theta)$ has increasing differences in (j, θ) since, for any $\theta, \theta' \in \theta$ such that $\theta' > \theta$, we have

$$v_2\left(\theta'\right) - v_2\left(\theta\right) = k_2\left(\theta' - \theta\right) > k_1\left(\theta' - \theta\right) = v_1\left(\theta'\right) - v_1\left(\theta\right) > 0 = v_0\left(\theta'\right) - v_0\left(\theta\right).$$

In particular, this means that, for any $j > j'$,

$$v_j\left(\theta\right) - v_{j'}\left(\theta\right) > v_j\left(\theta'\right) - v_{j'}\left(\theta'\right).$$

This observation, together with conditions (i) and (ii), means that

$$\begin{aligned}
\int_S \left[v_j\left(\theta\right) - v_{j'}\left(\theta\right)\right] \mathrm{d}F\left(\theta\right) &\geq \int_S \left[v_j\left(\widehat{\theta}\right) - v_{j'}\left(\widehat{\theta}\right)\right] \mathrm{d}F\left(\theta\right) \\
&= \left[v_j\left(\widehat{\theta}\right) - v_{j'}\left(\widehat{\theta}\right)\right] \ell \\
&= \int_{S'} \left[v_j\left(\widehat{\theta}\right) - v_{j'}\left(\widehat{\theta}\right)\right] \mathrm{d}F\left(\theta\right) \\
&\geq \int_{S'} \left[v_j\left(\theta\right) - v_{j'}\left(\theta\right)\right] \mathrm{d}F\left(\theta\right).
\end{aligned}$$

In fact, what we have shown is

$$\int_S v_j(\theta)\,\mathrm{d}F(\theta) + \int_{S'} v_{j'}(\theta)\,\mathrm{d}F(\theta) \geq \int_S v_{j'}(\theta)\,\mathrm{d}F(\theta) + \int_{S'} v_j(\theta)\,\mathrm{d}F(\theta);$$

i.e., the payoff from swapping assignments is greater. To see why the claim implies the result, observe that if φ is not monotone, then we can find $S, S' \subseteq \theta$ that satisfies the conditions of the claim. Now suppose that φ is monotone and involves randomisation, then there will exist measures of workers who are sometimes assigned to a higher position and otherwise assigned to a lower position. However, the firm can improve its payoff by assigning the workers to a higher position without changing the marginal distribution via the claim. ■

B Appendix: Some applications of the general solution approach

Our general approach to solving the base game is useful in settings where the associated auxiliary information design can be solved. For example, whenever the number of states is "small"—particularly when the state is binary—one can appeal to the concavification approach (Kamenica and Gentzkow, 2011) to solve the information design problem. Another case in which a well-established method exists is the "linear" case (Gentzkow and Kamenica, 2016; Kolotilin et al., 2017; Kolotilin, 2018; Dworczak and Martini, 2019; Arieli et al., 2023) in which the state is one-dimensional and posterior distributions are summarised by their mean. We briefly describe several applications that fall into one of these two settings.[47]

B.1 "Binary" signalling game

Consider the case in which the sets Θ, A_S and A_R are all binary: $\Theta := \{0,1\}$, $A_S := \{a_S^1, a_S^2\}$, and $A_R := \{a_R^1, a_R^2\}$. Let μ_0 denote the prior probability that the state is 1. For each $k \in \{1,2\}$, the Receiver's best response to a_S^k is either always taking one action or there is a cutoff $\mu^k \in (0,1)$ such that the Receiver is indifferent between the two actions. Consequently, $\omega(a_S^k, \cdot)$ is either constant or piecewise affine with a single discontinuity at μ^k. In this case, either the uninformative information structure is optimal (i.e., $\tau^* = \delta_{\mu_0}$) or the

[47] Our approach can easily be extended to study signalling games in which players take (non-signalling) actions after the receiver's action (e.g., Milgrom and Roberts, 1986).

Figure 11: Beer and Quiche game (Fudenberg and Tirole, 1991).

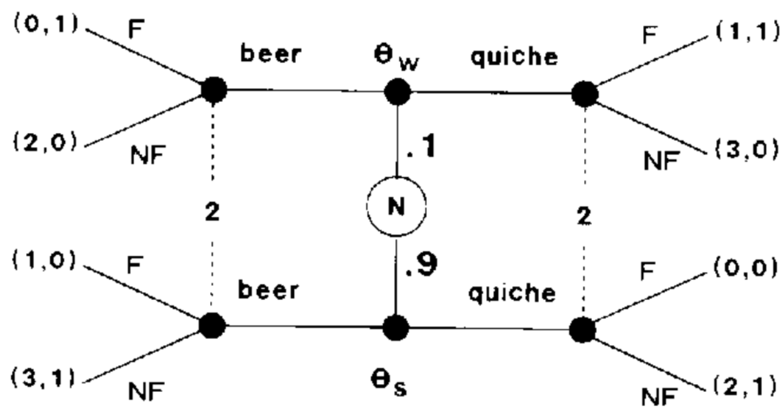

solution to the information design problem, τ^*, has support $\{\mu_\ell, \mu_r\}$ where $\mu_0 \in (\mu_\ell, \mu_r)$. In the first case, an optimal exPBE is one in which the Sender simply commits to taking an action that is optimal under the prior belief given that the Receiver best responds to that action also under the prior belief. In the second case, the Sender can attain the same payoff as in the auxiliary information design game in the signalling game if the Sender's optimal action given μ_ℓ and μ_r, denoted a_S^ℓ and a_S^r, are distinct. If this is the case, then there is a mixed signalling strategy σ in the signalling game that induces posterior beliefs at μ_ℓ after observing a_S^ℓ and μ_r after observing a_S^r pinned down by

$$\mu_\ell = \frac{\sigma\left(a_S^\ell|1\right)\mu_0}{\sigma\left(a_S^\ell|1\right)\mu_0 + \sigma\left(a_S^\ell|0\right)(1-\mu_0)},\ \mu_r = \frac{\left[1-\sigma\left(a_S^\ell|1\right)\right]\mu_0}{\left[1-\sigma\left(a_S^\ell|1\right)\right]\mu_0 + \left[1-\sigma\left(a_S^\ell|0\right)\right](1-\mu_0)}.$$

If $\mu_\ell, \mu_r \in \{0,1\}$ does not hold (i.e., τ^* is not fully revealing), the optimal signalling strategy is a semi-separating strategy (i.e., either $\mu_\ell = 0$ and $\mu_r < 1$, or $\mu_\ell > 0$ and $\mu_r = 1$).

As a particular example, we study the "Beer and Quiche game" (Cho and Kreps, 1987). Let $\theta := \{\theta_w, \theta_s\}$, $A_S := \{beer, quiche\}$, $A_R := \{F, N\}$, and denote the common prior belief that the sender is θ_s as $\mu_0 \in [0,1]$.[48] The payoffs are as given in the Figure 11. Let $\mu \in [0,1]$ denote the belief that $\theta = \theta_s$. In this game, the Receiver's best response contains F

[48]The ability for the Sender to commit might seem more appropriate if we interpreted the game as one of repeated entry deterrence: Sender's type representing the incumbent firm's type (*s*trong or *w*eak), *beer* (resp. *quiche*) representing the strong (resp. weak) incumbent's preferred action, and F and N respectively representing whether an entrant enters or not enter the market.

Figure 12: Sender's value function in the Beer and Quiche game.

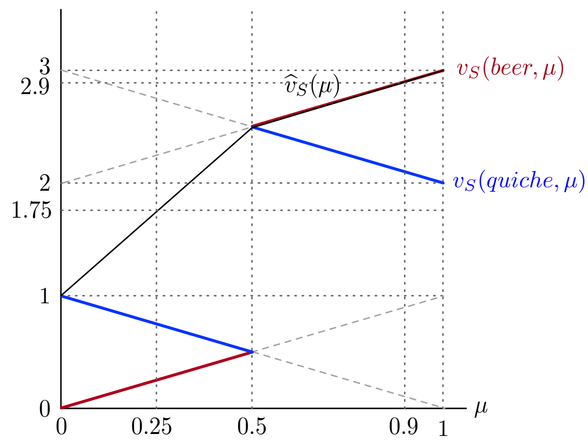

whenever $\mu \leq \frac{1}{2}$ and contains NF whenever $\mu \geq \frac{1}{2}$. Using this, one can then compute that

$$\omega(beer, \mu) = \begin{cases} \mu & \text{if } \mu < \frac{1}{2}, \\ 2+\mu & \text{if } \mu \geq \frac{1}{2}, \end{cases} \quad \omega(quiche, \mu) = \begin{cases} 1-\mu & \text{if } \mu < \frac{1}{2}, \\ 3-\mu & \text{if } \mu \geq \frac{1}{2}, \end{cases}$$

so that

$$w(\mu) = \max\{\omega(beer, \mu), \omega(quiche, \mu)\} = \begin{cases} 1-\mu & \text{if } \mu < \frac{1}{2}, \\ 2+\mu & \text{if } \mu \geq \frac{1}{2}. \end{cases}$$

Figure 12 below shows $\omega(beer, \mu)$, $\omega(quiche, \mu)$, $w(\mu)$ and $\widehat{w}(\mu)$.

Let us first consider the case when $\mu_0 = 0.9$ as studied by Cho and Kreps (1987). They show that the game then has two PBE: one in which the sender pools on *beer* and the other in which she pools on *quiche*. Cho and Kreps (1987) argue that the pooling equilibrium in which the sender chooses quiche is unreasonable based on their *intuitive criterion*. When the Sender can commit to signalling actions, her optimal signalling strategy involves the Sender pooling on *beer* but not on *quiche*. Thus, in this case, the Sender can attain her commitment payoff as a PBE that satisfies the intuitive criterion.

Now suppose $\mu_0 = 0.25$. In this case, there is a unique PBE given by

$$\sigma(\cdot|\theta) = \begin{cases} \delta_{beer} & \text{if } \theta = \theta_S, \\ \frac{1}{3}\delta_{beer} + \frac{2}{3}\delta_{quiche} & \text{if } \theta = \theta_w, \end{cases} \quad \rho(\cdot|a_S) = \begin{cases} \delta_F & \text{if } a_S = quiche, \\ \frac{1}{2}\delta_F + \frac{1}{2}\delta_N & \text{if } a_S = beer, \end{cases}$$

and $\mu(quiche) = 0$ and $\mu(beer) = \frac{1}{2}$,[49] and the Sender's payoff in the optimal PBE is 1.25. From the figure, optimal PBE also induces the same posterior beliefs as in the PBE above but the commitment payoff is higher at 1.75 (because the Sender can choose the Receiver's tie-breaking rule). In this example, the Sender's optimal commitment strategy coincides with her strategy in an optimal PBE; however, in general, the optimal strategies for the Sender will differ depending on whether she can commit to signalling actions.

B.2 Security Issuance

Consider the following example from Gibbons (1992), which itself is based on Myers and Majluf (1984). Consider an entrepreneur who requires an outside investment to undertake an attractive new investment. The profitability of the existing business, θ, can either be $h > 0$ or $\ell \in (0,h)$ and is distributed according to μ_0. An investment costs I and yields a payoff of Y. Let $r > 0$ denote the interest rate and suppose that $R > I(1+r)$; i.e., investment is always profitable. The entrepreneur offers an equity stake $a_S \in [0,1]$ to the investor who can either accept ($a_R = 1$) or reject ($a_R = 0$) the offer. If the investor rejects, the entrepreneur gets a payoff of π and the investor's payoff is $I(1+r)$. If the investor accepts, the entrepreneur gets a payoff of $(1-a_S)(\pi+Y)$ and the investor's payoff is $a_S(\pi+Y)$.

Suppose the entrepreneur can commit to an offer strategy $\sigma : \Theta \to \Delta([0,1])$. Denote the posterior mean of θ after observing an offer a_S as E, then

$$\omega\left(a_S, E\right) = \begin{cases} E & \text{if } a_S < \frac{I(1+r)}{E+Y}, \\ \left(1-a_S\right)\left(E+Y\right) & \text{if } a_S \geq \frac{I(1+r)}{E+Y}, \end{cases}$$

where we note that $\omega(\cdot, E)$ has a discontinuous jump up at $a_S^*(E) = \frac{I(1+r)}{E+Y}$ and decreases thereafter. Hence,

$$w\left(E\right) = \max_{a_S \in [0,1]} \omega_S\left(a_S, E\right) = \omega\left(\frac{I\left(1+r\right)}{E+Y}, E\right) = E + Y - I\left(1+r\right).$$

Since w is linear in E, the management is indifferent across all information structures. In particular, an ex ante optimal signalling strategy is to always offer $a_S^* = \frac{I(1+r)}{\mathbb{E}[\theta]+Y}$ in which case $W^* = \mathbb{E}[\theta] + Y - I(1+r)$.[50]

[49] Throughout, we use δ_x to denote a Dirac measure on a set X with mass on x.

[50] The result does not depend on the distribution of θ. We also note that if the entrepreneur is risk averse, then pooling is strictly optimal.

One can show that a pooling PBE exists only if

$$\frac{I(1+r)}{\mathbb{E}[\theta]+Y} \leq \frac{Y}{H+Y}$$

and the equilibrium stake offered is a_S^* as above. However, if the above inequality does not hold, then a pooling equilibrium does not exist but a separating equilibrium exists in which ℓ-type offers $s_\ell = \frac{I(1+r)}{\ell+r} < a_S^*$ and h-type offers $s_H < \frac{I(1+r)}{H+Y}$ and investor rejects.

The analysis above suggests that one way to avoid inefficiency associated with a separating equilibrium is to ensure that the entrepreneur can commit to offer strategies. We note that both the investor and the social planner have an incentive to ensure that the entrepreneur can commit.

B.3 Education as signalling

We apply our approach to study Spence (1973)'s job market signalling model. The worker is the Sender who chooses effort $a_S \in \mathbb{R}_+$ and the Receiver consists of firms that compete by making wage offers, $a_R \in \mathbb{R}_+$. Payoffs are: $u_S(\theta, a_S, a_R) = a_R - c(a_S, \theta)$, where $c_{a_S} > 0$, $c_{a_S a_S} > 0$, $c_\theta > 0$, $c_{a_S\theta} < 0$, and $u_R(\theta, a_S, a_R) = \theta(1+\gamma a_S) - a_R$. Observe that $\gamma = 0$ corresponds to the case in which effort is unproductive. Competition implies that $a_R^*(\mu) = (1+\gamma a_S)\mathbb{E}_\mu[\theta]$. We therefore have

$$w(\mu) = \max_{a_S} \mathbb{E}_\mu\left[(1+\gamma a_S)\theta - c(a_S, \theta)\right].$$

If Θ is binary, then the objective is essentially affine in μ. Hence, as a maximum of linear functions, w is convex so that fully revealing the state is optimal. Therefore, if different actions are optimal at the two extreme beliefs, we have $\widehat{W}^* = W^*$.

If c is affine in θ (e.g., $c(a_S, \theta) = (k-\theta)a_S^2$ for some $k > 0$), then the objective is linear in μ. As a maximum of linear functions, w is convex so that fully revealing the state is optimal. If a_S is finite, then w is piecewise linear (each piece corresponding to some $a_S \in A_S$) and it is also optimal to pool states "piecewise". Linearity means $B(a_S)$ is convex and (weakly) concave so that we have $\widehat{W}^* = W^*$. If a_S is continuous, w can be strictly convex but we would still have $\widehat{W}^* = W^*$ because $B(a_S)$ must be a singleton to generate strict convexity.

If c is mean-measurable so that $\mathbb{E}_\mu[c(a_S, \theta)] = c(a_S, \mathbb{E}_\mu[\theta])$, the objective, viewed as a function of $\mathbb{E}_\mu[\theta]$, is concave if (and only if) c is convex in θ. Since the price function

(Dworczak and Martini, 2019) must be convex, it can only intersect $c(a_S,\cdot)$ at a unique point (if at all) for each a_S. Hence, $\widehat{W}^* = W^*$.

B.4 Insurance market

We apply our approach to Rothschild and Stiglitz (n.d.). Suppose θ represents the probability of an accident that incurs a loss of $L > 0$ and that θ is distributed according to a full-support CDF F on $[0,1]$. A consumer has a strictly increasing and strictly concave utility function over money with initial wealth of $y > 0$ while firms are risk neutral and compete by offering full-insurance contracts at a price $a_R \in \mathbb{R}_+$. The consumer with belief μ buys the full-insurance contract at price a_R if

$$\begin{aligned} u(y-a_R) &\geq \int_{[0,1]} [\theta u(y-L) + (1-\theta)u(y)]\,\mathrm{d}\mu(\theta) \\ \Leftrightarrow u(y-a_R) &\geq [u(y-L)-u(y)]\,\mathbb{E}_\mu[\theta] + u(y) \\ \Leftrightarrow \mathbb{E}_\mu[\theta] &\geq \frac{u(y)-u(y-a_R)}{u(y)-u(y-L)}. \end{aligned}$$

Competition among firms with belief μ implies each firm offers the full-insurance contract for $a_R^*(\mu) = \mathbb{E}_\mu[\theta]L$. Letting $h(E) := \frac{u(y)-u(y-EL)}{u(y)-u(y-L)}$, the consumer's value function from inducing posterior mean $E \in [0,1]$ is

$$w(E) = \begin{cases} u(y-EL) & \text{if } E \geq h(E), \\ Eu(y-L) + (1-E)\,u(y) = u(y) - E\,[u(y) - u(y-L)] & \text{otherwise.} \end{cases}$$

Note that $h(0) = 0$, $h(1) = 1$, $h', h'' > 0$ so that $E > h(E)$ for all $E \in (0,1)$. That is, the consumer always purchases insurance in this case and so w is decreasing and concave. Hence, a fully uninformative information structure is optimal—of course, this is equivalent to the consumer committing to always purchasing in the signalling game. Thus, commitment payoff coincides with the pooling PBE.

If the consumer can choose a partial-coverage contract (B,p), where $B \in [0,L]$ is the coverage and p is the price, the consumer buys only if

$$\mathbb{E}_\mu[\theta] \geq \frac{u(y)-u(y-p)}{u(y)-u(y-p)+u(y-L+B-p)-u(y-L)}$$

and firms offer this contract for $p=\mathbb{E}_\mu[\theta]B$. Letting $h_B(E):=\frac{u(y)-u(y-EB)}{u(y)-u(y-EB)+u(y-L+B-EB)-u(y-L)}$, the signs of h' and h'' are now ambiguous (although $h(0)=0$ and $h(1)=1$ still). Thus, it is now possible that for some $E\in(0,1)$ the consumer does not buy a partial-coverage contract (B,p). However, observe that the consumer prefers $B=L$ and the firms are always indifferent to the level of B since they always make zero profits.

C Supplemental Material

C.1 Proof of Lemma 1

We first summarise results by Dworczak and Martini (2019) and Arieli et al. (2023) that are relevant in our proofs for convenience. The following set of facts follows directly from theorems 1 and 2 by Dworczak and Martini (2019).

Fact 2 (Dworczak and Martini, 2019)**.** *A CDF G is a solution to the firm's information design problem,* (2)*, if there exists a function* $p:\Theta\to\mathbb{R}$ *such that*

$$p\geq V \text{ and } p \text{ is convex,} \tag{21}$$

$$\operatorname{supp}(G)\subseteq\{E\in\Theta : V(E)=p(E)\}, \tag{22}$$

$$\int_\Theta p(E)\,\mathrm{d}G(E)=\int_\Theta p(\theta)\,\mathrm{d}F(\theta), \tag{23}$$

$$G\in\mathrm{MPC}(F). \tag{24}$$

Moreover, if V is Lipschitz continuous, then the converse also holds; i.e., if a CDF G solves (2)*, then there exists a function* $p:\Theta\to\mathbb{R}$ *that satisfies* (21)–(24).

Given a pair (G,p) that satisfies (21)–(24), we refer to p as a *support function of G*.[51]

Following Arieli et al. (2023), we call a mean-preserving contraction G of F, $G\in\mathrm{MPC}(F)$, a *bi-pooling* distribution (with respect to F) if G can be induced as a distribution of posterior means by an information structure that, for some collection (indexed by $\mathscr{I}$) of pairwise disjoint open intervals $\bigcup_{I\in\mathscr{I}}I\subseteq\Theta$: (i) partitions Θ into $\bigcup_{I\in\mathscr{I}}I$ and its complement; (ii) whenever $\theta\in\bigcup_{I\in\mathscr{I}}I$, the information structure reveals which interval the state is in and provides an additional binary signal; and (iii) the sender fully reveals the state whenever the state is in the complement of $\bigcup_{I\in\mathscr{I}}I$. If a nontrivial binary signal is provided

[51] Dworczak and Martini (2019) refer to p instead as a price function for reasons that are not relevant for our paper.

in an interval $I \in \mathscr{I}$, we call the interval I a *bi-pooling* interval; if only a trivial binary signal is provided, then we call I a *pooling* interval. We call an interval $I \in \mathscr{I}$ in which the sender fully reveals the state a *separating* interval. If all intervals are pooling intervals, we refer to G as a *pooling* distribution.

The following set of facts is a combination of Lemma 1, Theorem 1 and Proposition 3 by Arieli et al. (2023).

Fact 3 (Arieli et al., 2023)**.**

(i) A bi-pooling distribution G solves (2)*.*

(ii) Suppose V is upper semicontinuous and there exists a collection of $J \in \mathbb{N}$ points $\{\overline{\theta}_j\}_{j=-1}^{J}$, with $0 = \overline{\theta}_{-1} < \overline{\theta}_0 < \cdots < \overline{\theta}_J = 1$, such that V is either concave or convex in the interval $(\overline{\theta}_{j-1}, \overline{\theta}_j)$ for every $j \in \{0, 1, \ldots, J\}$. Let $p : \Theta \to \mathbb{R}$ denote the support function of an optimal bi-pooling distribution G. Then, there exists a collection of functions, $\{p_j : [\theta_{j-1}, \theta_j] \to \mathbb{R}\}_{j=0}^{J}$, where $0 = \theta_{-1} < \theta_0 < \ldots < \theta_m = 1$ for some $m \leq 2(J+1) - 1$,[52] such that $p|_{[\theta_{j-1},\theta_j]} = p_j$ for all $j \in \{0, 1, \ldots, J\}$, and each p_j satisfies exactly one of the following conditions:

(a) $p_j = V|_{[\theta_{j-1},\theta_j]}$;

(b) $p_j \geq V|_{[\theta_{j-1},\theta_j]}$, and p_j is affine and tangential to V at $\mathbb{E}[\theta|\theta_{j-1} < \theta < \theta_j]$; or

(c) $p_j \geq V|_{[\theta_{j-1},\theta_j]}$, and p_j affine and tangential to V at two distinct points $\underline{E}_j$ and $\overline{E}_j$ such that

$$\theta_{j-1} \leq \underline{E}_j \leq \mathbb{E}\left[\theta|\theta_{j-1} < \theta < \theta_j\right] \leq \overline{E}_j \leq \mathbb{E}\left[\theta|\theta'_{j-1} < \theta < \theta_j\right] \leq \theta_j, \quad (25)$$

where θ'_{j-1} is defined implicitly via $\mathbb{E}[\theta|\theta_{j-1} < \theta < \theta'_{j-1}] = \underline{E}_j$.

For each of the cases (a) to (c), we call the interval $[\theta_{j-1}, \theta_j]$ a *separating*, *pooling* and *bi-pooling* interval, respectively. Moreover, in case $[\theta_{j-1}, \theta_j]$ is a bi-pooling interval, we refer to p_j (extended to $\mathbb{R}$) as the *bi-tangent* of V on the interval $[\theta_{j-1}, \theta_j]$, and the points $\underline{E}_j$ and $\overline{E}_j$ as the bi-tangency points.

Observe that, for any tuple $(\theta_0, \theta_1, E_2) \in \mathscr{F}$, the CDF $G_{\theta_0,\theta_1,E_2}$ is a bi-pooling distribution in which $[0, \theta_0)$ is a pooling interval and

[52] While Proposition 3 by Arieli et al. (2023) states $m \leq 2(J+1)$ instead, the argument in their proof in fact implies $m \leq 2(J+1) - 1$.

- if $E_2 = \mathbb{E}[\theta|\theta > \theta_0]$ (in which case $\theta_0 = \theta_1$), $[\theta_0, 1]$ is a pooling interval;

- if $E_2 = \mathbb{E}[\theta|\theta \geq \theta_1]$, $[\theta_0, \theta_1)$ and $[\theta_1, 1]$ are pooling intervals;

- otherwise $[\theta_0, 1]$ is a bi-pooling interval because $\mathbb{E}[\theta|\theta_0 < \theta < \theta_1]$ and E_2 are bi-tagency points.[53]

The observation above means that simple CDFs, which are parameterised by a family of tuples $(\theta_0, \theta_1, E_2)$ in $\mathscr{F}$, is a class of bi-pooling distributions.

We first establish the necessary and sufficient conditions for V to have a bi-tangent on $[\overline{\theta}_0, 1]$. We note that, because V is constant and equals zero on $[0, \overline{\theta}_0]$, V always has a bi-tangent on any subinterval of $[0, \overline{\theta}_0]$.

Lemma 7. *V has a bi-tangent on $[\overline{\theta}_0, 1]$ if and only if (i) the tangent to V_1 at $\overline{\theta}_0$ lies above V_2 on $[\overline{\theta}_0, 1]$ and (ii) the tangent to V_2 at 1 lies above V_1 on $[\overline{\theta}_0, 1]$. Moreover, when it exists, the bi-tangent is unique, and it is tangent to V at some $E_1 \in [\overline{\theta}_0, \overline{\theta}_1)$ and $E_2 \in [\overline{\theta}_1, 1]$.*

Proof. Suppose a bi-tangent on $[\overline{\theta}_0, 1]$ exists; i.e., there exist $E_1, E_2 \in [\overline{\theta}_0, 1]$ with $E_2 > E_1$ such that the tangents to V at E_1 and E_2 coincide. Recall that

$$V(\theta) = \begin{cases} V_1(\theta) & \text{if } \theta \in \left[\overline{\theta}_0, \overline{\theta}_1\right], \\ V_2(\theta) & \text{if } \theta \in \left[\overline{\theta}_1, 1\right], \end{cases}$$

and that V_1 and V_2 are strictly concave. It cannot be that $E_1, E_2 \in [\overline{\theta}_0, \overline{\theta}_1]$: strict concavity means that each point on V_j has a unique tangent so that tangents to V_1 at two distinct points E_1 and E_2 cannot coincide. Similarly, it cannot be that $E_1, E_2 \in [\overline{\theta}_1, 1]$. Thus, we must have $E_1 \in [\overline{\theta}_0, \overline{\theta}_1]$ and $E_2 \in [\overline{\theta}_1, 1]$. For each $j \in \{1, 2\}$ and $E \in \Theta$, let $t_j(\cdot|E) : \Theta \to \mathbb{R}$ denote the tangent to V_j at point E; i.e.,

$$t_j(\theta|E) := V_j'(E)(\theta - E) + V_j(E).$$

Since E_1 and E_2 are bi-tangency points, we must have $\ell := V_1'(E_1) = V_2'(E_2)$ and $t_1(\cdot|E_1) =$

[53] To see this, observe that:

$$\theta_0 \leq \mathbb{E}[\theta|\theta_0 < \theta < \theta_1] \leq \underbrace{\mathbb{E}[\theta|\theta_0 < \theta < 1]}_{=\mathbb{E}[\theta|\theta > \theta_0]} \leq E_2 \leq \underbrace{\mathbb{E}[\theta|\theta_1 < \theta < 1]}_{=\mathbb{E}[\theta|\theta \geq \theta_1]} < 1.$$

$t_2(\cdot|E_2)$. Because V_1 is strictly concave and $\overline{\theta}_0 \leq E_1$,

$$t_1\left(\theta|\overline{\theta}_0\right) \geq t_1\left(\theta|E_1\right) = t_2\left(\theta|E_2\right)\ \forall\theta \geq E_1.$$

Because V_2 is strictly concave, V_2 must lie below the tangent line $t_1(\cdot|\overline{\theta}_0)$. Given that $V_2(\theta) \leq V_1(\theta)$ for any $\theta < \overline{\theta}_0$, it follows that $t_1(\cdot|\overline{\theta}_0)$ lies above V_2 on $[\overline{\theta}_0, 1]$. Analogously, because V_2 is strictly concave and $1 \geq E_2$,

$$t_2\left(\cdot|1\right) \geq t_2\left(\theta|E_2\right) = t_1\left(\theta|E_1\right)\ \forall\theta \leq E_2.$$

Because V_1 is strictly concave, V_1 must lie below the tangent line $t_2(\cdot|1)$. Given that $V_2(\theta) \geq V_1(\theta)$ for any $\theta \geq \overline{\theta}_1$, it follows that $t_2(\cdot|1)$ lies above V_1 on $[\overline{\theta}_0, 1]$.

To prove the converse, let $t_1(\cdot)$ and $t_2(\cdot)$ denote the tangent to V_1 at $\overline{\theta}_0$ and tangent to V_2 at 1, respectively. Suppose that t_j lies above $V_{j'}$ on $[\overline{\theta}_0, 1]$ for each distinct $j, j' \in \{1,2\}$. Then, $V_1'(\theta_1) > V_2'(1)$. By definition of $\overline{\theta}_1$, $V_1'(\overline{\theta}_1) < V_2'(\overline{\theta}_2)$. By strict concavity of V_1 and V_2, there must exist a unique bi-tangent. ■

Recall $c_E(\cdot)$ denotes the tangent to c at E. Then, condition (i) of Lemma 7 holds if and only if

$$c\left(\theta\right) - c_{\overline{\theta}_0}\left(\theta\right) \geq v_2\left(\theta\right) - v_1\left(\theta\right)\ \forall\theta \in \left[\overline{\theta}_1, 1\right], \tag{26}$$

where we note that $v_2(\cdot) - v_1(\cdot) \geq 0$ on $[\overline{\theta}_1, 1]$. Condition (ii) of Lemma 7 holds if and only if

$$v_1\left(\theta\right) - v_2\left(\theta\right) \geq c_1\left(\theta\right) - c\left(\theta\right)\ \forall\theta \in \left[\overline{\theta}_0, \overline{\theta}_1\right], \tag{27}$$

where we note that $v_1(\cdot) - v_2(\cdot) \geq 0$ on $[\overline{\theta}_0, \overline{\theta}_1]$. The first condition requires c to exhibit sufficient convexity, while the second condition limits how much convexity c can exhibit. Thus, a bi-tangent exists whenever c is sufficiently convex but not too convex; in other words, a bi-tangent exists whenever the increase in the marginal product of labour is not too little or too much.

We are now ready to prove Lemma 1.

Proof of Lemma 1. By Fact 3(i), a bi-pooling G solves (2). Because V is continuous, linear on $[\overline{\theta}_{-1} := 0, \overline{\theta}_0]$, strictly convex on the intervals $[\overline{\theta}_0, \overline{\theta}_1]$ and $[\overline{\theta}_1, \overline{\theta}_2 := 1]$, and has convex kinks at $\overline{\theta}_0$ and $\overline{\theta}_1$, by Fact 3(ii), any bi-pooling interval can be found by identifying the bi-tangents of V.

Consider first the case in which V has a bi-tangent on $[\overline{\theta}_0, 1]$. By Lemma 7, the bi-tangent is unique and we denote it as $V^b(\cdot)$. Let θ_0 be such that $V^b(\theta_0) = 0$. Given the assumptions on V, $\theta_0^b < \overline{\theta}_0$. Denote the two bi-tangency points as $E_1 \in [\overline{\theta}_0, \overline{\theta}_1)$ and $E_2 \in [\overline{\theta}_1, 1]$. If E_1 and E_2 satisfy (25), then $[\theta_0^b, 1]$ is a bi-pooling interval. Since V is constant and equals zero on $[0, \overline{\theta}_0]$, $\max\{V^b(\cdot), 0\}$ is function satisfying (21)–(24). Moreover, G that can be supported by such a p is associated with $(\theta_0, \theta_1, E_2)$, where θ_1 satisfies $\mathbb{E}[\theta|\theta_0 < \theta < \theta_1] = E_2$; i.e., G is simple.

If an optimal G is not bi-pooling (either because a bi-tangent does not exist or even if it exists, the tangency points, E_1 and E_2, do not satisfy (25)), then G must be a pooling distribution. By Fact 3(ii), a support function of an optimal distribution G, $p : \Theta \to \mathbb{R}$, is a concatenation of at most 5 affine functions; i.e., there exists a collection $\{p_j : [\theta_{j-1}, \theta_j] \to \mathbb{R}\}_{j=0}^4$, where

$$0 = \theta_{-1} < \theta_0 < \overline{\theta}_0 < \theta_1 < \theta_2 < \overline{\theta}_1 < \theta_3 < \theta_4 = 1, \tag{28}$$

such that $p|_{[\theta_{j-1}, \theta_j]} = p_j$ for each $j \in \{0, 1, \dots, 4\}$. Moreover, each p_j is such that either $p_j = V|_{[\theta_{j-1}, \theta_j]}$ and p_j is convex (i.e., $[\theta_{j-1}, \theta_j]$ is a separating interval), or $p_j \geq V_{|_{[\theta_{j-1}, \theta_j]}}$ and p_j is affine and tangent to V at $\mathbb{E}[\theta|\theta_{j-1} \leq \theta \leq \theta_j]$ (i.e., $[\theta_{j-1}, \theta_j]$ is a pooling interval). Since V is convex (in fact, linear) on $[0, \overline{\theta}_0]$, p_0 must be a constant function equal to zero. Furthermore, to ensure that p is convex and $p \geq V$, p_1 must be tangent to V at some $E_1 \in [\overline{\theta}_0, \overline{\theta}_1]$, and p_4 must be tangent to V at some $E_2 \in [\overline{\theta}_1, 1]$. However, because V is strictly concave on $[\overline{\theta}_0, \overline{\theta}_1]$ and on $[\overline{\theta}_1, 1]$, and has a convex kink at $\overline{\theta}_1$, p_2, if it exists, would lie strictly above V; contradicting that p_2 is tangential to V on $[\theta_1, \theta_2]$. Moreover, any point in which p_3 is tangential to V would be a concave kink; contradicting the convexity of p. Thus, it follows that p must be a concatenation of at most three affine functions. Moreover, any G that is pooling that can be supported by such a p is associated with $(\theta_0, \theta_1, \mathbb{E}[\theta|\theta \geq \theta_1])$ and so is simple. ■

C.2 Optimal simple G

As shown in Figure 13, for each $j \in \{1,2\}$, let $p_j^o : \Theta \to \mathbb{R}$ denote the smallest linear function that exceeds V_j, and let E_j^o denote the smallest $E \in \Theta$ such that $p_j^o(E) = V_j(E)$. Observe that while $p_2^o \geq p_1^o$ implies $\max V_2 > \max V_1$, the converse is not true.

Figure 13: p_j^o and E_j^o's.

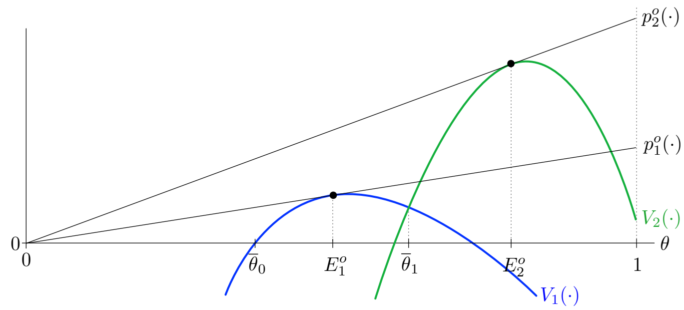

We first consider the case when $p_2^o \geq p_1^o$.

Proposition 10. *Suppose $p_2^o \geq p_1^o$ and $\mathbb{E}[\theta] \geq E_2^o$. Then, $G = F$ solves the firm's problem and assigning all workers to position $j = 2$ is optimal.*

Proof. Because V_2 is concave on $[E_2^o, 1] \subseteq [\overline{\theta}_1, 1]$, then $G = F$ (i.e., providing no information) is optimal for the firm because it can be supported by p given by the tangent to V at $\mathbb{E}[\theta]$. Since $\mathbb{E}[\theta] \in [\overline{\theta}_1, 1]$, the firm maximises its expected payoff by assigning all workers to position $j = 2$. ■

Given a simple G associated with $(\theta_0, \theta_1, E_2)$, we define the firm's expected payoff as:

$$\begin{aligned}\widehat{V}(\theta_0, \theta_1, E_2) &= \int_{\Theta} V(E)\, \mathrm{d}G_{\theta_0,\theta_1,E_2}(\theta) \\ &= (1 - F(\theta_0)) \left[V(\mathbb{E}[\theta|\theta_0 < \theta < \theta_1])\, \alpha_{\theta_0,\theta_1,E_2} + V(E_2)\left(1 - \alpha_{\theta_0,\theta_1,E_1}\right)\right],\end{aligned}$$

where

$$\alpha_{\theta_0,\theta_1,E_2} := \begin{cases} 0 & \text{if } \theta_1 = 1 \text{ and } E_2 = \mathbb{E}[\theta|\theta > \theta_0], \\ \frac{E_2 - \mathbb{E}[\theta|\theta > \theta_0]}{E_2 - \mathbb{E}[\theta|\theta_0 < \theta < \theta_1]} & \text{otherwise.} \end{cases}$$

The problem for the firm now, (3), is to maximise $\widehat{V}$ with respect to $(\theta_0, \theta_1, E_2) \in \mathscr{F}$. Since $\widehat{V}$ is differentiable, a necessary condition for optimality is that $(\theta_0, \theta_1, E_2)$ satisfies the first-order conditions (while accounting for boundary solutions). We establish sufficiency by constructing the support function for G associated with $(\theta_0, \theta_1, E_2)$ as an upper envelope of tangents to $\widehat{V}$ at $\mathbb{E}[\theta|\theta_0 < \theta < \theta_1]$ and E_2, while imposing an additional condition to ensure that the support function is convex. We say that a tuple in $\mathscr{F}$ is *nondegenerate* if it is not associated with F.

Lemma 8. *A nondegenerate tuple* $(\theta_0^*, \theta_1^*, E_2^*)$*in*$\mathscr{F}$ *solves the firm's problem,* (3)*, if and only if it satisfies the first-order conditions and* $V'(\mathbb{E}[\theta|\theta \geq \theta_1^*]) \geq V'(\mathbb{E}[\theta|\theta_0^* < \theta < \theta_1^*])$.

Proof. The idea is to construct a supporting function p^* that supports $G_{(\theta_0^*,\theta_1^*,E_2^*)}$ using the first-order conditions to (3). Then, by Fact 3, it would follow that $(\theta_0^*, \theta_1^*, E_2^*)$ must be optimal. To that end, let us first fix $\theta_0 \in (0, \overline{\theta}_0)$ and $\theta_1 \in (\overline{\theta}_0, 1)$ and solve

$$\max_{E_2 \in \Theta} \widehat{V}(\theta_0, \theta_1, E_2) \text{ s.t. } \mathbb{E}[\theta|\theta > \theta_0] \leq E_2 \leq \mathbb{E}[\theta|\theta \geq \theta_1].$$

We denote the maximiser as $E_2^*(\theta_0, \theta_1)$. The derivative of the objective with respect to E_2 is

$$\frac{\mathrm{d}\widehat{V}(\cdot)}{\mathrm{d}E_2} = (1 - F(\theta_0))\left(1 - \alpha_{\theta_0,\theta_1,E_2}\right)\left[V'(E_2) - \frac{V(E_2) - V(\mathbb{E}[\theta|\theta_0 < \theta < \theta_1])}{E_2 - \mathbb{E}[\theta|\theta_0 < \theta < \theta_1]}\right].$$

If $\frac{\mathrm{d}\widehat{V}(\cdot)}{\mathrm{d}E_2}|_{E_2=\mathbb{E}[\theta|\theta \geq \theta_1]} > 0$, we let $E_2^*(\theta_0, \theta_1) = \mathbb{E}[\theta|\theta \geq \theta_1]$ and if $\frac{\mathrm{d}\widehat{V}(\cdot)}{\mathrm{d}E_2}|_{E_2=\mathbb{E}[\theta|\theta > \theta_0]} < 0$, we let $E_2^*(\theta_0, \theta_1) = \mathbb{E}[\theta|\theta > \theta_0]$. Otherwise, we let $E_2^* \equiv E_2^*(\theta_0, \theta_1)$ be the solution to

$$V(E_2^*) - V(\mathbb{E}[\theta|\theta_0 < \theta < \theta_1]) = V'(E_2^*)(E_2^* - \mathbb{E}[\theta|\theta_0 < \theta < \theta_1]).$$

Next, consider the program:

$$\max_{\theta_0 \in [0,\overline{\theta}_0],\ \theta_1 \in [\overline{\theta}_0, 1]} \widehat{V}(\theta_0, \theta_1, E_2^*(\theta_0, \theta_1)) \text{ s.t. } \theta_1 \geq \theta_0.$$

Let θ_0^* and θ_1^* be the maximisers. There are three cases to consider:

(i) $\mathbb{E}[\theta|\theta > \theta_0^*] < E_2^*(\theta_0^*, \theta_1^*) < \mathbb{E}[\theta|\theta \geq \theta_1^*]$.

(ii) $E_2^*(\theta_0^*, \theta_1^*) = \mathbb{E}[\theta|\theta \geq \theta_1^*]$ because $\frac{\mathrm{d}\widehat{V}(\theta_0^*,\theta_1^*,\mathbb{E}[\theta|\theta \geq \theta_1^*])}{\mathrm{d}E_2} > 0$.

(iii) $E_2^*(\theta_0^*, \theta_1^*) = \mathbb{E}[\theta|\theta > \theta_0^*]$ because $\frac{\mathrm{d}\widehat{V}(\theta_0^*,\theta_1^*,\mathbb{E}[\theta|\theta \geq \theta_1^*])}{\mathrm{d}E_2} < 0$.

We consider case (i) first. Letting $E_2^* \equiv E_2^*(\theta_0, \theta_1)$ and $\alpha^*_{\theta_0,\theta_1} \equiv \alpha_{\theta_0,\theta_1,E_2^*(\theta_0,\theta_1)}$, the

Envelope theorem gives us that

$$\begin{aligned}
\frac{\mathrm{d}\widehat{V}\left(\theta_1,\theta_1,E_2^*\right)}{\mathrm{d}\theta_1} &= \left.\frac{\partial\widehat{V}\left(\theta_0,\theta_1,E_2\right)}{\partial\theta_1}\right|_{E_2=E_2^*} \\
&= f\left(\theta_1\right)\left(\theta_1-\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]\right)\frac{1-F\left(\theta_0\right)}{F\left(\theta_1\right)-F\left(\theta_0\right)}\alpha^*_{\theta_0,\theta_1} \\
&\quad\times\left[V'\left(\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]\right)-\frac{V\left(E_2^*\right)-V\left(\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]\right)}{E_2^*-\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]}\right]
\end{aligned}$$

and

$$\begin{aligned}
&\frac{\mathrm{d}\widehat{V}\left(\theta_0,\theta_1,E_2^*\right)}{\mathrm{d}\theta_0} = \left.\frac{\partial\widehat{V}\left(\theta_0,\theta_1,E_2\right)}{\partial\theta_0}\right|_{E_2=E_2^*} \\
&=\frac{f\left(\theta_0\right)}{f\left(\theta_1\right)}\frac{\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]-\theta_0}{\theta_1-\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]}\frac{\mathrm{d}\widehat{V}\left(\theta_0,\theta_1,E_2^*\left(\theta_0,\theta_1\right)\right)}{\mathrm{d}\theta_1} \\
&\quad+f\left(\theta_0\right)\left(\mathbb{E}\left[\theta|\theta>\theta_0\right]-\theta_0\right)\frac{V\left(E_2^*\right)-V\left(\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]\right)}{E_2^*-\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]} \\
&\quad-f\left(\theta_0\right)\left(\mathbb{E}\left[\theta|\theta>\theta_0\right]-\theta_0\right)\frac{V\left(\mathbb{E}\left[\theta|\theta_0<\theta<\theta_1\right]\right)\alpha^*_{\theta_0,\theta_1}+V\left(E_2^*\right)\left(1-\alpha^*_{\theta_0,\theta_1}\right)}{\mathbb{E}\left[\theta|\theta>\theta_0\right]-\theta_0}.
\end{aligned}$$

Thus, at any interior optimal $\theta_0^*\in(0,\theta_0)$ and $\theta_1^*\in(\theta_0^*,1)$, it must be that

$$V'\left(\mathbb{E}\left[\theta|\theta_0^*<\theta<\theta_1^*\right]\right)\left(E_2^*-\mathbb{E}\left[\theta|\theta_0^*<\theta<\theta_1^*\right]\right)=V\left(E_2^*\right)-V\left(\mathbb{E}\left[\theta|\theta_0^*<\theta<\theta_1^*\right]\right),$$

which ensures that the impact on the average profit from workers equals the gains from more workers being assigned to $j=1$ net of losses from fewer workers being assigned to $j=2$; and

$$\frac{V\left(E_2^*\right)-V\left(\mathbb{E}\left[\theta|\theta_0^*<\theta<\theta_1^*\right]\right)}{E_2^*-\mathbb{E}\left[\theta|\theta_0^*<\theta<\theta_1^*\right]}=\frac{V\left(\mathbb{E}\left[\theta|\theta_0^*<\theta<\theta_1^*\right]\right)\alpha^*_{\theta_0^*,\theta_1^*}+V\left(E_2^*\right)\left(1-\alpha^*_{\theta_0^*,\theta_1^*}\right)}{\mathbb{E}\left[\theta|\theta>\theta_0^*\right]-\theta_0^*},$$

which ensures that changes in profits due to changes in workers assigned to $j=1$ or 2 equals the changes in profits from retaining fewer workers. If an interior solution $(\theta_0^*,\theta_1^*,E_2^*)$ to the first-order condition exists, then we can construct a convex supporting function p^* as follows:

$$\begin{aligned}
p^*(\theta) &:= \left[V'\left(\mathbb{E}\left[\theta|\theta_0^* < \theta < \theta_1^*\right]\right)\left(\theta - \theta_0^*\right)\right]_+ \\
&= \begin{cases} V(\theta) = 0 & \text{if } \theta \in [0, \theta_0^*], \\ V'\left(\mathbb{E}\left[\theta|\theta_0^* < \theta < \theta_1^*\right]\right)\left(\theta - \theta_0^*\right) & \text{if } \theta \in (\theta_0^*, 1]. \end{cases}
\end{aligned}$$

By construction, p^* is convex and $p^* \geq V$. One can also verify that p^* and $G_{\theta_0^*,\theta_1^*,E_2(\theta_0^*,\theta_1^*)}$ satisfy (22)–(24). Hence, we conclude that $G_{\theta_0^*,\theta_1^*,E_2^*(\theta_0^*,\theta_1^*)}$ solves (2). Observe that this corresponds to the case in which the interval $(\theta_0^*, 1)$ is bi-pooled to $\mathbb{E}[\theta|\theta_0^* < \theta < \theta_1^*]$ and $E_2(\theta_0^*, \theta_1^*)$. Figure 14 depicts the typical example of this case.

Figure 14: Case 1: $\mathbb{E}[\theta|\theta \geq \theta_0^*] < E_2^* < \mathbb{E}[\theta|\theta \geq \theta_1^*]$.

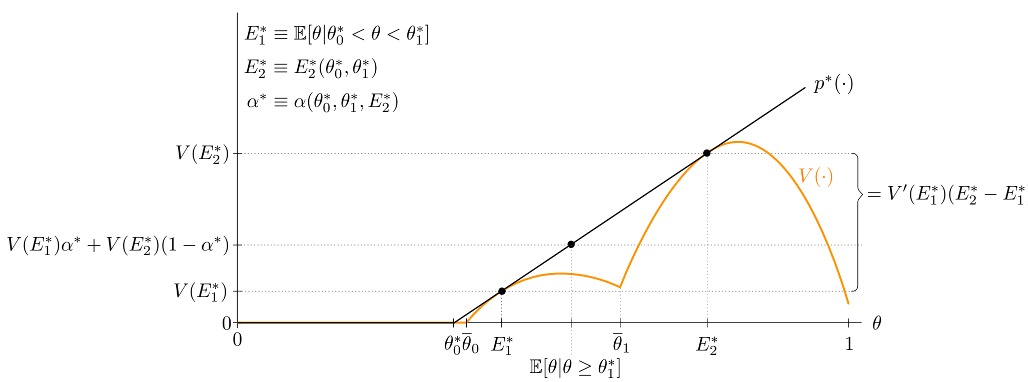

We now consider case (ii) in which $E_2^*(\theta_0^*, \theta_1^*) = \mathbb{E}[\theta|\theta \geq \theta_1^*]$. The expected payoff for the firm now is:

$$\begin{aligned}
\widehat{V}\left(\theta_0^*, \theta_1^*, \mathbb{E}\left[\theta|\theta \geq \theta_1^*\right]\right) = {} & V\left(\mathbb{E}\left[\theta|\theta_0^* < \theta < \theta_1^*\right]\right)\left(F\left(\theta_1^*\right) - F\left(\theta_0^*\right)\right) \\
& + V\left(\mathbb{E}\left[\theta|\theta \geq \theta_1^*\right]\right)\left(1 - F\left(\theta_1^*\right)\right)
\end{aligned}$$

Observe that

$$\begin{aligned}
&\frac{\mathrm{d}\widehat{V}\left(\theta_0, \theta_1, \mathbb{E}\left[\theta|\theta \geq \theta_1\right]\right)}{\mathrm{d}\theta_0} \\
&= f\left(\theta_0\right)\left[\mathbb{E}\left[\theta|\theta_0 < \theta < \theta_1\right] - \theta_0\right]\left(V'\left(\mathbb{E}\left[\theta|\theta_0 < \theta < \theta_1\right]\right) - \frac{V\left(\mathbb{E}\left[\theta|\theta_0 < \theta < \theta_1\right]\right)}{\mathbb{E}\left[\theta|\theta_0 < \theta < \theta_1\right] - \theta_0}\right),
\end{aligned}$$

and

$$
\begin{aligned}
&\frac{\mathrm{d}\widehat{V}\left(\theta_0, \theta_1, \mathbb{E}\left[\theta \mid \theta \geq \theta_1\right]\right)}{\mathrm{d}\theta_1} \\
&= f\left(\theta_1\right)\big(V\left(\mathbb{E}\left[\theta \mid \theta_0 < \theta < \theta_1\right]\right) + V'\left(\mathbb{E}\left[\theta \mid \theta_0 < \theta < \theta_1\right]\right)\left(\theta_1 - \mathbb{E}\left[\theta \mid \theta_0 < \theta < \theta_1\right]\right)\big) \\
&\qquad + f\left(\theta_1\right) V'\left(\mathbb{E}\left[\theta \mid \theta \geq \theta_1\right]\right)\left(\mathbb{E}\left[\theta \mid \theta \geq \theta_1\right] - \theta_1\right) \\
&\qquad - f\left(\theta_1\right)\left[V\left(\mathbb{E}\left[\theta \mid \theta \geq \theta_1\right]\right)\right].
\end{aligned}
$$

Suppose $\theta_0^* > 0$, then the first-order condition with respect to θ_0 means

$$
V'\left(\mathbb{E}\left[\theta \mid \theta_0^* < \theta < \theta_1^*\right]\right)\left(\mathbb{E}\left[\theta \mid \theta_0^* < \theta < \theta_1^*\right] - \theta_0^*\right) = V\left(\mathbb{E}\left[\theta \mid \theta_0^* < \theta < \theta_1^*\right]\right),
$$

but if $\theta_0^* = 0$, then

$$
\frac{\mathrm{d}\widehat{V}\left(0, \theta_1^*, \mathbb{E}\left[\theta \mid \theta \geq \theta_1^*\right]\right)}{\mathrm{d}\theta_0} = V'\left(\mathbb{E}\left[\theta \mid \theta < \theta_1^*\right]\right) - \frac{V\left(\mathbb{E}\left[\theta \mid \theta < \theta_1^*\right]\right)}{\mathbb{E}\left[\theta \mid \theta < \theta_1^*\right]} < 0.
$$

The first-order condition with respect to θ_1 is

$$
\begin{aligned}
V\left(\mathbb{E}\left[\theta \mid \theta \geq \theta_1^*\right]\right) = V&\left(\mathbb{E}\left[\theta \mid \theta_0^* < \theta < \theta_1^*\right]\right) \\
&+ V'\left(\mathbb{E}\left[\theta \mid \theta_0^* < \theta < \theta_1^*\right]\right)\left(\theta_1^* - \mathbb{E}\left[\theta \mid \theta_0^* < \theta < \theta_1^*\right]\right) \\
&+ V'\left(\mathbb{E}\left[\theta \mid \theta \geq \theta_1^*\right]\right)\left(\mathbb{E}\left[\theta \mid \theta \geq \theta_1^*\right] - \theta_1^*\right)
\end{aligned}
$$

Letting $E_1^* \equiv \mathbb{E}[\theta \mid \theta_0^* < \theta < \theta_1^*]$ and $E_2^* \equiv \mathbb{E}[\theta \mid \theta \geq \theta_1^*]$, construct p^* as follows:

$$
\begin{aligned}
p^*\left(\theta\right) &:= \max\Big\{\left[V\left(E_1^*\right) + V'\left(E_1^*\right)\left(\theta - E_1^*\right)\right]_+, V\left(E_2^*\right) + V'\left(E_2^*\right)\left(\theta - E_2^*\right)\Big\} \\
&= \begin{cases} \left[V\left(E_1^*\right) + V'\left(E_1^*\right)\left(\theta - E_1^*\right)\right]_+ & \text{if } \theta \in [0, \theta_1^*], \\ V\left(E_2^*\right) + V'\left(E_2^*\right)\left(\theta - E_2^*\right) & \text{if } \theta \in (\theta_1^*, 1]. \end{cases}
\end{aligned}
\tag{29}
$$

In case $\theta_0^* > 0$, the first-order condition with respect to θ_0 means that $V'(E_1^*) > 0$ so that p^* is convex on $[0, \theta_1^*)$. If $\theta_0^* = 0$, then p^* is affine (and thus convex) on $[0, \theta_1^*)$. To ensure that p^* is convex on $[0, 1]$, we further require that $V'(E_2^*) \geq V'(E_1^*)$. By construction, $p^* \geq V$. Moreover, p^* and $G_{\theta_0^*, \theta_1^*, E_2^*(\theta_0^*, \theta_1^*)}$ satisfy (22)–(24) and so we conclude that $G_{\theta_0^*, \theta_1^*, E_2^*(\theta_0^*, \theta_1^*)}$ solves (2). Observe that, in this case, the interval (θ_0^*, θ_1^*) is pooled to $E_1^*)$ and the interval $[\theta_1^*, 1]$ is pooled to E_2^*. Figure 15 shows the typical example of the case when $E_2^*(\theta_0^*, \theta_1^*) =$

$\mathbb{E}[\theta|\theta \geq \theta_1^*]$ and either $\theta_0^* \in (0, \overline{\theta}_0)$ (Figure 15a) or $\theta_0^* = 0$ (Figure 15b).

Figure 15: Case 2: $E_2^*(\theta_0^*, \theta_1^*) = \mathbb{E}[\theta|\theta \geq \theta_1^*]$.

(a) Case 2a: $\theta_0^* \in (0, \overline{\theta}_0)$.

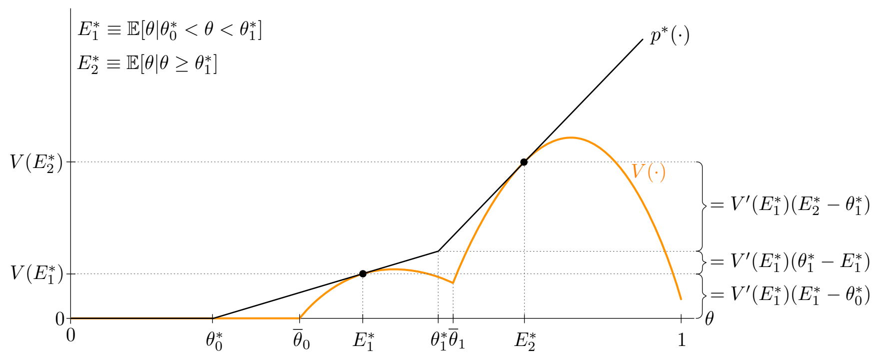

(b) Case 2b: $\theta_0^* = 0$.

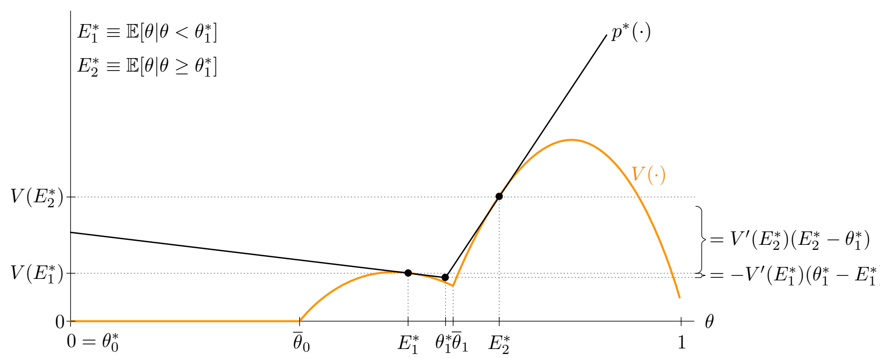

Finally, we consider case (iii) in which $E_2^*(\theta_0^*, \theta_1^*) = \mathbb{E}[\theta|\theta > \theta_0^*]$. The expected payoff for the firm is now:

$$\widehat{V}\left(\theta_0^*, \theta_1^*, \mathbb{E}\left[\theta|\theta > \theta_0^*\right]\right) = V\left(\mathbb{E}\left[\theta|\theta > \theta_0^*\right]\right)\left(1 - F\left(\theta_0^*\right)\right).$$

Since the payoff does not depend on θ_1, we set it to $\theta_1^* = 1$. Observe that

$$\frac{\mathrm{d}\widehat{V}\left(\theta_0, \theta_1, \mathbb{E}\left[\theta|\theta > \theta_0\right]\right)}{\mathrm{d}\theta_0} = f\left(\theta_0\right)\left[\mathbb{E}\left[\theta|\theta > \theta_0\right] - \theta_0\right]\left(V'\left(\mathbb{E}\left[\theta|\theta > \theta_0\right]\right) - \frac{V\left(\mathbb{E}\left[\theta|\theta > \theta_0\right]\right)}{\mathbb{E}\left[\theta|\theta > \theta_0\right] - \theta_0}\right).$$

Thus, the first-order condition then requires that, if $\theta_0^* \in (0, \theta_0^*)$ and $\theta_1^* \in (\theta_0^*, 1)$,

$$V'(\mathbb{E}[\theta|\theta > \theta_0^*])(\mathbb{E}[\theta|\theta > \theta_0^*] - \theta_0^*) = \frac{V\left(\mathbb{E}\left[\theta|\theta > \theta_0^*\right]\right)}{\mathbb{E}\left[\theta|\theta > \theta_0^*\right] - \theta_0^*};$$

i.e., the tangent to V at $\mathbb{E}[\theta|\theta > \theta_0^*]$ crosses the x-axis at θ_0^*. Construct p^* now as follows:

$$p^*(\theta) = \begin{cases} V(\theta) = 0 & \text{if } \theta \in [0, \theta_0^*], \\ V'\left(\mathbb{E}\left[\theta|\theta \geq \theta_0^*\right]\right)\left(\theta - \theta_0^*\right) & \text{if } \theta \in (\theta_0^*, 1]. \end{cases}.$$

Observe that p^* is convex, $p^* \geq V$. Moreover, p^* and $G_{\theta_0^*, \theta_1^*, E_2^*(\theta_0^*, \theta_1^*)}$ satisfy (22)—(24) so that $G_{\theta_0^*, \theta_1^*, E_2^*(\theta_0^*, \theta_1^*)}$ solves (2). In this case, the interval $(\theta_0^*, 1]$ is pooled to $\mathbb{E}[\theta|\theta > \theta_0^*]$.

Figure 16: Case 3: $E_2^*(\theta_0^*, \theta_1^*) = \mathbb{E}[\theta|\theta > \theta_0^*]$.

(a) Case 3a: $E_1^* \in (\overline{\theta}_0, \overline{\theta}_1)$.

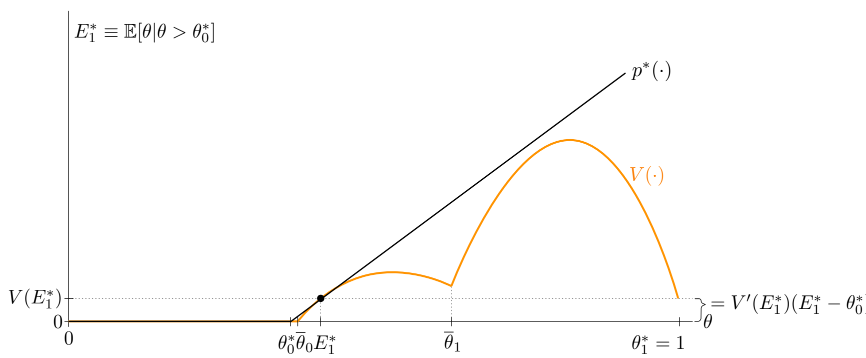

(b) Case 3b: $E_1^* \in (\overline{\theta}_1, 1)$.

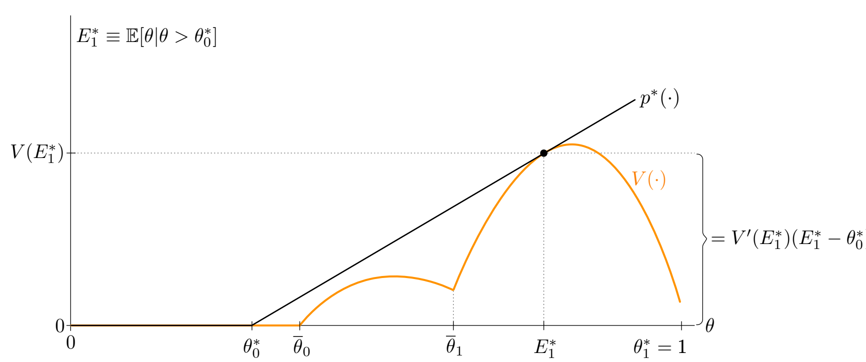

Figure 16 shows the typical examples of the case when $E_2^*(\theta_0^*, \theta_1^*) = \mathbb{E}[\theta|\theta > \theta_0^*]$. ■

We now characterise simple G that solves the firm's information design problem, (3). Towards this goal, in case V has a bi-tangent, we denote it as $V^b(\cdot)$, and let E_1^b and E_2^b denote the two tangency points, where $E_2^b > E_1^b$, and let θ_0^b and θ_1^b be points such that $V^b(\theta_0^b) = 0$ and $\mathbb{E}[\theta|\theta_0^b < \theta < \theta_1^b] = E_1^b$. Figure 17 shows how E_1^b, E_2^b, θ_0^b and θ_1^b are obtained.

Figure 17: V has a bi-tangent, V^b.

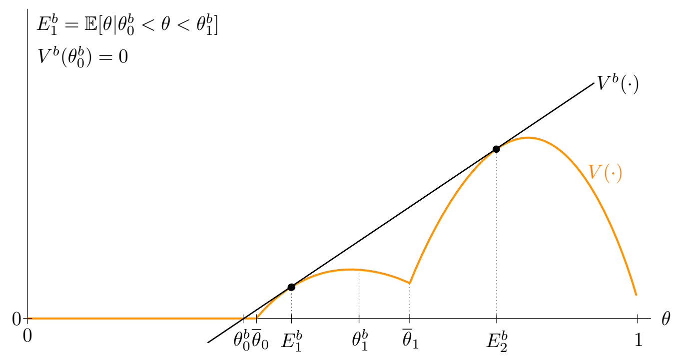

In case V does not have a bi-tangent, then at least one of (26) and (27) does not hold. If (26) does not hold, then let $E_2^b \in [\overline{\theta}_1, 1]$ be the point such that the tangent to V_2 at E_2^b goes through the point $(\overline{\theta}_0, 0)$. If (27) does not hold, then let $E_1^b \in [\overline{\theta}_0, \overline{\theta}_1]$ be the point such that the tangent to V_1 at E_1^b goes through the point $(1, V_2(1))$, and let θ_0^b be the point such that the tangents equals zero.

Figure 18: V does not have a bi-tagent.

(a) Case: (26) does not hold.

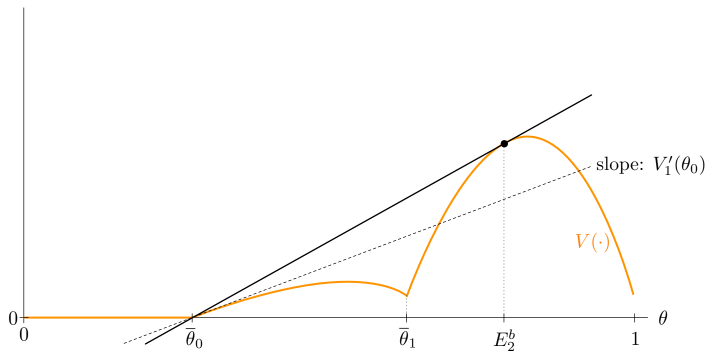

(b) Case: (27) does not hold

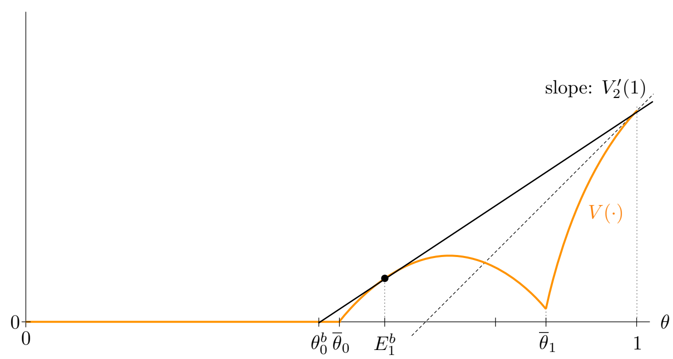

Theorem 5. *Suppose $p_2^o \geq p_1^o$ and $\mathbb{E}[\theta] < E_2^o$. Suppose a bi-tangent of V exists.*

(i) If

$$E_1^b \leq \mathbb{E}[\theta|\theta > \theta_0^b] \leq E_2^b \leq \mathbb{E}[\theta|\theta \geq \theta_1^b], \tag{30}$$

then an optimal G exists that bi-pools the interval $[\theta_0^b, 1]$ to E_1^b and E_2^b.

(ii) If (30) *does not hold and*

$$E_2^b > \mathbb{E}[\theta|\theta \geq \theta_1^b], \tag{31}$$

then an optimal G exists that pools the interval (θ_0^*, θ_1^*) to $\mathbb{E}[\theta|\theta_0^* < \theta < \theta_1^*] > E_1^b$, and pools the interval $[\theta_1^*, 1]$ to $E[\theta|\theta \geq \theta_1^*] < E_2^b$, where $\theta_0^* \in [0, \theta_0^b)$ and $\theta_1^* \in (\theta_0^*, 1)$.

(iii) If (30) does not hold and

$$E_1^b > \mathbb{E}[\theta|\theta > \theta_0^b], \tag{32}$$

then an optimal G exists that pools the interval $(\theta_0^*, 1]$ to $\mathbb{E}[\theta|\theta > \theta_0^*]$.

Suppose V does not have a bi-tangent: If (26) *does not hold, then optimal G has the same features as case (iii) above if $\mathbb{E}[\theta|\theta > \overline{\theta}_0] \geq E_2^b$ and otherwise G has the same features as case (ii).* (27) *does not hold, then optimal G has the same features as case (iii) above if $\theta_0^b < \mathbb{E}[\theta|\theta > \theta_0^b]$ and otherwise G has the same features as case (ii).*

Proof. (i) By Fact 3(ii), the interval $[\theta_0^b, 1]$ can be bi-pooled at E_1^b and E_2^b if and only if (30) holds. This case corresponds to case 1 from Lemma 8 with $\theta_0^* = \theta_0^b$, $\theta_1^* = \theta_1^b$ and $E_2^* = E_2^b$. (ii) This case includes the case in which (30) does not hold and $E_2^b > \mathbb{E}[\theta|\theta > \theta_0^b]$ because $\theta_1^b \geq \theta_0^b$. Consider first the case $E_2 \geq \overline{\theta}_1$. Letting denote $E_2 \equiv \mathbb{E}[\theta|\theta > \theta_0^b]$, we first argue that

$$V(E_2) < V(E_1^b) + V'(E_1^b)(\theta_1^b - E_1^b) + V'(E_2)(E_2 - \theta_1^b). \tag{33}$$

Toward a contradiction, suppose the inequality does not hold. Then, there must be a line through the point $(\theta_1^b, V(E_1^b) + V'(E_1^b)(\theta_1^b - E_1^b))$ that is tangent to V at $\mathbb{E}[\theta|\theta \geq \theta_1^b]$. However, by the choice of θ_1^b, V has a unique tangent that goes through the point $(\theta_1^b, V(E_1^b) + V'(E_1^b)(\theta_1^b - E_1^b))$ given by the tangent at E_2^b—a contradiction. Figure 19 depicts a typical example of this case.

Figure 19: Case (ii): Positive difference.

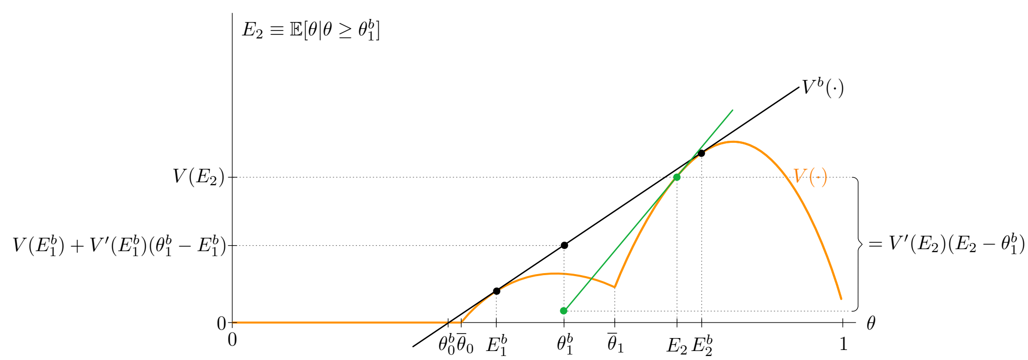

For any $\theta_0 \in [0, \theta_0^b]$, let $E_1(\theta_0)$ be given by

$$\theta_0 + V'(E_1(\theta_0))\, E_1(\theta_0) = V(E_1(\theta_0));$$

i.e., $E_1(\theta_0)$ is the point such that the line through $(\theta_0, 0)$ and $(E_1(\theta_0), V(E_1(\theta_0)))$ is tangent to V at $E_1(\theta_0)$. Because $E_1^b \in (\overline{\theta}_0, \overline{\theta}_1)$, concavity of V on $[\overline{\theta}_0, \overline{\theta}_1]$ means that $E_1(\theta_0)$ is continuous and strictly decreasing. Moreover, $\theta_1(\theta_0)$, defined implicitly by $\mathbb{E}[\theta|\theta_0 < \theta <$

$\theta_1(\theta_0)] = E_1(\theta_0)$, is also continuous and strictly decreasing. Now, let $p_1 : \theta \to \mathbb{R}$ denote the smallest linear function that exceeds V_1 and let $\widehat{E}_1$ denote the smallest $E \in \theta$ such that $p_1(E) = V_1(E)$. Let $\widehat{\theta}_1$ be such that $\mathbb{E}[\theta|\theta < \widehat{\theta}_1] = \widehat{E}_1$. Suppose, as shown in Figure 20,

$$V(\mathbb{E}[\theta|\theta \geq \widehat{\theta}_1]) > V(\widehat{E}_1) + V'(\widehat{E}_1)(\widehat{\theta}_1 - \widehat{E}_1) + V'(\mathbb{E}[\theta|\theta \geq \widehat{\theta}_1])(\mathbb{E}[\theta|\theta \geq \widehat{\theta}_1] - \widehat{\theta}_1). \quad (34)$$

Figure 20: Case (ii.a): Negative difference.

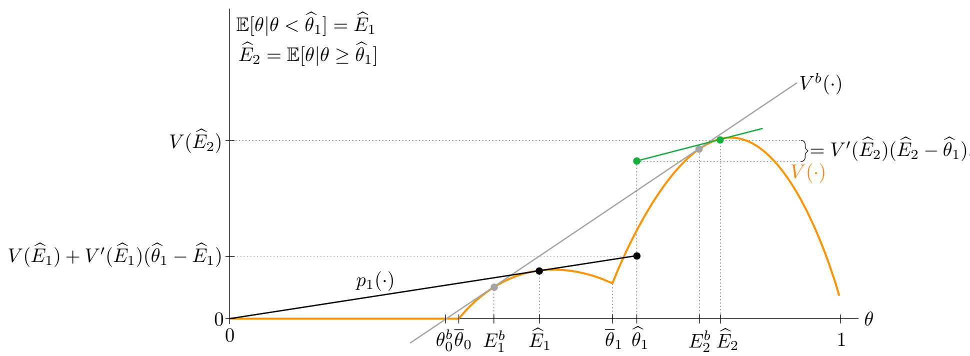

In this case, by the intermediate value theorem, there exists a $\theta_0 \in (0, \overline{\theta}_0)$ such that

$$\begin{aligned} V\left(\mathbb{E}\left[\theta|\theta \geq \theta_1\left(\theta_0\right)\right]\right) = V\left(E_1\left(\theta_0\right)\right) + V'\left(E_1\left(\theta_0\right)\right)\left(\theta_1\left(\theta_0\right) - E_1\left(\theta_0\right)\right) \\ + V'\left(\mathbb{E}\left[\theta|\theta \geq \theta_1\left(\theta_0\right)\right]\right)\left(\mathbb{E}\left[\theta|\theta \geq \theta_1\left(\theta_0\right)\right] - \theta_1\left(\theta_0\right)\right). \end{aligned}$$

Moreover, observe that such a θ_0 is unique due to the concavity of V on $[\overline{\theta}_0, \overline{\theta}_1]$ and on $[\overline{\theta}_1, 1]$. This case corresponds to the case 2a in the proof of Lemma 8.

Suppose now (34) does not hold as shown in Figure 21.

Figure 21: Case (ii.b): Positive difference.

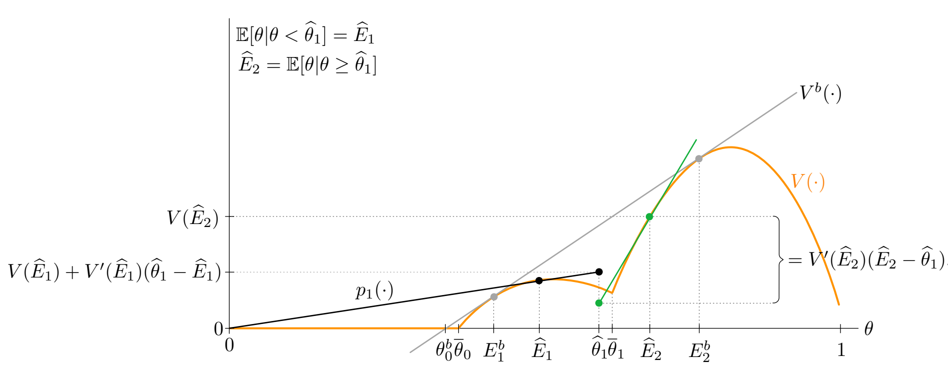

Now let $\widetilde{\theta}_1 > \widehat{\theta}_1$ be such that $p_1(\widetilde{\theta}_1|\widetilde{\theta}_1) = V(\widetilde{\theta}_1)$, where

$$p_1\left(\theta'|\theta_1\right) := V\left(\mathbb{E}\left[\theta|\theta < \theta_1\right]\right) + V'\left(\mathbb{E}\left[\theta|\theta < \theta_1\right]\right)\left(\theta' - \mathbb{E}\left[\theta|\theta < \theta_1\right]\right).$$

We will argue that (see Figure 22):

$$V(\mathbb{E}[\theta|\theta \geq \widetilde{\theta}_1]) > V(\mathbb{E}[\theta|\theta < \widetilde{\theta}_1]) + V'(\mathbb{E}[\theta|\theta < \widetilde{\theta}_1])(\widetilde{\theta}_1 - \mathbb{E}[\theta|\theta < \widetilde{\theta}_1])$$
$$+ V'(\mathbb{E}[\theta|\theta \geq \widetilde{\theta}_1])(\mathbb{E}[\theta|\theta \geq \widetilde{\theta}_1] - \widetilde{\theta}_1)$$

By construction, $\widetilde{\theta}_1 \in (\overline{\theta}_1, 1)$ is such that the tangent to V at $\mathbb{E}[\theta|\theta < \widetilde{\theta}_1] \in (\overline{\theta}_1, 1)$ intersects V at $\widetilde{\theta}_1$. Since V is concave on $[\overline{\theta}_1, 1]$, $V'(\cdot)$ is strictly decreasing on $(\widetilde{\theta}_1, 1]$, for any $E_2 \in (\widetilde{\theta}_1, E_2^b) \subseteq (\widetilde{\theta}_1, 1]$,

$$V'(E_2)(E_2 - \widetilde{\theta}_1) > V(E_2) - V(\tilde{\theta}_1).$$

Letting $E_2 = \mathbb{E}[\theta|\theta < \widetilde{\theta}_1]$ gives the desired inequality.

Figure 22: Case (ii.b): Negative difference.

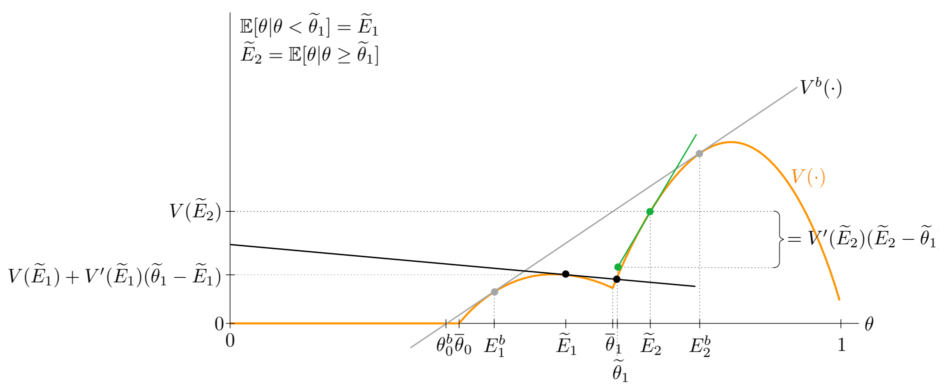

By the intermediate value theorem, there exists $\theta_1 \in (\widehat{\theta}_1, \widetilde{\theta}_1)$ such that

$$V\left(\mathbb{E}\left[\theta|\theta \geq \theta_1\right]\right) = V\left(\mathbb{E}\left[\theta|\theta < \theta_1\right]\right) + V'\left(\mathbb{E}\left[\theta|\theta < \theta_1\right]\right)\left(\theta_1 - \mathbb{E}\left[\theta|\theta < \theta_1\right]\right)$$
$$+ V'\left(\mathbb{E}\left[\theta|\theta \geq \theta_1\right]\right)\left(\mathbb{E}\left[\theta|\theta \geq \theta_1\right] - \theta_1\right).$$

Moreover, observe that such a θ_0 is unique due to the concavity of V on $[\overline{\theta}_0, \overline{\theta}_1]$ and on $[\overline{\theta}_1, 1]$. This case corresponds to case 2b in the proof of Lemma 8.

Now suppose that $\mathbb{E}[\theta|\theta \geq \theta_1^b] < \overline{\theta}_1$. Because V is concave on $[\overline{\theta}_0, \overline{\theta}_1]$, $\mathbb{E}[\theta|\theta \geq \theta_1^b] > E_1^b$ means that $V'(\mathbb{E}[\theta|\theta \geq \theta_1^b]) < V'(E_1^b)$. Hence, a function constructed as in (29) would not be convex. As already argued, choosing any $\theta_0 \in (0, \theta_0^b)$ means $E_1(\theta_0) > E_1^b$ and

$\theta_1(\theta_0) > \theta_1^b$ and $E_1(\theta_0) \geq \mathbb{E}[\theta|\theta \geq \theta_1^b]$. So long as $E_1(\theta_0) < \overline{\theta}_1$, a function constructed as in (29) would not be convex. If $E_1(\theta_0) \geq \overline{\theta}_1$ then we can proceed as above with θ_0 replacing $\overline{\theta}_0$ so we are once again in either case 2a or 2b of Lemma 8.

(iii) Let $\widetilde{\theta}_0$ be such that $\mathbb{E}[\theta|\theta > \widetilde{\theta}_0] = E_1^b$ and note that $\widetilde{\theta}_0 > \overline{\theta}_0$ because of condition (32). Define $\gamma\colon [\theta_0^b, \widetilde{\theta}_0] \to \mathbb{R}$ as

$$\gamma(\theta_0) := V'(\mathbb{E}[\theta|\theta > \theta_0]) - \frac{V(\mathbb{E}[\theta|\theta > \theta_0])}{\mathbb{E}[\theta|\theta > \theta_0] - \theta_0}.$$

Since the line segment connecting θ_0^b and E_1^b is tangent to V at E_1^b, $\widetilde{\theta}_0 > \theta_0^b$, and V is concave on $[\overline{\theta}_0, \overline{\theta}_1]$, the condition (32) implies that

$$\gamma(\theta_0^b) > 0 > \gamma(\widetilde{\theta}_0).$$

Since γ is continuous, by the intermediate value theorem, there exists $\theta_0 \in (\theta_0^b, \widetilde{\theta}_0)$ such that $\gamma(\theta_0^*) = 0$. Observe that this case corresponds to case (iii) in the proof of Lemma 8.

Now, suppose that a bi-tangent does not exist. Consider the first the case in which (26) does not hold. If $\mathbb{E}[\theta|\theta > \overline{\theta}_0] \geq E_2^b$. Since $\mathbb{E}[\theta] \leq \overline{E}$, there exists $\theta_0 \in [0, \overline{\theta}_0]$ such that

$$V\left(\mathbb{E}\left[\theta|\theta > \theta_0\right]\right) = V'\left(\mathbb{E}\left[\theta|\theta > \theta_0\right]\right)\left(\mathbb{E}\left[\theta|\theta > \theta_0\right] - \theta_0\right),$$

which is case (iii) in the proof of Lemma 8. Now suppose $\mathbb{E}[\theta|\theta > \overline{\theta}_0] < \overline{E}_0$. Observe that neither case (iii) nor case (i) in the proof of Lemma 8 is possible; i.e., only case (ii) can arise.

Finally, suppose (27) does not hold. If $\mathbb{E}[\theta|\theta > \theta_0^b] < E_1^b$, then any tangency to V on $[\overline{\theta}_0, 1]$ cannot be on $[\overline{\theta}_1, 1]$ so that this case corresponds to case (iii) in the proof of Lemma 8. If $\mathbb{E}[\theta|\theta > \theta_0^b] \geq E_1^b$, neither case (iii) nor case (i) in the proof of Lemma 8 is possible; i.e., only case (ii) can arise. ■

We now characterise the optimal G when $p_1^o \geq p_2^o$.

Theorem 6. *Suppose $p_1^o \geq p_2^o$. Then, V has a bi-tangent on $[\overline{\theta}_0, 1]$. Let E_1^b and E_2^b denote the two bi-tangency points.*

(i) If $\mathbb{E}[\theta] \in [0, E_1^o]$, then G associated with $(\theta_0, 1, 1)$ for some $\theta_0 \in [0, \overline{\theta}_0)$ solves the firm's problem. Moreover, no worker is ever assigned position $j = 2$.

(ii) If $\mathbb{E}[\theta] \in [E_1^o, E_1^b]$, then $G = F$ solves the firm's problem and assigning all workers to position $j = 1$ is optimal.

(iii) If $\mathbb{E}[\theta] \in [E_1^b, E_2^b]$...

(a) ... and (30) holds with $\theta_0^b = 0$,[54] then an optimal G bi-pools the Θ to E_1^b and E_2^b.

(b) ... but (30) does not hold, then an optimal G pools the interval $(0, \theta_1^*)$ to $\mathbb{E}[\theta|\theta < \theta_1^*] > E_1^b$, and pools the interval $[\theta_1^*, 1]$ to $E[\theta|\theta \geq \theta_1^*] < E_2^b$ for some $\theta_1^* \in \Theta$.

(iv) If $\mathbb{E}[\theta] \in [E_2^b, 1]$, then $G = F$ solves the firm's problem and assigning all workers to position $j = 2$ is optimal.

Proof. The result is immediate from previous arguments while observing that the slope of p_1^o must be greater than the tangent so that $\theta_1^b = 0$. ■

[54] As before, θ_1^b is given via $\mathbb{E}[\theta|\theta < \theta_1^b] = E_1^b$