

Career Concerns and Performance Reporting in Optimal Incentive Contracts*

Alexander Wolitzky

Microsoft Research and Stanford University

January 15, 2012

Abstract

In many contracting settings incentives are provided not only through transfers but also through performance reports to prospective employers. This paper studies a principal-agent model of this kind, in which the principal sends a non-verifiable report of output to a competitive labor market interested in the agent’s ability. It is assumed that the principal and agent write an enforceable contract over both payments and reports as a function of output, and that the contract terms they agree to cannot be observed by the market. When contracts are unrestricted, reports cannot affect the agent’s future wage in equilibrium, because the agent will always pay the principal to give a good report. Under limited liability, reports can affect future wages, but only by designating output as “good” or “bad.” This informative performance reporting benefits the principal, and may more than make up for the costs imposed by limited liability. The possibilities that the market may have some direct information about output, that the principal may have additional information about the agent’s ability, and that the principal may have psychological or reputational costs of misreporting output are also considered. (*JEL* D83, D86, J30, L20, M50)

1 Introduction

An important part of incentives in markets and organizations are provided by reports that today’s consumer or employer sends to prospective future consumers or employers—in practice, career concerns likely operate at least as much through such reports as through direct

*I thank Gabriel Carroll, Glenn Ellison, Bob Gibbons, Parag Pathak, Michael Powell, and Christian Ruzzier for helpful comments, and thank the NSF for financial support.

observation of past performance. Given that such reports tend to have a much larger impact on the future welfare of the employee than on that of the current employer, a major concern is that the employee and the current employer may collude through giving the employer favorable terms in the current relationship in exchange for a favorable report to prospective future employers. A leading example is provided by websites where consumers both search for and rate contractors, such as Angie’s List, UShip, and the services departments of eBay and Yelp. Consumers that find contractors on these sites can easily promise to write favorable reviews—or threaten to write negative ones—in exchange for improved performance or lower prices. The same is true for consumers who can make reports about firms to industry organizations like the Better Business Bureau. Towards understanding the economics of situations like these, this paper studies a simple but general principal-agent model with a competitive future labor market that learns about the agent’s ability through a non-verifiable report made by the principal, and derives stark implications for optimal incentive contracts and performance reporting policies.

The key ingredients in the model are that the principal (consumer) and agent (contractor) can write *secret* and *enforceable* contracts over both payments and performance reports as a function of output, so that the (binding) contract they ultimately agree to is not observed by the market (though the market correctly infers what contract is signed in equilibrium). The assumption that the contract is not observed by the market is the key departure from most of the existing literature, and it seems realistic in, for example, the case a homeowner contracting with a builder he found on Angie’s List. The assumption that the parties can enforceably contract over performance reports (i.e., that the principal can commit to an information disclosure policy) is standard in the literature, and provides a benchmark that seems reasonable in the motivating examples, where consumers are short-run players with little incentive to maintain reputations for accurately reporting performance to the market.¹ This combination of assumptions has novel implications. The benchmark result is that if

¹I do not mean to suggest that consumers literally write contracts specifying exactly what reports they will give about their contractors as a function of output, but rather than this model is a useful abstraction in the case where consumers have little incentive to renege on their promised reports. The literal contracting behavior underlying the model should be viewed as including making private promises such as, “I’ll pay an extra \$100 if you do a great job,” or “I’ll write a negative review if you do a bad job.” In many situations such promises are not incredible.

transfers between the principal and agent are unrestricted then the principal's report cannot affect the agent's future wage in equilibrium. The reason is simple but robust: if the market expects the principal to send a bad report after some realization of output, then the principal can offer a contract that instead sends a good report after that realization and charges the agent his valuation of receiving the good report. Since the contract offer is not observed by the market and the principal is willing to send the misleading report, this is a profitable deviation.²

The contracting problem is more interesting in the presence of a limited liability constraint that rules out payments from the agent to the principal, and the analysis of this contracting problem is the core of my analysis. In this case, the threat of giving a bad review is useful for providing incentives to the agent, and performance reports can therefore convey information that is relevant for wages. Nonetheless, quite generally the reporting policy can only partition output into two classes—"good news" and "bad news"—and furthermore the good report is weak, in that some agents receive good reports even if they do not receive payments.³ More precisely, the optimal contract consists of three parts: when output is low, the agent receives the bad report and no payment; when output is intermediate, the agent receives the good report but still no payment; and when output is high, the agent receives the good report and is paid. In addition, the principal is always better off in an equilibrium in which performance reports convey wage-relevant information to the market, because this makes performance reports useful instruments for providing incentives. This implies that the principal prefers to direct the agent to jobs which are informative about her ability. For example, if more difficult jobs are more informative about ability, then the principal directs the agent to inefficiently difficult jobs. Finally, the principal can even be better off in an informative equilibrium of this type than he is without the limited liability constraint. This yields conditions under which restricting transfers from contractors to employers may be efficient.

²However, the performance report can provide wage-irrelevant information in equilibrium, such as which of two tasks the agent is better at.

³Both the results that performance reports are not very informative and that the good report is very weak are consistent with evidence from feedback mechanisms on websites such as eBay. For example, Resnick et al (2006) write that, "The most striking feature of eBay feedback is that it is so positive. Sellers received negative feedback only 1% of the time, and buyers 2%."

The main results of the model are robust to three natural extensions. First, the form of the optimal contract and reporting policy are unchanged if the agent has additional career concerns due to direct observation of output by the market. Second, only information about the agent’s ability that is reflected in output can be conveyed to the market in equilibrium, and therefore all other information that the principal might have is irrelevant for the agent’s future wage. Finally, the main results are largely robust to the possibility that the principal might face “lying costs” if she fails to report output truthfully (which may be due to ethical constraints, reputational concerns, or any other reason).

The paper proceeds as follows. Section 2 discusses related literature. Section 3 introduces the model. Section 4 demonstrates the impossibility of wage-relevant communication with unrestricted contracts. Section 5, which contains the main analysis, characterizes optimal contracts with limited liability. Section 6 compares efficiency with and without the limited liability constraint. Section 7 briefly considers three extensions. Section 8 concludes. All proofs, except for the proof of main result of Section 5, are deferred to an appendix.

2 Related Literature

This paper studies a principal-agent model where the agent has career concerns that are mediated by an information disclosure rule chosen by the principal. It is therefore related to the literatures on career concerns, information disclosure in employment relationships, certification, and collusion in hierarchies (the latter because secret contracting between the principal and agent may be viewed as collusion).

Career Concerns: The literature on career concerns literature originated with Holmström (1982), who studies incentives generated by pure career concerns. The interaction of career concerns and static contracts is studied by Gibbons and Murphy (1992), who provide theory and evidence on the mix of career concerns and linear contracts, and Dewatripont, Jewitt, and Tirole (1999), who compare the role of information in career concerns and static contracts.⁴ These papers do not characterize optimal nonlinear contracts with both career

⁴My setup of career concerns with a general information structure in a two-period model is closely related to Dewatripont, Jewitt, and Tirole (1999).

concerns and explicit incentives, as this is in general intractable; however, the current paper shows that it is possible to describe optimal contracts when career concerns operate through non-verifiable performance reports.

Information Disclosure in Employment: A recent and rapidly growing literature on information disclosure in employment is closely related to my paper (Mukherjee, 2008a, 2008b; Koch and Peyrache, 2009; Albano and Leaver, 2005; Bar-Isaac, Jewitt, and Leaver, 2008; Calzolari and Pavan, 2006). The main difference is that I assume that contracts are not observed by prospective employers, which leads to very different results.⁵ My assumption seems reasonable for the case of consumers posting semi-anonymous reports about contractors online, while the alternative assumption seems better for large employers whose employment and reporting practices are well-known. Indeed, this difference in observability is the key difference between this paper and Mukherjee (2008a), who shows that full disclosure of the principal's information is often optimal. He also discusses how a version of my benchmark result on the impossibility of wage-relevant information transmission would hold in his model if contracts were unobservable.⁶ Mukherjee also studies the interaction of information disclosure and relational contracting in a second paper (Mukherjee, 2008b). Koch and Peyrache (2009) present a related model with two agents and limited liability and focus on how the disclosure policy can be used to transfer continuation utility between the agents. They show that full disclosure is generally suboptimal, and that relative performance evaluations can be superior to individual evaluations. Albano and Leaver (2005) study information disclosure in organizations with rigid wages (like many public sector organizations), and characterize when information disclosure makes recruiting and retaining employees more difficult for these organizations. Bar-Isaac, Jewitt, and Leaver (2008) build on the classic contributions of Waldman (1984) and Greenwald (1986) on adverse selection in the labor market by letting agents' abilities be firm-specific, and develop a general model of information disclosure and its consequences for the labor market. These models are somewhat less closely related to

⁵The assumption that the principal can commit to an information disclosure policy, however, is common to my paper and the literature.

⁶Another paper that explicitly discusses the problem that employees can bribe employers into giving positive reports is Karlan et al (2008). In their model, employers have incentives to report truthfully in order to maintain their reputations with other employers. Karlan et al do not focus on optimal contracting, but rather on what kind of jobs will be found through recommendations.

mine because there are no career concerns or moral hazard in these models, and also because in these models the initial employer can continue to employ the agent in the future, unlike in my paper. There is also no moral hazard in Calzolari and Pavan (2006), who consider a general model of sequential contracting with adverse selection (in that the agent’s type is her private information, in contrast to the symmetric information case of my paper and most of the career concerns literature). Their focus is on deriving conditions under which it is optimal for the upstream principal to conceal all information about the agent from the downstream principal.

Certification: Another related paper is Lizzeri’s (1999) model of certification intermediaries, in which an intermediary publicly commits to an information disclosure policy before testing a product and reporting the outcome of the test to prospective buyers (see also Albano and Lizzeri, 2001). As compared to Lizzeri, the agent being tested in my model also performs a job for the principal, and the chosen disclosure policy is observed by the agent but not by prospective buyers. These differences lead my analysis to be very different from his: for example, I focus on the contracting problem between the principal and agent, holding market beliefs fixed, while Lizzeri focuses on an unravelling argument that implies that all sellers go to the intermediary, which earns high profits. Nonetheless, issues raised by unobservable disclosure policies—which are the focus of the current paper—do seem relevant for the certification literature. For example, my results imply that certification intermediaries benefit from the presence of strong anti-bribery laws when disclosure policies are unobservable.⁷

Collusion in Hierarchies: Because I assume that the market does not observe the contract between the principal and agent, this paper is also related to the large literature on collusion in hierarchies, starting with Tirole (1986).⁸ The model would be fit into to this literature directly if prospective employers could provide incentives for the current employer to report

⁷I thank a referee for this point.

⁸See, for example, Laffort and Martimort (1997, 2000) for a general mechanism design approach to collusion, Baliga and Sjöström (1998) and Laffont and Martimort (1998) on the role of delegation in such models with moral hazard, Mookerherjee and Tsumagari (2004) on such models with adverse selection, and Baliga (1999) and Faure-Grimaud, Laffont, and Martimort (2003) on collusion with soft information. Further references may be found in Section 8.6 of Bolton and Dewatripont’s (2005) book. Strausz (2005) and Peyrache and Quesada (2011) study the possibility of collusion between sellers and intermediaries, connecting the literatures on collusion and certification.

accurately. This possibility seems likely to arise when, for example, the current employer is a long-run player in the market or when the current and future employers are part of the same organization, but does not seem as important in the motivating contracting examples.

Related Analysis: Finally, my analysis of optimal contracting with performance reporting and limited liability is related on a more technical level to Sherstyuk (2000), who analyzes a canonical principal-agent problem with limited liability and the ability of the principal to (inefficiently) punish the agent. In her model, there is one exogenous punishment that the principal may choose to inflict or not. Using a similar argument to mine, she shows that the principal punishes the agent for very bad outcomes, neither punishes nor rewards the agent for intermediate outcomes, and rewards the agent for good outcomes. In my model, giving a bad review is like “punishment” (in that this hurts the agent without helping the principal), where the size and nature of punishments are endogenous. Thus, Sherstyuk’s model could be interpreted as a “reduced form” of my model, where the consequences of bad reviews for the agent are exogenous.⁹ In terms of analysis, my approach to the principal-agent problem also draws heavily on Holmström (1979) and Jewitt, Kadan, and Swinkels (2008).

3 Model

I consider a two-period model consisting of a principal-agent problem followed by the agent’s hiring in a competitive labor market: in the main motivating example, the principal is a consumer hiring a contractor (the agent) to perform a job for her, before reporting to a website or other database.¹⁰ The periods are linked by the possibility that the principal may send a performance report to the market, which may influence the market’s assessment of the agent’s ability and thereby influence the wage the agent earns in the second period. Formally, the agent has underlying ability $\eta \sim G(\eta)$ with density $g(\eta)$ and expectation $\hat{\eta}$, which is initially unknown to all players. If an agent with ability η exerts effort a in period 1, period 1 output y is distributed according to cdf $F(y|a, \eta)$ with density $f(y|a, \eta)$, which I assume is continuously differentiable with respect to a . Exerting effort a costs the agent

⁹There is also uncertainty about the agent’s ability in my model, but not in Sherstyuk’s.

¹⁰Thus, the “labor market” is best interpreted as a market for short-term contracting, rather than long-term employment.

$c(a)$ utils. The principal observes realized output but not chosen effort. At the beginning of period 1, the principal offers the agent a contract $(s(y), r(y))$ consisting of a payment $s(y) \in \mathbb{R}$ and a report (real number) $r(y) \in \mathbb{R}$ conditional on y . If the agent rejects the offer, she obtains outside option \bar{u} , the principal receives 0, and the game ends. Otherwise, she chooses an effort level a , and then y , $s(y)$, and $r(y)$ are realized. I assume that the principal can fully commit to such a contract.¹¹

After output is realized, potential period 2 employers observe report r and rationally update their beliefs about the agent's ability η . Crucially, I assume that potential period 2 employers observe the principal's report r but do *not* observe the contract $(s(y), r(y))$ (though they correctly infer what contract is offered in equilibrium). For simplicity, I assume that the competitive period 2 labor market forces potential employers to pay the agent a second period wage that depends only on the expectation of the agent's ability conditional on the report. If the market's expectation of η conditional on report r is $\tilde{\eta}$, then the period 2 wage is denoted $w(\tilde{\eta})$. Assume that $w(\tilde{\eta})$ is increasing in $\tilde{\eta}$. Let $\hat{w} \equiv w(\hat{\eta})$, the period 2 wage when the market's posterior expectation of the agent's ability equals its prior expectation.¹²

There is no discounting between the periods, and all consumption takes place at the end of period 2. I assume that the agent has differentiable, increasing, and concave utility function $u(w)$ and that the principal is risk-neutral (though this can be easily relaxed). Therefore, given effort a , output y , contract $(s(y), r(y))$, and expected ability $\tilde{\eta}$, the principal's payoff is

$$y - s(y),$$

and the agent's payoff is

$$u(s(y) + w(\tilde{\eta})) - c(a).$$

¹¹Full commitment is a benchmark rather than an assumption that is intended to be perfectly realistic, as in reality there is always some scope for renegotiating the contract after output is realized. See the end of this section as well as Sections 5.1 and 7.3 for discussion of this assumption.

¹²The assumption that the period 2 wage depends only on expected ability is not needed. More generally, if the wage was a function $w(\tilde{G})$ of the entire posterior distribution of ability \tilde{G} , then all the results of the paper go through under the condition that if \tilde{G} is in the interior of the convex hull of a set of distributions \mathbb{G} and $\inf_{G \in \mathbb{G}} w(G) < \sup_{G \in \mathbb{G}} w(G)$, then $w(\tilde{G}) \in (\inf_{G \in \mathbb{G}} w(G), \sup_{G \in \mathbb{G}} w(G))$.

The goal is to characterize the pure-strategy perfect Bayesian equilibria (PBE) of this game, defined as follows: Given a contract $(s^*(y), r^*(y))$ and an effort level a^* , let $w^*(r) \equiv w(\tilde{\eta}(r))$, where $\tilde{\eta}(r)$ is the expectation of η conditional on r , given by

$$\tilde{\eta}(r) = \int_{\eta} \eta \left(\frac{g(\eta) \int_{y:r^*(y)=r} f(y|a^*, \eta) dy}{\int_{\eta'} g(\eta') \int_{y:r^*(y)=r} f(y|a^*, \eta') dy d\eta'} \right) d\eta$$

if $\int_{y:r^*(y)=r} \int_{\eta} g(\eta) f(y|a^*, \eta) dy d\eta > 0$, and given by $\tilde{\eta}(r) = \hat{\eta}$ otherwise.¹³ Then a (pure-strategy) PBE is a contract $(s^*(y), r^*(y))$ and an effort level a^* such that

$$a^* \in \arg \max_a \int_y \int_{\eta} u(s^*(y) + w^*(r^*(y))) g(\eta) f(y|a, \eta) d\eta dy - c(a)$$

and

$$\int_y \int_{\eta} u(s^*(y) + w^*(r^*(y))) g(\eta) f(y|a^*, \eta) d\eta dy - c(a^*) \geq \bar{u},$$

and there does not exist another contract $(s(y), r(y))$ and effort level a' satisfying

$$a' \in \arg \max_a \int_y \int_{\eta} u(s(y) + w^*(r(y))) g(\eta) f(y|a, \eta) d\eta dy - c(a),$$

$$\int_y \int_{\eta} u(s(y) + w^*(r(y))) g(\eta) f(y|a', \eta) d\eta dy - c(a') \geq \bar{u},$$

and

$$\int_y \int_{\eta} (y - s(y)) g(\eta) f(y|a', \eta) d\eta dy > \int_y \int_{\eta} (y - s^*(y)) g(\eta) f(y|a^*, \eta) d\eta dy.$$

Before proceeding, I briefly discuss the key assumption that the principal can fully commit to a contract $(s(y), r(y))$ and that this contract is not observable to the market. The assumption that the principal can commit to a monetary contract $s(y)$ is standard and is consistent with the motivating examples of spot-contracting. The assumption that the principal can commit to a reporting rule $r(y)$ is also standard in the literature.¹⁴ One

¹³That is, it is assumed that the market interprets out-of-equilibrium reports r as “trembles” by the principal, and thus as uninformative of the agent’s ability. This can be relaxed.

¹⁴For example, the principal commits to an information disclosure rule in Mukherjee (2008a, 2008b), Koch and Peyrache (2009), Albano and Leaver (2005), Bar-Isaac, Jewitt, and Leaver (2008), and Calzolari and

interpretation of the principal’s ability to commit to a reporting rule is that if he promises the agent that he will send report $r(y)$ after output y while the market expects that he will send report $r^*(y)$ after output y , then making report $r^*(y)$ is more psychologically costly to him than is making report $r(y)$; that is, the principal faces a higher cost of “breaking his promise” to the agent than of “misleading” the market. Since in many cases the principal interacts more directly with the agent than with prospective future employers, this interpretation is very consistent with evidence from the psychology and economics literature that face-to-face interactions tend to exhibit high levels of trust and trustworthiness (see, e.g., Valley, Moag, and Bazerman (1998) and references therein). This is especially true in the motivating example of a consumer hiring a contractor to provide a personal service and then posting a performance report to a large website. Nonetheless, in Section 7.3 I show that many of the implications of the model are robust to the possibility that the principal may face a cost of “lying” to the market but not to the agent, where the lying cost should be interpreted not only as a psychology cost but also as a reduced form for all sorts of costs of misleading the market, including in particular the cost of losing a reputation for truthful reporting (which could easily be derived explicitly in a repeated game model). However, if the principal cannot commit to a reporting rule at all, then he is indifferent among all reports ex post, and it is not clear what report he might be expected to send.¹⁵

Finally, the assumption that the contract is not observable to the market is the main difference between this paper and the existing literature on information disclosure in employment relationships, and it seems realistic in the case of a small employer whose employment and reporting practices are not well known by the market. Mukherjee (2008a) studies a model similar to mine under the assumption that the contract is observable to the market, and finds that full disclosure is often optimal, essentially because full disclosure is efficient and the principal can extract the resulting rents from the agent up front.

Pavan (2006). Of course, in reality a principal cannot commit to literally any reporting rule $r(y)$. But allowing him to do so in the model leads to the conclusion that the optimal reporting rule (both in the cases with and without limited liability) is quite simple and intuitive. This is similar to the situation in the classical principal-agent model, where the assumption that the principal can commit to literally any monetary contract $s(y)$ is not realistic but is still a natural and useful assumption.

¹⁵The possibility that the report is determined by renegotiation with the agent is discussed at the end of Section 5.1.

4 Unrestricted Contracts

This section presents the benchmark result that when transfers are unrestricted performance reporting cannot convey wage-relevant information about the agent’s ability. Formally, in any PBE the agent’s second-period wage equals \hat{w} , her wage in the absence of any performance reporting. The intuition for this result is simple: Suppose reports \bar{r} and \underline{r} induce second period wages $w^*(\bar{r}) > w^*(\underline{r})$. Then the principal can offer to always report \bar{r} and to pocket the difference $w^*(\bar{r}) - w^*(\underline{r})$ when output y is such that $r^*(y) = \underline{r}$, and the agent is willing to accept such a contract. The formal result is the following:

Proposition 1 *In any PBE, the agent’s period 2 wage equals \hat{w} with probability one.*

Proposition 1 is a very general result: if there are no restrictions on transfers between the principal and agent, then the principal’s performance report cannot affect the agent’s future wage. However, performance reports can still convey *some* information about the agent’s ability, in that the posterior distribution of the agent’s ability can differ from the prior. This information can potentially be socially valuable. For example, if ability is two-dimensional and consists of, say, ability to fix plumbing problems and ability to fix electrical problems, and if in period 2 the employee is paid the sum of her expected ability in each dimension, then the period 1 report can convey which dimension the employee is better at, and this could lead the employee to be hired for a more appropriate job in period 2 (while being paid the same wage as if this information had not been conveyed). In other words, it is as if the principal is biased toward increasing the agent’s period 2 wage, and she can only convey information “orthogonal” to this bias. Conceptually, this point is a special case of Chakraborty and Harbaugh’s (2010) general analysis of multidimensional communication with a sender with known bias.¹⁶

5 Limited Liability

This section presents the main part of the analysis, which characterizes the optimal incentive contract and performance reporting policy under limited liability. Section 5.1 contains the

¹⁶I thank a referee for pointing out this connection.

main characterization, and Section 5.2 discusses the consequences of informative reporting for the principal’s payoff and implications for job allocation.

5.1 Optimal Contracts

Assume now that the agent has limited liability, that is, that $s(y)$ must be non-negative for all y . As usual, this assumption can be justified in the presence of literal limited liability laws or when the agent is wealth and credit constrained. In the present context, it can also be viewed as a social norm under which employers are willing to accept “discounts” in exchange for positive reports, but not “bribes.” The analysis in this section is more involved than that in Section 4 and requires using the first-order approach to the principal-agent problem. I therefore impose the following conditions for the remainder of the paper:

Condition 1 (MLRP) $\frac{\int_{\eta} f_a(y|a,\eta)g(\eta)d\eta}{\int_{\eta} f(y|a,\eta)g(\eta)d\eta}$ is increasing in y .

Condition 2 (CDFC) $\int_{\eta} \frac{\partial^2 F(y|a,\eta)}{\partial a^2} g(\eta) d\eta \geq 0$.

Condition 1 says that the marginal distribution over η of $f(y|a,\eta)$ satisfies the monotone likelihood ratio property (MLRP). Condition 1 is related to requiring MLRP for each ability realization η but neither MLRP of the marginal distribution over η nor MLRP for each η imply the other. It is not hard to see that a sufficient, though certainly not a necessary, condition for Condition 1 is that $\frac{f_a(y|a,\eta)}{f(y|a,\eta)}$ is increasing in y for all η, η' —of course, this is also stronger than MLRP for each η . MLRP of the marginal distribution over η is analyzed in some detail in Dewatripont, Jewitt, and Tirole (1999).

Condition 2 says that the marginal distribution over η of $F(y|a,\eta)$ is convex in effort. This is the analog of the usual convexity of the distribution function (CDFC) property for the case where η is uncertain. Jewitt, Kadan, and Swinkels (2008) show that MLRP and CDFC are sufficient for the first-order approach the principal-agent problem to be valid when η is fixed, and it is immediate that their arguments apply when η is uncertain under Conditions 1 and 2. I thus use the first-order approach for the remainder of the paper without further comment.

The main result of this section is the following:¹⁷

Proposition 2 *Fix a PBE $(s(y), r(y), a)$, and let $\underline{w} \equiv \inf_y w^*(r(y))$ and $\bar{w} \equiv \sup_y w^*(r(y))$.*

Then there exist $\underline{y} < \bar{y}$ such that

1. *If $y < \underline{y}$, then $w^*(r(y)) = \underline{w}$ and $s(y) = 0$.*
2. *If $y \in (\underline{y}, \bar{y})$, then $w^*(r(y)) = \bar{w}$ and $s(y) = 0$.*
3. *If $y > \bar{y}$, then $w^*(r(y)) = \bar{w}$, $s(y) > 0$, and $s(y)$ is increasing in y .*

Furthermore, $w^(r(y)) = \bar{w}$ with positive probability.*

Proposition 2 says that, under MLRP and CDFC, the optimal contract consists of at most three regimes in any PBE. For the lowest realizations of output, the principal gives the worst possible report and does not pay the agent. For intermediate output, the principal gives the best possible report but still does not pay the agent. For the highest values of output, the principal gives the best possible report and also pays the agent.

The intuition for this result is fairly straightforward. As is standard in principal-agent models, MLRP implies that the principal wants to give the agent higher total compensation (performance report plus payment) for higher realizations of y . Increasing the agent's total compensation by giving a better report is free for the principal, while increasing the agent's total compensation through payment is costly, so the principal always gives the best possible report before she starts paying the agent. There can be an intermediate range of realizations of y for which the principal gives the best report but does not pay the agent. In addition, it is (with probability one) not optimal for the principal to give any report other than the best or worst one, as she always wants to either maximize or minimize the agent's period 2 wage.¹⁸ These considerations—worst report for low output, best report for medium and high, payment for high—pin down the form of the optimal contract.¹⁹

¹⁷To keep the statement of Proposition 2 concise, I ignore probability zero events. Technically, the conclusions about $w(\tilde{\eta}(r(y)))$ and $s(y)$ in points 1-3 in the statement of Proposition 2 hold with probability one, conditional on y being in the specified range.

¹⁸Formally, this follows because the principal's Lagrangian is linear in the agent's utility (as is standard in principal-agent models), and the principal's payoff does not directly depend on the agent's period 2 wage.

¹⁹Proposition 2 parallels Propositions 1 and 2 in Sherstyuk (2000).

Note that Proposition 2 states that $w^*(r(y)) = \bar{w}$ with positive probability, but not that $w^*(r(y)) = \underline{w}$ or $s(y) > 0$ with positive probability. Intuitively, this is because there may be no outputs that are sufficiently informative of low effort for the principal to want to give the bad report; and there may be no outputs that are sufficiently informative of high effort for the principal to want to make payments. However, the proof of Proposition 2 shows that all three regimes described in Proposition 2 exist if there are outputs that are sufficiently informative of both good and bad effort.

Proof of Proposition 2. *Step 1: Principal's Lagrangian and First-Order Condition.* The principal's problem is

$$\max_{s(y) \geq 0, r(y), a} \int_y \int_{\eta} (y - s(y)) g(\eta) f(y|a, \eta) d\eta dy$$

subject to

$$\int_y \int_{\eta} u(s(y) + w^*(r(y))) g(\eta) f(y|a, \eta) d\eta dy - c(a) \geq \bar{u}, \quad (\text{IR})$$

and

$$a \in \arg \max_{a'} \int_y \int_{\eta} u(s(y) + w^*(r(y))) g(\eta) f(y|a', \eta) d\eta dy - c(a'). \quad (\text{IC})$$

Using the first-order approach and letting λ be the multiplier on the IR constraint and μ the multiplier on the IC constraint, the principal's Lagrangian becomes

$$\begin{aligned} \mathcal{L} = & \int_y \int_{\eta} \{(y - s(y)) f(y|a, \eta) + \lambda [u(s(y) + w^*(r(y))) f(y|a, \eta) - c(a)] \\ & + \mu [u(s(y) + w^*(r(y))) f_a(y|a, \eta) d\eta dy - c'(a)]\} g(\eta) d\eta dy. \end{aligned}$$

The first order condition with respect to $s(y)$ is a version of the classic equation of Holmström (1979):

$$\frac{1}{u'(s(y) + w^*(r(y)))} = \lambda + \mu \frac{\int_{\eta} f_a(y|a, \eta) g(\eta) d\eta}{\int_{\eta} f(y|a, \eta) g(\eta) d\eta}. \quad (1)$$

This must hold whenever $s(y) > 0$.

Step 2: Optimality of Extremal $w^(r)$.* Note that the Lagrangian is strictly increasing in $w^*(r(y))$ if $\lambda \int_{\eta} f(y|a, \eta) g(\eta) d\eta > -\mu \int_{\eta} f_a(y|a, \eta) g(\eta) d\eta$, is strictly decreasing in

$w^*(r(y))$ if the reverse inequality holds, and does not depend on $w^*(r(y))$ if $\lambda \int_{\eta} f(y|a, \eta) g(\eta) d\eta = -\mu \int_{\eta} f_a(y|a, \eta) g(\eta) d\eta$. Assume for the remainder of the proof that $\int_{\eta} f(y|a, \eta) g(\eta) d\eta > 0$ (the argument for $\int_{\eta} f(y|a, \eta) g(\eta) d\eta < 0$ is symmetric). Then the principal chooses $r(y)$ to maximize $w^*(r(y))$ if

$$0 < \lambda + \mu \frac{\int_{\eta} f_a(y|a, \eta) g(\eta) d\eta}{\int_{\eta} f(y|a, \eta) g(\eta) d\eta}, \quad (2)$$

chooses $r(y)$ to minimize $w^*(r(y))$ if the reverse inequality holds, and is indifferent among all choices of $r(y)$ if there is equality in (2).

Step 3: Characterizing \underline{y} and \bar{y} . By MLRP, if \underline{y} satisfies

$$0 = \lambda + \mu \frac{\int_{\eta} f_a(\underline{y}|a, \eta) g(\eta) d\eta}{\int_{\eta} f(\underline{y}|a, \eta) g(\eta) d\eta},$$

then $0 < (>) \lambda + \mu \frac{\int_{\eta} f_a(y|a, \eta) g(\eta) d\eta}{\int_{\eta} f(y|a, \eta) g(\eta) d\eta}$ if $y > (<) \underline{y}$. Therefore, by (2) if $y < \underline{y}$ the principal chooses r to minimize $w^*(r)$, and if $y > \underline{y}$ the principal chooses r to maximize $w^*(r)$.²⁰

Note also that $\int_y \int_{\eta} f_a(y|a, \eta) g(\eta) d\eta dy = 0$, so MLRP implies that there is a set of outputs Y of positive measure such that $\int_{\eta} f_a(y|a, \eta) g(\eta) d\eta > 0$ for all $y \in Y$. Since λ and μ are non-negative (which follows by standard arguments), it follows that $w^*(r(y)) = \bar{w}$ for all $y \in Y$, and thus $w^*(r(y)) = \bar{w}$ with positive probability.

Finally, recall that (1) holds if $s(y) > 0$. So if $s(y) > 0$ then $y > \underline{y}$ (by MLRP), and hence $w^*(r(y)) = \bar{w}$. In particular, $w^*(r(y))$ is constant over all y such that $s(y) > 0$, which, combined with (1) and MLRP, implies that $s(y)$ is strictly increasing at any y such that $s(y) > 0$. Let \bar{y} be the infimum over y such that $s(y) > 0$. Then $\bar{y} > \underline{y}$, by MLRP and $u' > 0$. This completes the proof of Proposition 2. ■

Note that MLRP is needed only to obtain the cutoff form for equilibria in Proposition 2 and to justify the first-order approach to the principal-agent problem; it is far stronger than needed to obtain the result that the principal uses only the highest and lowest performance reports with probability one, under the maintained assumption that the first-order approach is valid. For this result, the following genericity condition is sufficient:

²⁰This shows that there exist reports \bar{r} and \underline{r} such that $w(\bar{r}) = \bar{w}$ and $w(\underline{r}) = \underline{w}$ in any PBE in which $y < \underline{y}$ and $y > \underline{y}$ with positive probability.

Condition 3 (Genericity) For any a and y , the probability that the output realization equals any y' (including y) such that

$$\frac{\int_{\eta} f_a(y|a, \eta) g(\eta) d\eta}{\int_{\eta} f(y|a, \eta) g(\eta) d\eta} = \frac{\int_{\eta} f_a(y'|a, \eta) g(\eta) d\eta}{\int_{\eta} f(y'|a, \eta) g(\eta) d\eta}$$

equals zero.

Genericity is sufficient for the result that the principal uses only the highest and lowest performance reports with probability one, because, by the proof of Proposition 2, only when

$$0 = \lambda + \mu \frac{\int_{\eta} f_a(y|a, \eta) g(\eta) d\eta}{\int_{\eta} f(y|a, \eta) g(\eta) d\eta}$$

is the principal willing to give a report r such that $w^*(r) \in (\underline{w}, \bar{w})$, and under Genericity the probability of this event is zero.

One potential concern about the real-world feasibility of the optimal contract described in Proposition 2 is that players might renegotiate after an output that was supposed to lead to a bad report under the original contract, with the agent offering payment or extra effort in exchange for a good report.²¹ However, with limited liability and unobserved effort it might be difficult for the agent to make this offer credible. More generally, it is not at all clear how one should think of this renegotiation proceeding: after the agent makes any payment or puts in any extra effort, the principal can always hold him up for more (and perhaps can even threaten to go back and change a good report to a bad one). Indeed, bad performance reports do get sent in reality, so there are clearly some frictions that can scuttle renegotiation. Thus, while studying renegotiation explicitly could certainly be interesting, the full commitment model seems at least like a reasonable benchmark.

A second, related, concern is that the agent could turn the tables by threatening to send a bad report about the principal if the principal sends a bad report about him. In particular, Dellarocas and Wood (2008) argue that fear of such negative reciprocity is an important factor behind the overwhelmingly positive feedback found on eBay. While this might be a problem in markets where buyers and sellers are both long-run players, it is less important

²¹I thank an anonymous referee for suggestions here.

in the leading example where buyers just need to get one job done but sellers get a large fraction of their business through the marketplace.

5.2 Profit Comparison and Job Allocation

This section studies implications of informative performance reporting for the principal's payoff and (thinking slightly outside the model) for the choice of job that the agent is hired for.

Since the principal's performance report is cheap talk, there always exists a babbling PBE in which it is ignored by the market, that is, in which $w^*(r) = w(\tilde{\eta})$ for all r . In this section, I show that if an informative PBE (i.e., a PBE in which $w^*(r(y)) \neq w(\tilde{\eta})$ with positive probability) exists, then it yields strictly higher profit for the principal than any babbling PBE, when the principal can commit to randomizing over performance reports. Thus, the principal benefits from wage-relevant performance reporting.

The intuition for this result is slightly subtle. Suppose that the principal induces effort a^* in a babbling PBE. Because there may not be a report r that the principal can send in an informative PBE such that $w^*(r) = w(\tilde{\eta})$, it is not clear that the principal can always induce a^* at as low a cost in the informative PBE as she could in the babbling PBE. If the principal can randomize over good and bad reports, however, she can always do so in such a way as to give the agent the same expected utility after every realization of y as in the babbling PBE, thus implementing a^* at the same cost as in the babbling PBE (as is verified in the proof). Formally, for this subsection only I allow the principal to offer a contract $(s(y), R(y))$, where $R(y)$ is a probability distribution over reports $r(y)$.²²

Proposition 3 *If the principal can commit to a randomized performance report $R(y)$, the principal receives a strictly higher expected profit in any informative PBE than in any babbling PBE.*

²²Allowing the principal to use fully random contracts throughout the paper would not affect any of the results, though it would substantially complicate notation. I also conjecture that Proposition 3 holds even without allowing random contracts, since the principal can simulate a random contract that reports \bar{r} with probability λ and reports \underline{r} with probability $1 - \lambda$ when output equals y by reporting \bar{r} for fraction λ of the outputs in a small neighborhood of y and reporting \underline{r} for the remaining outputs in the neighborhood, under the maintained assumption that $f(y|a, \eta)$ is continuous. Rather than going into the details of such an argument, I simply allow the principal to use randomized performance reports in this subsection.

An interesting consequence of Proposition 3 is that, all else equal, the principal prefers the agent’s job to be informative about the agent’s ability. For example, consider modifying the model by adding an ex ante stage in which the principal can assign the agent to one of two jobs, one of which is uninformative about the agent’s ability (i.e., $f(y|a, \eta) = f(y|a)$ for all y), while the other is informative. Suppose further that both jobs yield the same benefit to the principal, that job-assignment is observable to the market, and that an informative PBE exists when the agent is assigned to the informative job. For example, a homeowner may be approximately indifferent as to whether a contractor repairs his roof using a traditional method or a more high-tech method, and he may be able to document which method is used by posting photographs of the finished roof online. Proposition 3 then implies that the principal assigns the agent to the informative job (in the example, perhaps the high-tech method). Furthermore, since Proposition 3 shows that the principal’s payoff is strictly higher in an informative PBE, this result continues to hold if there is some cost of assigning the agent to the informative job, for example if the informative job is slightly less productive than the uninformative job. And, as is shown in Section 7.2, this continues to hold if the principal has additional information about the agent’s ability, as in the case where the principal knows the agent’s ability perfectly. As this makes clear, the principal benefits from assigning the agent to the informative job not because she herself learns more about the agent’s ability by doing so, but rather because this allows her to credibly convey information about the agent’s ability to the market, and thus to use performance reporting to provide incentives to the agent.

6 Comparing Unrestricted Contracts and Limited Liability

I have shown that limited liability allows for wage-relevant information to be conveyed by performance reports, and that informative performance reporting is good for the principal. This suggests that the principal may sometimes actually be better off in the presence of the limited liability constraint. This section investigates this possibility.

Whether the principal can be better off under limited liability turns out to be closely related to whether w is a linear function of $\tilde{\eta}$. If w is linear, then no performance reporting policy can increase the agent's ex ante expected period 2 wage, and it can be shown that limited liability makes the principal weakly worse off. If w is nonlinear, I show by example that limited liability can make the principal better off.

Before presenting the results, let me briefly remark on the condition of linearity of w . It seems plausible for w to be linear in cases where the agent is assigned to the same job in period 2 regardless of the market's belief about her ability: for example, when the agent is paid a piece-rate, and her ability is her expected hourly output. On the other hand, w is likely to be nonlinear when which job the agent does in period 2 job depends heavily on her ability. For example, Rosen (1981, 1982) presents two canonical theories of extreme nonlinearity of wages in ability: "superstar" fields (Rosen, 1981) and corporate hierarchies (Rosen, 1982). In these settings, it is natural to think that credible performance reporting is very important, and it is thus natural that previous employers can benefit from the limited liability constraint that makes this possible.

The following simple example can be thought of as a "superstar" job a la Rosen, where w is very nonlinear and both the principal and agent are better off in the presence of limited liability:²³

Example 1 *There be two ability levels, $\bar{\eta} > \underline{\eta}$, two effort levels, $\bar{e} > \underline{e}$, and two output levels, 1 and 0. The probability of each ability level is $\frac{1}{2}$, output is always 1 if both $\eta = \bar{\eta}$ and $e = \bar{e}$, and otherwise output is always 0. Let $\bar{u} = 0$, $c(\underline{e}) = 0$ and $c(\bar{e}) = 2$. Both the principal and agent are risk-neutral. Note that, in the absence of a second-period labor market, the optimal contract induces the first-best effort level \underline{e} , and both the principal and the agent receive payoff zero.*

Now suppose that wages in the second-period labor market are given by $w(\tilde{\eta}) = 10$ if $\tilde{\eta} = \bar{\eta}$, and $w(\tilde{\eta}) = 0$ if $\tilde{\eta} \in [\underline{\eta}, \bar{\eta})$; thus, the wage is 10 if the market is certain that the agent is high-ability, and otherwise it is 0. Without limited liability, Proposition 1 implies that $w(\tilde{\eta}) = \hat{w} = 0$ with probability one in any PBE. Therefore, the optimal contract without

²³To keep the example simple, I have assumed that ability, effort, and output are discrete, and that the agent is risk-neutral, so the example does not exactly fit the model. It should be clear that the example could be perturbed to fit the model exactly.

limited liability is identical to the optimal contract without any second-period labor market, and it yields payoff zero for both the principal and the agent.

However, with limited liability the following is a PBE: the principal offers the contract $r(1) = H$, $r(0) = L$ and $s(1) = s(0) = 0$; the agent chooses $e = \bar{e}$; and $w(H) = 10$ and $w(L) = 0$. This is a PBE because the principal receives her highest feasible payoff of $\frac{1}{2}$, and the agent receives payoff $\frac{1}{2}(10) + \frac{1}{2}(0) - 2 = 3$ from choosing effort \bar{e} but would receive payoff 0 from choosing effort \underline{e} .

In this example, the period 1 job is difficult and unimportant, and therefore not worth doing for its own sake, but is an excellent signal of the agent's ability. The agent's second period market wage is highly nonlinear in the agent's perceived ability; she is heavily rewarded if the market is certain that she is of high ability, and it otherwise paid nothing. For example, the second period labor market may consist of large companies looking for CEOs. Without limited liability, however, the principal cannot credibly convey the valuable information about the agent's ability contained in y , as in Proposition 1. With limited liability, the principal can now use honest reporting to induce effort, and this is anticipated by the market.

The logic behind the above example is compelling for many applications, but it depends crucially on the nonlinearity of $w(\tilde{\eta})$. In particular, the principal always does at least as well without limited liability if $w(\tilde{\eta})$ is linear. The intuition is that if $w(\tilde{\eta})$ is linear, then conveying information about the agent's ability to the market is not valuable to either the principal or the agent, so limited liability imposes only a cost and not a benefit. The proof proceeds by starting with an arbitrary PBE of the model with limited liability and showing that the principal's profit must be weakly higher in any PBE of the model with unrestricted contracts.

Proposition 4 *If $w(\tilde{\eta})$ is linear, then the principal's payoff in any PBE of the model with unrestricted contracts is at least as high as her payoff in any PBE of the model with limited liability.*

7 Extensions

This section considers three extensions of the model: the first two concern modifications to the information structure, and the third introduces costs to the principal of misreporting output. Section 7.1 allows the market to receive direct signals of output in addition to performance reports. Section 7.2 allows the principal to receive direct signals of the agent's ability in addition to output. And Section 7.3 allows the principal to face a cost of distorting the agent's period 2 wage.

7.1 Additional Career Concerns

In this section, I assume that the market directly observes an informative signal of output σ in addition to the principal's performance report, where σ is distributed $H(\sigma|y)$ with density $h(\sigma|y)$. The presence of such a signal may lead to "standard" career concerns for the agent acting directly through output rather than through the performance reporting, and this can affect the agent's behavior in equilibrium. Nonetheless, the main results on the form of optimal contracts and reporting policies from Sections 4 and 5 continue to apply. Formally, the following result establishes natural generalizations of Propositions 1 through 3 that allow for direct observations of output. The intuition is that the optimal use of performance reports for incentive provision depends only on the informativeness of different output realizations about the agent's effort, and this is independent of the fact that the agent may have additional incentives to exert effort via the effect of y on the distribution of σ .²⁴

Proposition 5 *With unrestricted contracts, in any PBE, $\int_{\sigma} u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) d\sigma = \max_r \int_{\sigma} u(s(y) + w^*(r, \sigma)) h(\sigma|y) d\sigma$ with probability one.*

With limited liability, in any PBE, there exist $\underline{y} < \bar{y}$ such that

1. *If $y < \underline{y}$, then $\int_{\sigma} u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) d\sigma = \min_r \int_{\sigma} u(s(y) + w^*(r, \sigma)) h(\sigma|y) d\sigma$ and $s(y) = 0$.*
2. *If $y \in (\underline{y}, \bar{y})$, then $\int_{\sigma} u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) d\sigma = \max_r \int_{\sigma} u(s(y) + w^*(r, \sigma)) h(\sigma|y) d\sigma$ and $s(y) = 0$.*

²⁴In Proposition 5, $w^*(r, \sigma)$ is the agent's period 2 wage given report r and signal σ , which generalizes $w^*(r)$ in the natural way.

3. If $y > \bar{y}$, then $\int_{\sigma} u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) d\sigma = \max_r \int_{\sigma} u(s(y) + w^*(r, \sigma)) h(\sigma|y) d\sigma$, $s(y) > 0$, and $s(y)$ is increasing in y .

Furthermore, the principal's expected profit is higher in any informative PBE than in any babbling PBE.

One difference between Proposition 5 and the earlier results is that the best performance report is no longer necessarily the report that maximizes $\int_{\sigma} w^*(r, \sigma) h(\sigma|y) d\sigma$, since the principal and agent care about the agent's expected utility rather than her expected wage, and these need not coincide now that the second period wage is not a deterministic function of the report. In this way, performance reporting now can convey wage-relevant information (for some distributions of σ), albeit only in a very special way. For example, if the agent's utility function displays decreasing absolute risk aversion, then it may be optimal for the principal to send a safe report that guarantees a moderate period 2 wage when the period 1 payment is low, and send a "risky" report that leads to either a low or high period 2 wage (depending on the realization of σ) when the period 1 payment is high.

7.2 Additional Signals of Ability

This section allows for the possibility that the principal may have information about the agent's ability beyond the information that is reflected in output, and asks whether such information can ever be conveyed in equilibrium. This possibility is of course quite natural even in the case of spot-contracting between the principal and agent. For example, a consumer likely learns about something a contractor's personality and professionalism in the course of discussing a job with her, and this information probably cannot be inferred from the quality of output.

Formally, returning to the baseline model suppose that the principal observes the realization of a random variable ρ , distributed $J(\rho|\eta)$ (in addition to observing y). For example, the principal could perfectly observe η . The result of this section is that in equilibrium the principal will essentially never use ρ to convey wage-relevant information about η , no matter how informative ρ is about η .

Proposition 6 *In any PBE with or without limited liability, $w^*(r(y, \rho)) = w^*(r(y, \rho'))$ for all ρ, ρ' with probability one.*²⁵

As is clear from the proof, Proposition 6 follows immediately from the earlier analysis. Nonetheless, this result has a possibly counterintuitive implication: the principal can only convey wage-relevant information about η via reports about a random variable that she personally cares about. This is because she is willing to convey information about η only to the extent that doing so provides incentives to the agent, and reporting on ρ directly has no effect on the agent's incentives. Hence, the fact that an agent's current employer may learn a lot about her ability in the course of their relationship does not imply that the employer is a good source of information about the agent's ability.

7.3 Lying Costs

As discussed in Section 3, it may be reasonable to assume that the principal would rather mislead the market about what wage the agent should receive than lie to the agent about what performance report he will send. But of course in some cases it may be the principal's cost of "lying" to the market that is more relevant. For example, this could be the case if the principal feels an ethical obligation to the fellow users of an online rating service, or if he has a reputation for honest reporting to protect. This section examines the robustness of Propositions 1 and 2 to this possibility, and provides simple conditions under which many of their conclusions are robust.

Since the model does not assign exogenous meanings to performance reports, I measure the extent to which a report is a "lie" by how much it distorts the agent's period 2 wage relative to reporting output truthfully. Formally, I assume that it costs the principal $l(w^*(r(y)) - \tilde{w}(y))$ to send report $r(y)$ when output is y , where $l : \mathbb{R} \rightarrow \mathbb{R}$ is a loss function representing the cost of lying and $\tilde{w}(y)$ is the wage that would prevail if the principal reporting output truthfully. The loss function l can be interpreted either as a direct psychological cost to lying or as a reduced form for a lost reputation for honest reporting (or for any other

²⁵As in the discussion following Proposition 2, the maintained assumptions of MLRP and CDFC can be replaced by Genericity for this result, under the assumption that the first-order approach is valid.

cost of misreporting output).²⁶

The first result of this section is that Proposition 1 continues to hold as long as distorting the agent’s future wage by one dollar always costs the principal less than one dollar in lying costs. This seems to be a fairly mild condition: for example, it holds whenever the principal can restore his good standard with a prospective employer by doing a favor for him worth 99 cents. The intuition for the result is that, under this condition, the principal can still gain more from “selling” a good report to the agent than she loses from the lying cost.

Proposition 7 *With unrestricted contracts and lying costs, if $l'(x) < 1$ for all $x \in \mathbb{R}$ then the agent’s period 2 wage equals \hat{w} with probability one.*

The second result of this section is that some of the conclusions of Proposition 2 also continue to hold with lying costs (maintaining the assumption that $l'(x) < 1$). In particular, it is still true that the best possible performance report accompanies any payment to the agent. However, it is no longer the case that the principal always gives either the best or the worst performance report, because the benefit of using intermediate reports to reduce lying costs may outweigh the benefit of using only extreme reports to provide incentives.

Proposition 8 *With limited liability and lying costs, if $l'(x) < 1$ for all $x \in \mathbb{R}$ then there exists \bar{y} such that*

1. *If $y < \bar{y}$, then $s(y) = 0$.*
2. *If $y > \bar{y}$, then $w^*(r(y)) = \bar{w}$, $s(y) > 0$, and $s(y)$ is strictly increasing in y .*

8 Conclusion

This paper has developed a model of performance reporting in the presence of career concerns when contracts are not observed by the market. The model has stark implications for optimal contracts and equilibrium performance reporting policies. In particular, when transfers

²⁶There are of course other ways of modelling lying costs. One appealing alternative is assuming that it costs the principal $l(w(r(y)) - w(r^*(y)))$ to send report $r(y)$ when the market *expects* him to send report $r^*(y)$, thus making the model a psychological game (Geanokoplos, Pearce, and Stacchetti, 1989). The robustness of Propositions 1 and 2 can also be established in this alternative model.

between the principal and agent are unrestricted, performance reports cannot convey wage-relevant information about the agent’s ability to the market in equilibrium. When transfers are restricted by limited liability, performance reports can convey wage-relevant information, but the level of detail they can contain is severely limited: essentially, reports can only be “good” or “bad.” In addition, the good report is sometimes given even if output is low enough that the agent does not receive monetary payment, which implies that performance reporting does not distinguish at all among relatively successful agents. Because performance reporting is useful to the principal as a means of providing incentives to the agent, the principal would like the agent to work on a job that is more informative about his ability than would otherwise be optimal. These results are robust to various extensions of the model, including the possibility that the principal may face psychological or reputational costs if he misreports the agent’s output to the market.

I have motivated the model with the example of websites where consumers search for and rate contractors, and the model has some potential implications for the design of feedback mechanisms on these websites. Most importantly, the model predicts that such feedback will be overwhelmingly positive (in that many contractors receive the highest rating) and uninformative (in that there can be at most two ratings) unless consumers can be given incentives for truthful reporting—and feedback on such websites indeed appears to have these properties. There are a number of ways in which one could imagine trying to provide incentives for truthful reporting on these websites, but it is far from clear how this should be done. Indeed, characterizing optimal mechanisms for inducing truthful reporting in this setting seems like an interesting direction for future work, and this could perhaps be done by building on the model of contracting developed in this paper.

Appendix: Omitted Proofs

Proof of Proposition 1. Suppose, towards a contradiction, that $(s^*(y), r^*(y), a^*)$ is a PBE in which $w^*(r) \neq \hat{w}$ with positive probability. Then there exist reports \bar{r} and \underline{r} such that $w^*(\bar{r}) > \hat{w} > w^*(\underline{r})$ and $\int_{y:r^*(y)=\bar{r}} \int_{\eta} g(\eta) f(y|a^*, \eta) dy d\eta > 0$ for $r \in \{\bar{r}, \underline{r}\}$. Replace $(s^*(y), r^*(y))$ with $(s^*(y) - w^*(\bar{r}) + w^*(\underline{r}), \bar{r})$ for all $y : r^*(y) = \underline{r}$, leaving the contract

unchanged elsewhere, and call the resulting contract $(s(y), r(y))$. Then for all $y : r^*(y) = \underline{r}$, the agent's payoff under $(s(y), r(y))$ is $s^*(y) - w^*(\bar{r}) + w^*(\underline{r}) + w^*(\bar{r}) = s^*(y) + w^*(r^*(y))$, and clearly the agent's payoff under $(s(y), r(y))$ and $(s^*(y), r^*(y))$ also coincide for all $y : r^*(y) \neq \underline{r}$, so a^* remains an optimal action under contract $(s(y), r(y))$. Finally, the principal's payoff under $(s(y), r(y))$ is strictly higher than under $(s^*(y), r^*(y))$ when the agent chooses action a^* , because the principal receives a weakly higher payoff under $(s(y), r(y))$ after every output, and receives a strictly higher payoff under $(s(y), r(y))$ whenever $r^*(y) = \underline{r}$, which occurs with positive probability when the agent chooses action a^* . Therefore, $(s^*(y), r^*(y), a^*)$ is not a PBE, a contradiction. ■

Proof of Proposition 3. First observe that, in any informative PBE, there exist \bar{r} and \underline{r} such that $w^*(\bar{r}) > \hat{w} > w^*(\underline{r})$ (since if the posterior $\tilde{\eta}$ is above $\hat{\eta}$ with positive probability then it must also be below $\hat{\eta}$ with positive probability, by the law of iterated expectation). Now suppose that $(s^*(y), R^*(y), a^*)$ is a babbling PBE. Fix an informative PBE, and define $\lambda(y) \in (0, 1)$ to be the solution to

$$\lambda(y) u(s^*(y) + w^*(\bar{r})) + (1 - \lambda(y)) u(s^*(y) + w^*(\underline{r})) = u(s^*(y) + \hat{w}). \quad (3)$$

for some such \bar{r} and \underline{r} . Consider the random contract $(s(y), R(y))$ given by

$$\begin{aligned} s(y) &= s^*(y) \\ R(y) &= \left\{ \begin{array}{l} \bar{r} \text{ with probability } \lambda(y) \\ \underline{r} \text{ with probability } 1 - \lambda(y) \end{array} \right\}. \end{aligned}$$

If the principal deviates by offering contract $(s(y), R(y))$ in the informative PBE, the agent's expected utility given y equals the left-hand side of (3). By (3), this is the same as her expected utility given y in the babbling PBE, so a^* is an optimal effort choice for the agent facing contract $(s(y), R(y))$. Since the principal's expected wage bill is the same when she offers $(s(y), R(y))$ or $(s^*(y), R^*(y))$ and $a = a^*$, she can guarantee herself at least her payoff from the babbling PBE by offering $(s(y), R(y))$. Furthermore, Proposition 2 (which continues to hold when the principal can commit to a randomized performance report, as can easily be shown) shows that $(s(y), R(y))$ is not an optimal contract in an informative

PBE, so the principal must receive a payoff strictly higher than this in any informative PBE.

■

Proof of Proposition 4. Suppose that the $(s^*(y), r^*(y), a^*)$ is a PBE in the model with limited liability. Let

$$s(y) \equiv s^*(y) + w^*(r^*(y)) - \hat{w}.$$

Fix any PBE of the model with unrestricted contracts, and suppose that the principal uses reporting rule $r(y)$. Suppose that, starting from this PBE of the model with unrestricted contracts, the principal deviates to offering contract $(s(y), r(y))$. If the agent accepts the contract, her payoff following output realization y is $s(y) + \hat{w}$ (by Proposition 1), which is precisely her payoff following output realization y when the principal offers contract $(s^*(y), r^*(y))$ in the given PBE of the model with limited liability. Hence, it is optimal for the agent to accept the contract and chooses effort a^* . The principal's payoff from this deviation is therefore

$$\begin{aligned} \int_y \int_\eta (y - s(y)) g(\eta) f(y|a, \eta) d\eta dy &= \int_y \int_\eta (y - s^*(y) - w^*(r^*(y)) + \hat{w}) g(\eta) f(y|a, \eta) d\eta dy \\ &= \int_y \int_\eta (y - s^*(y)) g(\eta) f(y|a, \eta) d\eta dy, \end{aligned} \quad (4)$$

because $\int_y \int_\eta w^*(r^*(y)) g(\eta) f(y|a, \eta) d\eta dy = \hat{w}$ by linearity of w . Now, by definition of PBE, the principal must attain a weakly higher payoff than (4) in the PBE of the model with unrestricted contracts, and (4) is exactly the principal's payoff in the PBE of the model with limited liability. ■

Proof of Proposition 5. The principal's problem is now

$$\max_{s(y), r(y), a} \int_y \int_\eta (y - s(y)) g(\eta) f(y|a, \eta) d\eta dy,$$

subject to

$$a \in \arg \max_{a'} \int_y \int_\eta \int_\sigma u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) g(\eta) f(y|a, \eta) d\sigma d\eta dy - c(a') \quad (\text{IC})$$

and

$$\int_y \int_\eta \int_\sigma u(s(y) + w^*(r(y))) h(\sigma|y) g(\eta) f(y|a, \eta) d\sigma d\eta dy - c(a) \geq \bar{u}. \quad (\text{IR})$$

This yields Lagrangian

$$\begin{aligned} \mathcal{L} = & \int_y \int_\eta \int_\sigma \{ (y - s(y)) f(y|a, \eta) + \lambda [u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) f(y|a, \eta) - c(a)] \\ & + \mu [u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) f_a(y|a, \eta) d\eta dy - c'(a)] \} d\sigma g(\eta) d\eta dy. \end{aligned}$$

Differentiating with respect to $s(y)$ yields

$$0 < \frac{1}{\int_\sigma u'(s(y) + w^*(r(y), \sigma)) h(\sigma|y) d\sigma} = \lambda + \mu \frac{\int_\eta f_a(y|a, \eta) g(\eta) d\eta}{\int_\eta f(y|a, \eta) g(\eta) d\eta}.$$

Linearity of the Lagrangian in $w^*(r(y), \sigma)$ now implies that the principal chooses $r(y)$ to maximize $\int_\sigma u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) d\sigma$ whenever

$$0 < \lambda + \mu \frac{\int_\eta f_a(y|a, \eta) g(\eta) d\eta}{\int_\eta f(y|a, \eta) g(\eta) d\eta},$$

and chooses $r(y)$ to minimize $\int_\sigma u(s(y) + w^*(r(y), \sigma)) h(\sigma|y) d\sigma$ whenever the opposite inequality holds. The rest of the argument follows immediately from the analysis in Sections 4 and 5. ■

Proof of Proposition 6. It is easy to check that allowing the principal to offer contracts of the form $(s(y, \rho), r(y, \rho))$ does not affect the validity of Propositions 1 or 2. Recall from the proof of Proposition 2 that the principal is willing to make a report r such that $w^*(r) \in (\underline{w}, \bar{w})$ only if

$$0 = \lambda + \mu \frac{\int_\eta f_a(y|a, \eta) g(\eta) d\eta}{\int_\eta f(y|a, \eta) g(\eta) d\eta}.$$

Under MLRP, this occurs with probability zero. Therefore, with probability one it is the case that either $w^*(r(y, \rho)) = \underline{w}$ for all ρ or $w^*(r(y, \rho)) = \bar{w}$ for all ρ . ■

Proof of Proposition 7. Suppose that $(s^*(y), r^*(y), a^*)$ is a PBE in which $w^*(r) \neq \hat{w}$ with positive probability. Consider the same alternative contract as the in the proof of Proposition 1; that is, let $s(y) = s^*(y) - w^*(\bar{r}) + w^*(\underline{r})$ for all $y : r^*(y) = \underline{r}$. After output

$y \in Y$, the principal receives payoff

$$y - s^*(y) + w^*(\bar{r}) - w^*(\underline{r}) - l(w^*(\bar{r}) - \tilde{w}(y))$$

under contract $(s(y), r(y))$, and receives payoff

$$y - s^*(y) - l(w^*(\underline{r}) - \tilde{w}(y))$$

under contract $(s^*(y), r^*(y))$. The difference between these payoffs equals

$$w^*(\bar{r}) - w^*(\underline{r}) - (l(w^*(\bar{r}) - \tilde{w}(y)) - l(w^*(\underline{r}) - \tilde{w}(y))).$$

Since $l'(x) < 1$ for all $x \in \mathbb{R}$, it follows that $l(w^*(\bar{r}) - \tilde{w}(y)) - l(w^*(\underline{r}) - \tilde{w}(y)) < w^*(\bar{r}) - w^*(\underline{r})$, so this difference is positive. Hence, deviating to contract $(s(y), r(y))$ remains profitable. The rest of the argument is as in the proof of Proposition 1. ■

Proof of Proposition 8. With lying costs, the Lagrangian for the principal's problem becomes

$$\begin{aligned} \mathcal{L} = & \int_y \int_\eta \{ (y - s(y) - l(w^*(r(y)) - \tilde{w}(y))) f(y|a, \eta) + \lambda [u(s(y) + w^*(r(y))) f(y|a, \eta) - c(a)] \\ & + \mu [u(s(y) + w^*(r(y))) f_a(y|a, \eta) d\eta dy - c'(a)] \} g(\eta) d\eta dy. \end{aligned}$$

Therefore, the first-order condition with respect to $r(y)$ continues to equal (1), and the principal chooses $r(y)$ to maximize $w^*(r(y))$ if

$$0 < -\frac{l'(w^*(r(y)) - \tilde{w}(y))}{u'(s(y) + w^*(r(y)))} + \lambda + \mu \frac{\int_\eta f_a(y|a, \eta) g(\eta) d\eta}{\int_\eta f(y|a, \eta) g(\eta) d\eta}.$$

It follows from (1) and the assumption that $l'(x) < 1$ for all $x \in \mathbb{R}$ that if $s(y) > 0$ then the principal chooses $r(y)$ to maximize $w^*(r(y))$. The rest of the argument is as in the proof of Proposition 2. ■

References

- [1] Albano, G. and Lizzeri, A. (2001), “Strategic Certification and Provision of Quality,” *International Economic Review*, 42, 267-283.
- [2] Albano, G. and Leaver, C. (2005), “Transparency, Recruitment and Retention in the Public Sector,” *mimeo*.
- [3] Baliga, S. (1999), “Monitoring and Collusion with “Soft” Information,” *Journal of Law, Economics, and Organization*, 15, 434-440.
- [4] Baliga, S. and Sjöström, T. (1998), “Decentralization and Collusion,” *Journal of Economic Theory*, 83, 196-232.
- [5] Baker, G.P., Gibbons, R., and Murphy, K.J. (1994), “Subjective Performance Measures in Optimal Incentive Contracts,” *Quarterly Journal of Economics*, 109, 1125-1156.
- [6] Bar-Isaac, H., Jewitt, I., and Leaver, C. (2008), “Information and Human Capital Management,” *mimeo*.
- [7] Bolton, P. and Dewatripont, M. (2005), *Contract Theory*, Cambridge: MIT Press.
- [8] Calzolari, G. and Pavan, A. (2006), “On the Optimality of Privacy in Sequential Contracting,” *Journal of Economic Theory*, 130, 168-204.
- [9] Chakraborty, A. and Harbaugh, R. (2010), “Persuasion by Cheap Talk,” *American Economic Review*, 100, 2361-2382.
- [10] Dellarocas, C. and Wood, C.A. (2008), “The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias,” *Management Science*, 54, 460-476.
- [11] Dewatripont, M., Jewitt, I., and Tirole, J. (1999), “The Economics of Career Concerns, Part 1: Comparing Information Structures,” *Review of Economic Studies*, 66, 183-198.
- [12] Faure-Grimaud, A., Laffont, J.J., and Martimort, D. (2003), “Collusion, Delegation and Supervision with Soft Information,” *Review of Economic Studies*, 70, 253-280.

- [13] Geanakoplos J., Pearce, D., and Stacchetti, E. (1989), “Psychological Games and Sequential Rationality,” *Games and Economic Behavior*, 1, 60-79.
- [14] Gibbons, R. and Murphy, K.J. (1992), “Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence,” *Journal of Political Economy*, 100, 468-505.
- [15] Greenwald, B. (1986), “Adverse Selection in the Labor Market,” *Review of Economic Studies*, 53, 325-347.
- [16] Holmström, B. (1979), “Moral Hazard and Observability,” *Bell Journal of Economics*, 10, 74-91.
- [17] Holmström, B. (1982), “Managerial Incentive Problems—A Dynamic Perspective,” in *Essays in Economics and Management in Honor of Lars Wahlbeck*, Helsinki: Swedish School of Economics. Reprinted in *Review of Economic Studies* (1999), 66, 169-182.
- [18] Holmström, B. and Milgrom, P. (1994), “The Firm as an Incentive System,” *American Economic Review*, 84, 972-991.
- [19] Jewitt, I., Kadan, O., and Swinkels, J. (2008), “Moral Hazard with Bounded Payments,” *Journal of Economic Theory*, 143, 59-82.
- [20] Karlan, D., Möbius, M., Rosenblat, T., and Szeidl, A. (2008), “Trust and Social Collateral,” forthcoming *Quarterly Journal of Economics*.
- [21] Koch, A., and Peyrache, E. (2009), “Aligning Ambition and Incentives,” *Journal of Law, Economics and Organization*, forthcoming.
- [22] Laffont, J.J. and Martimort, D. (1997), “Collusion under Asymmetric Information,” *Econometrica*, 65, 875-911.
- [23] Laffont, J.J. and Martimort, D. (1998), “Collusion and Delegation,” *RAND Journal of Economics*, 29, 280-305.
- [24] Laffont, J.J. and Martimort, D. (2000), “Mechanism Design with Collusion and Correlation,” *Econometrica*, 68, 309-342.

- [25] Lizzeri, A. (1999), “Information Revelation and Certification Intermediaries,” *RAND Journal of Economics*, 30, 214-231.
- [26] Mookherjee, D. and Tsumagari, M. (2004), “The Organization of Supplier Networks: Effects of Delegation and Intermediation,” *Econometrica*, 72, 1179-1219.
- [27] Mukherjee, A. (2008a), “Career Concerns, Matching, and Optimal Disclosure Policy,” *International Economic Review*, 49, 1211-1250.
- [28] Mukherjee, A. (2008b), “Sustaining Implicit Contracts when Agents have Career Concerns: The Role of Information Disclosure,” *RAND Journal of Economics*, 39, 469-490.
- [29] Peyrache, E. and Quesada, L. (2011), “Intermediaries, Credibility and Incentives to Collude,” *Journal of Economics and Management Strategy*, 20, 1099-1133.
- [30] Resnick, P., Zeckhauser, R., Swanson, J., and Lockwood, K. (2006), “The Value of Reputation on eBay: A Controlled Experiment,” *Experimental Economics*, 9, 79-101.
- [31] Rosen, S. (1981), “The Economics of Superstars,” *American Economics Review*, 71, 845-858.
- [32] Rosen, S. (1982), “Authority, Control, and the Distribution of Earnings,” *Bell Journal of Economics*, 13, 311-323.
- [33] Sherstyuk, K. (2000), “Performance Standards and Incentive Pay in Agency Contracts,” *Scandinavian Journal of Economics*, 102, 725-736.
- [34] Strausz, R. (2005), “Honest Certification and the Threat of Capture,” *International Journal of Industrial Organization*, 23, 45-62.
- [35] Tirole, J. (1986), “Hierarchies and Bureaucracies: On the Role of Collusion in Organizations,” *Journal of Law, Economics, and Organization*, 2, 181-214.
- [36] Valley, K.L., Moag, J., and Bazerman, M.H. (1998), “A Matter of Trust: Effects of Communication on the Efficiency and Distribution of Outcomes,” *Journal of Economic Behavior and Organization*, 34, 211-238.

- [37] Waldman, M. (1984), "Job Assignments, Signalling, and Efficiency," *RAND Journal of Economics*, 15, 255-267.