Games and Economic Behavior 67 (2009) 665–681

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Discretionary rewards as a feedback mechanism

Anton Suvorov [a,b], Jeroen van de Ven [c,d,*]

[a] CEFIR, Moscow, Russia
[b] New Economic School, Moscow, Russia
[c] University of Amsterdam, Amsterdam, The Netherlands
[d] Amsterdam Center for Law and Economics, Amsterdam, The Netherlands

ARTICLE INFO

Article history:
Received 7 March 2006
Available online 19 March 2009

JEL classification:
D82
J33
M50

Keywords:
Discretionary rewards
Feedback
Self confidence
Subjective performance
Moral hazard

ABSTRACT

This paper studies the use of discretionary rewards in a finitely repeated principal–agent relationship with moral hazard. The key aspect is that rewards have informational content. When the principal obtains a private subjective signal about the agent's performance, she may pay discretionary bonuses to provide credible feedback to the agent. In accordance with the often observed compression of ratings, we show that in equilibrium the principal communicates the agent's interim performance imperfectly, i.e., she does not fully differentiate good and bad performance. Furthermore, we show that small rewards can have a large impact on the agent's effort, provided that the principal's stake in the project is small.

1. Introduction

Incentive problems are part of most relationships. Much of the literature on incentive contracts is focused on how verifiable performance measures can be used to mitigate these problems. However, most people do not work in jobs where such measures are available for all important dimensions of performance (Prendergast, 1999). Instead, many firms extensively use discretionary rewards based on subjective, non-contractible performance measures.

In this paper, we study how the principal can use a subjective performance measure to motivate the agent. We show that, even in a finitely repeated game, the principal may have an incentive to give discretionary rewards despite their non-contractibility. This happens because she has private information about the agent's performance, and wants to communicate favorable information credibly to the agent to boost his motivation in future periods. The key factors that make the agent react to positive feedback are his initial uncertainty about his ability, and complementarity between ability and effort.

Since the principal is tempted to overstate the agent's performance, cheap talk would be ineffective in this situation.[1] A monetary reward ensures credibility and, in equilibrium, gives positive feedback about past performance. This result, establishing the informational content of a reward, is in line with the literature on motivation both in psychology and economics (see e.g., Deci and Ryan, 1985; Bénabou and Tirole, 2003). In equilibrium, the bonus is proportional to the principal's

* Corresponding author at: University of Amsterdam, ACLE, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands. Fax: +31 (0) 20 5255318.
E-mail addresses: asuvorov@nes.ru (A. Suvorov), j.vandeven@uva.nl (J. van de Ven).

[1] One can check that only the babbling, uninformative equilibrium would be possible if costly signaling were replaced by a cheap talk game in our model. See also the discussion in Section 3.2.

0899-8256/$ – see front matter
doi:10.1016/j.geb.2009.03.002

payoff from the project's success. An interesting implication of this is that even a small bonus can have a substantial effect on the agent's subsequent motivation.

Our model furthermore predicts that good performance is not fully separated from bad performance by the principal. This result provides a new insight into the reluctance of managers to rate workers differently. The resulting compression of performance ratings is well documented in the literature (Prendergast, 1999). Another important result is that the prospect of getting a discretionary reward in case of success improves the agent's incentives to exert effort in the first place, even though the principal cannot *commit* to giving the reward. This improvement in incentives stems from the *information effect* (the agent gains from the opportunity to make a better decision concerning future efforts) and the *direct incentive effect* (the higher chance to get the reward at the interim stage if he works hard in the beginning).

The driving assumption behind all these results is that the principal has private information about the agent's performance. The widespread use of performance feedback systems shows the relevance of this assumption (cf. Gibbs, 1991). It seems especially relevant when workers are in their learning phase, produce a complex good, or contribute to a project that involves many individual tasks. The manager's experience enables her to form a better judgement of an employee's performance.

The model is closely related to those of Bénabou and Tirole (2003) and Suvorov (2003), which also focus on the informational content of rewards. In those papers, the principal's *promise* of a higher performance-contingent bonus sends a negative signal, whereas in our model a higher *discretionary* bonus gives positive feedback. The key difference stems from the assumption made in these papers: performance measures are supposed to be objective so that performance-contingent bonuses become contractible. If a principal observes that the agent has low ability, she also expects him to have low self-confidence. She, therefore, proposes a high reward in case of success to motivate the agent. In our case, in which the principal cannot commit to a contract, the returns to giving a high reward are higher after better interim performance.

The existing literature on subjective performance measures is mostly focussed on infinitely repeated games (e.g., MacLeod and Malcomson, 1989; Levin, 2003).[2] In that setup, non-contractible bonuses can be sustained in equilibrium through, for instance, the threat of termination of the relationship. Given that such reputation mechanisms are relatively well understood, we focus on a finitely repeated game. While our model can be extended to an infinite number of periods, we believe that a finite horizon is often realistic. This is certainly true for fixed-term contracts, but even under permanent contracts the principal and agent ultimately retire. Whenever there is some known end date, reputation mechanisms cease to be credible, and under standard assumptions, no self-enforcing discretionary bonus based on reputation can exist in equilibrium.[3]

Some other papers study subjective performance measures in finitely repeated games. MacLeod (2003) studies a one-shot game where there is no balanced budget between what the principal gives as a bonus and what the agent receives. The principal can promise to reward the agent after good performance, and to give the bonus away to a third party after bad performance (or "burn the money"). If she commits to spend the money in any event, she has no incentive to renege on giving it to the agent after good performance. Real life examples of such mechanisms do exist (see Fuchs, 2007), but it is unclear how important this practice is in reality. In this paper we focus instead on the balanced budget.

Lizzeri et al. (2003) and Fuchs (2007) also study the role of interim feedback about performance within a model with a finite horizon. In Lizzeri et al. (2003) the principal privately observes the agent's interim performance and his ability. Effort by a more able agent is more valuable to the principal, but does not increase the likelihood of success. The agent is interested in getting feedback because it allows him to adjust future efforts, given that his payoff depends on the entire history of performance and his ability (both of which are assumed to be ultimately verifiable). They show that the principal gains from revealing early the agent's ability, but not performance. One crucial difference with our approach, is that Lizzeri et al. (2003) assume that the principal can ex ante commit to information disclosure rules, so credibility is not an issue. Fuchs (2007) considers a repeated moral hazard problem without uncertainty about the agent's ability: his work is an extension of MacLeod (2003) to a dynamic environment. The principal is privately informed about the agent's performance. Since the agent's payoff depends on the entire history of payoffs, he is interested in receiving early feedback. However, Fuchs shows that the principal prefers to avoid giving informative interim feedback in order to exploit "reusability of punishments." Moreover, in the version of the model with a finite horizon, as in MacLeod (2003), there is no equilibrium with a positive reward unless one can break the balanced budget.

Finally, there is a complementary strand of literature that explains the use of discretionary bonuses with social preferences, which is distinct from our focus on the role of information. For instance, Fehr et al. (1997) argue that reciprocity can induce enforceability.

The paper is organized as follows. Section 2 presents the model and discusses the main assumptions. In Section 3 the equilibria are derived. In Section 4 we relax some of the assumptions and consider several extensions. Finally, Section 5 provides concluding remarks.

[2] See also Bull (1987), Baker et al. (1994), Pearce and Stacchetti (1998).

[3] In the presence of irrational types or types with social preferences and incomplete information, reputation building can be part of the equilibrium (cf. Kreps et al., 1982).

2. The model

2.1. Preliminaries

We consider a two-period model, where both the principal (she) and the agent (he) are risk-neutral. The agent faces a sequence of two identical projects and, at the beginning of each period $t = 1, 2$, he decides whether to exert effort or not: $e_t \in \{0, 1\}$. The cost of effort is $c(0) = 0$ and $c(1) = c > 0$.

The agent's ability for the task, θ, can be either high, $\theta = \theta_H$, or low, $\theta = \theta_L < \theta_H$. The outcome of each project y_t can be a success or a failure, $y_t \in \{S, F\}$. The probability of success depends on the agent's ability and effort:

$$\Pr\{y_t = S \mid e_t\} = e_t\theta. \tag{1}$$

Thus, ability and effort are complements. Not exerting effort induces failure with certainty. We assume that the agent's effort decision is not observed by any third parties and is therefore non-contractible. However, to simplify the exposition we assume that effort is observed by the principal.[4] In case of success, the agent receives payoff $V > 0$ and the principal gets $W > 0$. Failure yields nothing to both. The agent's reservation utility is normalized to 0. The discount factor is normalized to 1.

At the end of each period, the principal can give a *non-contractible* discretionary bonus $b_t \geqslant 0$. We assume that the budget is balanced, so that the reward paid by the principal equals the reward obtained by the agent.

2.2. The main assumptions

We assume that neither the agent himself, nor the principal, precisely knows the agent's ability. It generally takes time to learn one's ability, which depends partly on the specifics of the task, the available technology, and the input of colleagues involved.[5] Our mechanism relies on the fact that first-period performance is indicative of future performance, hence some uncertainty about the agent's characteristics is required. Like Holmstrom (1999), we believe uncertainty about talent is a natural candidate. Closely related is the idea of many search models in labor economics that jobs have a match-specific component which can be learned only through experience (e.g., Jovanovic, 1979).

The prior probability that the agent has high ability is given by ρ, and this is common knowledge. The agent's initial estimate of the probability to succeed if he exerts effort is $\hat{\theta} \equiv E[\theta] = \rho\theta_H + (1-\rho)\theta_L$. This estimate, akin to self-confidence, is increasing in ρ. With some abuse of language, we equate ρ with the agent's *initial self-confidence*.

Before deciding on the bonus, the principal observes performance y_t in each period $t = 1, 2$. However, performance cannot be observed by the agent himself or by any third party.[6] In many circumstances, the principal is in a better position to evaluate performance, a fact that is illustrated by the ubiquity of performance feedback systems in organizations (Gibbs, 1991). A manager often knows better how useful an employee's work was for the firm, a professor can better judge the quality of a student's paper, etc. This is particularly relevant for complex jobs or for agents who are in their learning phase, e.g., at school or after a promotion to a new position.

Nevertheless, the agent forms some impression of his performance, even if less accurate than what the principal observes: he receives a signal $\sigma \in [0, 1)$ that is imperfectly informative about the first-period outcome y_1. Let $G(\sigma|y_1) = G_{y_1}(\sigma)$ be the distribution function of σ conditional on $y_1 \in \{S, F\}$, and $g(\sigma|y_1) = g_{y_1}(\sigma)$ the conditional density function, assumed to be continuous and positive on $[0, 1)$. The realized signal σ is private information to the agent, but the conditional distribution functions are common knowledge. A high signal brings good news about the outcome. Formally:

Assumption 1. The likelihood ratio $l(\sigma) \equiv g_S(\sigma)/g_F(\sigma)$ is continuous in σ and assumes all values in $[0, +\infty)$. Furthermore, the monotone likelihood ratio property (MLRP) is satisfied: $l(\sigma)$ is everywhere increasing in σ.

An important assumption we make is that the signal σ is the agent's *private information*. The agent's perception of his performance is (partly) unobservable to the principal for several reasons. The principal cannot continually monitor the agent, so she will not be aware of all the information the agent receives about his performance. The agent might get feedback from other parties, such as colleagues in the workplace, other academics at conferences, or parents in a schooling context. Even if the principal knew all the information the agent gets, she might simply not know how the agent interprets it. For instance, the agent's interpretation will be based on his past experience with similar tasks that the principal has not observed.

Our paper is not unique in assuming that the receiver in a signaling game has private information unavailable to the sender. In a similar model, but with contractible bonuses, Bénabou and Tirole (2003) also assume that the less informed

[4] These simplifying assumptions are not crucial for our main results. See Section 4 for a brief analysis of more realistic cases where the principal does not observe the agent's effort or the agent's low effort does not imply that the project is certain to fail.

[5] From a different perspective, a person can at times manipulate self-confidence to boost her own motivation in struggling with self-control problems, as in Carrillo and Mariotti (2000) or Bénabou and Tirole (2002) or for hedonic reasons as in Köszegi (2006).

[6] Of course, since the agent gets payoff V if his project succeeds and 0 otherwise, he eventually learns about his performance. We assume that the agent's intrinsic payoffs from the first- and second-period projects (V or 0) are realized after his interaction with the principal is over.

party (the agent) receives a private signal. This signal plays a crucial role in their analysis, and creates what they call the *trust effect*. In our model, it means that a principal who has good news about the agent's ability θ expects the agent to have received a high private signal σ, and may consequently expect him to be highly motivated without any additional feedback. In a different context, Feltovich et al. (2002) investigate how private information by a receiver in a signaling game can lead to *countersignaling*, that is, the sender's action becoming non-monotonic in his "type."

In our setting, the existence of a private signal is not crucial but it helps in selecting a unique equilibrium. Because of the important role the private signal plays in the equilibrium selection, we also consider the cases where the agent gets no signal and where the signal is publicly observable before the principal chooses her policy. In Section 4.3 we show this yields qualitatively similar results.

To make things interesting, we assume that the payoff V that the agent receives in case of success is sufficiently high to motivate him to work even if he expects no bonus, provided that his self-confidence is high enough. The payoff V might represent intrinsic benefits from good performance (Deci and Ryan, 1985), or future rewards from success, such as job prospects. We do not model this explicitly.

Assumption 2. Were the agent perfectly informed about his type, he would exert effort without a bonus if and only if he had high ability: $\theta_L V < c \leqslant \theta_H V$.

Timing. We omit the agent's signal about his second-period performance as it has no impact on the analysis. Otherwise, the timing of the game is the same in both periods: first, the agent makes the effort decision; then, the principal observes performance and the agent his signal; finally, the principal pays a bonus (if she wishes to).

3. Equilibrium analysis

In this section we analyze perfect Bayesian equilibria (PBE) of the game (see Fudenberg and Tirole, 1991, for a standard formal definition). We first characterize equilibria of continuation "subgames" that start with the principal's choice of bonus b_1, and then of the whole game. We use the equilibrium refinement D1 to identify a unique equilibrium outcome. All proofs not given in the main text are in Appendix A.

3.1. Continuation equilibria

Note, first, that *no bonus is paid in period* 2: paying a bonus after all decisions have been made is costly and meaningless since it has no impact on the agent's behavior. Since only b_1 is non-trivial, in the remainder we simply write b for b_1.

We specify the belief functions for the agent as follows. His belief that success occurred in the first period, $y_1 = S$, given that he exerted effort e_1, received bonus b from the principal and observed signal σ, is denoted by $\lambda(e_1, b, \sigma)$, with:

$$\lambda(e_1, b, \sigma) = \frac{g_S(\sigma)\tilde{\lambda}(e_1, b)}{g_S(\sigma)\tilde{\lambda}(e_1, b) + g_F(\sigma)(1 - \tilde{\lambda}(e_1, b))} \tag{2}$$

where $\tilde{\lambda}(e_1, b)$ is the common belief that $y_1 = S$ of different types of agent independent of signal σ (see Fudenberg and Tirole, 1991). The agent's effort decision in period 2 is based on his beliefs about his ability. Let $\mu(e_1, b, \sigma)$ be the agent's belief that $\theta = \theta_H$:

$$\mu(e_1, b, \sigma) = \lambda(e_1, b, \sigma)\mu_S + \big(1 - \lambda(e_1, b, \sigma)\big)\mu_F, \tag{3}$$

where $\mu_y = \Pr\{\theta = \theta_H | y_1 = y\}$, $y \in \{S, F\}$.

The agent only exerts effort if his expected intrinsic payoff is high enough (since $b_2 = 0$). Assuming throughout that he exerts effort if indifferent, $e_2 = 1$ if and only if:

$$\big[\mu(e_1, b, \sigma)\theta_H + \big(1 - \mu(e_1, b, \sigma)\big)\theta_L\big]V \geqslant c. \tag{4}$$

Lemma 1. *If the agent exerts no effort in period* 1, $e_1 = 0$, *the principal does not pay a bonus*, $b = 0$, *and the agent chooses to exert effort in period* 2 *if and only if* $\rho \geqslant \bar{\rho}$, *where* $\bar{\rho}$ *is determined by*

$$\big[\bar{\rho}\theta_H + (1 - \bar{\rho})\theta_L\big]V = c. \tag{5}$$

Proof. If $e_1 = 0$, the agent is certain to fail, so no player gets any new information about ability θ. The agent then keeps the prior beliefs about his ability after any bonus b: $\mu(0, b, \sigma) = \rho$. Indeed, for the equilibrium bonus(es) b, this claim follows from Bayesian updating and the fact that the principal herself gets no information. For out-of-equilibrium bonuses this follows from the same fact and the "no signaling what you do not know" condition of the PBE (Fudenberg and Tirole, 1991). The agent chooses $e_2 = 1$, if his expected payoff from working is non-negative, i.e., $\rho \geqslant \bar{\rho}$. Since a bonus has no impact on period 2 effort, in equilibrium $b = 0$ if $e_1 = 0$. □

If the agent exerted effort in period 1, his posterior beliefs μ can depend on the new information: his signal σ about y_1, and the bonus b, that potentially conveys (some of) the principal's private information about performance y_1.

Lemma 2. *Let $e_1 = 1$. There exists a threshold level $\tilde{\sigma}(b) \in [0,1]$ such that the agent chooses $e_2 = 1$ if and only if $\sigma \geqslant \tilde{\sigma}(b)$. Moreover, there exist threshold levels of the initial self-confidence $\tilde{\rho}_S$ and $\tilde{\rho}_F > \tilde{\rho}_S$ such that in any continuation equilibrium*:

(i) *If $\rho < \tilde{\rho}_S$, $e_2 = 0$ after any signal σ and bonus b (i.e., $\tilde{\sigma}(b) = 1$);*
(ii) *If $\rho \geqslant \tilde{\rho}_F$, $e_2 = 1$ after any signal σ and bonus b (i.e., $\tilde{\sigma}(b) = 0$);*
(iii) *If $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$, e_2 is sensitive to information: $e_2 = 1$ if the agent is sure of success, $\tilde{\lambda}(1,b) = 1$, or if $0 < \tilde{\lambda}(1,b) < 1$ and $\sigma \geqslant \tilde{\sigma}(b)$; $e_2 = 0$ otherwise.*

Proof. Assume $e_1 = 1$. From (2) and (3) it follows immediately that $\mu(\cdot)$ is (weakly) increasing in σ; therefore, for any b there is a cut-off level $\tilde{\sigma}(b) \in [0,1]$ such that $e_2 = 1$ iff $\sigma \geqslant \tilde{\sigma}(b)$. Using (2) and (3) and the fact that by definition $\mu_S = \rho\theta_H/\hat{\theta}$, $\mu_F = \rho(1-\theta_H)/(1-\hat{\theta})$, where $\hat{\theta} = \rho\theta_H + (1-\rho)\theta_L$, inequality (4) with $e_1 = 1$ can be rewritten as:

$$l(\sigma)\tilde{\lambda}(1,b)(1-\hat{\theta})\big(\rho\theta_H - (1-\rho)\theta_L\phi\big) \geqslant \big(1-\tilde{\lambda}(1,b)\big)\hat{\theta}\big((1-\rho)(1-\theta_L)\phi - \rho(1-\theta_H)\big), \tag{6}$$

where $l(\sigma) \equiv g_S(\sigma)/g_F(\sigma)$ is the likelihood ratio of receiving signal σ, and $\phi \equiv \frac{c-\theta_L V}{\theta_H V - c}$ is the ratio of the expected loss from choosing $e = 1$ for the low ability agent, to the expected gain from choosing $e = 1$ for the high-ability agent. If $\rho < \tilde{\rho}_S$ determined by

$$\frac{\tilde{\rho}_S}{1-\tilde{\rho}_S} = \frac{\theta_L}{\theta_H}\phi, \tag{7}$$

the non-positive left-hand side (LHS) of (6) is strictly smaller than the non-negative right-hand side (RHS), so $e_2 = 0$, irrespective of signal σ and beliefs $\tilde{\lambda}(1,b)$; this proves claim (i). If $\rho \geqslant \tilde{\rho}_F$, where:

$$\frac{\tilde{\rho}_F}{1-\tilde{\rho}_F} = \left(\frac{1-\theta_L}{1-\theta_H}\right)\phi, \tag{8}$$

the (non-negative) LHS of (6) is larger than the (non-positive) RHS and $e_2 = 1$, irrespective of signal σ and beliefs $\tilde{\lambda}(1,b)$; this proves claim (ii). For $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$ and $\tilde{\lambda}(1,b) > 0$, simple transformations show that (6) is equivalent to

$$l(\sigma) \geqslant \left(\frac{1-\tilde{\lambda}(1,b)}{\tilde{\lambda}(1,b)}\right)\left(\frac{\hat{\theta}}{1-\hat{\theta}}\right)A, \tag{9}$$

where the compound parameter A is defined as

$$A = \frac{\Pr\{y_1 = F | e_1 = 1\} E[c - \theta V | y_1 = F,\ e_1 = 1]}{\Pr\{y_1 = S | e_1 = 1\} E[\theta V - c | y_1 = S,\ e_1 = 1]} = \frac{(1-\rho)(1-\theta_L)\phi - \rho(1-\theta_H)}{\rho\theta_H - (1-\rho)\theta_L\phi}. \tag{10}$$

Claim (iii) now follows from Assumption 1 that $l(\sigma)$ is a continuous increasing function that assumes all values in $[0, +\infty)$; $\tilde{\sigma}(b)$ is the unique value of σ that turns (9) into equality. If $\tilde{\lambda}(1,b) = 0$, $e_2 = 0$ from (6). If $\rho = \tilde{\rho}_S$, the LHS of (6) equals 0 and the RHS is positive and $\tilde{\sigma}(b) = 1$, unless $\tilde{\lambda}(1,b) = 1$; in the latter case $\tilde{\sigma}(b) = 0$. □

Lemma 2 is intuitive. If his initial self-confidence is sufficiently low, $\rho < \tilde{\rho}_S$, not exerting effort is optimal for the agent even if he is sure that the first-period project was successful: even the best news is insufficient to compensate for the initial pessimism. Similarly, for sufficiently high initial self-confidence, the agent exerts effort in period 2 even if he is sure that a failure occurred.

For intermediate levels of initial self-confidence $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$, were the agent to know the first-period outcome, it would be optimal for him to work after success and to shirk after failure. Therefore, within this range of parameters, the agent's reaction is sensitive to the news he gets: only if sufficiently good news arrives, he works in period 2. From the proof of Lemma 2 it follows that the agent's expected second-period effort is higher (i.e., threshold $\tilde{\sigma}(b)$ is lower) if, other things being fixed, his initial self-confidence ρ is higher or his interpretation of bonus b is more optimistic ($\tilde{\lambda}(1,b)$ is higher).

Consider now the principal's behavior in period 1. For $\rho \notin [\tilde{\rho}_S, \tilde{\rho}_F)$, the principal gives no bonus in period 1, since she is not able to influence the agent's behavior in period 2. From here on, we focus on the range of initial self-confidence $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$, where the agent's second-period behavior is sensitive to information about the outcome of the first period, and, therefore, she may try to signal through a bonus that the agent was successful.

The following lemma shows that in equilibrium a higher bonus increases the likelihood of effort, and is helpful in deriving the equilibrium bonus.

Lemma 3. *For any bonuses $b > b'$ offered in equilibrium with positive probability, it must be that $\tilde{\sigma}(b) < \tilde{\sigma}(b')$.*

Proof. Suppose $b > b'$ are equilibrium bonuses but $\tilde{\sigma}(b) \geqslant \tilde{\sigma}(b')$. The lower bonus b' (weakly) increases the likelihood of effort and thus makes the principal strictly better off, so b cannot be an equilibrium bonus. □

Without further restrictions on beliefs, there is generally a continuum of PBE for $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$. To identify a unique equilibrium outcome and eliminate PBE sustained by "unreasonable" out-of-equilibrium beliefs, we apply the equilibrium refinement D1 as defined in Cho and Kreps (1987) and based on the idea of Divinity of Banks and Sobel (1987). Our game is not a standard signaling game for which D1 was originally developed: first, because the signal space is a continuum, and, second, because the agent receives a private signal. We thus adapt D1 to our game in the spirit of the original definition. For these reasons, we cannot apply any of the results derived for standard signaling games, and derive results (such as equilibrium existence and uniqueness) independently for the current model.[7]

We call "type-y_1 principal" a principal who observed $y_1 \in \{S, F\}$. For any PBE of the game let $u_P^*(y_1)$ be the equilibrium expected second-period payoff of type-y_1 principal net of first-period bonus. Let $MBR(b)$ be the set of the agent's mixed strategies such that each is a best response reaction to a bonus b given some beliefs $\tilde{\lambda}(e_1, b)$. For an out-of-equilibrium bonus b let $D(y_1, b) \subseteq MBR(b)$ be the set of the agent's reactions that would make a type-y_1 principal strictly prefer to deviate to b from her equilibrium strategy. Similarly, let $D^o(y_1, b) \subseteq MBR(b)$ be the set of the agent's reactions that would make a principal of type y_1 indifferent.

Definition 1 *(D1).* A PBE in our model satisfies D1 if the agent's beliefs after out-of-equilibrium bonus b are such that the agent assigns probability 0 to the principal having type y_1 if $D(y_1, b) \cup D^o(y_1, b) \subseteq D(y_1', b)$ for some $y_1' \neq y_1$ such that $D(y_1', b) \neq \emptyset$.

Thus, after an out-of-equilibrium bonus b, beliefs are restricted to $\tilde{\lambda}(e_1, b) = 0$ whenever $D(S, b) \cup D^o(S, b) \subseteq D(F, b)$ and $D(F, b) \neq \emptyset$, and to $\tilde{\lambda}(e_1, b) = 1$ whenever $D(F, b) \cup D^o(F, b) \subseteq D(S, b)$ and $D(S, b) \neq \emptyset$. The logic of D1 refinement is that the agent should interpret b as coming from type-y_1' principal if she gains from this deviation for a strictly larger set of the agent's reactions than type-y_1 principal, and is consequently "more likely" to gain from this deviation.

According to Lemma 2, the agent's best response reaction to a bonus b can be fully described by a threshold $\tilde{\sigma}(b) \in [0, 1]$. From the proof of that lemma it also follows that when $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$, each $\tilde{\sigma} \in [0, 1]$ determines a best-response reaction for some beliefs $\tilde{\lambda}(e_1, b)$. Hence, $MBR(b) = [0, 1]$. Denote by $\hat{\theta}_F$ and $\hat{\theta}_S$ the probability of success after effort, conditional on failure or success in the first period: $\hat{\theta}_y = E[\theta \mid e_1 = 1,\ y_1 = y] = \mu_y \theta_H + (1 - \mu_y)\theta_L$ for $y \in \{S, F\}$. Then

$$u_P^*(y_1) = \hat{\theta}_{y_1}\big(1 - G_{y_1}\big(\tilde{\sigma}(\tilde{b})\big)\big)W - \tilde{b},$$

where $\tilde{b}$ is paid after outcome y_1 with positive probability in equilibrium.

For an out-of-equilibrium bonus b paid after outcome y_1, two cases are possible. In the first case there exists $\tilde{\sigma}(b) \in [0, 1)$ that makes the principal indifferent between deviating and not deviating, $u_P^*(y_1) = \hat{\theta}_{y_1}(1 - G_{y_1}(\tilde{\sigma}(b)))W - b$, in which case $D^o(y_1, b) = \{\tilde{\sigma}(b)\}$. Since our game is "monotonic" in that the principal's expected payoff is strictly decreasing in threshold $\tilde{\sigma}$ that determines the agent's best-response, in this case $D(y_1, b) = [0, \tilde{\sigma}(b))$. In the second case no such $\tilde{\sigma}(b)$ exists: the bonus is too large and $D^o(y_1, b) = D(y_1, b) = \emptyset$. Note that the case $D^o(y_1, b) = \{1\}$ and $D(y_1, b) = [0, 1)$ is impossible: bonus b would then dominate the equilibrium bonus.

The preceding analysis implies that in our model the condition $D(y_1, b) \cup D^o(y_1, b) \subseteq D(y_1', b)$ in the definition of D1, is equivalent to $D^o(y_1, b) \subseteq D(y_1', b)$, which corresponds to *Never a Weak Best Response* (NWBR) refinement for signaling games (Cho and Kreps, 1987). It is this simpler condition that we check in the proofs.[8]

Under D1, the continuation equilibrium outcome is (generically) unique: it is either pooling or semi-separating. To gain intuition, we now derive these two equilibria in detail before stating Proposition 1 that classifies all continuation equilibria for the case $e_1 = 1$.

3.1.1. Pooling equilibrium

We first look for a *pooling* equilibrium, in which the principal gives the same bonus $\tilde{b}$ regardless of y_1. Given this bonus, the agent only works for signals exceeding threshold $\tilde{\sigma} > 0$ that satisfies $l(\tilde{\sigma}) = A$ (see (9) satisfied as equality, where beliefs must be $\tilde{\lambda}(1, \tilde{b}) = \hat{\theta}$ according to Bayes' rule). Consider a deviation to bonus $\hat{b} = \tilde{b} + \varepsilon$ with $\varepsilon > 0$. For $\tilde{b}$ to be the equilibrium bonus, such a deviation should not be profitable. A necessary condition for this is that the agent not have beliefs $\tilde{\lambda}(1, \hat{b}) = 1$ for ε small enough (otherwise, the principal would achieve an upward jump in the probability of effort at an infinitesimal cost by giving $\hat{b}$ instead of $\tilde{b}$).

Let $\hat{\sigma}$ be the agent's reaction to bonus $\hat{b}$ that makes the principal indifferent between deviating or not after failure (it exists for ε small enough):

$$\hat{\theta}_F\big(1 - G_F(\tilde{\sigma})\big)W - \tilde{b} = \hat{\theta}_F\big(1 - G_F(\hat{\sigma})\big)W - \hat{b}. \tag{11}$$

[7] Van Damme (1987), Feltovich et al. (2002), Bénabou and Tirole (2003), and Bernheim (1994) are among other papers that use similar refinements for "non-standard" signaling games. Cho and Sobel (1990), in their analysis of refinements for "monotonic" signaling games, allow for infinite strategy spaces, although they do not consider games with two-sided information asymmetry such as ours and some of the aforementioned.

[8] Our definition of NWBR takes into account the fact that the principal has only two possible types. Otherwise type–action pair (y_1, b) would be "deleted" if $D^o(y_1, b) \subseteq \bigcup_{y_1' \neq y_1} D(y_1', b)$. Equivalence between D1 and NWBR for a general class of monotonic signaling games (which does not, however, include our game) is proved in Cho and Sobel (1990).

From (11) and $\hat{b} > \tilde{b}$, it follows that $\tilde{\sigma} > \hat{\sigma}$: the agent works with a higher probability after the larger bonus $\hat{b}$. Beliefs $\tilde{\lambda}(1, \hat{b}) < 1$, are consistent with D1 if $D^o(F, \hat{b}) \nsubseteq D(S, \hat{b})$, so given the agent's reaction $\hat{\sigma}$ the principal should not gain by deviating to $\hat{b}$ after success:

$$\hat{\theta}_S\big(1 - G_S(\tilde{\sigma})\big)W - \tilde{b} \geqslant \hat{\theta}_S\big(1 - G_S(\hat{\sigma})\big)W - \hat{b}. \tag{12}$$

Combining (11) and (12) yields $\hat{\theta}_S(G_S(\tilde{\sigma}) - G_S(\hat{\sigma})) \leqslant \hat{\theta}_F(G_F(\tilde{\sigma}) - G_F(\hat{\sigma}))$. Dividing both sides by $\tilde{\sigma} - \hat{\sigma} > 0$, and taking the limit $\varepsilon \to +0$, one gets $\hat{\theta}_S g_S(\tilde{\sigma}) \leqslant \hat{\theta}_F g_F(\tilde{\sigma})$ or:

$$l(\tilde{\sigma}) \leqslant \frac{\hat{\theta}_F}{\hat{\theta}_S}. \tag{13}$$

Conversely, assume that (13) is satisfied and consider a possible deviation $\hat{b} = \tilde{b} + \varepsilon$ with *arbitrary* (not necessarily small) $\varepsilon > 0$. If $D^o(S, \hat{b}) = \emptyset$, beliefs $\tilde{\lambda}(1, \hat{b}) = 0$ are consistent with D1. Otherwise, if $D^o(S, \hat{b}) = \{\hat{\sigma}\}$ for some $\hat{\sigma} < \tilde{\sigma}$, the MLRP implies that:

$$\frac{G_S(\tilde{\sigma}) - G_S(\hat{\sigma})}{G_F(\tilde{\sigma}) - G_F(\hat{\sigma})} < l(\tilde{\sigma}).^9 \tag{14}$$

From condition (13) it then follows that $\hat{\theta}_S(G_S(\tilde{\sigma}) - G_S(\hat{\sigma})) < \hat{\theta}_F(G_F(\tilde{\sigma}) - G_F(\hat{\sigma}))$ and $D^o(S, \hat{b}) \subseteq D(F, \hat{b})$, so beliefs $\tilde{\lambda}(1, \hat{b}) = 0$ satisfy D1. Given these beliefs the principal has no incentive to deviate to a higher bonus, so they are consistent with the equilibrium. Similarly one shows that the principal does not want to deviate to *any* lower bonus (more precisely, $\tilde{\lambda}(1, \hat{b}) = 0$ for $\hat{b} < \tilde{b}$) if and only if:

$$l(\tilde{\sigma}) \geqslant \frac{\hat{\theta}_F}{\hat{\theta}_S}. \tag{15}$$

There are two cases. Either inequalities (13) and (15) are satisfied simultaneously, which implies $A = \hat{\theta}_F / \hat{\theta}_S$. In this case for any $\tilde{b} \in [0, \tilde{b}_P]$ with $\tilde{b}_P = \hat{\theta}_F(1 - G_F(\tilde{\sigma}(\tilde{b}_P)))W$ there exists a pooling continuation equilibrium in which the principal gives $\tilde{b}$, and the agent works if $\sigma \geqslant \tilde{\sigma}$ (with $l(\tilde{\sigma}) = A$) after the equilibrium bonus. In the generic case, if $A \neq \hat{\theta}_F / \hat{\theta}_S$, the only candidate for the pooling equilibrium is the one with $\tilde{b} = 0$. Limited liability prevents the principal from downward deviations. For the upward deviations to be unprofitable, (13) must be satisfied, which implies a restriction on parameters:

$$A \leqslant \frac{\hat{\theta}_F}{\hat{\theta}_S}. \tag{16}$$

Our analysis implies that this condition is both necessary and sufficient for the existence of some pooling equilibrium. The equilibrium is sustained by beliefs $\tilde{\lambda}(1, \hat{b}) = 0$ for any $\hat{b} \neq \tilde{b}$ and these beliefs satisfy D1.

3.1.2. Semi-separating equilibrium

In a *semi-separating* equilibrium, the principal offers $\tilde{b}_S$ after success, and randomizes between $\tilde{b}_S$ and $\tilde{b}_F$ after failure (with probabilities $\tilde{x}_F$ and $1 - \tilde{x}_F$). Bonus $\tilde{b}_F$ reveals failure in this case. It follows that $\tilde{b}_F = 0$, since there is no reason to incur a cost for conveying a negative signal. After $\tilde{b}_F = 0$, the agent does not work (recall that $\hat{\theta}_F V < c$ for $\rho \in [\rho_S, \rho_F)$) so the payoff for the principal is zero. Let $\tilde{\sigma}_S$ be the threshold signal such that the agent works after bonus $\tilde{b}_S$ if $\sigma \geqslant \tilde{\sigma}_S$. The principal must be indifferent between $\tilde{b}_S$ and $\tilde{b}_F$ after failure (to be willing to mix):

$$\hat{\theta}_F\big(1 - G_F(\tilde{\sigma}_S)\big)W - \tilde{b}_S = 0. \tag{17}$$

Condition (17) determines bonus $\tilde{b}_S$. For the principal not to be willing to deviate to bonus $\hat{b} = \tilde{b}_S - \varepsilon$ with $\varepsilon > 0$, the agent should have beliefs $\tilde{\lambda}(1, \hat{b}) < 1$. D1 implies that a necessary condition for this is that the agent's reaction $\hat{\sigma}$ that makes the principal indifferent to deviation to $\hat{b}$ after failure,

$$\hat{\theta}_F\big(1 - G_F(\tilde{\sigma}_S)\big)W - \tilde{b}_S = \hat{\theta}_F\big(1 - G_F(\hat{\sigma})\big)W - \hat{b}, \tag{18}$$

does not make deviation to $\hat{b}$ attractive after success:

$$\hat{\theta}_S\big(1 - G_S(\tilde{\sigma}_S)\big)W - \tilde{b}_S \geqslant \hat{\theta}_S\big(1 - G_S(\hat{\sigma})\big)W - \hat{b}. \tag{19}$$

As in the analysis of pooling equilibrium, for $\varepsilon \to +0$ conditions (18) and (19) imply (using MLRP):

$$l(\tilde{\sigma}_S) \geqslant \frac{\hat{\theta}_F}{\hat{\theta}_S}. \tag{20}$$

[9] To prove this (following Milgrom, 1981): suppose $l(x) < l(z)\ \forall z \in [x, y]$. Then, since $l(z) \equiv g_S(z)/g_F(z)$, $g_F(z)l(x) < g_S(z)$. This implies $\int_x^y g_F(z)l(x)\,dz < \int_x^y g_S(z)\,dz$. Integration yields: $l(x) < [G_S(y) - G_S(x)]/[G_F(y) - G_F(x)]$. Similarly, for $l(z) < l(y)\ \forall z \in [x, y]$, it follows that $[G_S(y) - G_S(x)]/[G_F(y) - G_F(x)] < l(y)$.

Condition (20) implies that $\tilde{\sigma}_S > 0$ (otherwise, if $\tilde{\sigma}_S = 0$, then $l(0) \geqslant \hat{\theta}_F/\hat{\theta}_S > 0$: a contradiction to Assumption 1 which stipulates $l(0) = 0$): the agent does not always work after bonus $\tilde{b}_S$. The principal should also not be willing to deviate to a (slightly) higher bonus $\hat{b} = \tilde{b}_S + \varepsilon$ to separate the success outcome (which would induce the agent to work for any σ). This implies:

$$l(\tilde{\sigma}_S) \leqslant \frac{\hat{\theta}_F}{\hat{\theta}_S}. \tag{21}$$

Hence, combining (20) and (21) gives as the only possibility:

$$l(\tilde{\sigma}_S) = \frac{\hat{\theta}_F}{\hat{\theta}_S}. \tag{22}$$

As in the case of pooling, if (22) holds, then $\tilde{\lambda}(1, \hat{b}) = 0$, for any out-of-equilibrium bonus $\hat{b}$, satisfies D1 and supports the equilibrium. By Bayes' rule, $\tilde{\lambda}(1, \tilde{b}_S) = \hat{\theta}/(\hat{\theta} + (1 - \hat{\theta})\tilde{x}_F)$, thus from (9), the agent's reaction is given by

$$l(\tilde{\sigma}_S) = \tilde{x}_F A. \tag{23}$$

Condition (22) determines $\tilde{\sigma}_S$, (23) gives $\tilde{x}_F$ and (17) determines $\tilde{b}_S$. Finally note that conditions (22) and (23) imply that $A \geqslant \hat{\theta}_F/\hat{\theta}_S$ as $\tilde{x}_F \leqslant 1$.

3.1.3. Rewards, self-confidence, and motivation

The following proposition describes the continuation equilibrium outcomes after $e_1 = 1$; its complete proof is given in Appendix A.

Proposition 1. *For $e_1 = 1$ there always exists a (generically) unique*[10] *continuation equilibrium outcome supported by beliefs satisfying the D1 criterion. Furthermore, there exists a unique value of initial self-confidence ρ^*, defined as the (unique) solution to $A = \hat{\theta}_F/\hat{\theta}_S$, such that:*

(i) *If $\rho \in [\tilde{\rho}_S, \rho^*)$, the unique continuation equilibrium outcome is semi-separating: in period 1, the principal always gives a bonus $\tilde{b}_S = \hat{\theta}_F(1 - G_F(\tilde{\sigma}_S))W$ after success, and randomizes between $\tilde{b}_S$ and $\tilde{b}_F = 0$ after failure with probabilities $\tilde{x}_F = l(\tilde{\sigma})/A$ and $1 - \tilde{x}_F$ respectively.*[11] *After receiving bonus $\tilde{b}_S$ the agent works if his signal σ exceeds the threshold $\tilde{\sigma}_S$, determined by $l(\tilde{\sigma}_S) = \hat{\theta}_F/\hat{\theta}_S$; after getting no bonus the agent does not work: $\tilde{\sigma}_F = 1$.*
(ii) *If $\rho \in (\rho^*, \tilde{\rho}_F)$, the unique continuation equilibrium outcome is pooling and no bonus is ever given by the principal. The agent works if his signal exceeds the threshold $\tilde{\sigma}$ determined by $l(\tilde{\sigma}) = A$.*
(iii) *If $\rho \notin [\tilde{\rho}_S, \tilde{\rho}_F)$, the unique continuation equilibrium outcome is pooling: the principal pays no bonus; if $\rho \in (0, \tilde{\rho}_S)$ the agent exerts no effort after any signal σ; if $\rho \in [\tilde{\rho}_F, 1)$, the agent exerts effort after any signal σ.*

Remark 1. Proposition 1 can be re-stated in terms of disutility of effort: there exists a threshold value c^* defined as the unique solution to $A = \hat{\theta}_F/\hat{\theta}_S$ such that the unique continuation equilibrium outcome is pooling for $c \in (\tilde{c}_F, c^*)$ and semi-separating for $c \in (c^*, \tilde{c}_S]$, with equilibrium strategies being the same as those described in parts (ii) and (i) of Proposition 1 respectively.

That the unique continuation equilibrium outcome for sufficiently high initial self-confidence or low disutility of effort is a pooling equilibrium with no bonus is quite intuitive. Indeed, for such parameters the threshold signal $\tilde{\sigma}$, which makes the agent indifferent between working or not in the second period in the pooling equilibrium, is relatively low. Low signals are more likely after failure (due to MLRP), so it is after failure that the principal gains more from a marginal decrease in $\tilde{\sigma}$. According to the logic of D1, this makes the agent interpret an (out-of-equilibrium) increase in the bonus as coming from the principal who observed a failure, thus undermining the principal's incentives to increase the bonus. In fact, a *decrease* in bonus would signal to the agent that the principal observed a success, but the agent's limited liability prevents negative bonuses.

A more interesting result of the proposition is the existence of the region where the principal does give a positive bonus in (continuation) equilibrium, and this bonus increases the agent's self-confidence. In this region, the agent is relatively unlikely to exert effort in period 2, so, were the principal to play a pooling strategy, the threshold $\tilde{\sigma}$ would be high. High signals are more likely after success, and in this case it is a principal who observed a success, who would gain more from a marginal decrease in $\tilde{\sigma}$. This time the agent would interpret an (out-of-equilibrium) increase in bonus as a signal of success, thus destroying the pooling equilibrium. By paying a positive bonus, the principal sends a costly credible signal of success.

[10] The equilibrium outcome is unique unless $\hat{\theta}_F/\hat{\theta}_S = A$, in which case there is a continuum of pooling equilibria.

[11] To be more precise, when $\rho = \tilde{\rho}_S$, the equilibrium outcome is separating: $\tilde{x}_F = 0$; $\tilde{\sigma}_S$ and $\tilde{b}_S$ are given by the same formulae as for the semi-separating case.

Proposition 1 also rules out equilibria where good performance is completely separated from bad performance. The proof given in Appendix A shows that generic inexistence of separating equilibria relies on the assumption that "really bad" signals exist, i.e., $l(0) = 0$. It may seem counterintuitive that even bad performance gets rewarded, but there is considerable evidence that this often occurs in practice (Prendergast, 1999). Supervisors are reluctant to differentiate good from bad performance, resulting in a well documented compression of ratings. Prendergast (1999) conjectured that a possible reason could be that the supervisor avoids discouraging the agent by revealing poor performance to him. A similar view is shared by Beer (1990) who notes that many people are rated on the high side because managers "…do not want to damage an employee's self-esteem, thereby demotivating the employee.…"[12] This interpretation fits well our result.[13]

Another important point is that for levels of initial self-confidence below the threshold level ρ^*, the size of the reward $\tilde{b}_S$ is proportional to W. Intuitively, if the principal derives higher benefits from a success, a bonus of a given size is relatively less costly, and she needs to increase the bonus to keep it credible in equilibrium. This means that the agent can be very glad to get a seemingly negligible reward, provided that it is given by the principal who does not have too much interest in the agent's performance. A small reward from a division manager with small stakes, or a distinguished professor's willingness to discuss one's work may be very encouraging. In these cases it is not the reward *per se* that motivates, but foremost its informational content.

Comparative statics with respect to ρ When the initial self-confidence ρ crosses the threshold value ρ^*, the equilibrium switches from the semi-separating to the pooling regime. In the pooling regime the only relevant equilibrium parameter – the probability that the agent works in the second period – increases with ρ.

In the semi-separating regime, the impact of ρ on the size and frequency of the positive bonus is ambiguous. This is not surprising: the equilibrium strategy is determined by the ratio of the agent's ability estimates, $r(\rho) = \hat{\theta}_F/\hat{\theta}_S$, which itself varies non-monotonically in ρ. At $\rho = 0$ and $\rho = 1$ this ratio equals 1 since there is no uncertainty about his ability and the interim outcome does not bring new information. For intermediate values of ρ, the principal's estimate of the agent's ability is lower after observing a failure, and consequently, $r(\rho) < 1$. It can be easily checked that there exists $\rho_r \in (0, 1)$ such that $r(\rho)$ is decreasing on $(0, \rho_r)$ and increasing on $(\rho_r, 1)$. Whether $r(\rho)$ is decreasing, increasing, or non-monotonic for $\rho \in [\tilde{\rho}_S, \rho^*)$, depends on the specific parametrization. If $r'(\rho) > 0$, the threshold $\tilde{\sigma}$ and the probability of a positive bonus after failure $\tilde{x}_F$ increase in ρ. The impact on $\tilde{b}_S$ is ambiguous. If $r'(\rho) < 0$, $\tilde{\sigma}$ decreases in ρ, $\tilde{b}_S$ increases, but the impact on $\tilde{x}_F$ is ambiguous.

Comparative statics with respect to c The ratio $\hat{\theta}_F/\hat{\theta}_S$, does not depend on c, while A increases in c. First, this implies that in the pooling equilibrium the probability that the agent works period 2 decreases in c. Second, an increase in c causes a switch in equilibrium regimes when it crosses the threshold c^*. In this case the equilibrium bonus rises discontinuously from 0 to $b_S = \hat{\theta}_F(1 - G_F(\tilde{\sigma}_S))W$. Third, in the semi-separating regime the equilibrium value of the positive bonus, b_S, and the agent's strategy, $\tilde{\sigma}$, do not depend on c, but the probability of a positive bonus after a failure, x_F, decreases. The agent is less likely to work in the case of the absence of feedback, so that the bonus must provide a stronger positive signal.

3.2. The first-period choice of effort

To fully characterize the equilibria of the whole game, it remains to specify the agent's choice of effort in the first period given the continuation equilibria. Fix all parameters except for ρ, and denote by $\mathcal{R}_1 \subset [0, 1]$ the set of values of ρ for which there exists a PBE (satisfying D1) in which the agent chooses $e_1 = 1$, and similarly $\mathcal{R}_0 \subset [0, 1]$ the set of values of ρ for which there exists an equilibrium with $e_1 = 0$.

Proposition 2.

(i) *An equilibrium of the whole game (satisfying D1) exists and is generically unique: $\mathcal{R}_1 \Delta \mathcal{R}_0 = (\mathcal{R}_1 \backslash \mathcal{R}_0) \cup (\mathcal{R}_0 \backslash \mathcal{R}_1)$ is an open set everywhere dense on $(0, 1)$.*
(ii) *If the agent has sufficiently high initial self-confidence, in equilibrium he chooses $e_1 = 1$: $[\bar{\rho}, 1) \subset \mathcal{R}_1 \backslash \mathcal{R}_0$ with $\bar{\rho}$ defined in (5); if the agent has sufficiently low initial self-confidence, in equilibrium he chooses $e_1 = 0$: $(0, \tilde{\rho}_S) \subset \mathcal{R}_0 \backslash \mathcal{R}_1$.*
(iii) *The agent's choice of effort for the intermediate range of self-confidence depends on other parameters of the model. Under any parametrization the set $\mathcal{R}_1 \cap [\tilde{\rho}_S, \rho^*)$ has positive measure: there exists a range of values of self-confidence such that in equilibrium the agent exerts effort in period 1, and rewards are given with positive probability.*

Since, as we argued before, the expected size of the bonus and the probability that it is given after failure may be non-monotonic in ρ, the equilibrium effort decision also need not be monotonic with respect to ρ (nor with respect to c[14]).

[12] Quoted in Gibbs (1991, p. 7).
[13] Other reasons are plausible, such as the one mentioned by Prendergast (1999, p. 30) that it is simply an unpleasant task to offer poor ratings to workers.
[14] A statement, similar to Proposition 2, could be formulated in terms of cost of effort c; recall also Remark 1 concerning the classification of continuation equilibria.

The second part of Proposition 2 cannot be strengthened for that reason. The third claim shows that there is a range of parameters where the agent chooses to exert effort in the first period and gets rewards with positive probability.

We can now emphasize the dual role that rewards play in our model (for those values of parameters under which the continuation equilibrium is semi-separating). First of all, a reward provides the agent with positive feedback and encourages him to continue working hard. This *information effect* increases the agent's incentive to exert effort in period 1 because information acquisition is valuable: it allows the agent to make a better decision concerning second-period effort. Second, after exerting effort, the agent foresees a bonus with positive probability, which gives him an additional stimulus to work hard. This is the traditional *direct incentive effect*.

The traditional incentive effect provides a rationale for giving feedback through the use of rewards that are useful to the agent. Indeed, in the model we assume that discretionary rewards are monetary transfers from the principal to the agent. However, the principal could give feedback to the agent in many other ways, including burning money or wasting time, as long as such actions were costly to her. But actions that directly benefit the agent create an additional incentive effect, so their use for providing feedback is more attractive (from an *ex ante* point of view). There is an interesting interaction between cheap talk communication and costly signaling ("burning money") in the standard Crawford–Sobel (1982) setting, investigated in Austen-Smith and Banks (2000) (with an important correction by Kartik, 2007). In our model, however, cheap talk cannot be effective (or "influential" in Austen-Smith and Banks' terminology) even in the presence of costly signaling because the principal would *always* want to send the same message inducing the highest expected effort.

The information and the incentive effects are closely interconnected. Given that performance is not verifiable, it is the principal's superior information that makes her promises to pay discretionary rewards credible. If, by contrast, both parties are equally well informed about interim performance, there is no information to be conveyed to the agent and the promise of a reward is not credible.

What if the principal can *ex ante* endow the agent with a technology to perfectly monitor the interim performance?[15] One natural interpretation is that the principal could share with the agent her knowledge on how to interpret various imperfect unverifiable signals that can be available at the interim stage. In addition, she could also increase the number of available signals by setting up appropriate accounting systems, employing coworkers as peer reviewers, etc. If the agent chooses to exert effort in the first period, how the principal's payoffs compare under the two regimes depends on the specific parameters: in the relevant range of parameters, $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$, under full transparency, the agent works with a larger probability after success (i.e., with probability 1) and with a smaller probability after failure (i.e., with probability 0).

Making performance fully transparent has two effects on the agent's first-period incentives. First, the principal loses the opportunity to pay discretionary bonuses, so the direct incentive effect disappears. On the other hand, since the equilibrium in the original setting is only *semi*-separating at best, the information effect is amplified. Proposition 3 shows that which effect is stronger depends crucially on the principal's stake in the project because of the impact the stake has on the bonus. A preference for not committing to share information has some empirical relevance, reflected in the common practice of using monitoring systems that provide little information even though more informative performance measures are available (Gibbs, 1991).

Proposition 3. *The range of initial self-confidence levels under which the principal strictly prefers not to commit ex ante to making performance observable has a positive measure, and it* (*weakly*) *expands when W increases.*

4. Extensions to the basic model

4.1. Unobservable effort

We next discuss some generalizations and extensions. First, we show that the main results also hold if we make the more standard assumption that the agent's first-period effort is not observed by the principal. Indeed, the results for the continuation equilibria remain valid, with the qualification that they apply to the continuation equilibria that occur on the equilibrium path, where the principal holds correct beliefs about the agent's first-period choice of effort. However, the analysis must be also extended to out-of-equilibrium situations, in which the agent deviates and chooses the first-period effort that is different from that expected by the principal.

When the continuation equilibrium is semi-separating, shirking in the first period becomes (weakly) more attractive than in the case with observable effort. Indeed, if the principal believes that the agent exerted effort, the agent may still count on getting a positive bonus with some probability even if he shirks, since the principal does not observe the deviation. Similarly, if the principal believes that the agent shirks in the first period, but the agent deviates and chooses to supply effort, such a deviation will not be detected and rewarded by the principal unless the project succeeds (in which case the principal updates her beliefs about the agent's strategy). As a result of this complementarity between the principal's

[15] We still assume in this section that a third party cannot verify performance, even if the principal and the agent are symmetrically informed about it (cf. for instance Levin, 2003, Section III). This is a standard assumption in the subjective evaluation literature, and MacLeod (2003) gives some examples. That said, we admit that making performance more transparent will often also make it verifiable, in which case performance-contingent contracts become credible.

expectations and the agent's behavior, in the model with unobservable effort, there is a set of parameters of a positive measure for which there are multiple equilibria. That is, for a set of values of $\rho \in [\tilde{\rho}_S, \rho^*)$ of a positive measure, there exist equilibria with $e_1 = 0$ and $e_1 = 1$, as well as a mixed-strategy equilibrium. When effort is unobservable, statement (iii) of Proposition 2 remains valid, and statement (ii) as well with the qualifier that the described equilibria are not unique (i.e., $\mathcal{R}_1 \backslash \mathcal{R}_0$ in the statement should be replaced by $\mathcal{R}_1$, and $\mathcal{R}_0 \backslash \mathcal{R}_1$ by $\mathcal{R}_0$). Proposition 3 remains valid as well.

4.2. Positive probability of success under shirking

In many circumstances even low effort may leave some chance for success. So, suppose that the model described in Section 2 is modified in one aspect: the probability of the agent's success after $e = 0$, is now given by θk with some $k \in (0, 1)$; we also assume that after $e_1 = 0$, the agent gets signal σ that has the same distribution conditional on success/failure as after $e_1 = 1$. To check the robustness of our findings, we are particularly interested in the limit case of $k \to 0$.

The continuation subgame that follows $e_1 = 0$, is now qualitatively similar to the subgame that follows $e_1 = 1$. More specifically, Lemma 2 is now true for the case $e_1 = 0$ with thresholds $\tilde{\rho}_S$ and $\tilde{\rho}_F$ replaced by appropriate $\tilde{\rho}_{S0}$ and $\tilde{\rho}_{F0}$; one can check that $\tilde{\rho}_{S0} = \tilde{\rho}_S$ and $\tilde{\rho}_{F0} \in (\tilde{\rho}_S, \tilde{\rho}_F)$. Proposition 1 applies to the continuation subgame after $e_1 = 0$, with $A, \hat{\theta}_F, \hat{\theta}_S$ and ρ^* replaced by $A_0, \hat{\theta}_{F0}, \hat{\theta}_{S0}$ and ρ_0^*. Once again we observe that the role of bonuses in our framework is to give credibility to communication, rather than to reward the agent's effort. Even if the agent shirks, success is good news about ability which makes it worthwhile for the principal to communicate this news to the agent via a costly reward.

It can be checked that in the limit, for k small enough, all the statements of Propositions 2 and 3 are valid. Indeed, $\tilde{\rho}_{F0} \to \bar{\rho}$ when $k \to 0$; hence, for all $\rho > \bar{\rho}$, in the limit when $k \to 0$ the continuation equilibria after the agent chooses $e_1 = 0$ are pooling with no bonus offered. In contrast, for all $\rho \in (\tilde{\rho}_S, \bar{\rho})$, the continuation equilibria after the agent chooses $e_1 = 0$, are almost separating (i.e., $\tilde{x}_{F0} \to 0$ when $k \to 0$). Hence, for $\rho \in (\tilde{\rho}_S, \bar{\rho})$ the shirking option may be quite attractive for the agent: he gets almost perfect information from the principal. Thus, $\mathcal{R}_1$, the set of values of ρ for which in equilibrium $e_1 = 1$, contracts, while $\mathcal{R}_0$, the set of values of ρ for which in equilibrium $e_1 = 0$, expands, compared to the basic model with $k = 0$.

However, if k is small, the probability that the agent gets a bonus after shirking is negligible. Thus, when $\rho \in (\tilde{\rho}_S, \bar{\rho})$, he faces the following trade-off: to get (almost) perfect information by shirking or to get a bonus from the principal with non-negligible probability by working. Clearly, the optimal behavior depends on the equilibrium size of the reward: if W is small enough, the shirking option becomes more attractive, if W is large, the agent prefers to work in the equilibrium.

4.3. No private signal

We now suppose that the agent receives no private signal. We consider two cases. In the first case ("no signal"), the agent receives no signal at all (apart from the reward). In the second case ("public signal") he receives a noisy signal, but it is also observed by the principal before she selects the reward. The case of a private signal that we considered before is an intermediate case between these two extremes.

We consider only the case where $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$: outside this interval, the unique equilibrium outcome is pooling with no bonus. The following is a direct counterpart of Proposition 1 for the *no signal* case.

Proposition 4. *Suppose the agent receives no signal. For $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$, if $e_1 = 1$, there always exists a continuation equilibrium outcome (supported by some beliefs satisfying D1).*

(i) *If $\rho \in [\tilde{\rho}_S, \bar{\rho})$, the continuation equilibrium outcome is either*
  (a) *separating with bonus $\tilde{b}_S = \hat{\theta}_F W$ paid after success and no bonus paid after failure, and the agent choosing $e_2 = 1$ after getting $\tilde{b}_S$, and $e_2 = 0$ after getting no bonus, or*
  (b) *semi-separating, with the principal always paying bonus $\tilde{b}_S = \hat{\theta}_F W$ after success and randomizing between $\tilde{b}_S$ and $\tilde{b}_F = 0$ after failure, with probabilities $\tilde{x}_F \leqslant \frac{1}{A}$ and $1 - \tilde{x}_F$ respectively, and the agent choosing $e_2 = 1$ after getting $\tilde{b}_S$, and $e_2 = 0$ after getting no bonus.*

(ii) *If $\rho \in [\bar{\rho}, \tilde{\rho}_F)$, the continuation equilibrium outcome is either separating or semi-separating as in* (i)*, or pooling with first period bonus $\tilde{b}_P \leqslant \hat{\theta}_F W$ always given by the principal and the agent choosing $e_2 = 1$.*

In our main model with the agent's private information, the richer information structure, combined with the forward-induction logic of the D1 refinement, created a tight link between the principal's out-of-equilibrium deviations and the agent's plausible reactions to these deviations, which allowed us to (generically) restrict the set of PBE to a unique outcome. Without this private information for the agent, the set of equilibrium outcomes expands. Nevertheless, results are qualitatively similar. It is still true that for a range of parameters all equilibria are informative: rewards are given with positive probability, and provide positive feedback. For another range, informative equilibria coexist with uninformative pooling equilibria.

Very similar remarks apply for the case of a *public signal*. As in the case with no signal, there is no uncertainty about the agent's information. Define $\bar{\sigma}$ as the solution to $l(\bar{\sigma}) \equiv A(\rho)$. Proposition 4 can now be restated with two modifications.[16] First, conditions on initial self-confidence are replaced by conditions on the level of public signal: $\rho \in [\tilde{\rho}_S, \bar{\rho})$ in claim (i) is replaced by $\sigma < \bar{\sigma}$, while $\rho \in [\bar{\rho}, \tilde{\rho}_F)$ in claim (ii) is replaced by $\sigma \geqslant \bar{\sigma}$. Second, for semi-separating equilibria there is now a constraint $\tilde{x}_F \leqslant l(\sigma)/A$.

Overall, we conclude that when the agent does not have a private noisy signal, informative equilibria where rewards give positive feedback, are sustainable for a wider range of parameters, and there still exist cases in which only informative equilibria exist.[17]

4.4. Substitutable effort and ability

For some tasks, effort and ability are substitutes, rather than complements. For example, this is the case when a student takes an exam for which no grades are given except "pass" or "fail." As a higher self-confidence now lowers the agent's effort, it is no longer in the interest of the principal to provide positive feedback. *Vice versa*, she has an incentive to give negative feedback to encourage the agent to exert more effort. To make communication credible, the principal again has to spend resources. It is unlikely, though, that the principal will choose to give a bonus to the agent: in the current context, this form of signaling is unrealistic and gives perverse *ex ante* incentives. It is more likely that the principal will spend resources to convince the agent that he is not doing well enough, or help him, or start monitoring more closely or even spend resources to punish the agent. Any such intervention works as a short-term reinforcer but has a long-term negative effect on the agent's self-confidence.

To formally illustrate our argument, assume that our basic model is modified in three ways. First, the technology is now different (ability and effort are substitutes): the agent now succeeds if $\theta + e \geqslant \psi$, where $\theta_L < \psi < \theta_H < 1$. That is, if the agent is smart, he may master the task without much effort; if he is less gifted, he can compensate the lack of ability by working harder. Second, the agent gets no private signal σ. And, third, Assumption 2 is now replaced by $(1-\rho)V < c < V$.

In equilibrium, the agent does not exert effort in the first period (this is the optimal strategy whether he expects interim feedback from the principal or not). If the principal observes the agent's failure in the first period, she has an incentive to communicate this and encourage the agent to exert effort in the second period. Any action with some positive cost to the principal will do – were the agent to succeed in the first period, the principal would not want to spend any resources since she would be sure the agent is strong and effort is not needed.[18]

The negative effect of "crowding out" of intrinsic motivation can also be achieved even with positive rewards: see, in particular, closely related models of Bénabou and Tirole (2003) and Suvorov (2003) (previously discussed in the introduction), and Deci and Ryan (1985) for a review of psychological explanations. The crowding-out effect is found in many experiments in psychology and economics. In experiments where rewards were not made explicitly contingent on performance in advance, crowding out was not found, in line with our previous results (Deci and Ryan, 1985).[19] The next section extends our model to three outcomes and shows that in that case, even positive rewards can decrease self-confidence in the short run.

4.5. More than two outcomes

The main model with two possible outcomes is enough to get the key insights. Bonuses playing a signaling role will still occur in a more general model with more outcomes, and also, other interesting aspects come into play. In particular, the principal's equilibrium policy may be non-monotonic in outcomes: she may give a reward after moderate performance levels, but neither after a poor nor after a successful performance. The complete analysis of a general model is beyond the scope of this paper, but we give an illustrative example.[20]

We introduce a third outcome, "moderate performance": $y_t \in \{F, M, S\}$. The probability of outcome y conditional on ability $\alpha \in \{H, L\}$ is given by $e\theta_\alpha^y$. Outcome y yields a payoff V_y to the agent and W_y to the principal, with $V_S > V_M > V_F = 0$ and $W_S > W_M > W_F = 0$. The probability of better outcomes increases with ability (in the MLRP sense):

$$\frac{\theta_H^S}{\theta_L^S} > \frac{\theta_H^M}{\theta_L^M} > \frac{\theta_H^F}{\theta_L^F}.$$

As in our main model, the agent will then exert a (weakly) higher effort level when he is more confident in his ability. Assume (for simplicity) there are only two possible realizations of the agent's private signal: σ_1 and σ_2. After a successful

[16] A brief sketch of the proof is included in the proof of Proposition 4 in Appendix A.

[17] We have not explicitly analyzed the agent's behavior in period 1. Such analysis is ambiguous due to the multiplicity of continuation equilibria. For any equilibrium selection, though, a proposition similar to Proposition 2 can be proved.

[18] D1 does not allow to pin down a unique equilibrium here due to the open set problem: any action with positive costs can occur in equilibrium.

[19] Deci et al. (1999) confirm the existence of crowding out in a meta-analysis. See also Deci and Ryan (1985) and Frey (1997).

[20] The line of reasoning where only medium (and not high) types use signaling to distinguish themselves from low types is developed in an interesting contribution by Feltovich et al. (2002). They call it "countersignaling". Their analysis shows that restrictive assumptions are needed to get specific results in a more general setup.

performance, the agent always receives σ_1, after a failure always σ_2, and after moderate performance either signal with probability 1/2.

Let the parameters be such that it is optimal for the agent to exert effort in the second period if he knows that the intermediate outcome was M or S, and no effort if he failed. Assume further, that in the absence of any feedback from the principal, it is optimal for the agent to exert high effort after observing σ_1 and no effort after observing σ_2. Then, in the absence of feedback the agent chooses not to work after moderate interim performance with probability 1/2. To prevent this, the principal may want to give a bonus $b_M > 0$ after $y_1 = M$. There is no need for the principal to separate success from the other outcomes by paying a bonus: a successful agent that receives no bonus *knows* he is successful as he receives signal σ_1.

Thus, for certain parameters there exist PBE (satisfying D1) where moderate performance is separated from failure, while successful and failing outcomes are pooled: $b_S = b_F = 0$ and $b_M > 0$. Receiving a bonus is *motivating* for the agent who achieved intermediate performance but received pessimistic signal σ_2, but *demotivating* for the agent who achieved intermediate performance but received optimistic signal σ_1. The assumption that the principal is uncertain about the agent's signal when she chooses the bonus is crucial for this "non-monotonicity" result.[21]

5. Conclusions

In this paper we focus on the role of information given by rewards, and give a new explanation for the use of discretionary rewards, emphasizing their role in stimulating subsequent motivation. Although we frame our model as one where the agent is uncertain about his performance, uncertainty about either his own payoff or cost of effort yields similar results.

The informational aspect of rewards is the main distinction from other approaches in the literature, having several empirically testable implications. One important implication is that small rewards can have a large impact on motivation. In contrast, if rewards have no informational content, they need to be high to have a large effect. In practice, merit-based compensation systems often show very little differences in pay, by themselves not giving strong incentives unless employees are extremely sensitive to small variations (Milgrom and Roberts, 1992, p. 406). More generally, our model predicts that the bonus will be proportional to the principal's stake even though the agent has no explicit bargaining power. Another prediction that follows from our theory is that, by increasing self-confidence, a reward given for one task can result in more effort on a related task that requires similar abilities. In a standard model, rewarding one task does not lead to increased efforts in another one.

The focus on a finite horizon is another difference from other approaches on subjective performance measures. Other approaches mostly rely on infinitely repeated relationships based on reputation, where the threat of punishment sustains the use of discretionary rewards.

The relevance of our model is limited to situations where the principal is better informed. If the agent becomes familiar with the task, or the job is standardized and has objective performance measures, less frequent use of discretionary rewards should be expected. Based on empirical evidence, MacLeod and Parent (1999) find that *ex post* discretionary rewards are more likely to be used in complex environments. Although consistent with our model, this evidence is open to alternative interpretations. A more direct empirical test of our theory is desirable, ideally involving manipulation of the agent's information.[22]

This paper is only one of the first few attempts to formally study the interaction between rewards and self-confidence. One possible extension is to consider more periods. This would shed light on the dynamics of rewards and self-confidence similar to Suvorov (2003). Another extension would be to use an environment where both contractually specified and discretionary rewards might be used, and to then analyze the optimal mix, as in Baker et al. (1994) and Pearce and Stacchetti (1998). It would also be interesting to consider a model with several agents, where agents can make inferences from each others' rewards. We hope to report on this in future work.

Acknowledgments

We are grateful to Michèle Belot, Roland Bénabou, V. Bhaskar, Roberta Dessi, Armin Falk, Sergei Guriev, Thomas Mariotti, Jean-Charles Rochet, Sergey Stepanov, Theo van de Klundert, seminar participants at various conferences and universities, and especially Jean Tirole, an anonymous editorial board member and referee for very valuable comments. Jeroen van de Ven gratefully acknowledges support from the Netherlands Organisation for Scientific Research (NWO), grant Veni 451-05-005.

[21] In a more general model, whether the equilibrium bonuses are monotonic in outcomes or not depends crucially on the information structure. For example, if the agent gets signal σ_1 with probability p_y after outcome y with $p_S > p_M > p_F = 1/2$, and gets signal σ_2 with probability $1 - p_y$, then, if the agent's signal is sufficiently uninformative (p_S close enough to 1/2), only monotonic equilibria are possible.

[22] A recurrent theme in psychology is that rewards have informational content. It is a key element of cognitive evaluation theory (Deci and Ryan, 1985). In a leading textbook on work psychology, Arnold (2007, p. 337) concludes that "Financial rewards tend to enhance performance, especially when they are seen as fair and providing accurate feedback about how well the person is doing." Also in that field, direct evidence is scarce and not conclusive.

Appendix A

Proof of Proposition 1. Claim (iii), the existence of equilibria and the (generic) uniqueness in the class of pooling and semi-separating equilibria are proved in Section 3.1. It remains to prove that there exists a unique ρ^* such that $A(\rho^*) = r(\rho^*)$ and that there are no other continuation equilibria satisfying D1 (we consider parameters A and $r = \hat{\theta}_F/\hat{\theta}_S$ as functions of ρ: $A(\rho)$ and $r(\rho)$).

Lemma 4. *On $[\tilde{\rho}_S, \tilde{\rho}_F)$ there exists a unique ρ^* such that $A(\rho^*) = r(\rho^*)$. If $\rho \in [\tilde{\rho}_S, \rho^*)$, then $A(\rho) > r(\rho)$; if $\rho \in (\rho^*, \tilde{\rho}_F)$, then $A(\rho) < r(\rho)$.*

Proof. With $\alpha = \frac{\rho}{1-\rho}$, equation $A(\rho) = r(\rho)$ is equivalent to:

$$\begin{aligned}&\big((1-\theta_L)\phi - \alpha(1-\theta_H)\big)\big(\alpha(1-\theta_H) + (1-\theta_L)\big)\big(\alpha\theta_H^2 + \theta_L^2\big)\\ &\quad = \big(\alpha\theta_H - \theta_L\phi\big)\big(\alpha(1-\theta_H)\theta_H + (1-\theta_L)\theta_L\big)(\alpha\theta_H + \theta_L).\end{aligned} \tag{24}$$

Denote by $Q_1(\alpha)$ the LHS of this equation which is a cubic polynomial in α with 3 real roots $\alpha_1^L = -\frac{1-\theta_L}{1-\theta_H}$, $\alpha_2^L = -\frac{\theta_L^2}{\theta_H^2}$, $\alpha_3^L = \frac{1-\theta_L}{1-\theta_H}\phi$, it tends to $-\infty$ when $\alpha \to +\infty$ and tends to $+\infty$ when $\alpha \to -\infty$. The RHS (which we denote by $Q_2(\alpha)$) is also a cubic polynomial in α with 3 real roots $\alpha_1^R = -(\frac{1-\theta_L}{1-\theta_H})(\frac{\theta_L}{\theta_H})$, $\alpha_2^R = -\frac{\theta_L}{\theta_H}$, $\alpha_3^R = \frac{\theta_L}{\theta_H}\phi$; it tends to $+\infty$ when $\alpha \to +\infty$ and tends to $-\infty$ when $\alpha \to -\infty$. From the order of the roots ($\alpha_1^L < \alpha_1^R < \alpha_2^R < \alpha_2^L < \alpha_3^R < \alpha_3^L$), it can be easily seen that $Q_1(\alpha) - Q_2(\alpha)$ changes sign on each of the intervals (α_1^L, α_1^R), (α_2^R, α_2^L) and (α_3^R, α_3^L), and thus Eq. (24) has a solution on each of the intervals. Since $\alpha_3^R = \frac{\tilde{\rho}_S}{1-\tilde{\rho}_S}$ and $\alpha_3^L = \frac{\tilde{\rho}_F}{1-\tilde{\rho}_F}$, there exists $\rho^* \in [\tilde{\rho}_S, \rho^*)$ such that $A(\rho^*) = r(\rho^*)$. Since (24) has at most three real roots, this also proofs uniqueness of ρ^*. □

Lemma 5. *If bonus $\tilde{b}$ is given in equilibrium with probability $\tilde{x}_S > 0$ after success and with $\tilde{x}_F > 0$ after failure, and $\tilde{\sigma}$ is the agent's best response, then $l(\tilde{\sigma}) = \frac{\tilde{x}_F}{\tilde{x}_S}A$ and either: 1) $\tilde{b} > 0$ and $l(\tilde{\sigma}) = \hat{\theta}_F/\hat{\theta}_S$, or 2) $\tilde{b} = 0$ and $l(\tilde{\sigma}) \leqslant \hat{\theta}_F/\hat{\theta}_S$.*

Proof. The agent's optimal reaction $\tilde{\sigma}$ is determined by

$$l(\tilde{\sigma}) = \frac{\tilde{x}_F}{\tilde{x}_S}A$$

(this follows from (9) with $\tilde{\lambda}(1, \tilde{b}) = \tilde{x}_S\hat{\theta}/(\tilde{x}_S\hat{\theta} + \tilde{x}_F(1-\hat{\theta}))$ by Bayes' law, $0 < \tilde{\lambda}(1, \tilde{b}) < 1$). When $\tilde{b} > 0$, for the agent to have beliefs $\tilde{\lambda}(1, \hat{b}) < 1$ for $\hat{b} = \tilde{b} \pm \varepsilon$ with $\varepsilon > 0$ (a necessary condition for $\tilde{b}$ to be given an equilibrium bonus given that $\tilde{\sigma} > 0$) a necessary and sufficient condition is $l(\tilde{\sigma}) = r(\rho)$ (the proof is analogous to that in Section 3.1.1). When $\tilde{b} = 0$, only deviations to $\tilde{b} + \varepsilon$ are relevant so the requirement reduces to $l(\tilde{\sigma}) \leqslant \hat{\theta}_F/\hat{\theta}_S$. □

Lemma 6. *In equilibrium only one bonus is offered after success.*

Proof. Assume that b and $b' > b$ are offered after success with positive probability, and $\tilde{\sigma}$ and $\tilde{\sigma}'$ are the corresponding agent's reactions ($\tilde{\sigma} > \tilde{\sigma}'$ by Lemma 3). The smaller bonus, b, must be offered after failure with a positive probability (otherwise the agent would always work after b and the principal would never give the larger one, b'). For the principal not to be willing to separate the successful outcome by offering $b + \varepsilon$, it must be that $l(\tilde{\sigma}) \leqslant \hat{\theta}_F/\hat{\theta}_S$. Then,

$$b' - b = \big[\hat{\theta}_S\big(G_S(\tilde{\sigma}) - G_S(\tilde{\sigma}')\big)\big]W < \big[\hat{\theta}_F\big(G_F(\tilde{\sigma}) - G_F(\tilde{\sigma}')\big)\big]W. \tag{25}$$

The equality in (25) follows from indifference between b and b' after success; the inequality follows from $l(\tilde{\sigma}_1) \leqslant \hat{\theta}_F/\hat{\theta}_S$ and MLRP,[23] it implies that the principal strictly prefers b' to b after a failure – a contradiction. □

Corollary 1. *At most two different bonuses are offered with positive probability in equilibrium. There are three potential types of equilibrium*: A) *pooling – the same bonus offered to both types*; B) *semi-separating – the principal always gives $\tilde{b}_S$ after success and randomizes between $\tilde{b}_S$ and $\tilde{b}_F \neq \tilde{b}_S$ after failure*; C) *separating – the principal always gives $\tilde{b}_S$ after success and $\tilde{b}_F \neq \tilde{b}_S$ after failure.*

Lemma 7. *There are no separating continuation equilibria when $\rho \in (\tilde{\rho}_S, \tilde{\rho}_F)$.*

---

[23] See footnote 14.

Proof. In a separating equilibrium the principal gives $\tilde{b}_F = 0$ after a failure (there is no sense to incur any cost to send a negative signal) and $\tilde{b}_S > 0$ after success. The agent always works after $\tilde{b}_S$ and never works after $\tilde{b}_F$. For this to be an equilibrium, the principal should not strictly prefer to give $\tilde{b}_S$ after failure: $\hat{\theta}_F W - \tilde{b}_S \leqslant 0$. If this were a strict inequality, she would prefer to reduce $\tilde{b}_S$ by a small ε to $\hat{b} = \tilde{b}_S - \varepsilon$ so that the inequality would still be satisfied: according to D1, $\tilde{\lambda}(1, \hat{b}) = 1$ since $D^o(F, \hat{b}) = \emptyset$, $D^o(S, \hat{b}) \neq \emptyset$. Hence, $\tilde{b}_S = \hat{\theta}_F W$. She must also not want to separate the success outcome by a bonus lower than $\tilde{b}_S$. The necessary and sufficient condition for this is:

$$l(0) \geqslant \frac{\hat{\theta}_F}{\hat{\theta}_S}. \tag{26}$$

To see this, assume that the agent's reaction $\hat{\sigma}$ to an out-of-equilibrium bonus $\hat{b}$ is such that the principal is indifferent between deviating to $\hat{b}$ or not after a failure, $\hat{\theta}_F(1 - G_F(\hat{\sigma}))W - \hat{b} = 0$. To have $\tilde{\lambda}(1, \hat{b}) < 1$ satisfying D1, she must not want to deviate after success, $\hat{\theta}_S(1 - G_S(\hat{\sigma}))W - \hat{b} < \hat{\theta}_S W - \tilde{b}_S$. Using $\tilde{b}_S = \hat{\theta}_F W$, this implies $\hat{\theta}_S G_S(\hat{\sigma}) > \hat{\theta}_F G_F(\hat{\sigma})$. For this inequality to be satisfied for all $\varepsilon \in (0, \hat{\theta}_F W)$, i.e., for all $\hat{\sigma} \in (0, 1)$, a necessary and sufficient condition is (26). Since $l(0) = 0$, this can be ruled out. □

The proof of Proposition 1 is now complete. □

Proof of Proposition 2. Let the continuation equilibrium (satisfying D1) be characterized by bonus $\tilde{b}_S$ that the principal pays after success and, with probability $\tilde{x}_F$, after failure, and by the agent's reaction $\tilde{\sigma}$ to this bonus. Note that both pooling equilibria (with $\tilde{x}_F = 1$) and semi-separating equilibria (with $\tilde{x}_F < 1$) fit this characterization; equilibrium values of these parameters are given in Proposition 1. The agent weakly prefers to work in the first period if

$$\begin{aligned}
&\rho\big[(\theta_H V - c)\big(1 + \theta_H\big(1 - G_S(\tilde{\sigma})\big) + (1 - \theta_H)\tilde{x}_F\big(1 - G_F(\tilde{\sigma})\big)\big) + b_S\big(\theta_H + \tilde{x}_F(1 - \theta_H)\big)\big] \\
&\quad \geqslant (1 - \rho)\big[(c - \theta_L V)\big(1 + \theta_L\big(1 - G_S(\tilde{\sigma})\big) + (1 - \theta_L)\tilde{x}_F\big(1 - G_F(\tilde{\sigma})\big)\big) - b_S\big(\theta_L + \tilde{x}_F(1 - \theta_L)\big)\big],
\end{aligned} \tag{27}$$

and is indifferent if (27) is satisfied as an equality. If $\rho > \rho^*$, the unique continuation equilibrium satisfying D1 is pooling: $\tilde{b}_S = 0$, $\tilde{x}_F = 1$. The agent's reaction $\tilde{\sigma}$ is characterized by $l(\tilde{\sigma}) = A$ when $\rho \in (\rho^*, \tilde{\rho}_F)$ and $\tilde{\sigma} = 0$ for $\rho \in [\tilde{\rho}_F, 1]$. The definition of $\bar{\rho}$, (5), is equivalent to $\bar{\rho}/(1 - \bar{\rho}) = \phi$, and also to $A(\bar{\rho}) = 1$. Then,

$$\frac{\rho}{1 - \rho} > \phi \tag{28}$$

for all $\rho \geqslant \bar{\rho}$. It is easy to verify that MLRP implies:

$$\frac{1 + \theta_H(1 - G_S(\tilde{\sigma})) + (1 - \theta_H)\tilde{x}_F(1 - G_F(\tilde{\sigma}))}{1 + \theta_L(1 - G_S(\tilde{\sigma})) + (1 - \theta_L)\tilde{x}_F(1 - G_F(\tilde{\sigma}))} > 1. \tag{29}$$

From (28) and (29) it follows that (27) is satisfied for all $\rho \geqslant \bar{\rho}$, so $[\bar{\rho}, 1) \subset \mathcal{R}_1 \backslash \mathcal{R}_0$. If $\rho < \tilde{\rho}_S$, the unique continuation equilibrium is pooling with no bonus offered and the agent exerting no effort in the second period. Then, (27) reduces to

$$\frac{\rho}{1 - \rho} \geqslant \phi$$

which is not satisfied for any $\rho < \tilde{\rho}_S$ since, as can be easily seen, $\tilde{\rho}_S < \bar{\rho}$. Hence, $(0, \tilde{\rho}_S) \subset \mathcal{R}_0 \backslash \mathcal{R}_1$ which completes the proof of part (ii) of the proposition.[24]

Note that $A(\rho^*) = r(\rho^*) < 1$, therefore $\rho^* > \bar{\rho}$, which proves part (iii).

Existence of equilibrium has been proved by construction: when (27) is satisfied as a strict inequality, there exists an equilibrium with $e_1 = 1$; if (27) is not satisfied, there exists an equilibrium with $e_1 = 0$; finally, if (27) is satisfied as equality, both equilibria exist: one with $e_1 = 0$ and the other with $e_1 = 1$. Also note that the equilibrium outcome (satisfying D1) will be unique, unless $\rho \in \{\tilde{\rho}_S, \rho^*, \tilde{\rho}_F\}$ or (27) is satisfied as equality. Parameter values for which none of these exceptions takes place, constitute an open and everywhere dense set in the manifold of admissible parameter values, so part (i) is proven. □

Proof of Proposition 3. Let $\mathcal{R}^T$ be the set of values of ρ such that there exist a D1 equilibrium in which the agent exerts effort in period 1 when interim performance is fully transparent to the agent.

First, $(0, \tilde{\rho}_S) \cap \mathcal{R}^T = \emptyset$ and $[\tilde{\rho}_F, 1) \subset \mathcal{R}^T$ – in these regions the agent does not care about information about y_1. When $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$, if $e_1 = 1$ under transparency, the optimal strategy is to exert effort in period 2 if $y_1 = S$ and to shirk if $y_1 = F$. Therefore, for these intermediate values of ρ, $e_1 = 1$ is optimal if $\rho \geqslant \bar{\rho}^T$, where

$$\frac{\bar{\rho}^T}{1 - \bar{\rho}^T} = \frac{1 + \theta_L}{1 + \theta_H}\phi.$$

[24] When $\rho = \rho^*$, there is a continuum of pooling continuation equilibria, but the proof goes through for each of them.

Thus, $\mathcal{R}^T = [\bar{\rho}^T, 1)$, and $\tilde{\rho}_S < \bar{\rho}^T < \bar{\rho} < \rho^*$. Note first that for $\rho \in \mathcal{R}_1 \cap \mathcal{R}^T$ the principal's choice whether to make interim performance fully transparent does not depend on W: her expected payoff in both cases is proportional to W. Also, note that for $\rho = \tilde{\rho}_F - \varepsilon$ with small enough $\varepsilon > 0$ it is not optimal to introduce full transparency: in the pooling equilibrium of the semi-transparent model the agent works in the second period both after success and failure with probability close to 1 (since $\tilde{\sigma} \approx 0$ in this case), while under full transparency the agent works only after success. This proves the first claim of the proposition. From the proof of Proposition 2 it is clear that $\mathcal{R}_1$ (weakly) expands with W, and so does $\mathcal{R}_1 \backslash \mathcal{R}^T$. This proves the second claim, because for any $\rho \in \mathcal{R}_1 \backslash \mathcal{R}^T$ the regime of limited transparency is preferable (the principal's expected payoff is positive under this regime, while under full transparency it is zero). □

Proof of Proposition 4. A detailed proof follows many of the steps of the proof of Proposition 1. We, therefore only give a sketch here. In particular, Corollary 1 applies, so we can restrict attention to pooling, separating and semi-separating equilibria.

Pooling Equilibria (PE). Since the agent acquires no new information, he works if and only if $\rho \in [\bar{\rho}, \tilde{\rho}_F)$. For $\rho \in [\tilde{\rho}_S, \bar{\rho})$ effort is zero so the only PE possible would have $\tilde{b}_P = 0$. However, a deviation to $\hat{b}$ with $\hat{\theta}_S W > \hat{b} \geqslant \hat{\theta}_F W$ makes the principal better off after success if in response the agent chooses $e_2 = 1$ with a large enough probability, while making her worse off after failure for any agent's response. Hence, according to D1, $\tilde{\lambda}(1, \hat{b}) = 1$, upsetting the equilibrium.

For $\rho \in [\bar{\rho}, \tilde{\rho}_F)$, effort $e_2 = 1$. We must have $\tilde{b}_P \leqslant \hat{\theta}_F W$ – otherwise $b = 0$ is strictly better after failure. Let $q \in [0, 1]$ denote the probability that the agent chooses $e_2 = 1$ after getting an out-of-equilibrium bonus $\hat{b}$ (this is the agent's MBR to a bonus given some beliefs $\tilde{\lambda}(1, \hat{b})$). Clearly no principal wants to deviate to $\hat{b} > \tilde{b}_P$ for any agent's reaction. Any downward deviation from $\tilde{b}_P > 0$ to $\hat{b} = \tilde{b}_P - \varepsilon$ would make the principal who observed success indifferent if $\hat{\theta}_S W - \tilde{b}_P = \hat{\theta}_S W q - \hat{b}$, i.e., if $\tilde{b}_P - \hat{b} = \hat{\theta}_S W (1 - q)$. But then $(\hat{\theta}_F W q - \hat{b}) - (\hat{\theta}_F W - \tilde{b}_P) = (\hat{\theta}_S - \hat{\theta}_F) W (1 - q) > 0$ so $D^o(S, \hat{b}) \subseteq D(F, \hat{b})$ and, by D1, $\tilde{\lambda}(1, \hat{b}) = 0$: a deviation is not profitable.

Separating equilibria (SE). It is straightforward to show that the only candidate SE is one where $\tilde{b}_S = \hat{\theta}_F W$ and $\tilde{b}_F = 0$. Any deviation to $\hat{b} \in (0, \hat{\theta}_F W)$ would make the principal who observed success indifferent if $\tilde{b}_S - \hat{b} = \hat{\theta}_S W (1 - q)$, but then $\hat{\theta}_F W q - \hat{b} = (\hat{\theta}_S - \hat{\theta}_F) W (1 - q) > 0$. Hence $D^o(S, \hat{b}) \subseteq D(F, \hat{b})$ and by D1 $\tilde{\lambda}(1, \hat{b}) = 0$: no deviation is profitable.

Semi-separating equilibria (SSE). It is straightforward to show that the only candidate SSE is one where the principal who observed success gives $\tilde{b}_S = \hat{\theta}_F W$ and the principal who observed failure randomizes between $\tilde{b}_S$ with probability $\tilde{x}_F$ and $\tilde{b}_F = 0$. In any SSE, the agent works after $\tilde{b}_S$ (otherwise, a deviation to $\hat{b} = \tilde{b}_S + \varepsilon$ would imply $\tilde{\lambda}(1, \hat{b}) = 1$ by D1 and would be profitable) which requires $\tilde{x}_F \leqslant 1/A$. As in SE, any deviation to $\hat{b} \in (0, \hat{\theta}_F W)$ is interpreted by the agent as coming from the principal who observed failure and no principal will want to deviate. For any $\rho \in [\tilde{\rho}_S, \bar{\rho})$, a SSE can be sustained for a sufficiently small $\tilde{x}_F$. For $\rho \in [\bar{\rho}, \tilde{\rho}_F)$, $A < 1$. In that case, a SSE exists for any $\tilde{x}_F \leqslant 1$, which completes the proof.

Public signal case. We briefly discuss how the proof changes if the agent receives a public signal. For every $\rho \in [\tilde{\rho}_S, \tilde{\rho}_F)$, there always exists a signal σ such that the agent chooses to work if he receives it. When the agent receives no information from the principal, then this threshold signal is given by $\bar{\sigma}$, determined by $l(\bar{\sigma}) = A(\rho)$. In any SSE, the agent should work after getting $\tilde{b}_S$, which requires $\tilde{x}_F \leqslant l(\sigma)/A$. The rest of the proof is easily obtained from the proof for the no signal case, with $\bar{\sigma}$ taking over the role of $\bar{\rho}$. □

References

Arnold, J., 2007. Work Psychology: Understanding Human Behaviour in the Workplace. Prentice Hall, Harlow.
Austen-Smith, D., Banks, J., 2000. Cheap talk and burned money. J. Econ. Theory 91, 1–16.
Baker, G., Gibbons, R., Murphy, K., 1994. Subjective performance measures in optimal incentive contracts. Quart. J. Econ. 109 (439), 1125–1156.
Banks, J., Sobel, J., 1987. Equilibrium selection in signaling games. Econometrica 55 (3), 647–661.
Beer, M., 1990. Unpublished memorandum on performance appraisal. Harvard Business School.
Bénabou, R., Tirole, J., 2002. Self-confidence and personal motivation. Quart. J. Econ. 117 (3), 871–915.
Bénabou, R., Tirole, J., 2003. Intrinsic and extrinsic motivation. Rev. Econ. Stud. 70 (3), 489–520.
Bernheim, D., 1994. A theory of conformity. J. Polit. Economy 94, 841–877.
Bull, C., 1987. The existence of self-enforcing implicit contracts. Quart. J. Econ. 102 (1), 147–159.
Carrillo, J., Mariotti, T., 2000. Strategic ignorance as a self-disciplining device. Rev. Econ. Stud. 67, 529–544.
Cho, I.K, Kreps, D., 1987. Signaling games and stable equilibria. Quart. J. Econ. 102, 179–221.
Cho, I.K., Sobel, J., 1990. Strategic stability and uniqueness in signaling games. J. Econ. Theory 50 (2), 381–413.
Crawford, V., Sobel, J., 1982. Strategic information transmission. Econometrica 50, 1431–1451.
Deci, E., Ryan, R., 1985. Intrinsic Motivation and Self-Determination in Human Behavior. Plenum Press, New York.
Deci, E., Koestner, R., Ryan, R., 1999. A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. Psychol. Bull. 125 (6), 627–668.
Fehr, E., Gächter, S., Kirchsteiger, G., 1997. Reciprocity as a contract enforcement device. Econometrica 65, 833–860.
Feltovich, N., Harbaugh, R., To, T., 2002. Too cool for school? Signaling and countersignalling. RAND J. Econ. 33 (4), 630–649.
Frey, B., 1997. Not Just for the Money: An Economic Theory of Personal Motivation. Edward Elgar, Brookfield.
Fuchs, W., 2007. Contracting with repeated moral hazard and private evaluations. Amer. Econ. Rev. 97 (4), 1432–1448.
Fudenberg, D., Tirole, J., 1991. Perfect Bayesian equilibrium and sequential equilibrium. J. Econ. Theory 53, 236–260.
Gibbs, M., 1991. An economic approach to process in pay and performance appraisals. Mimeo, Harvard Business School.
Holmstrom, B., 1999. Managerial incentive problems: A dynamic perspective. Rev. Econ. Stud. 66 (1), 169–182.
Jovanovic, B., 1979. Job matching and the theory of turnover. J. Polit. Economy 87 (5), 972–990.
Kartik, N., 2007. A note on cheap talk and burned money. J. Econ. Theory 136, 749–758.

Köszegi, B., 2006. Ego utility, overconfidence, and task choice. J. Europ. Econ. Assoc. 4, 673–707.
Kreps, D., Milgrom, P., Roberts, J., Wilson, R., 1982. Rational cooperation in the finitely repeated prisoners' dilemma. J. Econ. Theory 27, 245–252.
Levin, J., 2003. Relational incentive contracts. Amer. Econ. Rev. 93 (3), 835–857.
Lizzeri, A., Meyer, M., Persico, N., 2003. The incentive effect of interim performance evaluations. Mimeo.
MacLeod, B., 2003. Optimal contracting with subjective evaluation. Amer. Econ. Rev. 93 (1), 216–240.
MacLeod, B., Malcomson, J., 1989. Implicit contracts, incentive compatibility, and involuntary unemployment. Econometrica 57 (2), 447–480.
MacLeod, W., Parent, D., 1999. Job characteristics, wages, and the employment contract. Rev. Fed. Res. Bank of St. Louis, 13–27.
Milgrom, P., 1981. Good news and bad news: Representation theorems and applications. Bell J. Econ. 12 (2), 380–391.
Milgrom, P., Roberts, J., 1992. Economics, Organization and Management. Prentice Hall, Englewood Cliffs.
Pearce, D., Stacchetti, E., 1998. The interaction of implicit and explicit contracts in a repeated agency. Games Econ. Behav. 23, 75–96.
Prendergast, C., 1999. The provision of incentives in firms. J. Econ. Lit. 37, 7–63.
Suvorov, A., 2003. Addiction to rewards. Mimeo.
Van Damme, E., 1987. Stability and Perfection of Nash Equilibria. Springer-Verlag, Berlin.