Information Design and Career Concerns*

David Rodina†

September 12, 2017

Abstract

This paper studies the interplay between information and incentives in principal-agent relationships with career concerns, that is when the agent wants to be perceived as of high ability. I derive conditions for when more precise information about performance or more uncertainty about the agent's ability lead to stronger incentives due to career concerns. A key condition for deriving these comparative statics is how effort changes the informativeness of performance signals regarding ability. An indirect, yet tractable representation of information structures enables a full pure strategy equilibrium analysis without ad-hoc restrictions on the set of information structures. Moreover, I show that more sophisticated information revelation technologies that are implicitly ruled out in the literature overturn commonly held assertions regarding information design and career concerns.

*I am particularly grateful to Eddie Dekel and Asher Wolinsky for their valuable insights and continual encouragement. This work also benefited from discussions with Jeff Ely, John Farragut, George Georgiadis, Yingni Guo, Alessandro Pavan, Mikhail Safronov, Bruno Strulovici and Jean Tirole.

†University of Bonn. E-mail: david.rodina@u.northwestern.edu

Contents

1 Introduction

The focus of this paper is a principal-agent relationship where there is uncertainty about the agent's ability, and the agent has a preference for being perceived as of high ability. This motive is often called career concerns, and provides incentives beyond those derived from any existing wage contract, as by exerting effort the agent can manipulate the belief about his ability. The strength of this incentive is affected by two types of information, namely observable performance measures, and prior uncertainty about the agent's ability. The overall aim is to provide a better understanding of the interplay between information and incentives in such an environment.

The main contribution of this paper is a derivation of substantive conditions in which more or less information about performance or ability leads to stronger incentives related to career concerns. If there is a measure of performance such as realized output, when does transparency maximize career-concerns incentives? When are incentives depressed when more information about the agent's ability becomes available? These questions are relevant to the extent that career concerns are associated with situations where explicit incentives are weak and effort is undersupplied. For each of these comparative statics, interpretable assumptions are derived which make them true without ad-hoc restrictions on the set of information structures or equilibria.

Consider a situation where an agent exerts effort once, after which a performance measure such as output is realized. Output is informative about the agent's ability, and observed by a player called the market. The market pays the agent a wage equal to the perceived ability, thereby generating career concerns. Effort is unobserved, so the market belief is based on a conjecture about the agent's behavior. The agent then can try to manipulate the market belief, which he does by trading off the expected change in the perception of his ability against the effort cost. In equilibrium, the agent finds it not worthwhile to influence the market belief any further, and typically a positive effort is then sustained.

To introduce information design, suppose there is a principal who to some extent controls the market's information. First, consider the informativeness of performance measures. Output is the most informative signal of performance, but the principal can commit to garble it in an unrestricted way. The market no longer observes output, but whatever signal is chosen by the principal. Second, to study the effect of uncertainty regarding the agent's ability, suppose that the principal cannot garble output. What she can do is to disclose or withhold a signal of the agent's ability, potentially observable by the market together with output. Given the question of what type of information maximizes career-concerns incentives, the principal is endowed with the objective of maximizing effort.

It seems intuitively reasonable that more information about performance and uncertainty about ability both increase career-concerns incentives. After all, if the market does not observe any signal of performance or if it knows the agent's ability, it is impossible for the agent to manipulate its belief and career-concerns incentives disappear. Beyond these extreme cases, the effect of information is less obvious when no ad-hoc restrictions on signals are imposed. One might argue that garbling information about performance leads to less variability in terms of the wage paid, and therefore weakens incentives to exert more effort. This turns out to be true, but only in a relatively weak sense.

As far as information about performance is concerned, suppose that effort increases the distribution of output (in some appropriate sense), and that output is a favorable signal of ability. It turns out that these natural ordering conditions are not sufficient for effort to be maximized by full revelation of output. Under the additional assumption that higher effort increases the informativeness of output about ability, such a result can be derived. In fact, if one can find two effort levels such that for the higher effort output is less informative, then (conditional on the ordering assumptions) there exists a cost of effort function such that career-concerns incentives are maximized by a noisy signal of performance.

As for the prior uncertainty about ability, I derive conditions under which career-concerns incentives are reduced when the market observes a signal of ability (in addition to output). Conditional on every output level, the signal needs to be indicative of high ability, and higher effort leads to an adverse inference about the ability estimate. Also, it is still required that effort increases the informativeness of output. While the ordering conditions appear innocuous, I present a natural example where they can be violated and then indeed career-concerns incentives are stronger with less uncertainty about ability.

In the main exposition, incentives are derived only from career concerns. The comparative statics are shown to remain true if in addition the agent is motivated by an explicit wage contract. On the other hand, relaxing risk neutrality with respect to the market belief renders either result untrue, and several other extensions delineate their validity. Perhaps surprisingly, once the agent randomizes over effort there exist more sophisticated, and somewhat non-obvious information revelation technologies that endogenously create asymmetric information between the agent and the market. These overturn the result that under appropriate conditions, full revelation of performance measures maximizes incentives due to career concerns.

Given the strategic interaction between agent and market, standard techniques from the literature on Bayesian Persuasion cannot be applied. The equilibrium nature of the model together with an absence of ad-hoc restrictions on the set of information structures makes a direct analysis in terms of signals intractable. To that end, an alternative representation of the payoff effect due to career concerns is developed. This provides an indirect means of representing information structures that is more conducive to performing comparative statics with respect to the amount of information available. The techniques developed here might be applicable in other information design problems with multiple "receivers" when the interaction among them satisfies similar assumptions on timing and observability.

The considerations at the core of this paper are relevant in the following examples. An employee is subjected to a performance review, and one can ask what information this review should incorporate. In a multi-agent environment, should the employee be evaluated based on the performance of co-workers, given that they are subject to similar productivity shocks? A promotion or tenure decision is based on the perception of the agent's ability. Should the decision maker commit ex-ante to use or disregard certain demographic information that is correlated with ability? The findings made here provide a foundation for transparency in organizations with respect to measures of performance, and opaqueness regarding certain information about the agent's ability, such as demographics. That said, the goal of this paper is not to provide the details of an optimal information disclosure policy for these examples, but address more generally what type of information is conducive to career-concerns incentives.

1.1 Literature

The directly related literature is discussed here, and an overview of applied theory work featuring information design questions in the presence of career concerns can be found in Section 6.

Career concerns. The idea of career concerns was introduced formally by Holmström (1999).[1] In a dynamic situation where a principal is not able to offer long term contracts and output contingent wages, incentives can still be generated endogenously. Due to competition for the agent, future wages depend on the belief about the agent's ability, providing incentives today. Holmström (1999) identifies a tractable way of modeling career concerns through the normal-linear model, on which most applications of career concerns are based. While the focus is on the dynamic effort profile for a given information structure, one can perform comparative statics with respect to the noisiness of the performance signal and the uncertainty about the agent's ability and show that the former is detrimental, and the latter beneficial for incentives. Here I explore how general these results are, and show what substantive assumptions are implicit in the normal-linear model that generate them.

Information design. Through examples, Dewatripont, Jewitt and Tirole (1999*a*) show

[1] The paper was circulated in 1982.

that a more accurate performance measure might not raise career-concerns incentives. Like me, they ask when full revelation maximizes incentives in a (quasi-) static model. They show that output needs to be a favorable signal about ability, and its distribution must be increased by effort. These conditions are derived through a partial equilibrium argument wherein the adjustment of the market conjecture about effort is ignored, and they interpret this as local comparative statics around an equilibrium. But I show by example that their qualitative conclusion does not apply in a full equilibrium model, and derive conditions that make it true. Also, the effect of uncertainty regarding the agent's ability is examined here, for which they provide no systematic analysis.

2 Model

There are three players: the principal (she), the agent (he), and a market (it).

Actions. The agent exerts effort e that affects the joint distribution of his ability θ, a signal of ability s_θ called "ability estimate", and output q. This joint distribution is denoted by $\hat{F}(\theta, s_\theta, q \,|\, e)$. Let $\Theta \subset \mathbb{R}$ be the set of abilities, $S_\theta \subset \mathbb{R}$ be a finite set of possible realizations of the ability estimate, and $Q = \{q_1, ..., q_M\} \subset \mathbb{R}$ be a finite set of outputs with $q_1 < ... < q_M$. The set of effort levels E is compact. The marginal distribution over output given effort, denoted $F(q \,|\, e)$, is continuous in e and has full support. The market pays a wage $w \in \mathbb{R}$ to the agent.

The principal decides what information the market gets to observe in form of a signal s, through an information structure H.

- In Section 3.2, she garbles output. Formally, the principal selects an information structure H that specifies for each q a distribution over the signal $s \in \mathbb{R}$, so $H = \{H(s \,|\, q)\}_{q \in Q}$.[2] The ability estimate is irrelevant, in that no player observes it.
- In Section 3.3, she cannot garble output so the market observes q. What she can decide is whether the market observes only output, or output and the ability estimate. Formally, H specifies whether $s = q$ or $s = (q, s_\theta)$.

Information and timing. The timing is summarized as follows.

1. The principal publicly commits to an information structure H.

[2] The signal space is taken to be $\mathbb{R}$ for convenience, and in fact without loss of generality as any equilibrium can be implemented through a direct recommendation mechanism where the principal recommends a wage to the market, and the recommendation is incentive compatible. At any rate, the results are unaffected by the choice of the signal space as long as it is rich enough to fully reveal output.

2. The agent exerts effort e, unobserved by the principal and market.

3. Ability θ, ability estimate s_θ, output q and signal s are realized, and signal s is disclosed to the market.

4. The market pays wage w to the agent.

Since effort can potentially affect the distribution of ability through investment or depreciation, ability is only realized after effort is exerted. This implies in particular that when deciding on the level of effort, the agent has no private information about his ability, a typical assumption in career concern models that is relaxed in Section 5.2.

Payoffs. The principal wants to maximize (the expectation of some increasing function of) effort, but any stated result will be true if the objective is to maximize expected output. The agent maximizes the expectation of $u_A(w, e) = w - c(e)$, where $c(e)$ is nondecreasing and continuous.[3] The market pays the agent his expected ability, perhaps due to Bertrand competition, so formally it maximizes the expectation of $-(w-\theta)^2$ conditional on its information which comprises the signal and a conjecture about the agent's behavior.

Equilibrium. The principal's strategy specifies a randomization over an information structure, and its realization is mapped by the agent's strategy into a randomization over effort. The market's strategy specifies for each information structure and signal realization a wage payment. An equilibrium is a Perfect Bayesian Equilibrium (PBE) that in addition induces a PBE in each subgame following a choice of H. In line with the principal-agent literature, attention will be restricted to the principal's preferred equilibrium.

Since the information structure is publicly observed, it is without loss of optimality to look at equilibria where the principal plays a pure strategy.[4] In a mixed strategy equilibrium, the agent draws effort from $\sigma \in \Delta(E)$ on-path. A pure strategy equilibrium is a mixed strategy equilibrium where σ assigns probability one to some e on-path. Except in Section 4, attention is restricted to pure strategy equilibria.

In most of the analysis I consider the subgame following an arbitrary H, and loosely refer to the PBE induced by the continuation strategies as "equilibrium".

Interpretation of the model. Taken literally, the proposed model represents a two period interaction where in each period, a principal offers a short term contract specifying an upfront (and hence output independent) wage. Ability θ represents the surplus from hiring the agent in the second period, and said surplus is divided through a linear sharing

[3] Assuming convexity of $c(\cdot)$ would be without loss of generality as one can redefine effort to be the cost, which is a monotone transformation of effort. Any assumptions that are imposed on $\hat{F}(\theta, s_\theta, q \mid e)$ are preserved under a monotone transformation of e, unless otherwise stated.

[4] This is even true if only the principal's randomization is publicly observed.

rule between the agent and second period principal (with Bertrand competition as a special case). Since the first period principal cannot offer an output contingent contract, she wants to maximize effort (as it maximizes expected output under an assumption made later). She does so through a disclosure policy about information regarding the agent.

Comment on the production technology. The production side of the economy is specified by $\hat{F}(\theta, s_\theta, q \mid e)$. A more concrete primitive would be a stochastic production function, $\Gamma(q \mid \theta, s_\theta, e)$, and the distribution of ability and ability estimate, $\Phi(\theta, s_\theta \mid e)$. One can then derive $\hat{F}(\theta, s_\theta, q \mid e)$ from $\Gamma(q \mid \theta, s_\theta, e)$ and $\Phi(\theta, s_\theta \mid e)$.

3 Analysis

Section 3.1 provides a short discussion of the normal-linear specification of career concerns, deriving some commonly believed results about how information affects incentives. Sections 3.2 and 3.3 present the main results regarding the informativeness of performance measures and prior uncertainty about ability, respectively.

3.1 The normal-linear example

In this example only, it is assumed that $Q = \mathbb{R}$. Output is additive in ability and effort, so

$$q = \theta + e.$$

When garbling output, the principal can only add normally distributed noise ϵ to the realized output and select its variance. Simple information structures such as partitions of q are not allowed. The market observes the signal

$$s = q + \epsilon,$$

where $\epsilon \sim N(0, \sigma_\epsilon^2)$ and σ_ϵ^2 is controlled by the principal. Ability is normally distributed, $\theta \sim N(\mu_\theta, \sigma_\theta^2)$, and independent of ϵ. Straightforward calculations show that if the market expects effort e^*, the wage satisfies (the information structure is written as $H_{\sigma_\epsilon^2}$)

$$w(s \mid H_{\sigma_\epsilon^2}, e^*) = \alpha(\sigma_\theta^2, \sigma_\epsilon^2, e^*) + \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\epsilon^2} \cdot s,$$

where $\alpha(\sigma_\theta^2, \sigma_\epsilon^2, e^*)$ is independent of the actual effort exerted. Given that $E[s \mid e] = \mu_\theta + e$,

e	$\delta(e)$	$f(q_H \mid e)$	$c(e)$
e_L	1	1/3	0
e_M	2	1/2	$1/6 + 2\epsilon$
e_H	1	2/3	$1/2 + \epsilon$

Table 1: Details for Example 3.1

when exerting e the agent receives an expected payoff

$$E[u_A \mid e] = \alpha(\sigma_\theta^2, \sigma_\epsilon^2, e^*) + \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\epsilon^2} \cdot (\mu_\theta + e) - c(e).$$

Note that marginal incentives to exert effort, $\frac{\sigma_\theta^2}{\sigma_\theta^2+\sigma_\epsilon^2}$, do not depend on the market's conjecture e^*. With a differentiable strictly convex cost function, an interior equilibrium effort is characterized by the first order condition

$$c'(e^*) = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\epsilon^2}.$$

It is immediate that less noise about output (lower σ_ϵ^2) and more uncertainty about ability (higher σ_θ^2) lead to a higher equilibrium effort. As for disclosing the ability estimate, suppose (θ, s_θ) are jointly normal. Disclosing s_θ to the market amounts to reducing σ_θ^2, and is therefore detrimental.

Beyond the normal-linear example, it is not clear under what conditions career-concerns incentives are raised by fully revealing performance measures or increasing uncertainty about the agent's ability. A general analysis is not straightforward for two reasons:

1. The principal optimizes over the set of all information structures.

2. For a given garbling of output H, it is not transparent how $w(s \mid H, \hat{e})$ depends on the market conjecture $\hat{e}$.

In the normal-linear example, one restricts attention to a parametric class of information structures so that one can derive a closed form solution for the wage schedule, which eliminates either problem.

3.2 Information about performance

In the following example, output is a favorable signal of ability and effort increases output, yet a noisy disclosure of output maximizes equilibrium effort.

Example 3.1 *Let* $E = \{e_L, e_M, e_H\}$ *and* $Q = \{q_L, q_H\}$*. Since output is binary, all that matters for incentives given the market's conjecture* $\hat{e}$ *under full revelation is* $\delta(\hat{e}) := E(\theta \,|\, q_H, \hat{e}) - E(\theta \,|\, q_L, \hat{e})$*, since up to a constant the agent then maximizes* $f(q_H \,|\, e)\delta(\hat{e}) - c(e)$*. The primitives are specified in Table 1, where* ϵ *is a positive number sufficiently close to zero. Under full revelation, the unique equilibrium is* e_L*. If the market conjecture is* e_M*, the agent wants to deviate to* e_H*, and if* e_H *is conjectured the agent wants to deviate to* e_L*.*

Yet e_M *can be sustained by a noisy information structure* H *that maps* q_L *into signal* s_L*, and* q_H *into a uniform randomization over* s_L *and* s_H*. Given that* $\delta(e_M) = 2$*, one can verify that* $E_s[w(s \,|\, H, e_M) \,|\, q_H] - E_s[w(s \,|\, H, e_M) \,|\, q_L] = 3/2$*. This sufficiently decreases incentives to deviate to* e_H*, yet* e_M *remains preferable to* e_L *thus implementing* e_M*.*

In Example 3.1, better information leads to stronger incentives *given the market's conjecture*. So in a partial equilibrium sense, full revelation maximizes incentives. Proposition 3.2 below provides conditions for when this is true *in equilibrium*. First, say that $F(q \,|\, e)$ is ordered according to the monotone likelihood ratio property (MLRP) if for any e_1, e_2 with $e_2 > e_1$, $f(q \,|\, e_2)/f(q \,|\, e_1)$ is nondecreasing in q.

Proposition 3.2 *Assume that*

- $F(q \,|\, e)$ *is ordered according to the MLRP,*
- $E(\theta \,|\, q, e)$ *is nondecreasing in* q*,*
- $E(\theta \,|\, q, e)$ *is supermodular in* (q, e)*.*

A best pure strategy equilibrium exists, and is achieved by fully revealing output.

The first two assumptions are natural ordering conditions. The relevant concept of effort increasing the distribution of output is the MLRP, and output is a favorable signal of ability as $E(\theta \,|\, q, e)$ is nondecreasing in q. It will follow from the sketch of the proof below that supermodularity of $E(\theta \,|\, q, e)$ in (q, e) is a *sufficient* condition for effort to increase the informativeness of output about ability. In fact, it can be weakened to the *exact* condition that effort increases informativeness, but for expositional reasons the discussion is in terms of supermodularity.

Whether effort makes output a more informative signal of ability will depend on the application, as the following example illustrates. Suppose effort and output are binary, so the agent can either work or shirk, and succeed or fail. First consider a task that is difficult in that the agent always fails if he shirks, but succeeds with a positive probability that is increasing in ability when working. Under shirking the outcome is completely uninformative about ability, but not so when the agent works. The opposite is true for a mundane task where the agent always succeeds when working, yet there is some chance of failure when shirking that is decreasing in the agent's ability. For such a task, informativeness of the outcome is decreasing in effort and $E(\theta \mid q, e)$ is (strictly) submodular. Indeed, the "intuitive" result that full revelation of output maximizes effort then fails in that a noisy signal of output maximizes effort for some cost of effort function.

Necessity of the conditions. Given the somewhat abstract environment, it is natural to ask to what extent the conditions in Proposition 3.2 are necessary. Attention is restricted to primitives $\hat{F}(\theta, s_\theta, q \mid e)$ that satisfy the first two assumptions, and these are called *ordered statistical environments*. So the question is, to what extent is it necessary that effort increases the informativeness of output. Fixing the statistical environment $\hat{F}(\theta, s_\theta, q \mid e)$, the only remaining primitive is $c(\cdot)$.

Proposition 3.3 *Take any ordered statistical environment. If there exist e_1, e_2 with $e_2 > e_1$ such that $E(\theta \mid q, e)$ is strictly submodular*[5] *on $Q \times \{e_1, e_2\}$, and $F(\cdot \mid e_1) \neq F(\cdot \mid e_2)$, then there exists a cost function such that the best pure strategy equilibrium is achieved through an information structure that is not fully revealing.*

Propositions 3.2 and 3.3 do not provide a tight characterization in that $E(\theta \mid q, e)$ can be neither sub- nor supermodular. While the proof of either proposition uses weaker conditions, there remains a gap which is related to the incompleteness of the Blackwell ordering of informativeness.

The following discussion develops the techniques used to show Proposition 3.2, and provides a sketch of the proof. A fully revealing information structure is denoted by H_{FR}.

A representation of information structures. The wage paid by the market depends on the realized signal s, the information structure H, and market conjecture $\hat{e}$. Let $t(q \mid H, \hat{e})$ denote the expected wage the agent receives given $(H, \hat{e})$ if output q is realized,

$$t(q \mid H, \hat{e}) = E_s[w(s \mid H, \hat{e}) \mid q].$$

$t(\cdot \mid H, \hat{e})$ will be referred to as the *transfer schedule* induced by $(H, \hat{e})$. After each output realization the agent faces a lottery over wages, which can be identified with its mean under risk neutrality.

If output is garbled, one would expect that roughly speaking the agent is made better off when realized output indicates low ability, and the opposite when it indicates high ability.

[5]Strictly submodular means that $\forall q_2 > q_1$, $E(\theta \mid q_2, e_2) - E(\theta \mid q_1, e_2) < E(\theta \mid q_2, e_1) - E(\theta \mid q_1, e_1)$.

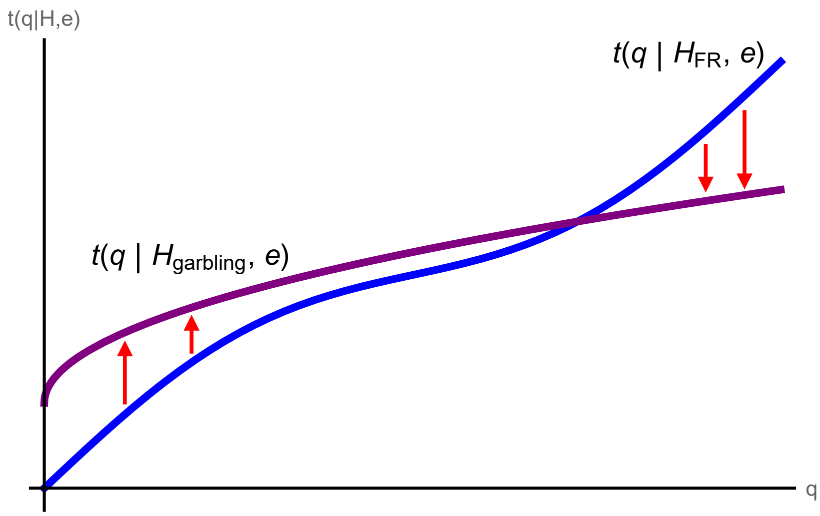

Figure 1: Effect of a garbling

E.g. if output is binary, any nontrivial garbling will sometimes lead to a signal that leaves the market to some extent uncertain about whether output is high or low. Conditional on the low output being realized, the garbling is preferable to full revelation for the agent.

This suggests that $t(q \mid H, \hat{e})$ lies above (below) $t(q \mid H_{\mathrm{FR}}, \hat{e})$ for low (high) values of q, as illustrated Figure 1. In what follows, I define the sense in which $t(\cdot \mid \cdot, \hat{e})$ becomes "flatter" as one garbles information. The significance of this is that flatter transfer schedules depress incentives, albeit in a relatively weak sense.

Implementability: a necessary condition. For the fully revealing information structure, the transfer schedule satisfies $t(q \mid H_{\mathrm{FR}}, \hat{e}) = E(\theta \mid q, \hat{e})$, $\forall q$, given conjecture $\hat{e}$. For any other information structure H, one can say that output q gets "subsidized" if $t(q \mid H_{\mathrm{FR}}, \hat{e}) > E(\theta \mid q, \hat{e})$, and "taxed" if the inequality is reversed. As discussed, it should be the case that lower outputs tend to be subsidized, and the opposite for higher outputs. This is formalized as follows.

Lemma 3.4 *Fix any information structure H and conjecture $\hat{e}$. If $E(\theta \mid q, \hat{e})$ is nondecreasing in q, then*

$$\sum_{i=1}^{j} t(q_i \mid H, \hat{e}) f(q_i \mid \hat{e}) \geq \sum_{i=1}^{j} t(q_i \mid H_{FR}, \hat{e}) f(q_i \mid \hat{e}) \qquad (1)$$
$$\forall j = 1, ..., M, \text{ with equality for } j = M.$$

The result says that for any garbling, the "cumulative subsidy up to j", $\sum_{i=1}^{j} [t(q_i \mid H, \hat{e}) - t(q_i \mid H_{\mathrm{FR}}, \hat{e})] f(q_i \mid \hat{e})$, has to be non-negative, and vanish for the highest possible output. The "with equality for $j = M$" condition has to be satisfied since the expected wage has to

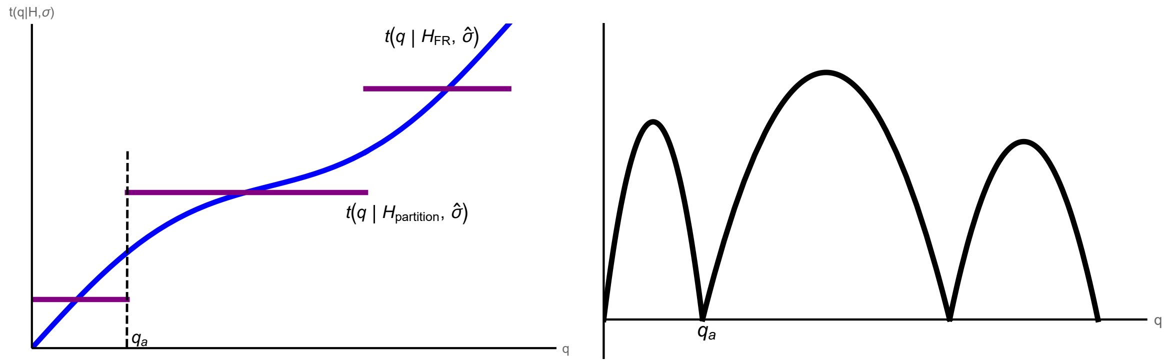

(a) $t(q \mid H_{\text{partition}}, \hat{\sigma})$ and $t(q \mid H_{\text{FR}}, \hat{\sigma})$ (b) Cumulative subsidy

Figure 2: Illustration of the cumulative subsidy condition

coincide with the expected ability. From now on, Equation (1) will be called the "cumulative subsidy condition".

Figure 2 illustrates the result of Lemma 3.4. Information structure $H_{\text{partition}}$ corresponds to an "interval" partition, where each element of the partition is an adjacent set of scores. In this special case, the cumulative subsidy is zero at each of the break points such as q_a. The market never confounds outputs weakly below q_a with those above it, so it pays a "fair" wage conditional on q being weakly below q_a.

Moving beyond partitions, consider first the case of binary output. After any signal realization, the market is to some extent uncertain as to which output generated it, and pays an intermediate wage. This subsidizes (taxes) the low (high) output, and when weighted by the relative frequency of the two outputs, the subsidy and tax are of the same magnitude. The same reasoning applies if output is not binary, yet every signal is generated by at most two outputs. For a more general information structure where every signal can be generated by an arbitrary subset of outputs, consider any signal s'. It turns out that one can replace signal s' with a finite number of new signals, where each new signal is generated by two outputs only, and leads to exactly the same wage as signal s'.[6] The bottom line is that after the translation to such an information structure that has the same joint distribution over wages and outputs, each signal generates a subsidy and tax in a way that satisfies the cumulative subsidy condition.

Implementability: a sufficient condition. The next step is to ask which transfer schedules can be implemented by some information structure. Satisfying the cumulative

[6]Given the posterior distribution over output conditional on s', start out by pooling the highest and lowest output into signal s^1 in such a way that s^1 leads to $w(s')$. This can be done until the mass of one of the two outputs is exhausted, say the highest output. Now start pooling the second highest output with the lowest one into signal s^2, again generating $w(s')$. This procedure can be continued until one exhausts the probability mass of all outputs, indeed it terminates after at most M steps.

subsidy condition of Lemma 3.4 turns out not to be sufficient.

Lemma 3.5 *Fix a conjecture* $\hat{e}$ *and assume that* $E(\theta \,|\, q, \hat{e})$ *is nondecreasing in* q*. Suppose that for some nondecreasing function* $g : Q \to \mathbb{R}$*,*

$$\sum_{i=1}^{j} g(q_i) f(q_i \,|\, \hat{e}) \geq \sum_{i=1}^{j} t(q_i \,|\, H_{FR}, \hat{e}) f(q_i \,|\, \hat{e})$$
$$\forall j = 1, ..., M, \text{ with equality for } j = M.$$

Then there exists an information structure H *with* $t(\cdot \,|\, H, \hat{e}) = g(\cdot)$*.*

The added condition in Lemma 3.5 is that the transfer schedule is nondecreasing, under which one can construct an information structure that implements it. It is fairly immediate that a strictly decreasing transfer schedule is not implementable, since given that output is a favorable signal of ability, it would imply that the market is systematically fooled in that it pays a lower wage for outputs which indicate high ability. But the following example shows that a locally decreasing transfer schedule might be implementable. Take $Q = \{q_1, q_2, q_3\}$, and an information structure $\tilde{H}$ that reveals q_2 and pools $\{q_1, q_3\}$ into the same signal. For a given conjecture $\hat{e}$, one typically has $E(\theta \,|\, q_2, \hat{e}) \neq E(\theta \,|\, q \in \{q_1, q_3\}, \hat{e})$, which leads to a nonmonotone transfer schedule as either $t(q_2 \,|\, \tilde{H}, \hat{e}) > t(q_1 \,|\, \tilde{H}, \hat{e}) = t(q_3 \,|\, \tilde{H}, \hat{e})$ or $t(q_2 \,|\, \tilde{H}, \hat{e}) < t(q_1 \,|\, \tilde{H}, \hat{e}) = t(q_3 \,|\, \tilde{H}, \hat{e})$. Be that as it may, such transfer schedules are not optimal for the principal since they discourage effort by (sometimes) penalizing high output.

Information and incentives: partial equilibrium. The cumulative subsidy condition of Lemma 3.4 defines a sense in which garbling output makes a transfer schedule "flatter". The effect on incentives is the content of the following lemma.

Lemma 3.6 *Assume that* $E(\theta \,|\, q, e)$ *is nondecreasing in* q *and* $F(q \,|\, e)$ *is ordered according to the MLRP. Fix* $\hat{e}$ *and* H*. Then* $\forall e \leq \hat{e}$*,*

$$E_q[u_A(t(q \,|\, H_{FR}, \hat{e}), \hat{e}) \,|\, \hat{e}] \;-\; E_q[u_A(t(q \,|\, H_{FR}, \hat{e}), e) \,|\, e] \geq$$
$$E_q[u_A(t(q \,|\, H, \hat{e}), \hat{e}) \,|\, \hat{e}] \;-\; E_q[u_A(t(q \,|\, H, \hat{e}), e) \,|\, e].$$

Lemma 3.6 states that for a given market conjecture $\hat{e}$, moving from an arbitrary garbling to full revelation makes $\hat{e}$ more attractive compared to any lower effort. Better information about output and effort are complementary, albeit in a weak sense. It is for example *not* true that for two arbitrary effort levels $e_H > e_L$, the higher effort e_H becomes more attractive when moving to full revelation, unless $e_H = \hat{e}$.

Nevertheless, Lemma 3.6 can be seen as a partial equilibrium version of the result that under appropriate ordering conditions, noisy disclosure of output reduces effort due to career-concerns incentives. For if some e^* is an equilibrium effort under a garbling H, then Lemma 3.6 implies that the agent wants to deviate (weakly) upwards if output gets fully revealed, yet the market's conjecture remains at e^*. Below, it will be shown how this can be developed into a full equilibrium argument under the supermodularity assumption on $E(\theta \mid q, e)$.

That $F(q \mid e)$ has to satisfy the MLRP in Lemma 3.6 is due to the absence of restrictions on how output can be garbled. If one were to restrict attention to information structures H such that the transfer schedule corresponding to full revelation is everywhere steeper, then the concept of first order stochastic dominance (FOSD) would be sufficient to deliver the same conclusion.[7] Yet this would rule out even simple information structures like partitions.

Proof of Proposition 3.2: a sketch. The idea behind Proposition 3.2 will be explained in a series of steps, which are illustrated in Figure 3. The result is shown by contradiction, starting with an arbitrary information structure H_1 and equilibrium effort e_1. It will be shown that either e_1 can be also implemented by full revelation, or there exists an alternative (possibly noisy) information structure that implements an effort $e_2 > e_1$. Therefore, the best effort across all pure strategy equilibria (which can be shown to exist) must be achieved by full revelation. The sketch is as follows:

(a) Start with the equilibrium that involves a garbled information structure H_1 and effort level e_1. This leads to transfer schedule $t(\cdot \mid H_1, e_1)$.

(b) Suppose that the principal moves to full revelation of output, but the market still conjectures that the agent exerts e_1, so that the transfer schedule changes to $t(\cdot \mid H_{\text{FR}}, e_1)$. By Lemma 3.6, the agent wants to deviate weakly upwards to some effort level $e_2 \geq e_1$. If $e_2 = e_1$, this implies that H_{FR} also implements e_1. Going forward, the case $e_2 > e_1$ is considered.

(c) The goal is to show that e_2 can be implemented by some information structure. Observe that H_{FR} provides incentives that are too strong, since under supermodularity $t(q \mid H_{\text{FR}}, e_2) = E(\theta \mid q, e_2)$ is steeper than $t(\cdot \mid H_{\text{FR}}, e_1) = E(\theta \mid q, e_1)$, and hence might induce even more effort.

(d) But by Lemma 3.5 there exists some information structure H_2 that induces a transfer

[7] $t(\cdot \mid H_{\text{FR}}, \hat{e})$ is everywhere steeper than $t(\cdot \mid H, \hat{e})$ if $\forall q_2 > q_1$, $t(q_2 \mid H_{\text{FR}}, \hat{e}) - t(q_1 \mid H_{\text{FR}}, \hat{e}) \geq t(q_2 \mid H, \hat{e}) - t(q_1 \mid H, \hat{e})$. See Lemma A.5 for the argument. In the normal-linear example, all transfer schedules are linear and hence can be ordered by steepness, but in any case the normal distribution satisfies the MLRP. Another example is when output is binary (and the set of information structures is unrestricted), yet in this case FOSD implies the MLRP.

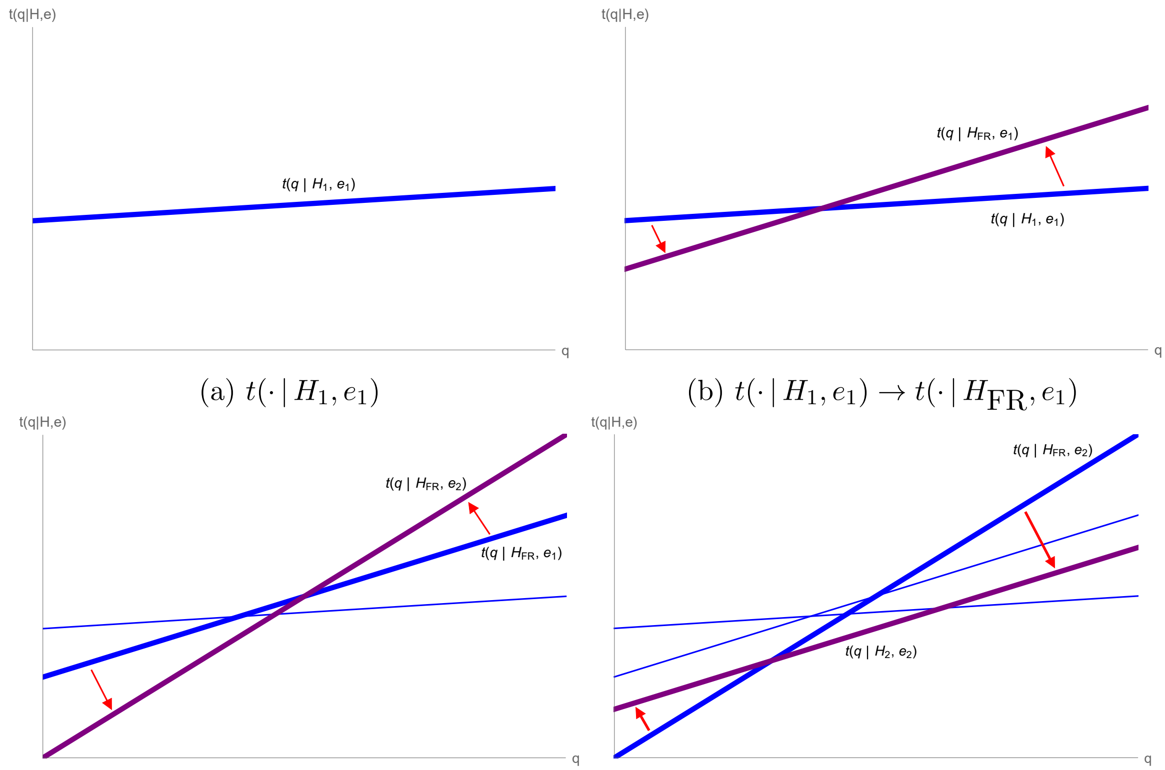

(a) $t(\cdot \,|\, H_1, e_1)$ (b) $t(\cdot \,|\, H_1, e_1) \to t(\cdot \,|\, H_{\text{FR}}, e_1)$

(c) $t(\cdot \,|\, H_{\text{FR}}, e_1) \to t(\cdot \,|\, H_{\text{FR}}, e_2)$ (d) $t(\cdot \,|\, H_{\text{FR}}, e_2) \to t(\cdot \,|\, H_2, e_2)$

Figure 3: Illustration of the proof of Proposition 3.2

schedule that coincides with $t(\cdot \,|\, H_{\text{FR}}, e_1)$ up to a constant. Since the constant does not affect incentives, e_2 is a best response to $t(\cdot \,|\, H_2, e_2)$.

Supermodularity of $E(\theta \,|\, q, e)$ is used in steps (c) and (d). What one really needs is that when the market conjectures e_2, then $t(\cdot \,|\, H_{FR}, e_1)$ can be induced by some information structure H_2 up to an appropriately defined constant. The proof of Proposition 3.2 formalizes this by Assumption A.3, and its meaning is that effort increases the informativeness of output about ability in the Blackwell sense. It is implied by supermodularity which is used for expositional reasons.[8]

On the cardinality of Q. Proposition 3.2 holds for any measurable $Q \subset \mathbb{R}$, as shown in Proposition A.8 in Appendix A.1 under a mild regularity condition.[9] In particular, one can verify that the normal-linear example satisfies the assumptions of Proposition 3.2. The

[8]Strictly speaking, the interpretation regarding informativeness is only accurate when effort does not affect the marginal distribution of ability. Otherwise, there is a second effect in that the underlying uncertainty about ability changes. Given that the goal of this section is to perform comparative statics with respect to uncertainty about performance, the assumption will be referred to as effort increasing informativeness of output even though Proposition 3.2 applies also to situations where effort affects ability.

[9]There is a slight gap in that it is only shown under supermodularity of $E(\theta \,|\, q, e)$, but not under the weaker Assumption A.3.

MLRP holds since $F(q \,|\, e) = N(\mu_\theta + e, \sigma_\theta^2)$, and $E(\theta \,|\, q, e) = q - e$ which is increasing in q and (trivially) supermodular.

Corollary 3.7 *In the normal-linear example, full revelation leads to the highest effort among all pure strategy equilibria.*

3.3 Information about ability

To study the effect of uncertainty about ability on incentives, it is assumed that output is not garbled but the principal can decide whether to disclose *ability estimate* s_θ to the market. To fix ideas, suppose that s_θ is *exogenous* in that its marginal distribution, $G(s_\theta \,|\, e)$, does not depend on effort. This is a natural condition if the marginal distribution of ability is independent of effort, and s_θ is a garbling of θ. Yet it can be also satisfied if effort enhances or depreciates ability.[10] In line with the previous analysis, the agent does not observe s_θ when taking effort. Section 5.2 explores the opposite case.

Intuitively one would expect that there is less of an incentive to try to fool the market, as its inference will partially rely on the exogenous signal s_θ. The following example shows that this is not true in general.

Example 3.8 *Let $\Theta = \{0, 1\}$, $E = \{e_L, e_H\}$, and $Q = \{q_L, q_H\}$. The task is one of two "types", τ_0 and τ_1, and the agent does not know the task type. The probability of high ability $\theta = 1$ is 2/3, as is the probability of facing task τ_1. Ability and task type are independently distributed, and unaffected by effort. Effort e_L leads to q_L, for any θ and τ.[11] When e_H is exerted, then q_H is generated if $(\theta, \tau) \in \{(0, \tau_0), (1, \tau_1)\}$, that is when ability and task "match", while the remaining combinations of (θ, τ) lead to q_L.*

Suppose that the market conjectures $\hat{e} = e_H$. When exerting e_H, the agent receives an expected wage of $Pr(\theta = 1) = 2/3$. If he deviates to e_L, output q_L is realized and he receives

$$Pr(\theta = 1 \,|\, q = q_L, \hat{e} = e_H) = \frac{Pr(\theta = 1) \cdot Pr(\tau = \tau_0)}{Pr(\theta = 1) \cdot Pr(\tau = \tau_0) + Pr(\theta = 0) \cdot Pr(\tau = \tau_1)} = \frac{1}{2}.$$

If $c(e_H) - c(e_L) > \frac{2}{3} - \frac{1}{2} = \frac{1}{6}$, then e_H is not an equilibrium.

Now suppose that the actual task gets revealed to the market through signal $s_\theta \in \{\tau_0, \tau_1\}$, together with output (the agent still does not know which type of task he is facing). Given conjecture $\hat{e} = e_H$, the agent still receives an expected wage of $Pr(\theta = 1) = 2/3$ when exerting

[10] For example, consider the case where ability is a deterministic function ϕ of some random variable γ and effort so that $\theta = \phi(\gamma, e)$, and the marginal distribution of γ is exogenous. If s_θ is a garbling of γ this fits the setup as well.

[11] The violation of the full support assumption on $F(q \,|\, e)$ is not critical for the conclusion.

e_H. If the market observes q_L and $s_\theta = \tau_0$ ($s_\theta = \tau_1$), given its conjecture of $\hat{e} = e_H$ it believes the agent is of ability $\theta = 1$ ($\theta = 0$). So the agent receives an expected wage of $Pr(\tau = \tau_0) = \frac{1}{3}$ when deviating to e_L. If $c(e_H) - c(e_L) \leq \frac{2}{3} - \frac{1}{3} = \frac{1}{3}$, effort e_H is indeed an equilibrium.

In summary, if $\frac{1}{6} < c(e_H) - c(e_L) \leq \frac{1}{3}$, strictly higher effort is generated if the market observes the task.

Proposition 3.9 below spells out conditions when disclosing the ability estimate hurts incentives. An important component will be the distribution of s_θ conditional on (q, e), denoted as $H(s_\theta \,|\, q, e)$. Even when the marginal distribution of s_θ does not depend on e, this is in general not the case once one also conditions on q. This is since both effort and ability estimate affect the distribution of output.

Proposition 3.9 *Assume that*

- *$F(q \,|\, e)$ is increasing in e according to FOSD,*
- *$E(\theta \,|\, q, e)$ is supermodular in (q, e),*
- *$E(\theta \,|\, q, s_\theta, e)$ is nondecreasing in s_θ,*
- *$H(s_\theta \,|\, q, e)$ is decreasing in e according to FOSD.*

Equilibrium effort is higher if the ability estimate s_θ is not disclosed to the market.

The first two assumptions are already familiar from Proposition 3.2, except that the MLRP can be weakened to FOSD as output is not garbled. That $E(\theta \,|\, q, s_\theta, e)$ is nondecreasing in s_θ is an ordering condition on the ability estimate. The remaining condition says that conditional on output, higher effort causes a negative inference about s_θ. To interpret this, suppose s_θ is exogenous (the result does not rely on this). Roughly speaking the assumption then means that output is increasing in effort and ability estimate.

In Example 3.8, the ordering condition that $E(\theta \,|\, q, s_\theta, e)$ is nondecreasing in s_θ is violated. Given a market conjecture of e_H, $s_\theta = \tau_1$ is a favorable signal of ability if q_H is realized, but bad news conditional on q_L.

Notice that under the conditions of Proposition 3.2 it is optimal to fully reveal output if s_θ is not disclosed, and Proposition 5.7 below derives conditions when this is also optimal if s_θ is disclosed. Under the assumptions of Propositions 3.2, 3.9, and 5.7, the principal benefits from the absence of s_θ, even if she can garble q. Yet this is not true if the principal can garble (q, s_θ) directly, as will be discussed in detail at the end of Section 5.3.1.

4 Random effort

In this section only, the agent is allowed to randomize over effort. Section 4.1 studies mixed strategy equilibria, while Section 4.2 allows for a more general information revelation technology that becomes available when the agent randomizes. Either section considers a principal garbling output, so s_θ is suppressed.

4.1 Mixed strategy equilibrium

Notice that for a given information structure, a pure strategy equilibrium might not exist. Such garblings were implicitly ruled out above. The following can be said once the agent is allowed to randomize.

Proposition 4.1 *Assume that*

- $F(q \mid e)$ *is ordered according to the MLRP,*
- $E(\theta \mid q, \sigma)$ *is nondecreasing in* q $\forall \sigma \in \Delta(E)$,
- $E(\theta \mid q, e)$ *is supermodular in* (q, e),
- $E(\theta \mid q, e)$ *is nonincreasing in* e.

Then:

1. *If* e_P *is the best pure strategy equilibrium, then for any mixed equilibrium* σ *under full revelation,*
$$e_P \geq sup(support(\sigma)).$$
2. *If either* E *or* Q *is binary, the best mixed equilibrium is pure and is induced by full revelation.*

One complication is that on the set of mixed strategies, the MLRP cannot be satisfied in general.[12] Lemma 3.6 still applies in that starting from an arbitrary garbling H and market conjecture σ, any e that dominates (is dominated by) σ according to the MLRP becomes more (less) attractive than σ after moving to full revelation. But since some effort levels are not comparable to σ, one cannot conclude that the agent deviates upwards.

Consider a non-degenerate mixed strategy equilibrium under full revelation. Given the MLRP of $F(q \mid e)$, high output is indicative of high effort, and whether this effect leads to

[12] More precisely, if $F(q \mid e)$ satisfies the MLRP given the natural order on E, then unless E is binary there is in general no complete order on $\Delta(E)$ such that the MLRP is satisfied with respect to that order.

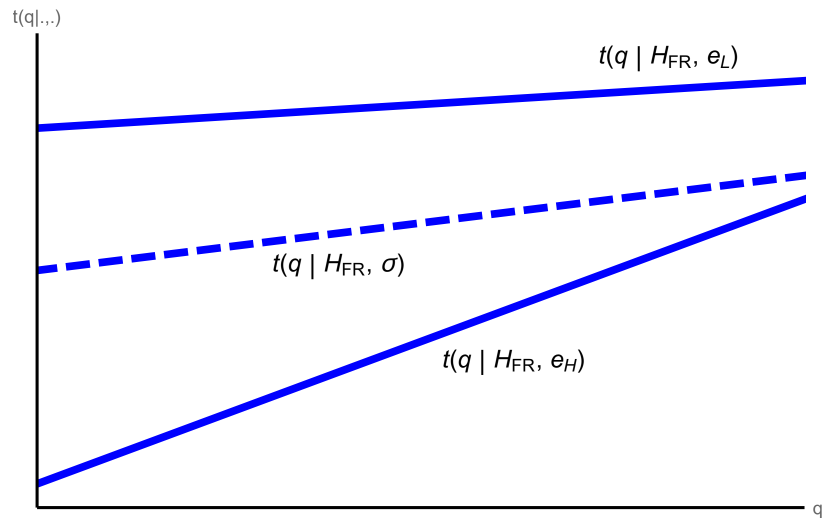

Figure 4: Effect of mixing on the transfer schedule: $t(q \mid H_{FR}, \sigma)$ lies in between $t(q \mid H_{FR}, e_L)$ and $t(q \mid H_{FR}, e_H)$, and under MLRP a higher q is more indicative of high effort.

a favorable inference about ability or not depends on how $E(\theta \mid q, e)$ varies with e. This is illustrated in Figure 4, where σ is a randomization over (e_L, e_H). $E(\theta \mid q, e)$ is decreasing in e so that there is a "bad news" effect of high output about ability, in that for higher output it is more likely that the agent exerted e_H which leads to an adverse inference about θ. This dampens incentives, as formalized in the following lemma.

Lemma 4.2 *Fix a mixed strategy σ. Let $\bar{e} = sup\ (support(\sigma))$. Assume that $E(\theta \mid q, e)$ is supermodular in (q, e) and nonincreasing in e, and that $F(q \mid e)$ is ordered according to the MLRP. Then for any $q_2 > q_1$,*

$$E(\theta \mid q_2, \bar{e}) - E(\theta \mid q_1, \bar{e}) \geq E(\theta \mid q_2, \sigma) - E(\theta \mid q_1, \sigma).$$

Under full revelation, the transfer schedule becomes steeper when the market conjecture changes from mixed strategy σ to the highest effort in the support of σ, called $\bar{e}$. This is due to supermodularity and the bad news effect of high output when $E(\theta \mid q, e)$ is nonincreasing in e. The significance of Lemma 4.2 is that if σ is an equilibrium, then $\bar{e}$ has to be a best response (by continuity). But adjusting the conjecture to $\bar{e}$, incentives become stronger. The best pure strategy equilibrium has to be at least $\bar{e}$ by the arguments leading to Proposition 3.2, explaining part 1 of Proposition 4.1.

When E is binary, full revelation leads to stronger incentives since $F(q \mid e)$ satisfies the MLRP on $\Delta(E)$. That random effort cannot be optimal is due to the bad news effect of higher output (as $E(\theta \mid q, e)$ is nonincreasing in e). This is also true when Q is binary, where incentives for effort are summarized by $\delta(H, \sigma) := t(q_2 \mid H, \sigma) - t(q_1 \mid H, \sigma)$. More information then leads to unambiguously stronger incentives by increasing $\delta(\cdot,\cdot)$.

To my knowledge, only Dewatripont, Jewitt and Tirole (1999*b*) considered mixed strategies in the presence of career concerns, but the context is different. Effort is multidimensional, and the agent randomizes over the task he exerts effort on. The overall amount of effort is constant, and the randomization across tasks makes output a less accurate signal of ability. Here, the agent randomizes over the overall amount of effort, and the effect of this on the market's inference is different.

4.2 Mediated information structures

This section allows for a generalized communication protocol between the principal and agent. Instead of just garbling output, the principal also commits to make a secret effort recommendation to the agent. The market only observes the realized signal, but not the effort recommendation. This endogenously creates asymmetric information between the agent and market, and is potentially beneficial for the principal.

Formally, an information structure consists of $(\mathcal{M}, \alpha_{\mathcal{M}}, H)$, where $\mathcal{M}$ is a message space, $\alpha_{\mathcal{M}} \in \Delta(\mathcal{M})$ is a distribution over the message space and H specifies for each combination of $q \in Q$ and $m \in \mathcal{M}$ a signal distribution $H(\cdot \mid m, q) \in \Delta(\mathbb{R})$. The timing is now as follows.

1. The principal publicly commits to an information structure $(\mathcal{M}, \alpha_{\mathcal{M}}, H)$.
2. The principal draws a realization from $\alpha_{\mathcal{M}}$ and privately communicates it to the agent.
3. The agent exerts effort e, unobserved by the principal and market.
4. Ability θ, output q and signal s are realized, and signal s is disclosed to the market.
5. The market pays wage w to the agent.

Any information structure in this communication protocol will be called a *mediated*, while the standard mechanisms up to now are called *conventional*. Since the principal can commit to communicate the realization of $\alpha_{\mathcal{M}}$ truthfully, it is without loss of generality to look at equilibria where the agent is obedient. That is, $\mathcal{M} = E$ and $\alpha_{\mathcal{M}} = \sigma$, where σ is the equilibrium distribution over effort.

Clearly, a mediated mechanism can improve upon a conventional one only when the agent plays a randomized strategy. Compared to a mixed strategy equilibrium induced by a conventional mechanism, one difference is that the agent need not be indifferent between all effort levels in the support of his strategy after each message he receives, but possibly strictly prefers the recommended effort.

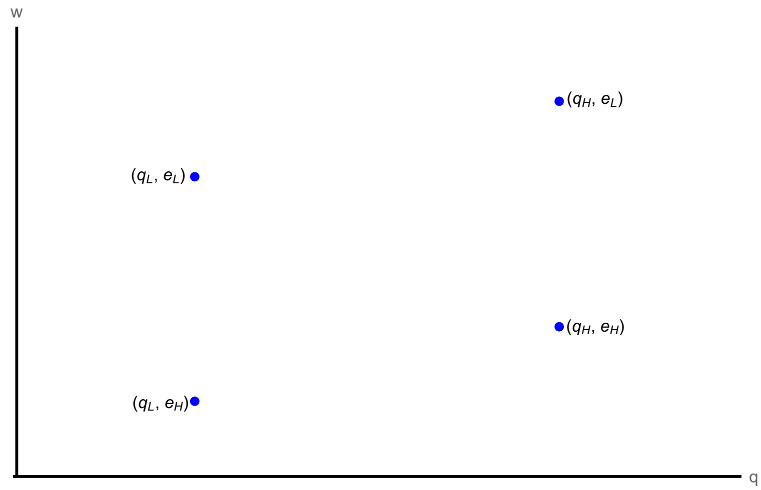

(a) Wages under full revelation

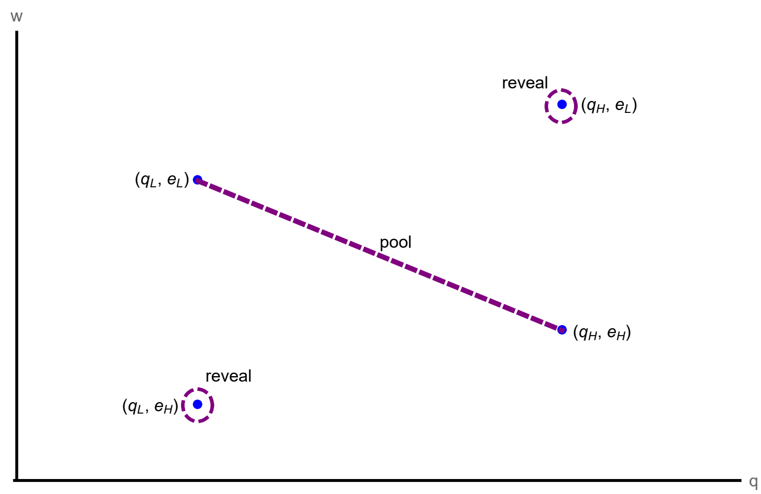

(b) Mediated information structure

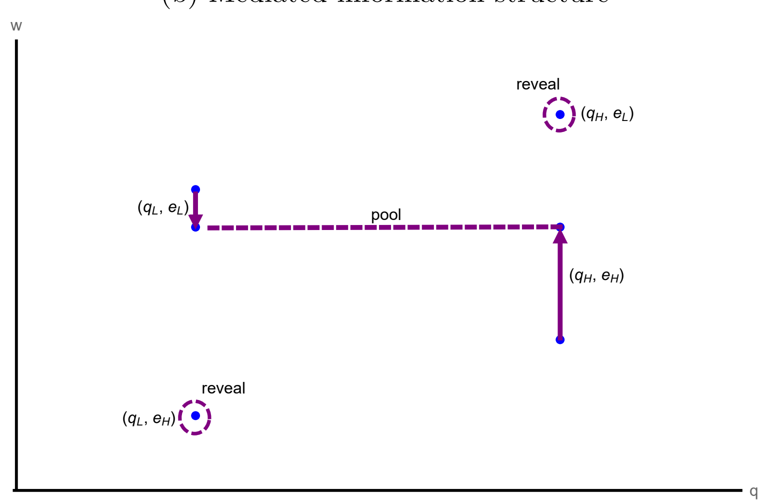

(c) Wages under mediated information structure

Figure 5: Illustration of a mediated information structure

e	$E(\theta \mid q_L, e)$	$E(\theta \mid q_H, e)$	$f(q_H \mid e)$	$c(e)$
e_L	4	5	1/3	0
e_H	0	2	2/3	1

Table 2: Details for Example 4.3

signal s	outcomes inducing s	wage $w(s)$
s_1	(q_L, e_H)	0
s_2	$\{(q_L, e_L), (q_H, e_H)\}$	3
s_3	(q_H, e_L)	5

Table 3: Mediated equilibrium in Example 4.3: $\sigma(e_L) = \sigma(e_H) = 1/2$

A mediated information structure allows the principal to correlate the effort and signal, *conditional on* q. The following example illustrates how the principal can benefit from this.

Example 4.3 *Let* $E = \{e_L, e_H\}$ *and* $Q = \{q_L, q_H\}$*, the primitives are specified in Table 2. Incentives are driven by the difference in expected wage across outcomes* q_L *and* q_H*. In order to sustain* e_H*, this difference needs to be at least 3. Under full revelation,* e_L *is the unique equilibrium as conditional on* e_H*, the wage difference is* $2 - 0 = 2$*. Since the conditions of Proposition 4.1 apply, this is the best outcome that can be achieved through conventional mechanisms even when considering mixed equilibria. Table 3 summarizes a mediated equilibrium where the agent randomizes with equal probability between* e_L *and* e_H*. Using* (q, e) *to denote a combination of an output realization and an effort recommendation, then* (q_L, e_H) *and* (q_H, e_L) *get revealed (leading to wages of 0 and 5, respectively), while* $\{(q_L, e_L), (q_H, e_H)\}$ *get pooled into the same signal (leading to a wage of 3). The point is that after recommendation* e_L*, the agent's difference in wages across* q_L *and* q_H *equals* $5 - 3 = 2$*, while after recommendation* e_H*, the wage difference equals* $3 - 0 = 3$*, making the respective recommendations incentive compatible.*

Figure 5 illustrates the mechanics of the example. In equilibrium, the agent randomizes between e_L and e_H, and panel (a) plots the wages the agent would get under full revelation and either effort conjecture. Panel (b) describes how the mediated information structure reveals different output realizations and effort recommendations. Since $E(\theta \mid q_L, e_L) > E(\theta \mid q_H, e_H)$, pooling $\{(q_L, e_L), (q_H, e_H)\}$ rewards q_H after recommendation e_H, and punishes q_L after recommendation e_L, compared to full revelation. By increasing the difference in the wage across the two outputs, incentives are raised after either recommendation as illustrated in panel (c). The example specifies a cost function that makes the effort recommendations incentive compatible, yet can only sustain e_L with conventional mechanisms.

When a mediated mechanism improves upon the best conventional one, it will do so by not fully revealing output. Ruling out the optimality of mediated mechanisms appears elusive, but should not be confused with the difficulty of establishing a full revelation result in the class of mixed strategy equilibria (with conventional mechanisms).[13] There, one can show that garbling output "flattens" the transfer schedule in the sense of the cumulative subsidy condition. But since there is no complete ordering on $\Delta(E)$ satisfying the MLRP, the agent does not necessarily want to deviate upwards when output gets fully revealed.

In contrast, a mediated mechanism can violate the cumulative subsidy condition, as happens in Example 4.3. This is true both in terms of the overall strategy σ, but also conditional on every effort recommendation. Mediated information structures make this possible because they allow to correlate the signal with effort, *conditional on* q. This is discussed in more detail in Section 6.

5 Extensions

Several extensions are considered in this section, such as the presence of explicit wage contracts (Section 5.1), the agent having information about his ability (Section 5.2), and various other extensions (Section 5.3).

5.1 Wage contracts

So far, incentives were only derived from career concerns. Suppose now that there is also an explicit wage contract $\hat{w}$, offered by the principal. The agent maximizes the expectation of $\hat{u}_A(\hat{w}, w, e) = \hat{w} + w - c(e)$. The question is to what extent the full revelation result of Proposition 3.2 still applies.

For now, a wage contract $\hat{w} : Q \to \mathbb{R}$ is taken as given. Garbling output raises a conceptual issue as to what wage the agent receives. One possibility is that the principal can achieve a separation between payments and performance evaluation. This requires that the market does not observe the wage payment $\hat{w}(q)$, and the principal still observes the actual output realization even though it is garbled. In this case, the wage contract is said to be *unaffected by the disclosure policy.*

Alternatively, the information about the actual output realization might get lost and the wage payment has to be based on the signal realization. It is then assumed that the agent receives the wage he "deserves" given the available information, that is $\hat{w}(s \mid H, e^*) =$

[13] However, one can show that if $E(\theta \mid q, e)$ is independent of e (say because there is no residual uncertainty about θ given q), then there is no benefit from mediated mechanisms.

$E[\hat{w}(q)\,|\,s, H, e^*]$. The wage contract is said to be *affected by the disclosure policy* since the information structure has a direct effect on the distribution of payments received from it.

Proposition 5.1 *Suppose that all assumptions of Proposition 3.2 are satisfied and that in addition to career concerns, the agent is rewarded through a wage contract.*

- *If the wage contract is unaffected by the disclosure policy, effort is maximized by full disclosure.*
- *If the wage contract is affected by the disclosure policy and $\hat{w}(q)$ is nondecreasing, effort is maximized by full disclosure.*

If the explicit wage schedule is unaffected by the information structure, then the incentives due to it are independent of the garbling. This is not true in the second case, and a nondecreasing explicit wage schedule boosts incentives in addition to those from career concerns. Noisy information about performance depresses both types of incentives.

Instead of taking the wage contract $\hat{w} : Q \to \mathbb{R}$ as given, suppose the principal designs $\hat{w}$ and H to maximize profits subject to a limited liability constraint $\hat{w} \geq 0$. Depending on whether output information gets lost under a garbling, $\hat{w}$ might map either realized output or signal into a wage. For either case, Proposition 5.2 gives conditions when full revelation will be part of the principal's solution, under the following condition. Say there are *decreasing returns to signal jamming* if given any $\hat{e}$, $E_q[E(\theta\,|\,q, \hat{e})\,|\,e] - c(e)$ is concave in e for $e \geq \hat{e}$.

Proposition 5.2 *Assume that $E = [\underline{e}, \overline{e}]$, $f(q_M|e)$ is concave, there are decreasing returns to signal jamming, and all conditions of Proposition 3.2 are satisfied. If $(H, \hat{w})$ implements effort $\tilde{e}$, then under full revelation of output there exists an alternative wage contract $\hat{w}'$ that implements $\tilde{e}' \geq \tilde{e}$ and leads to a lower expected wage payment, that is $E[\hat{w}'\,|\,H_{FR}, \tilde{e}'] \leq E[\hat{w}\,|\,H, \tilde{e}]$.*

Career concerns deliver "free incentives" for the principal, and providing incentives through a wage contract is costly at the margin. Fully revealing output and only rewarding the highest output (which is optimal due to the MLRP) generates more effort (and hence output) at a lower cost.

5.2 Informed agent

Up to now, the agent did not observe the ability estimate. This is appealing if s_θ represents a test whose outcome is not known when taking effort. On the other hand, some signals of ability such as demographic information are known to the agent, but perhaps not to the

$\Pr(\theta = 1)$	$\Pr(q_H \mid \theta = 0, e_H)$	$\Pr(q_H \mid \theta = 1, e_H)$	ρ	$c(e_H) - c(e_L)$
0.4	0.7	0.8	0.5	0.092

Table 4: Parameters for Example 5.3

market in that it can commit to ignore such information. This section studies the effect of revealing s_θ to the market if the agent exerts effort knowing its realization.[14] The following example shows that Proposition 3.9 can be overturned.

Example 5.3 *Let $\Theta = \{0, 1\}$, $E = \{e_L, e_H\}$, and $Q = \{q_L, q_H\}$. Ability is unaffected by effort. If the agent exerts e_L, q_L is generated for sure.*[15] *If e_H is exerted, then q_H is generated according to $Pr(q_H \mid \theta, e_H)$ which is increasing in θ. The signal s_θ takes two values $\{s_L, s_H\}$, and $Pr(s_L|\theta = 0) = \rho$, $Pr(s_H|\theta = 0) = 1 - \rho$, and $Pr(s_H \mid \theta = 1) = 1$. Under the parameter values in Table 4, the following statements can be verified.*

Uninformed agent: *The conditions of Proposition 3.9 are satisfied. When s_θ is not disclosed, e_H can be sustained in equilibrium, but not when s_θ is disclosed to the market.*

Informed agent: *When s_θ is not disclosed, $e(s_L) = e(s_H) = e_L$ is the unique pure strategy equilibrium. When s_θ is disclosed to the market, there exists an equilibrium with $e(s_L) = e_L$ and $e(s_H) = e_H$.*

Under the conditions of Proposition 3.9, s_θ is a favorable signal of ability conditional on (q, e). But as the agent's strategy now depends on s_θ, it is possible that conditional on q a higher realization of s_θ leads to an adverse inference about ability. The following definition formalizes this.

Definition 5.4 *In an equilibrium with effort strategy $e^*(\cdot)$, expected ability is nondecreasing (nonincreasing) in the ability estimate if $\forall q$,*

$$E(\theta \mid q, s_\theta, e^*(s_\theta))$$

is nondecreasing (nonincreasing) in s_θ.

Whether expected ability is nondecreasing or nonincreasing in equilibrium has a crucial impact on whether disclosing the ability estimate to the market leads to more or less incentives. Define $E(\theta \mid q, e(\cdot))$ as the expected ability if q is realized, s_θ is not disclosed, and the agent's strategy is conjectured to be $e(\cdot)$. Similarly, $H(s_\theta \mid q, e(\cdot))$ is the distribution of s_θ when the agent's strategy is $e(\cdot)$ and q is realized.

[14] This requires a slight change in timing in that s_θ has to be realized before e is exerted, so s_θ needs to be exogenous.

[15] The violation of the full support assumption on $F(q \mid e)$ is not critical for the conclusion.

Proposition 5.5 *Suppose that the agent observes ability estimate s_θ. Assume that*

O-1 *$E(\theta \mid q, s_\theta, e)$ is supermodular in (q, e) and*
$H(s_\theta \mid q, e)$ is decreasing in e according to FOSD,

O-2 *$E(\theta \mid q, e(\cdot))$ is supermodular in $(q, e(\cdot))$*[16] *and*
$H(s_\theta \mid q, e(\cdot))$ is decreasing in $e(\cdot)$ according to FOSD,[17]

M *$F(q \mid s_\theta, e)$ is increasing in e according to FOSD,*

S *$E(\theta \mid q, s_\theta, e)$ is nondecreasing in s_θ,*

When s_θ is disclosed to the market and the equilibrium is such that expected ability is nondecreasing in the ability estimate, then under assumptions $\mathbf{O-2, M, S}$ there exists an equilibrium with higher effort (pointwise) when s_θ is not disclosed.

When s_θ is not disclosed to the market and the equilibrium is such that expected ability is nonincreasing in the ability estimate, then under assumptions $\mathbf{O-1, M, S}$ there exists an equilibrium with higher effort (pointwise) when s_θ is disclosed.

Remember from Proposition 3.9 that when the agent does not observe the ability estimate, disclosing it to the market is harmful if s_θ is favorable information about ability, and effort causes an adverse inference about s_θ conditional on q. With an informed agent, these ordering conditions are still satisfied when expected ability is nondecreasing in the ability estimate. If on the other hand s_θ becomes bad news through the agent's strategy, yet effort still causes an adverse inference about s_θ conditional on q, it is better to not disclose it.

Assumptions $\mathbf{O-1}$ and $\mathbf{O-2}$ in Proposition 5.5 are notions in which effort makes output more informative about ability, and according to which effort is bad news about ability conditional on output. In Example 5.3, $E(\theta \mid q, s_\theta, e)$ is supermodular in (q, e), but $E(\theta \mid q, e(\cdot))$ is not supermodular in $(q, e(\cdot))$.

Example 5.3, continued *When the agent is informed, then given market conjecture $e(s_L) = e_L$, $e(s_H) = e_H$ he wants to deviate to playing e_H after either signal. But due to a failure of supermodularity in the sense of Assumption $\mathbf{O-2}$, he wants to deviate downwards given market conjecture $e(s_L) = e(s_H) = e_H$.*

In Example 5.6 below, expected ability is nonincreasing in the ability estimate once taking the equilibrium strategy into account, and disclosure of s_θ leads to more effort.

[16] $E(\theta \mid q, e(\cdot))$ is supermodular in $(q, e(\cdot))$ if $\forall q_2 > q_1$, and $\forall e_2(\cdot) \geq e_1(\cdot)$, it is the case that $E(\theta \mid q_2, e_2(\cdot)) - E(\theta \mid q_1, e_2(\cdot)) \geq E(\theta \mid q_2, e_1(\cdot)) - E(\theta \mid q_1, e_1(\cdot))$.

[17] $H(s_\theta \mid q, e(\cdot))$ is decreasing in $e(\cdot)$ according to FOSD if for $\forall e_2(\cdot) \geq e_1(\cdot)$, $H(s_\theta \mid q, e_1(\cdot))$ dominates $H(s_\theta \mid q, e_2(\cdot))$ according to FOSD.

Example 5.6 *Let* $\Theta = \{0,1\}$*,* $E = \{e_L, e_H\}$*, and* $Q = \{q_L, q_H\}$*. Ability is uniformly distributed and unaffected by effort. If the agent is of low ability* $\theta = 0$ *or exerts* e_L*, output* q_L *is generated for sure.*[18] *If the agent is of high ability* $\theta = 1$ *and* e_H *is exerted, then output* q_H *is generated for sure. The ability estimate* s_θ *takes two values* $\{s_L, s_H\}$*, and* $Pr(s_L \,|\, \theta = 0) = Pr(s_H \,|\, \theta = 1) = \rho$*. For given* $c(e_H) > c(e_L)$*,* $e(s_L) = e_L$ *in any equilibrium if* ρ *is sufficiently large. The question is whether* $e(s_H) = e_H$ *can be sustained in equilibrium.*

s_θ not disclosed: *Given market conjecture* $e(s_L) = e_L, e(s_H) = e_H$*, the market belief is*

$$Pr(\theta = 1 \,|\, q_H, e(\cdot)) = 1, \quad Pr(\theta = 1 \,|\, q_L, e(\cdot)) = 1 - \rho.$$

Type s_H *exerts* e_H *if*

$$1 - c(e_H) \geq 1 - \rho - c(e_L).$$

s_θ disclosed: *When the market observes* s_H*, it believes that* e_H *was exerted. Its belief is*

$$Pr(\theta = 1 \,|\, q_H, s_H, e_H) = 1, \quad Pr(\theta = 1 \,|\, q_L, s_H, e_H) = 0.$$

Type s_H *exerts* e_H *if*

$$1 - c(e_H) \geq -c(e_L).$$

So if $1 \geq c(e_H) - c(e_L) > \rho$*, the best pure strategy equilibrium effort is higher when* s_θ *is disclosed to the market.*

If the agent's strategy is $e(s_L) = e_L, e(s_H) = e_H$, then *conditional on* q_L, signal s_H becomes bad news about ability. This explains why disclosure of s_θ leads to more effort.

5.3 Other extensions

In this section I examine whether various results apply if the principal is forced to disclose the ability estimate, output is observed with noise, the agent is not risk neutral with respect to the market belief, and the full support assumption on output is violated.

5.3.1 Exogenous information

In Section 3.2, the principal can garble q, while s_θ is not disclosed. When the market observes s_θ, the following result is analogous to Proposition 3.2.

Proposition 5.7 *Assume that*

[18] The violation of the full support assumption on $F(q \,|\, e)$ is not critical for the conclusion.

- *$F(q \mid s_\theta, e)$ is ordered according to the MLRP,*
- *$E(\theta \mid q, s_\theta, e)$ is nondecreasing in q,*
- *$E(\theta \mid q, s_\theta, e)$ is supermodular in (q, e).*

When s_θ is disclosed to the market, effort is maximized by full revelation of output.

In fact, full revelation is optimal even when the principal can make the garbling of q dependent on the realization of s_θ. Proposition 5.7 is not a corollary of Proposition 3.2, as in order to apply it one would need to find a complete and transitive order on the set of possible realizations of (q, s_θ) that satisfies all ordering assumptions. There are circumstances where this is not possible, yet Proposition 5.7 applies.

It is worth emphasizing that the principal cannot control the signal s_θ. If she can commit to garble (q, s_θ) directly, it might be optimal to do so even though it is not optimal to garble q when s_θ has to be disclosed. As an example, suppose $s_\theta = \theta$, so ability gets revealed. If the marginal of ability is independent of effort, equilibrium effort is zero. But if the principal can garble (q, s_θ) directly, one feasible signal is to disclose whether the sum of output and ability is above a threshold. One can find examples where such a signal sustains a nontrivial amount of effort.

The idea is that s_θ is a favorable signal about ability, but unaffected by e. On the other hand, e affects q, yet it provides no information about θ conditional on s_θ. By disclosing an appropriately chosen statistic of (q, s_θ), one creates a favorable signal of ability that the agent can manipulate.

5.3.2 Imperfectly observed output

Output is a baseline signal that represents the maximal available amount of information regarding performance. If only a noisy signal of output, $\tilde{q}$, is available, one can verify whether it satisfies the conditions of Proposition 3.2.

When Q has three or more values, there always exists a garbling of q such that the conditions of Proposition 3.2 are not satisfied with respect to such a signal. In particular, making the garbling less informative can boost incentives. So under the conditions of Proposition 3.2, it is *not* the case that a more informative garbling of output induces higher effort. Rather, fully revealing output maximizes incentives across all garblings (see Figure 7b for a schematic illustration).

5.3.3 General risk attitudes

Propositions 3.2 and 3.9 rely critically on the agent being risk neutral with respect to the market belief. E.g. to illustrate why Proposition 3.2 fails, suppose the agent maximizes the expectation of $u_A(w, e) = \mathbb{1}\{w \geq w_t\} - c(e)$. The agent is only concerned whether the belief is above a threshold w_t, say because the reward is a promotion and hence indivisible. Once Q has three or more elements, full revelation fails to be optimal in general. Suppose there is an output $\tilde{q}$ such that effort increases (decreases) the probability of outputs $q > \tilde{q}$ ($q < \tilde{q}$). Under full revelation of q, there will be a threshold level q^* such that the agent gets promoted only if output is above q^*. If say q^* is above $\tilde{q}$, the principal can typically induce more effort by recommending a promotion for output levels above q^*, but also for some adjacent output levels below q^*. This is incentive compatible as the market is unaware of the output realization that induced the recommendation.

More generally, a garbling of output decreases the dispersion of the marginal distribution of wages according to second order stochastic dominance (SOSD). While the marginal distribution over wages does not pin down the transfer schedule and hence incentives, one can show that under risk neutrality even the "steepest" transfer schedule that can be derived from the new marginal distribution over wages depresses incentives, in the sense of Lemma 3.6. The idea is that when wages are less dispersed, there is less leeway to reward good over bad performance.

With general risk attitudes, transfer schedules should be defined to specify the expected utility the agent receives for each output. But a garbling of output does not cause an SOSD shift in the distribution of utilities. For example, if the agent is risk averse then pooling information insures the agent. By selectively pooling high output realizations, the principal might raise the reward associated with high output through the insurance effect, thereby generating stronger incentives.

5.3.4 Support of q depends on e

In this part only, the support of output is allowed to depend on effort. Let $Q(e) \subset Q$ be the support given e. Relaxing the full support assumption raises one conceptual issue. Up to now, in any equilibrium beliefs were pinned down by Bayes' rule, and the agent could never induce an off-path signal by deviating. If on the other hand some combinations of effort and output are technologically not possible, then for certain information structures equilibrium analysis will require specifying off-path beliefs. This is relevant since off-path beliefs determine on-path behavior.

The equilibrium definition now includes a sequence of strategies $\beta = \{\beta_n\}$, where $\forall n$, $\beta_n \in$

e	$f_n(\boldsymbol{q}\,|\,e)$	$E_n(\theta\,|\,\boldsymbol{q},e)$	$c_n(e)$
e_L	(1-2ϵ_n , ϵ_n , ϵ_n)	(0 , 0 , 0)	0
e_M	(ϵ_n , 2/3 , 1/3-ϵ_n)	(0 , 1 , 2)	1
e_H	(ϵ_n , 1/3 , 2/3-ϵ_n)	(-1 , 0 , 1)	1.1

Table 5: Full support production technology in Example 5.8

e	$f(\boldsymbol{q}\,|\,e)$	$E(\theta\,|\,\boldsymbol{q},e)$	$c(e)$
e_L	(1 , 0 , 0)	(0 , . , .)	0
e_M	(0 , 2/3 , 1/3)	(. , 1 , 2)	1
e_H	(0 , 1/3 , 2/3)	(. , 0 , 1)	1.1

Table 6: Limit production technology in Example 5.8

$\Delta(E)$ has full support on Q, and $\lim_{n\to\infty}\beta_n = e^*$. Formally, a *consistent* transfer schedule $\tilde{t}$ also has the sequence β as an argument, and is written as $\tilde{t}(\cdot\,|\,H,\beta,e^*)$ with $\lim_{n\to\infty}\beta_n = e^*$. In the spirit of sequential equilibrium, I require that $\forall q$, $\tilde{t}(q\,|\,H,\beta,e^*) = \lim_{n\to\infty} t(q\,|\,H,\beta_n)$, where $t(q\,|\,H,\beta_n)$ can be derived from Bayes' rule since β_n has full support.

An output q is said to be fully revealed if $\tilde{t}(q\,|\,H,\beta,e^*) = \lim_{n\to\infty} E(\theta\,|\,q,\beta_n)$. Perhaps one would not expect full revelation to be optimal now because by pooling outputs below $Q(e^*)$, deviations inducing such outputs can be punished more severely through off-path beliefs. On the other hand, one can still show that if one moves to full revelation of on-path outputs, the agent wants to deviate upwards in the partial equilibrium sense of Lemma 3.6. Also, it is true that for any approximation $F_n(q\,|\,e), E_n(\theta\,|\,q,e)$ of the statistical environment $F(q\,|\,e), E(\theta\,|\,q,e)$ that satisfies full support and all assumptions of Proposition 3.2, full revelation of output maximizes effort. Why this is not necessarily true in the limit is illustrated by Example 5.8.

Example 5.8 *Let $E = \{e_L, e_M, e_H\}$, $Q = \{q_L, q_M, q_H\}$. The primitives are specified in Tables 5 and 6, where ϵ_n is a sequence of positive numbers converging to zero. First notice that $\forall n$ large enough, $F_n(q\,|\,e), E_n(\theta\,|\,q,e)$ satisfy the assumptions of Proposition 3.2, so full revelation is optimal. The important feature is that in the limit statistical environment $F(q\,|\,e), E(\theta\,|\,q,e)$, no signal can ever generate a negative wage under the sequential equilibrium-type refinement since $E(\theta\,|\,q,e) \geq 0$, $\forall(q,e)$ with $q \in Q(e)$. In particular, for any possible conjecture e and associated sequence β converging to e, $\tilde{t}(q_L\,|\,H,\beta,e) \geq 0 > -1 = E_n(\theta\,|\,q_L,e_H)$ $\forall(H,\beta,e)$. This limits the extent to which output q_L can be punished, and one can show that this prevents e_H from being implemented. Now even considering information structures that at least reveal on-path outputs, these are not optimal by the following argument. Any information structure that fully reveals on-path outputs conditional on e_M*

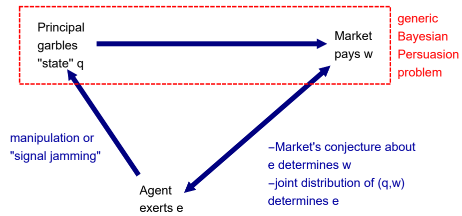

Figure 6: Relationship with Bayesian Persuasion

leads to a deviation to either e_L *or* e_H*, for any sequence of fully mixed strategies approximating* e_M*. Similarly, the agent wants to deviate to* e_L *if* e_H *is conjectured by the market. The information structure that pools all outputs leads to* e_L *and hence fully reveals the on-path output* q_L*. This is the highest effort that can be induced by information structures that fully reveal on-path outputs. However* e_M *can be sustained by an information structure that pools* q_M *and* q_H *and reveals* q_L*.*

The example makes two points. First, that the set of equilibria that satisfy the sequential equilibrium-type refinement is not upper hemicontinuous in the primitives $F(q \mid e), E(\theta \mid q, e)$. If beliefs are not constrained by the refinement, one can sustain e_H as an equilibrium with a fully revealing info structure that assigns a sufficiently low wage after q_L. While strictly speaking this constitutes a perfect Bayesian equilibrium, it is an unreasonable one. The second point is that even information structures that only reveal on-path scores are in general not optimal.

6 Discussion

This section attempts to put some results of this paper into perspective.

On the cumulative subsidy condition. Consider a joint distribution of three real valued random variables (θ, q, s), where θ and s are independent conditional on q. The cumulative subsidy condition in Lemma 3.4 states a relationship between $E_s[E(\theta \mid s) \mid q]$, $E(\theta \mid q)$, and the marginal distribution of q. This should not be confused with the result that the marginal distribution of $E(\theta \mid s)$ is a mean-preserving contraction of the marginal distribution of $E(\theta \mid q)$ (see Blackwell (1953)). This result, applied by Gentzkow and Kamenica (2016) to the Bayesian Persuasion literature, is not sufficient in the present context.

Fixing the agent's effort, any *marginal distribution* over wages that is a mean-preserving contraction relative to that generated by full revelation can be induced by some information structure. However, the market behavior cannot be treated in isolation. Two different information structures can induce the same marginal distribution over wages, yet still lead to different transfer schedules and therefore different incentives for the agent. Figure 6 highlights the difference with the literature on Bayesian Persuasion focusing on how a garbling affects market behavior. Here, the agent can manipulate the distribution of the "state" q, and the behaviors of market and agent are intertwined. This also explains why direct recommendation mechanisms (in terms of the market wage) are not very useful and were eschewed - whether a direct recommendation mechanism is incentive compatible or not depends on the conjecture held by the market.

Comparison with standard principal-agent problems. Information design as considered here is irrelevant in a standard principal-agent model where the agent is motivated by wages only. Disregarding available information about performance amounts to constraining the set of feasible contracts, and more accurate information about the agent's ability (keeping the agent's information fixed) is beneficial for a profit-maximizing principal.[19]

In the standard model, let $T \subset \mathbb{R}^Q$ be the set of feasible transfer schedules. If the principal wishes to implement effort e^*, she has to select a transfer schedule $t(\cdot) \in T$ such that $e^* \in \text{argmax } E_q[u_A(t(q), e) \,|\, e]$. Here, let $T(\hat{e})$ be the set of transfer schedules implementable by some H, given conjecture $\hat{e}$. For e^* to be an equilibrium, it must be that $e^* \in \text{argmax } E_q[u_A(t(q \,|\, H, e^*), e) \,|\, e]$ for some $t(\cdot \,|\, H, e^*) \in T(e^*)$. The normal-linear example circumvents this fixed point nature, because effectively $T(\hat{e})$ is independent of $\hat{e}$. Transfer schedules are linear and can be identified with their slope, which is independent of $\hat{e}$.

Lemma 3.4 shows that there exists an "extremal" transfer schedule with respect to the cumulative subsidy condition within $T(\hat{e})$ that maximizes incentives. The supermodularity condition on $E(\theta \,|\, q, e)$ guarantees that $T(\hat{e})$ expands as one increases $\hat{e}$.[20] The cumulative subsidy condition is a binding constraint in that higher incentives can be generated if it is violated, and full revelation of output induces the extremal transfer schedule.

Partial vs full equilibrium. Dewatripont, Jewitt and Tirole (1999*a*) show in their Proposition 5.2 that for a given market conjecture, inducing noise locally decreases incentives. Lemma 3.6 can be seen as a non-local analogue in that *any* effort below $\hat{e}$ is shown to become less attractive.[21] Since their result holds for any market conjecture, they interpret it as a local comparative static around equilibrium. Example A.7 in Appendix A.1 shows that under their assumptions, effort might be maximized by an information structure that is not fully

[19] This is not true if better information is more costly, as is the case in Li and Yang (2016).

[20] Given Lemma 3.5, this is strictly speaking only true for the subset of nondecreasing transfer schedules, yet one can restrict attention to these as far as maximal incentives for effort are concerned.

[21] Besides the MLRP, Dewatripont, Jewitt and Tirole (1999*a*) assume that θ and q are affiliated conditional on e, which is stronger than $E(\theta \,|\, q, e)$ being nondecreasing in q.

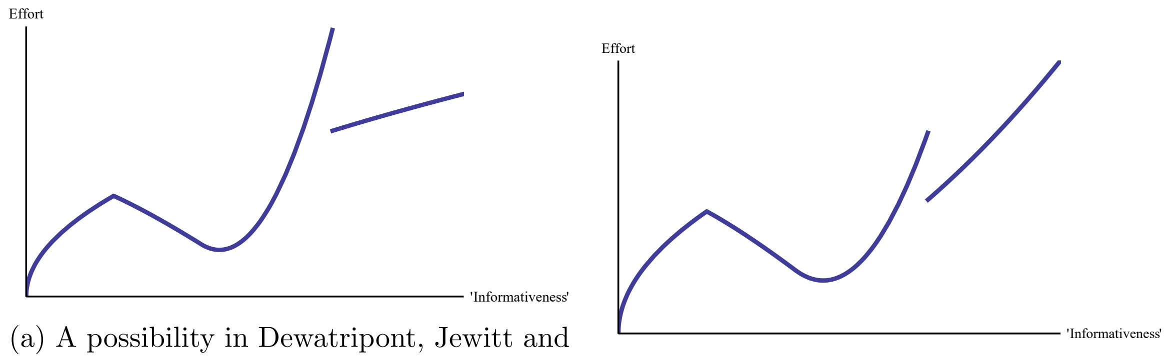

(a) A possibility in Dewatripont, Jewitt and Tirole (1999*a*)

(b) Result of Proposition 3.2

Figure 7: Comparison with Dewatripont, Jewitt and Tirole (1999*a*)

revealing. This does not occur under the conditions of Proposition 3.2, although it is possible that increasing informativeness reduces incentives when starting from a noisy information structure. Also, in general the equilibrium correspondence is not lower hemicontinuous in the information structure, explaining why the *maximal* effort can be discontinuous. Figure 7 illustrates the difference in results schematically.

Mediated equilibria. Regarding mediated information structures, the way the principal correlates the wage distribution with the effort recommendation is reminiscent of the mediated contracts found in Rahman and Obara (2010). At an abstract level, in both papers the principal commits to a randomization over mechanisms. The randomization is publicly observed, and the principal discloses differential information about the realized mechanism. In Rahman and Obara (2010), this makes it easier to identify the identity of a deviator. Here, it allows the principal to violate the cumulative subsidy condition.

In a conventional mechanism with associated mixed strategy equilibrium σ, signal s provides no information about ability *conditional on* q, that is, $E(\theta \,|\, q, s, \sigma) = E(\theta \,|\, q, \sigma)$. The reason is that q is a sufficient statistic for (q, s) when estimating the realization from σ. In a mediated mechanism this is not true because conditional on q, e is correlated with s. Unless $E(\theta \,|\, q, e)$ is constant in e, q is not a sufficient statistic for (q, s) when estimating θ. Indeed, one can show that if $E(\theta \,|\, q, e)$ is constant in e, there is no benefit from mediated mechanisms.

At first sight, mediated information structures might appear unrealistic because of the seemingly artificial implementation involving randomized strategies. At an abstract level, however, there is nothing implausible about the idea that within organizations, one can make recommendations that are opaque to an outsider.

Signal jamming and noisy signaling. Career concerns are also referred to as *signal jamming* incentives, in that the agent tries to interfere with the market's inference. This incentive is distinct from *signaling*, where an agent takes costly action to transmit information credibly (in a separating equilibrium). A typical assumption in career-concerns models is that the agent has no private information about ability, in order to abstract from signaling incentives. When in Section 5.2 the agent observes s_θ but not the market, this can be seen as a hybrid model that features both signal jamming and signaling incentives.[22]

Section 5.2 shows that in the hybrid model, the signaling effect can be detrimental for career-concerns incentives. In that case, the ability estimate should be disclosed to the market. Somewhat related to this is Chen (2015), where effort affects the riskiness of output. She fixes the information held by the market (no information) and studies the effect of giving more information to the agent about his ability. Both effort and output are observed, and the signaling incentives when the agent has private information lead to the riskier action being taken (which is more informative about ability).

Literature. A brief summary of the applied theory literature that features both career-concerns incentives and information design questions follows.

Ratchet effect. In a dynamic principal-agent relationship with uncertainty about the agent's ability, career concerns arise when the principal cannot commit to long-term contracts. Future contracts are influenced by current performance, and a forward-looking agent takes this into account which is described as the ratchet effect. A tractable formalization was introduced by Gibbons and Murphy (1992), where the relevant uncertainty is normally distributed, the agent has CARA preferences, and the principal offers linear contracts. This has been used to study the effects of relative performance comparisons across agents (Meyer and Vickers (1997)), optimal job design (Meyer (1995) and Meyer, Olsen and Torsvik (1996)), and optimal duration of employment (Auriol, Friebel and Pechlivanos (2002)). Each of the preceding examples can be posed as a question of optimal information disclosure about performance and ability in the presence of career concerns.

Fully dynamic models. Most career-concerns applications are static in that effort is exerted once. Hörner and Lambert (2016) consider information design within a dynamic version of the normal-linear model. They determine the optimal way to disclose past and present performance, and compare the amount of information revealed when the market has access to previously disclosed information and when it does not. Moav and Neeman (2010) argue that more informative performance measures, while initially conducive to inducing output, remove uncertainty about ability in the future. This is detrimental for incentives when ability is persistent.

[22] Compared to the signaling model of Spence (1973), there is only a noisy signal of effort. The fact that output is informative about ability *conditional on effort* leads to career-concerns incentives that are absent in signaling models. An example of this can be found in de Haan, Offerman and Sloof (2011).

Other work. In Bar-Isaac (2007), an agent can acquire an asset with unknown productivity. This provides him with an incentive device to exert effort if the market is not able to ascertain whether output was produced because of effort or because of the high productivity of the asset. To maintain these incentives, the agent does not wish to reveal information about the source of high output. In Jeon (1996), a principal has to assign workers to teams. Uncertainty about ability differs across agents, and team composition affects the informativeness of output regarding the ability of the team members. Dewatripont, Jewitt and Tirole (1999*b*) study multi-tasking with career-concerns, in particular the effect of transparency about the agent's focus on incentives. Bar-Isaac and Ganuza (2008) consider how different recruitment and training policies affect the career concerns of agents. Recruitment policies affect the type uncertainty, and training policies affect the informativeness of output, hence these are questions of information design. The effect of information on career-concerns incentives has also been studied in relational contract settings (Mukherjee, 2008*b*, 2010) and in matching markets (Mukherjee (2008*a*)). The setting in Wolitzky (2012) is similar to that in Section 5.1, in particular Proposition 5.2, except that the principal's wage contract and disclosure policy are unobserved by the market. In Rodina and Farragut (2016), an agent wishes to create the impression that he has produced high output, which can be seen as a special case of the current model by defining ability to equate to output. The transparency result of Proposition 3.2 does not apply in this setting, and some qualitative features of the optimal garbling are derived.

7 Conclusion

This paper derives substantive conditions under which full revelation of performance measures and more uncertainty about the agent's ability lead to stronger career-concerns incentives. A key assumption for deriving these comparative statics is that effort increases the informativeness of output regarding ability. The results provide a foundation for two principles regarding the evaluation of agents within organizations: transparency with respect to measures of performance, and opaqueness regarding certain information about the agent's ability. The following three issues have not been addressed, however.

First, when the conditions for full revelation of performance measures are violated, it is natural to ask what features the optimal garbling has. The same can be said for mediated mechanisms, because they are not fully revealing when improving upon conventional disclosure policies.

Second, comparative statics with respect to information about performance or ability are considered separately. Certain instruments such as relative performance evaluation across agents might affect both types of information simultaneously. Plausibly, a tradeoff arises between the forces identified here. Conceptually more interesting is the case where the principal can control information about performance and ability *simultaneously*, that is, she can garble (q, s_θ).

Third, it would be desirable to understand the effect of information beyond the (quasi-) static model considered here. One complication in a (fully) dynamic model is the presence of asymmetric information between the principal and the agent, at least off the equilibrium path (which one has to solve in order to determine on-path behavior).[23] More directly related to information design is that whatever information is released today also provides information in future periods when the agent's ability is persistent. A natural conjecture is that it remains optimal to withhold information about ability when the signal structure in the stage game satisfies appropriate conditions as derived here, yet full revelation of performance measures requires stronger assumptions.

References

Auriol, Emmanuelle, Guido Friebel, and Lambros Pechlivanos. 2002. "Career Concerns in Teams." *Journal of Labor Economics*, 20(2): 289–307.

Bar-Isaac, Heski. 2007. "Something to Prove: Reputation in Teams." *The RAND Journal of Economics*, 38(2): 495–511.

Bar-Isaac, Heski, and Juan-José Ganuza. 2008. "Recruitment, Training, and Career Concerns." *Journal of Economics & Management Strategy*, 17(4): 839–864.

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics*, 24(2): 265–272.

Chen, Ying. 2015. "Career Concerns and Excessive Risk Taking." *Journal of Economics & Management Strategy*, 24(1): 110–130.

Cisternas, Gonzalo. 2016. "Two-Sided Learning and Moral Hazard." Working paper.

de Haan, Thomas, Theo Offerman, and Randolph Sloof. 2011. "Noisy Signaling: Theory and Experiment." *Games and Economic Behavior*, 73(2): 402–428.

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999*a*. "The Economics of Career Concerns, Part I: Comparing Information Structures." *The Review of Economic Studies*, 66(1): 183–198.

[23] Some progress on this has been made by Cisternas (2016).

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999*b*. "The Economics of Career Concerns, Part II: Application to Missions and Accountability of Government Agencies." *The Review of Economic Studies*, 66(1): 199–217.

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

Gibbons, Robert, and Kevin J Murphy. 1992. "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence." *The Journal of Political Economy*, 100(3): 468–505.

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *The Review of Economic Studies*, 66(1): 169–182.

Hörner, Johannes, and Nicolas S Lambert. 2016. "Motivational Ratings." Working paper.

Jeon, Seonghoon. 1996. "Moral Hazard and Reputational Concerns in Teams: Implications for Organizational Choice." *International Journal of Industrial Organization*, 14(3): 297–315.

Li, Anqi, and Ming Yang. 2016. "Optimal Incentive Contract with Costly and Flexible Monitoring." Working paper.

Meyer, Margaret A. 1995. "Cooperation and Competition in Organizations: A Dynamic Perspective." *European Economic Review*, 39(3): 709–722.

Meyer, Margaret A, and John Vickers. 1997. "Performance Comparisons and Dynamic Incentives." *Journal of Political Economy*, 105(3): 547–581.

Meyer, Margaret A, Trond E Olsen, and Gaute Torsvik. 1996. "Limited Intertemporal Commitment and Job Design." *Journal of Economic Behavior & Organization*, 31(3): 401–417.

Moav, Omer, and Zvika Neeman. 2010. "The Quality of Information and Incentives for Effort." *The Journal of Industrial Economics*, 58(3): 642–660.

Mukherjee, Arijit. 2008*a*. "Career Concerns, Matching, and Optimal Disclosure Policy." *International Economic Review*, 49(4): 1211–1250.

Mukherjee, Arijit. 2008*b*. "Sustaining Implicit Contracts when Agents have Career Concerns: The Role of Information Disclosure." *The RAND Journal of Economics*, 39(2): 469–490.

Mukherjee, Arijit. 2010. "The Optimal Disclosure Policy when Firms offer Implicit Contracts." *The RAND Journal of Economics*, 41(3): 549–573.

Rahman, David, and Ichiro Obara. 2010. "Mediated Partnerships." *Econometrica*, 78(1): 285–308.

Rodina, David, and John Farragut. 2016. "Inducing Effort through Grades." Working paper.

Spence, Michael. 1973. "Job Market Signaling." *The Quarterly Journal of Economics*, 355–374.

Tarski, Alfred. 1955. "A Lattice-Theoretical Fixpoint Theorem and its Applications." *Pacific Journal of Mathematics*, 5(2): 285–309.

Wolitzky, Alexander. 2012. "Career Concerns and Performance Reporting in Optimal Incentive Contracts." *The BE Journal of Theoretical Economics*, 12(1).

A Appendix

The following result will be used repeatedly.

Result A.1 (Tarski's fixed point theorem) *Let $(L, \geq)$ be a complete lattice. Suppose $f : L \to L$ is nondecreasing. Then the set of all fixed points of f is a complete lattice with respect to $\geq$ (Tarski (1955)). So the set of fixed points is nonempty and has a maximum.*

In particular, let $N \in \mathbb{N}$ and $\geq$ be the partial order on $\mathbb{R}^N$. For a given $\hat{e}(\cdot) \in E^N$, let $\hat{E} := \{e(\cdot) \in E^N : e(\cdot) \geq \hat{e}(\cdot)\}$, and suppose $\xi : \hat{E} \to \hat{E}$ is nondecreasing. Then $e^*(\cdot) = \xi(e^*(\cdot))$ for some $e^*(\cdot) \in \hat{E}$.

A.1 Proofs for Section 3

Proof of Lemma 3.4 Fix a possibly mixed conjecture $\hat{\sigma}$. In inequality j in Equation (1),

$$F(q_j \,|\, \hat{\sigma}) E_q[E_s[E_\theta[\theta \,|\, s, \hat{\sigma}] \,|\, q] \,|\, q \leq q_j, \hat{\sigma}]$$

is compared with

$$F(q_j \mid \hat{\sigma}) E_q[E_\theta[\theta \mid q, \hat{\sigma}] \mid q \leq q_j, \hat{\sigma}].$$

Notice that $E_\theta[\theta \mid s, \hat{\sigma}] = E_q[E_\theta[\theta \mid q, s, \hat{\sigma}] \mid s, \hat{\sigma}] = E_q[E_\theta[\theta \mid q, \hat{\sigma}] \mid s, \hat{\sigma}]$, where the first equality follows from the law of iterated expectations, and the second from s being a garbling of q. So one can write

$$\begin{aligned}
&E_q[E_s[E_\theta[\theta \mid s, \hat{\sigma}] \mid q] \mid q \leq q_j, \hat{\sigma}] \\
&= E_q[E_s[E_q[E_\theta[\theta \mid q, \hat{\sigma}] \mid s, \hat{\sigma}] \mid q] \mid q \leq q_j, \hat{\sigma}] \\
&= E_s[E_q[E_\theta[\theta \mid q, \hat{\sigma}] \mid s, \hat{\sigma}] \mid q \leq q_j, \hat{\sigma}].
\end{aligned}$$

On the other hand,

$$\begin{aligned}
&E_q[E_\theta[\theta \mid q, \hat{\sigma}] \mid q \leq q_j, \hat{\sigma}] \\
&= E_s[E_q[E_\theta[\theta \mid q, \hat{\sigma}] \mid s, q \leq q_j, \hat{\sigma}] \mid q \leq q_j, \hat{\sigma}].
\end{aligned}$$

Since $E_\theta[\theta \mid q, \hat{\sigma}]$ is nondecreasing in q,

$$E_q[E_\theta[\theta \mid q, \hat{\sigma}] \mid s, \hat{\sigma}] \geq E_q[E_\theta[\theta \mid q, \hat{\sigma}] \mid s, q \leq q_j, \hat{\sigma}],$$

with equality for $j = M$. ∎

Proof of Lemma 3.5 Fix a possibly mixed conjecture $\hat{\sigma}$. Since $\hat{\sigma}$ will be kept fixed throughout, the following shorthands are introduced: $f_i = f(q_i \mid \hat{\sigma}), t_i = t(q_i \mid H_{FR}, \hat{\sigma})$, and $g_i = g(q_i)$. Define $\Delta_i = g_i - t_i$, and let $J = \{j : \Delta_j > 0\}$ be the set of all output realizations after which the expected wage under the transfer schedule $g(\cdot)$ is higher than under the fully informative signal. If J is empty, the supposition of the lemma implies that $\forall i$, $g(q_i) = t(q_i \mid H_{FR}, \hat{\sigma})$ so that the fully revealing information structure implements $g(\cdot)$.

Step 1: A building block in the construction of the information structure that implements $g(\cdot)$ will be a collection of vectors satisfying certain properties summarized as follows. There exists a collection of vectors $\{\alpha^j\}_{j \in J}$, where each α^j is an $M \times 1$ vector, and satisfies the following properties:

- α^j is nonnegative,
- $\alpha_i^j = 0$ for $i < j$ and $\alpha_j^j = 1$,
- if $i > j$ and $\alpha_i^j > 0$, then $\Delta_i \leq 0$,
- $\sum_{i=1}^M \alpha_i^j f_i \Delta_i = 0$,

- the collection of α^j's satisfies $\sum_{j\in J}\alpha^j = e_M$.[24]

Since each α^j is nonnegative, the last property implies that $\forall i,j,\ \alpha_i^j \in [0,1]$. Indeed, the elements of α^j will correspond to certain probabilities in the information structure.

The following algorithm is used to show that a collection of vectors $\{\alpha^j\}_{j\in J}$ with the desired properties exists. The algorithm consists of $|J|$ steps and its starting point is an $M\times 1$ vector of "budgets" $b^1 = e_M$ that at each step k gets updated from b^k to b^{k+1}. By construction, b^k will be nonnegative at each step k. The algorithm moves from the highest to the lowest element in J, so let $j(k)$ be the k-th highest element in J, that is $j(1) > ... > j(|J|)$. Let $\xi(b^k, j(k))$ be the largest element such that

$$\sum_{i=j(k)}^{\gamma} b_i^k f_i \Delta_i \geq 0,\ \forall\gamma = j(k), ..., \xi(b^k, j(k)),$$

which exists since $\Delta_{j(k)} > 0$. So either $\xi(b^k, j(k)) = M$, or $\sum_{i=j(k)}^{\xi(b^k,j(k))+1}\Delta_i f_i b_i^k < 0$.

Step k in the algorithm starts with the nonnegative vector of budgets b^k, and uses it to define the vector $\alpha^{j(k)}$ as

$$\alpha_i^{j(k)} = \begin{cases} 0 & \text{for } i < j(k), \\ b_i^k & \text{for } j(k) \leq i \leq \xi(b^k, j(k)), \\ \frac{\sum_{h=j(k)}^{\xi(b^k,j(k))} b_h^k f_h \Delta_h}{-f_{\xi(b^k,j(k))+1}\Delta_{\xi(b^k,j(k))+1}} & \text{for } i = \xi(b^k, j(k)) + 1, \\ 0 & \text{for } i > \xi(b^k, j(k)) + 1. \end{cases}$$

If $\xi(b^k, j(k)) = M$, only the first two cases apply. At the end of step k, b^{k+1} is defined as $b^k - \alpha^{j(k)}$. For b^{k+1} to be nonnegative, it needs to be shown that $\alpha^{j(k)} \leq b^k$. This is obviously true for all elements except $i = \xi(b^k, j(k))+1$, which is relevant when $\xi(b^k, j(k)) < M$. Then by the definition of $\xi(b^k, j(k))$

$$\sum_{i=j(k)}^{\xi(b^k,j(k))+1} b_i^k f_i \Delta_i < 0,$$

and $f_{\xi(b^k,j(k))+1}\Delta_{\xi(b^k,j(k))+1} b^k_{\xi(b^k,j(k))+1} < 0$. This implies that $\alpha^{j(k)}_{\xi(b^k,j(k))+1} \leq b^k_{\xi(b^k,j(k))+1}$.

Now all the desired properties of the collection $\{\alpha^j\}_{j\in J}$ will be verified. At step k, $\alpha^{j(k)}$ gets defined and is obviously nonnegative for $i \neq \xi(b^k, j(k)) + 1$. For $i = \xi(b^k, j(k)) + 1$ (which is relevant when $\xi(b^k, j(k)) < M$), this follows from $\sum_{h=j(k)}^{\xi(b^k,j(k))} b_h^k f_h \Delta_h \geq 0$ and $f_{\xi(b^k,j(k))+1}\Delta_{\xi(b^k,j(k))+1} < 0$. By construction $\alpha_i^j = 0$ for $i < j$, and $\alpha_j^j = 1$ since the algorithm moves from the highest to the lowest element in J so that at each step k, $b^k_{j(k)} = 1 = \alpha^{j(k)}_{j(k)}$.

[24] e_M is an $M\times 1$ vector of 1's.

This implies that at step k_1 which defines $\alpha^{j(k_1)}$, $b^{k_1}_{j_2} = 0$ for $j_2 > j(k_1)$. Therefore if $i > j$ and $\alpha^j_i > 0$, then $\Delta_i \leq 0$. The next property is that $\sum_{i=1}^M \alpha^j_i f_i \Delta_i = 0$, which for any step k with $\xi(b^k, j(k)) < M$ is satisfied as can be seen from the construction of $\alpha^{j(k)}_{\xi(b^k,j(k))+1}$. When at some step k it is that $\xi(b^k, j(k)) = M$, a violation of $\sum_{i=1}^M \alpha^j_i f_i \Delta_i = 0$ would imply that $\sum_{i=j(k)}^M b^k_i f_i \Delta_i > 0$. But this leads to a contradiction as for the first $\tilde{k}$ that such a violation can be found it is that

$$\sum_{k=1}^{\tilde{k}} \sum_{i=j(k)}^{M} b^k_i f_i \Delta_i = \sum_{i=j(\tilde{k})}^{M} f_i \Delta_i > 0.$$

This would contradict the supposition of the lemma that

$$\sum_{i=1}^{j} g_i f_i \geq \sum_{i=1}^{j} t_i f_i \qquad \forall j = 1, ..., M, \text{ with equality for } j = M.$$

The final property is that $\sum_{j \in J} \alpha^j = e_M$. Given that $b^{|J|+1} = e_M - \sum_{j \in J} \alpha^j$, this is equivalent to showing that at the the end of the final step $k = |J|$, $b^{|J|+1} = 0$. A violation of this would imply that at the final step $k = |J|$, $\xi(b^k, j(k)) < M$ and $b^{|J|+1}_i > 0$ for (possibly several) i with $\Delta_i < 0$. On the one hand, it was previously shown that $\forall j$, $\sum_{i=1}^M \alpha^j_i f_i \Delta_i = 0$, so

$$\sum_{k=1}^{|J|} \left(\sum_{i=1}^{M} \alpha^{j(k)}_i f_i \Delta_i \right) = 0.$$

But this generates a contradiction, since using that $b^{|J|+1} = e_M - \sum_{j \in J} \alpha^j$,

$$\sum_{k=1}^{|J|} \sum_{i=1}^{M} \alpha^{j(k)}_i f_i \Delta_i = \sum_{i=1}^{M} \left(\sum_{k=1}^{|J|} \alpha^{j(k)}_i \right) f_i \Delta_i = \sum_{i=1}^{M} \left(1 - b^{|J|+1}_i \right) f_i \Delta_i = -\sum_{i=1}^{M} b^{|J|+1}_i f_i \Delta_i > 0.$$

Step 2: The information structure will be supported on $|J| \cdot M$ signals, so define the set of signals $\{s^j_i\}_{\{j \in J, i=1,...,M\}}$ consisting of arbitrary numbers that are all distinct.

For $j \in J$ and $i \notin J$, set $\Pr(s^j_i \,|\, q_i) = \alpha^j_i$, and $\Pr(s^j_i \,|\, q_j) = \beta^j_i$, where $\beta^j_i = 0$ if $\alpha^j_i = 0$, and if $\alpha^j_i > 0$ then β^j_i is given by

$$\beta^j_i = \frac{\alpha^j_i f_i (t_i - g_i)}{f_j (g_i - t_j)}. \tag{2}$$

$\beta^j_i \geq 0$ since $\alpha^j_i > 0$ implies both that $\Delta_i = g_i - t_i \leq 0$ and $i > j$, so that $g_i \geq g_j > t_j$.

For $j \in J$, set $\Pr(s_j^j \,|\, q_j) = 1 - \sum_{i \neq j} \beta_i^j$. This is nonnegative since for $i > j$, β_i^j satisfies

$$\beta_i^j f_j(g_i - t_j) + \alpha_i^j f_i \Delta_i = 0.$$

Summing over $i > j$ gives

$$\sum_{i>j} \beta_i^j f_j(g_i - t_j) + \sum_{i>j} \alpha_i^j f_i \Delta_i = 0.$$

By the construction from step 1, $\sum_{i>j} \alpha_i^j f_i \Delta_i = -f_j \Delta_j$. So

$$f_j \Delta_j = \sum_{i>j} \beta_i^j f_j(g_i - t_j) \geq \sum_{i>j} \beta_i^j f_j \Delta_j,$$

where the inequality follows from g being nondecreasing. Since $f_j \Delta_j > 0$, $\sum_{i>j} \beta_i^j = \sum_{i \neq j} \beta_i^j \leq 1$.

All probabilities are specified, except for two cases. The first is $\Pr(s_i^{j_1} \,|\, q_{j_2})$ where $j_1, j_2 \in J$, $j_1 \neq j_2$, and the second is $\Pr(s_{i_1}^j \,|\, q_{i_2})$ where $i_1, i_2 \notin J$, $i_1 \neq i_2$, and $j \in J$. All of these are set to zero. For this to constitute a valid information structure, it remains to show that every output is mapped with probability 1 into some set of signals. For $j \in J$, it was just shown that

$$\sum_{s \in S} \Pr(s \,|\, q_j) = \sum_{i=1}^{M} \Pr(s_i^j \,|\, q_j) = \Pr(s_j^j \,|\, q_j) + \sum_{i \neq j} \Pr(s_i^j \,|\, q_j) = 1 - \sum_{i \neq j} \beta_i^j + \sum_{i \neq j} \beta_i^j = 1.$$

For $i \notin J$,

$$\sum_{s \in S} \Pr(s \,|\, q_i) = \sum_{j \in J} \Pr(s_i^j \,|\, q_i) = \sum_{j \in J} \alpha_i^j = 1.$$

This completes the specification of H.

Step 3: It now remains to show that H implements the transfer schedule $g(\cdot)$. There are two types of signals $s \in S$. Either $s = s_i^j$ for $i \notin J$ and $j \in J$, or $s = s_j^j$ for $j \in J$. Some signal $s' \notin S$ might get never realized, more formally $\Pr(s' \,|\, q) = 0$, $\forall q \in Q$. Any such signal plays no role under the assumption that the marginal distribution over q has full support for all effort levels, as no deviation by the agent can induce such a signal realization.

Any signal of the form s_i^j for $i \notin J$ and $j \in J$ gets only induced by outputs q_i and q_j. The wage associated with such a signal is (suppressing notation for the information structure and effort conjecture $\hat{\sigma}$

$$w(s_i^j) = \frac{f_i \alpha_i^j t_i + f_j \beta_i^j t_j}{f_i \alpha_i^j + f_j \beta_i^j} = g_i,$$

as can be seen from the definition of β_i^j in Equation (2). Any remaining signal has to be of the form s_j^j for $j \in J$. Such a signal realization can only be induced by q_j, so output gets revealed and $w(s_j^j) = t_j$.

To show that transfer schedule $g(\cdot)$ is implemented, first outputs of the form $q = q_i$ with $i \notin J$ are considered. For any such output

$$E[w(s) \,|\, q_i] = \sum_{j \in J} \alpha_i^j w(s_i^j) = \sum_{j \in J} \alpha_i^j g_i = g_i.$$

For any output of the form $q = q_j$ with $j \in J$,

$$\begin{aligned} E[w(s) \,|\, q_j] = & \quad \textstyle\sum_{i \notin J} \beta_i^j w(s_i^j) + \left(1 - \sum_{i \notin J} \beta_i^j\right) w(s_j^j) \\ = & \quad \textstyle\sum_{i \notin J} \beta_i^j g_i + \left(1 - \sum_{i \notin J} \beta_i^j\right) t_j \\ = & \quad \textstyle t_j + \sum_{i \notin J} \beta_i^j \left(g_i - t_j\right) \\ = & \quad \textstyle t_j + \sum_{i \notin J} \alpha_i^j \frac{f_i}{f_j} \left(t_i - g_i\right), \end{aligned}$$

where the final equality follows from the definition of β_i^j in Equation (2). After subtracting g_j from both sides,

$$(E[w(s) \,|\, q_j] - g_j) f_j = -\sum_{i \notin J} \alpha_i^j f_i \Delta_i - f_j \Delta_j,$$

and since the right hand side is zero by construction of α^j (remember that $\alpha_j^j = 1$) it follows that $E[w(s) \,|\, q_j] = g_j$. ∎

Let $u_A(w, \sigma) := E_\sigma[u_A(w, e)]$, define the notation $U(t, \sigma) := E_q[u_A(t(q), \sigma) \,|\, \sigma]$ as the expected payoff the agent gets when being rewarded according to transfer schedule t and playing the mixed strategy σ. Also, $c(\sigma) = E_\sigma[c(e)]$.

Lemma A.2 *Fix a mixed strategy σ, and suppose that two transfer schedules t_1, t_2 (not necessarily implemented by an information structure) satisfy*

$$\sum_{i=1}^{j} t_2(q_i) f(q_i \,|\, \sigma) \geq \sum_{i=1}^{j} t_1(q_i) f(q_i \,|\, \sigma) \qquad (3)$$
$$\forall j = 1, ..., M, \text{ with equality for } j = M.$$

For any $\tilde{e}$ that is dominated by σ according to the MLRP,

$$U(t_1, \sigma) - U(t_1, \tilde{e}) \geq U(t_2, \sigma) - U(t_2, \tilde{e}).$$

For any $\tilde{e}$ that dominates σ according to the MLRP,

$$U(t_1, \sigma) - U(t_1, \tilde{e}) \leq U(t_2, \sigma) - U(t_2, \tilde{e}).$$

Proof of Lemma A.2 First consider the case where $\tilde{e}$ is dominated by σ according to the MLRP.

$$U(t_2, \sigma) - U(t_2, \tilde{e}) = \sum_{i=1}^{M} t_2(q_i)[f(q_i \,|\, \sigma) - f(q_i \,|\, \tilde{e})] - [c(\sigma) - c(\tilde{e})].$$

Defining $L_i = f(q_i \,|\, \tilde{e})/f(q_i \,|\, \sigma)$, one can write $f(q_i \,|\, \sigma) - f(q_i \,|\, \tilde{e}) = f(q_i \,|\, \sigma)(1 - L_i)$. By the MLRP supposition, L_i is weakly decreasing.

$$\begin{aligned}
&\sum_{i=1}^{M} t_2(q_i)[f(q_i \,|\, \sigma) - f(q_i \,|\, \tilde{e})] - [c(\sigma) - c(\tilde{e})] \\
&= \sum_{i=1}^{M} t_2(q_i) f(q_i \,|\, \sigma)(1 - L_i) - [c(\sigma) - c(\tilde{e})] \\
&= \sum_{i=1}^{M} t_2(q_i) f(q_i \,|\, \sigma)(1 - L_M + \sum_{j=i+1}^{M} L_j - L_{j-1}) - [c(\sigma) - c(\tilde{e})] \\
&= \sum_{i=1}^{M} t_2(q_i) f(q_i \,|\, \sigma)(1 - L_M) + \sum_{i=1}^{M} \sum_{j=i+1}^{M} t_2(q_i) f(q_i \,|\, \sigma)(L_j - L_{j-1}) - [c(\sigma) - c(\tilde{e})].
\end{aligned}$$

But the first term is $E[t_2(q) \,|\, \sigma](1 - L_M) = E[t_1(q) \,|\, \sigma](1 - L_M)$ by Equation (3). The only term where the transfer schedule appears is

$$\begin{aligned}
&\sum_{i=1}^{M} \sum_{j=i+1}^{M} t_2(q_i) f(q_i \,|\, \sigma)(L_j - L_{j-1}) \\
&= \sum_{j=2}^{M} (L_j - L_{j-1}) \sum_{i=1}^{j-1} t_2(q_i) f(q_i \,|\, \sigma).
\end{aligned}$$

Since $L_j - L_{j-1} \leq 0$, it follows from Equation (3) that

$$U(t_1, \sigma) - U(t_1, \tilde{e}) \geq U(t_2, \sigma) - U(t_2, \tilde{e}).$$

In the other case where $\tilde{e}$ MLRP dominates σ, the proof can be taken verbatim except that L_i is weakly increasing, and this is why the inequality is reversed. ∎

Proof of Lemma 3.6 This follows from Lemma A.2. Fix an arbitrary e^* and information structure H. Let $\sigma = e^*$ and $\tilde{e} = e$. Since $e \leq e^*$ the MLRP condition is satisfied. Define

$t_2(\cdot)$ as $t(\cdot\,|\,H, e^*)$, and $t_1(\cdot)$ as $t(\cdot\,|\,H_{\text{FR}}, e^*)$. By Lemma 3.4, the supposition in Equation (3) is satisfied. ∎

Proof of Proposition 3.2 First, it is shown that the highest pure strategy equilibrium is achieved by full revelation. It is maintained that $F(q\,|\,e)$ is ordered according to the MLRP, and that $E(\theta\,|\,q, e)$ is nondecreasing in q, but the supermodularity condition on $E(\theta\,|\,q, e)$ is replaced by Assumption A.3.

Assumption A.3 *For any $e_1, e_2 \in E$ with $e_2 > e_1$, $\forall j = 1, ..., M$*

$$\sum_{i=1}^{j} \{E(\theta\,|\,q_i, e_1) - E(\theta\,|\,q_i, e_2) + \kappa\}\, f(q_i\,|\,e_2) \geq 0,$$

where $\kappa = E(\theta\,|\,e_2) - E_q[E(\theta\,|\,q, e_1)\,|\,e_2]$.

Assumption A.3 is implied by supermodularity of $E(\theta\,|\,q, e)$, since then the term in curly brackets is decreasing in q_i and the sum equals zero for $j = M$.

Define $\text{BR}(t(\cdot))$ as the set of utility maximizing effort levels for the agent when he faces transfer schedule $t(\cdot)$. Since $F(q\,|\,e)$ and $c(e)$ are continuous, the agent's objective function is continuous in $t(\cdot)$ and e, so that by the Maximum Theorem $\text{BR}(t)$ is nonempty, compact-valued and upper hemicontinuous.

Lemma A.4 *Suppose that e_1 is induced by some H_1. Then there exists H_2 that induces e_2, and either $e_2 > e_1$ or $e_2 = e_1$ and $H_2 = H_{FR}$.*

Proof of Lemma A.4 Since e_1 is induced by H_1, one has $e_1 \in \text{BR}(t(\cdot\,|\,H_1, e_1))$. Moving to full revelation, Lemma 3.6 implies that $e_2 := \max \text{BR}(t(\cdot\,|\,H_{FR}, e_1)) \geq e_1$. The final step is to show that $e_2 \in \text{BR}(t(\cdot\,|\,H_2, e_2))$ for some information structure H_2. When $e_2 = e_1$ this is satisfied since $e_2 \in \text{BR}(t(\cdot\,|\,H_{FR}, e_2))$. If $e_2 > e_1$, one can find an information structure H_2 so that $t(\cdot\,|\,H_2, e_2) = t(\cdot\,|\,H_{FR}, e_1) + \kappa$, where κ is the constant described in Assumption A.3. Since $t(q\,|\,H_{FR}, e_1) = E(\theta\,|\,q, e_1)$ is nondecreasing, Assumption A.3 in conjunction with Lemma 3.5 imply that this is possible. e_2 is implemented since

$$e_2 \in \text{BR}(t(\cdot\,|\,H_{FR}, e_1)) = \text{BR}(t(\cdot\,|\,H_{FR}, e_1) + \kappa) = \text{BR}(t(\cdot\,|\,H_2, e_2)).$$

∎

In Lemma A.6 I show that a best pure strategy equilibrium exists, strengthening Assumption A.3 to supermodularity and using the following Lemma A.5.

Lemma A.5 *Suppose $\exists e' \geq \hat{e}$ with $e' \in BR(t(\cdot \,|\, H_{FR}, \hat{e}))$. If $F(q \,|\, e)$ is ordered according to FOSD (so a fortiori, if it satisfies the MLRP), and $E(\theta \,|\, q, e)$ is supermodular, $\exists e^* \geq \hat{e}$ such that $e^* \in BR(t(\cdot \,|\, H_{FR}, e^*))$, that is e^* is an equilibrium under full revelation. Also, there exists a highest pure strategy equilibrium under full revelation.*

Proof of Lemma A.5 Let $\hat{E} := \{e \in E : e \geq \hat{e}\}$, and let $\xi : \hat{E} \to \hat{E}$ be defined as $\xi(e) := \max(\mathrm{BR}(t(\cdot \,|\, H_{\mathrm{FR}}, e)))$. The maximum of $\mathrm{BR}(\cdot)$ exists by the Maximum Theorem, and that $\xi(e) \in \hat{E}$ follows from $\xi(\hat{e}) \geq \hat{e}$ and $\xi(\cdot)$ being nondecreasing. To show that $\xi(\cdot)$ is nondecreasing, take any $e_2 > e_1$. Observe that

$$U(t(\cdot \,|\, H_{\mathrm{FR}}, e_2), e) - U(t(\cdot \,|\, H_{\mathrm{FR}}, e_1), e) = \sum_{i=1}^{M} [E(\theta \,|\, q_i, e_2) - E(\theta \,|\, q_i, e_1)] f(q_i \,|\, e). \tag{4}$$

By supermodularity, $E(\theta \,|\, q, e_2) - E(\theta \,|\, q, e_1)$ is nondecreasing in q. Under the FOSD ordering, Equation 4 is nondecreasing in e.

By Result A.1, $e^* = \xi(e^*)$ for some $e^* \in \hat{E}$. Also, $\xi(\cdot)$ has a highest fixed point so there exists a highest pure strategy equilibrium under full revelation. ∎

Lemma A.6 *The set of pure strategy effort levels across all information structures has a maximum.*

Proof of Lemma A.6 The set of all possible pure strategy effort levels, E^*, is non-empty since for the information structure that reveals no information has an equilibrium with $e = \min E$. Let $\bar{e} = \sup E^*$, which is finite by compactness of E. There exists a sequence of information structures $\{H_n\}$ and associated equilibrium effort levels $\{e_n\}$ with $e_n \to \bar{e}$. By Lemma A.5, one can find another sequence $\{\tilde{e}_n\}$ with $\tilde{e}_n \geq e_n$, and each $\tilde{e}_n$ is induced by H_{FR}. Since a highest equilibrium exists under H_{FR}, it must coincide with $\bar{e}$. ∎

Proof of Proposition 3.3 The submodularity assumption is relaxed. Suppose there exists $e_2 > e_1$ with

$$\sum_{i=1}^{j} \{E(\theta \,|\, q_i, e_2) - E(\theta \,|\, q_i, e_1) + \kappa\} \, f(q_i \,|\, e_1) > 0, \; \forall j < M - 1, \tag{5}$$

where $\kappa = E(\theta \,|\, e_1) - E_q[E(\theta \,|\, q, e_2) \,|\, e_1]$. This is a violation of Assumption A.3, and is implied by strict submodularity. Given that $F(q \,|\, e)$ is not constant and ordered according to the MLRP, it follows from Lemma 3.6 that

$$0 < E_q[E(\theta \,|\, q, e_2) \,|\, e_2] - E_q[E(\theta \,|\, q, e_2) \,|\, e_1] < E_q[E(\theta \,|\, q, e_1) \,|\, e_2] - E_q[E(\theta \,|\, q, e_1) \,|\, e_1]. \tag{6}$$

To get the conclusion of Equation (6), one can alternatively assume that the inequalities in Equation (5) are weak, and strict for only some j as long as the ratio $f(q \mid e_2)/f(q \mid e_1)$ is strictly increasing in q.

At any rate, consider the following cost function

$$c(e) = \begin{cases} 0 & \text{if } e \leq e_1, \\ \xi & \text{if } e_1 < e \leq e_2, \\ \infty & \text{if } e > e_2, \end{cases}$$

which is not continuous, and where $\xi > 0$ satisfies

$$\xi \in (E_q[E(\theta \mid q, e_2) \mid e_2] - E_q[E(\theta \mid q, e_2) \mid e_1], E_q[E(\theta \mid q, e_1) \mid e_2] - E_q[E(\theta \mid q, e_1) \mid e_1]).$$

Given this cost function, under full revelation a pure equilibrium can only involve e_1 or e_2. However, given that the market conjectures e_1, the agent wants to deviate to e_2, and given a conjecture of e_2, the agent wants to deviate to e_1. On the other hand, if the market conjectures e_1, then by Equation (5) and Lemma 3.5 one can implement a transfer schedule that coincides with $E(\theta \mid q, e_2)$ up to a constant, and therefore makes e_1 incentive compatible. Since $F(q \mid e)$ is continuous in e, the same conclusion can be reached with a continuous cost function that approximates the previously stated one.

Perhaps one is more concerned with the fact that under full revelation, no equilibrium exists. Yet if one slightly strengthens the assumption (for simplicity stated in terms of supermodularity instead of the condition of Equation (5)), an equilibrium under full revelation exists yet is inferior. The assumption is that $E(\theta \mid q, e)$ is strictly supermodular on $Q \times \{e_0, e_1, e_2\}$ given the order $e_0 \triangleleft e_2 \triangleleft e_1$ on $\{e_0, e_1, e_2\}$, for some e_0, e_1, e_2 with $e_0 < e_1 < e_2$. Then one can find a cost function under which H_{FR} implements only e_0, and $\exists H$ that implements e_1. ∎

Example A.7 *In the following example, all assumptions of Proposition 5.2 Dewatripont, Jewitt and Tirole (1999*a*) are satisfied, yet full revelation does not maximize effort.*

Let θ and η be two random variables that are independently distributed, and have a standard normal distribution, that is $\theta \sim N(0,1)$ and $\eta \sim N(0,1)$. The production function satisfies

$$q = e + \theta + \alpha(e)e\theta + \beta(e)\eta,$$

where $\alpha(e)$ and $\beta(e)$ are functions that will be specified later, and satisfy

$$(1 + \alpha(e)e)^2 + \beta(e)^2 = k,$$

for some constant k*. As long as* $\forall e,\ 1+\alpha(e)e > 0$*, one has that* θ *and* q *are affiliated conditional on* e*. Also, the marginal over output is* $q \mid e \sim N(e,k)$*, so the MLRP is satisfied. Consider the simple class of signals where* $s = q + \epsilon$*, with* $\epsilon \sim N(0,\sigma_\epsilon^2)$ *that is independent of* (θ,η)*. The information structure will be denoted as* $H_{\sigma_\epsilon^2}$*.*

Up to a term that is independent of s *and therefore irrelevant,*

$$E(\theta|H_{\sigma_\epsilon^2}, s, e) = s \cdot \frac{1+\alpha(e)e}{k+\sigma_\epsilon^2}.$$

Given a market conjecture of $\hat{e}$*, the agent maximizes*

$$e \cdot \frac{1+\alpha(\hat{e})\hat{e}}{k+\sigma_\epsilon^2} - c(e).$$

Now consider the following parametrization. $E = \mathbb{R}_+$*,* $c(e) = e^2/2, k = 4$*. Also,* $\alpha(e) = 0$ *for* $e \le 0.25+\delta$ *or* $e \ge 0.25+3\delta$*, while* $\alpha(e) = 1$ *for* $e \in (0.25+\delta, 0.25+3\delta)$*, for some* $\delta > 0$ *sufficiently small. Under full revelation of* q*, that is when* $\sigma_\epsilon^2 = 0$*, there is a unique pure strategy equilibrium with* $e_{FR} = 0.25$*. On the other hand, if* $\sigma_\epsilon^2 = \frac{0.25-6\delta}{0.25+2\delta}$*, there exists an equilibrium with* $e = 0.25+2\delta$*.*

*While all assumptions from Dewatripont, Jewitt and Tirole (1999*a*) are satisfied, the example might be considered artificial because of the discontinuity in* $\alpha(\cdot)$*. An similar example where* $\alpha(\cdot)$ *is continuous can generate the same conclusion with a more complicated (but still continuous and increasing) cost function.*

Proposition A.8 *Suppose all assumptions in Proposition 3.2 are satisfied,* $Q \subset \mathbb{R}$ *is measurable, and* $F(q \mid e)$ *has continuous density. Restrict attention to information structures that for any conjecture induce a transfer schedule with a countable number of discontinuities. Then full revelation maximizes effort among all pure strategy equilibria.*

Proof of Proposition A.8 Start with information structure H and equilibrium $\hat{e}$. Lemma 3.4 also holds in that $\forall \hat{q}$, $\int_{-\infty}^{\hat{q}} t(q \mid H, \hat{e})\ \mathrm{d}F(q \mid \hat{e}) \ge \int_{-\infty}^{\hat{q}} t(q \mid H_{\mathrm{FR}}, \hat{e})\ \mathrm{d}F(q \mid \hat{e})$, with equality for $\hat{q} = +\infty$. As argued momentarily, one can show that Lemma 3.6 applies so that $e' := \max \mathrm{BR}(t(\cdot \mid H_{\mathrm{FR}}, \hat{e})) \ge \hat{e}$, and by the argument in Lemma A.5 there exists an equilibrium $e^* \ge \hat{e}$ under full revelation.

It remains to show that Lemma 3.6 applies. Suppose that for some $\hat{e}$ and two transfer schedules with countably many discontinuities, t_1 and t_2, satisfy $\forall \hat{q}$, $\int_{-\infty}^{\hat{q}} t_2(q)\ \mathrm{d}F(q \mid \hat{e}) \ge \int_{-\infty}^{\hat{q}} t_1(q)\ \mathrm{d}F(q \mid \hat{e})$, with equality for $\hat{q} = +\infty$. It needs to be shown that for any $e \le \hat{e}$,

$$\delta := [U(t_1,\hat{e}) - U(t_1,e)] - [U(t_2,\hat{e}) - U(t_2,e)] \ge 0.$$

Let $L(q) = f(q \,|\, e)/f(q \,|\, \hat{e})$, which is nonincreasing, and define $\Delta(q) := t_2(q) - t_1(q)$. Straightforward calculations show that

$$\delta = \int_{-\infty}^{+\infty} \Delta(q)L(q) \; \mathrm{d}F(q \,|\, \hat{e}) = \sum_{i=1}^{+\infty} \int_{I_i} \Delta(q)L(q) \; \mathrm{d}F(q \,|\, \hat{e}),$$

where I_i is an interval on which $\Delta(q)$ is continuous, and I_i lies below I_{i+1}. To show that $\sum_{i=1}^{+\infty} \int_{I_i} \Delta(q)L(q) \; \mathrm{d}F(q \,|\, \hat{e})$ is nonnegative, it is claimed that $\forall j$,

$$\sum_{i=1}^{j} \int_{I_i} \Delta(q)L(q) \; \mathrm{d}F(q \,|\, \hat{e}) \geq 0. \tag{7}$$

Observe that Equation (7) is equivalent to

$$\sum_{i=1}^{j} \frac{L(q_i)}{L(q_j)} \int_{I_i} \Delta(q) \; \mathrm{d}F(q \,|\, \hat{e}) \geq 0,$$

where $q_i \in I_i$. It is now claimed that

$$\sum_{i=1}^{j} \frac{L(q_i)}{L(q_j)} \int_{I_i} \Delta(q) \; \mathrm{d}F(q \,|\, \hat{e}) \geq \sum_{i=1}^{j} \int_{I_i} \Delta(q) \; \mathrm{d}F(q \,|\, \hat{e}) \geq 0,$$

where the second inequality follows from the supposition. For $j = 1$, the first inequality is satisfied. If it is satisfied up to k, then for $j = k + 1$,

$$\begin{aligned} \sum_{i=1}^{k+1} \frac{L(q_i)}{L(q_{k+1})} \int_{I_i} \Delta(q) \; \mathrm{d}F(q \,|\, \hat{e}) = & \quad \textstyle\sum_{i=1}^{k} \frac{L(q_i)}{L(q_{k+1})} \int_{I_i} \Delta(q) \; \mathrm{d}F(q \,|\, \hat{e}) + \int_{I_{k+1}} \Delta(q) \; \mathrm{d}F(q \,|\, \hat{e}) \\ \geq & \quad \textstyle\sum_{i=1}^{k} \int_{I_i} \Delta(q)L(q) \; \mathrm{d}F(q \,|\, \hat{e}) + \int_{I_{k+1}} \Delta(q) \; \mathrm{d}F(q \,|\, \hat{e}) \geq 0. \end{aligned}$$

The first inequality follows from $L(q_{k+1}) \leq L(q_k)$. ∎

Proof of Proposition 3.9 First consider the case where the market observes (q, s_θ), with some equilibrium effort $\hat{e}$. The wage after each realization of (q, s_θ) satisfies $w(q, s_\theta, \hat{e}) = E(\theta \,|\, q, s_\theta, \hat{e})$. Since the distribution of (q, s_θ) can be written as $f(q, s_\theta \,|\, e) = f(q \,|\, s_\theta, e)g(s_\theta)$, the agent's payoff is

$$\begin{aligned} & E_{(q,s_\theta)}[E(\theta \,|\, q, s_\theta, \hat{e}) \,|\, e] - c(e) \\ = \; & E_q[E_{s_\theta}[E(\theta \,|\, q, s_\theta, \hat{e}) \,|\, q, e] \,|\, e] - c(e). \end{aligned}$$

Now compare this to a situation where only q gets disclosed to the market, and its conjecture is $\hat{e}$. Since the wage satisfies $w(q,\hat{e}) = E(\theta \,|\, q,\hat{e})$, the agent maximizes

$$\begin{aligned} & E_q[E(\theta \,|\, q,\hat{e}) \,|\, e] - c(e) \\ = \quad & E_q[E_{s_\theta}[E(\theta \,|\, q, s_\theta, \hat{e}) \,|\, q, \hat{e}] \,|\, e] - c(e). \end{aligned}$$

Now for $e < \hat{e}$,

$$E_{s_\theta}[E(\theta \,|\, q, s_\theta, \hat{e}) \,|\, q, \hat{e}] \le E_{s_\theta}[E(\theta \,|\, q, s_\theta, \hat{e}) \,|\, q, e]$$

since $E(\theta \,|\, q, s_\theta, \hat{e})$ is nondecreasing in s_θ and $H(s_\theta \,|\, q, e)$ dominates $H(s_\theta \,|\, q, \hat{e})$ according to FOSD. This means that in a situation where only q is disclosed and the market's conjecture is $\hat{e}$, the agent's highest best response is weakly above $\hat{e}$. Lemma A.5 then implies that there exists an equilibrium $e^* \ge \hat{e}$. ∎

A.2 Proofs for Section 4

Proof of Lemma 4.2 For a generic mixed strategy σ, define $\bar{e} = \sup(\text{support}(\sigma))$. Fix an arbitrary $q_2 > q_1$, and define the likelihood ratio $L_1^2(e) := \frac{f(q_2 \,|\, e)}{f(q_1 \,|\, e)}$, which is nondecreasing in e by the MLRP property. Then

$$\begin{aligned} E(\theta \,|\, q_2, \sigma) - E(\theta \,|\, q_1, \sigma) &= \frac{E_\sigma[f(q_2 \,|\, e) E(\theta \,|\, q_2, e)]}{E_\sigma[f(q_2 \,|\, e)]} - \frac{E_\sigma[f(q_1 \,|\, e) E(\theta \,|\, q_1, e)]}{E_\sigma[f(q_1 \,|\, e)]} \\ &= \frac{E_\sigma[f(q_1 \,|\, e) L_1^2(e) E(\theta \,|\, q_2, e)]}{E_\sigma[f(q_1 \,|\, e) L_1^2(e)]} - \frac{E_\sigma[f(q_1 \,|\, e) E(\theta \,|\, q_1, e)]}{E_\sigma[f(q_1 \,|\, e)]} \\ &= \frac{E_\sigma[f(q_1 \,|\, e) L_1^2(e) \{E(\theta \,|\, q_2, e) - E(\theta \,|\, q_1, e)\}]}{E_\sigma[f(q_1 \,|\, e) L_1^2(e)]} \\ &+ \frac{E_\sigma[f(q_1 \,|\, e) L_1^2(e) E(\theta \,|\, q_1, e)]}{E_\sigma[f(q_1 \,|\, e) LR(e)]} - \frac{E_\sigma[f(q_1 \,|\, e) E(\theta \,|\, q_1, e)]}{E_\sigma[f(q_1 \,|\, e)]}. \end{aligned}$$

By supermodularity,

$$E(\theta \,|\, q_2, \bar{e}) - E(\theta \,|\, q_1, \bar{e}) \ge \frac{E_\sigma[f(q_1 \,|\, e) L_1^2(e) \{E(\theta \,|\, q_2, e) - E(\theta \,|\, q_1, e)\}]}{E_\sigma[f(q_1 \,|\, e) L_1^2(e)]}.$$

Since $L_1^2(e)$ is nondecreasing and $E(\theta \,|\, q_1, e)$ is nonincreasing in e, we have

$$0 \ge \frac{E_\sigma[f(q_1 \,|\, e) L_1^2(e) E(\theta \,|\, q_1, e)]}{E_\sigma[f(q_1 \,|\, e) L_1^2(e)]} - \frac{E_\sigma[f(q_1 \,|\, e) E(\theta \,|\, q_1, e)]}{E_\sigma[f(q_1 \,|\, e)]}.$$

The bottom line is that $E(\theta \,|\, q_2, \bar{e}) - E(\theta \,|\, q_1, \bar{e}) \ge E(\theta \,|\, q_2, \sigma) - E(\theta \,|\, q_1, \sigma)$. ∎

Proof of Proposition 4.1 For a mixed strategy σ, define $\bar{e} = \sup(\text{support}(\sigma))$.

1. Start with any mixed strategy equilibrium σ under the fully revealing information structure. By continuity, $\bar{e}$ is a best response for the agent. By Lemma 4.2 and the argument in Lemma A.4, there exists an information structure H' such that $t(\cdot \mid H', \bar{e})$ coincides with $t(\cdot \mid H_{\text{FR}}, \sigma^*)$ up to a constant, implementing $\bar{e}$. Since Proposition 3.2 applies under the assumptions of Proposition 4.1, the conclusion follows.

2. This part is a corollary of the following result. For any finite E, if there is an information structure H with mixed strategy equilibrium σ that is supported on two adjacent effort levels e_a and e_b with $e_a < e_b$, then under full revelation there is a pure equilibrium that generates weakly higher effort than e_b. For this, the steps of the proof of Proposition 3.2 are reproduced. When moving from $t(\cdot \mid H, \sigma)$ to $t(\cdot \mid H_{\text{FR}}, \sigma)$, there is a best response weakly higher than e_b by Lemma A.2, called e_c. Here it is used that σ is supported on adjacent effort levels, as any pure strategy can be compared to σ according to the MLRP. By supermodularity and Lemma 4.2, $t(\cdot \mid H_{\text{FR}}, \sigma)$ can be sustained under some H' up to a constant if the market's conjecture is e_c. Hence the optimal information structure cannot induce mixing between two adjacent effort levels.

3. With $Q = \{q_L, q_H\}$, define $\delta(H, \sigma) = t(q_H \mid H, \sigma) - t(q_L \mid H, \sigma)$. Starting out with information structure H and equilibrium σ, it is the case that

$$\delta(H, \sigma) \leq \delta(H_{\text{FR}}, \sigma) \leq \delta(H_{\text{FR}}, \bar{e}).$$

   The first inequality is due to Lemma 3.4, and the second due to Lemma 4.2. By Lemma 3.5 there exists an information structure H' such that $\delta(H, \sigma) = \delta(H', \bar{e})$, so that under H' one can sustain $\bar{e}$ as an equilibrium. ∎

A.3 Proofs for Section 5

Proof of Proposition 5.1 There are two cases to consider.

Explicit incentives are unaffected by the disclosure policy. For a given information structure and market conjecture, after output realization q the agent can expect a wage of

$$\hat{t}(q \mid H, e, \hat{w}) = \hat{w}(q) + t(q \mid H, e),$$

which consists of wage contract $\hat{w}(q)$ and the career-concerns term $t(q \mid H, e)$. From this point, one can verify that the steps in the proof of Proposition 3.2 still apply, so that full revelation maximizes effort.

Explicit incentives are affected by the disclosure policy. Define $\xi(q,e)$ as

$$\xi(q,e)=\hat{w}(q)+E(\theta\,|\,q,e).$$

As $E(\theta\,|\,q,e)$ is supermodular, so is $\xi(q,e)$, and since $\hat{w}(q)$ is nondecreasing in q, this is also true for $\xi(q,e)$. Replacing $E(\theta\,|\,q,e)$ with $\xi(q,e)$ in the proof of Proposition 3.2, one can reproduce the full revelation result. ∎

Proof of Proposition 5.2 It is enough to show the result only for the case where the wage can depend on q, even if it gets garbled. For if the wage has to depend on the signal, this just induces an additional constraint on the principal when she deviates from full revelation. Take any $(H,\tilde{w})$ that induces effort $\tilde{e}$, and let $w_e=E[\tilde{w}\,|\,\tilde{e}]\geq 0$ be the expected wage paid to the agent. Consider contracts of the form that pay a wage of zero for all outputs, except a wage r for the highest output q_M. Denote any such contract $w_M(\cdot\,|\,r):Q\to\mathbb{R}$.

Let $\mathrm{BR}(H,w,e)$ be the agent's best response under information structure H, wage contract w and market conjecture e. By construction of $w_M(q\,|\,w_e/f(q_M\,|\,\tilde{e}))$ and the limited liability constraint $\tilde{w}(q)\geq 0$, $\forall j=1,...,M$,

$$\sum_{i=1}^{j}\tilde{w}(q_i)f(q_i\,|\,\tilde{e})\geq\sum_{i=1}^{j}w_M(q_i\,|\,w_e/f(q_M\,|\,\tilde{e}))f(q_i\,|\,\tilde{e}),$$

with equality for $j=M$. In Lemma A.2, letting $\sigma=\tilde{e}$ and

$$t_1(q)=w_M(q\,|\,w_e/f(q_M\,|\,\tilde{e}))+t(q\,|\,H_{\mathrm{FR}},\tilde{e}),\quad t_2(q)=\tilde{w}(q)+t(q\,|\,H,\tilde{e}),$$

it follows that $\mathrm{BR}(H_{\mathrm{FR}},w_M(w_e/f(q_M\,|\,\tilde{e})),\tilde{e})\geq\tilde{e}$.

If $\mathrm{BR}(H_{\mathrm{FR}},w_M(0),\tilde{e})\geq\tilde{e}$, then full revelation absent any wage payment induces some effort $\tilde{e}'\geq\tilde{e}$.

If $\mathrm{BR}(H_{\mathrm{FR}},w_M(0),\tilde{e})<\tilde{e}$, it follows from the concavity of the agent's objective that any

$$e\in[\mathrm{BR}(H_{\mathrm{FR}},w_M(0),\tilde{e}),\mathrm{BR}(H_{\mathrm{FR}},w_M(w_e/f(q_M\,|\,\tilde{e})),\tilde{e})]$$

can be made optimal by some contract $w_M(t)$ with $t\in[0,w_e/f(q_M\,|\,\tilde{e})]$. So there exists contract $w_M(\tilde{t})$ that implements $\tilde{e}$ at cost $f(q_M\,|\,\tilde{e})\cdot\tilde{t}\leq w_e$. That the agent's objective is concave follows from the concavity of $f(q_M\,|\,e)$ and the decreasing returns to signal jamming assumption. ∎

Proof of Proposition 5.5 Suppose s_θ is disclosed to the market, and its conjecture is $\hat{e}(\cdot)$.

From an ex-ante perspective, the agent chooses strategy $e(\cdot) : S_\theta \to E$ to maximize

$$\begin{aligned} u_D(e(\cdot)\,|\,\hat{e}(\cdot)) := & \quad E_{(q,s_\theta)}[E(\theta\,|\,q, s_\theta, \hat{e}(s_\theta))\,|\,e(\cdot)] - E_{s_\theta}[c(e(s_\theta))] \\ = & \quad E_q[E_{s_\theta}[E(\theta\,|\,q, s_\theta, \hat{e}(s_\theta))\,|\,q, e(\cdot)]\,|\,e(\cdot)] - E_{s_\theta}[c(e(s_\theta))] \\ = & \quad E_{s_\theta}\{E_q[E_{s'_\theta}[E(\theta\,|\,q, s'_\theta, \hat{e}(s'_\theta))\,|\,q, e(\cdot)]\,|\,s_\theta, e(s_\theta)] - c(e(s_\theta))\}. \end{aligned}$$

If s_θ is not disclosed, then under conjecture $\hat{e}(s_\theta)$ his payoff is

$$\begin{aligned} u_N(e(\cdot)\,|\,\hat{e}(\cdot)) := & \quad E_{(q,s_\theta)}[E(\theta\,|\,q, \hat{e}(\cdot))\,|\,e(\cdot)] - E_{s_\theta}[c(e(s_\theta))] \\ = & \quad E_q[E_{s_\theta}[E(\theta\,|\,q, s_\theta, \hat{e}(s_\theta))\,|\,q, \hat{e}(\cdot)]\,|\,e(\cdot)] - E_{s_\theta}[c(e(s_\theta))] \\ = & \quad E_{s_\theta}\{E_q[E_{s'_\theta}[E(\theta\,|\,q, s'_\theta, \hat{e}(s'_\theta))\,|\,q, \hat{e}(\cdot)]\,|\,s_\theta, e(s_\theta)] - c(e(s_\theta))\}. \end{aligned}$$

Now consider the following two cases.

Expected ability is nondecreasing in the ability estimate. Suppose $\hat{e}(\cdot)$ is an equilibrium when s_θ is disclosed to the market, and $E(\theta\,|\,q, s_\theta, \hat{e}(s_\theta))$ is nondecreasing in s_θ. I claim that if s_θ is not disclosed and the market conjecture is $\hat{e}(\cdot)$, then there exists a best response $\tilde{e}(\cdot)$ that satisfies $\tilde{e}(\cdot) \geq \hat{e}(\cdot)$.

For any s_θ and $e \in E$, define strategy $\hat{e}^{(s_\theta,e)}(\cdot)$ that coincides with $\hat{e}(\cdot)$, except at s_θ where it prescribes e. Given any s_θ, if $e < \hat{e}(s_\theta)$ then

$$u_N(\hat{e}^{(s_\theta,e)}(\cdot)\,|\,\hat{e}(\cdot)) \leq u_D(\hat{e}^{(s_\theta,e)}(\cdot)\,|\,\hat{e}(\cdot)) \leq u_D(\hat{e}(\cdot)\,|\,\hat{e}(\cdot)) = u_N(\hat{e}(\cdot)\,|\,\hat{e}(\cdot)).$$

The first inequality follows from $\hat{e}^{(s_\theta,e)}(\cdot) \leq \hat{e}(\cdot)$. As then, $H(\cdot\,|\,q, \hat{e}^{(s_\theta,e)}(\cdot))$ dominates $H(\cdot\,|\,q, \hat{e}(\cdot))$ according to FOSD, and remember that $E(\theta\,|\,q, s_\theta, \hat{e}(s_\theta))$ is nondecreasing in s_θ. The second inequality follows from $\hat{e}(\cdot)$ being an equilibrium when the ability estimate is disclosed. The equality follows from the definition of u_D and u_N.

That under no disclosure of the ability estimate there exists a best response $\tilde{e}(\cdot) \geq \hat{e}(\cdot)$ follows from the fact that u_N can be written as $u_N(e(\cdot)\,|\,\hat{e}(\cdot)) = \sum_{s_\theta} \alpha(e(s_\theta)\,|\,\hat{e}(\cdot))$, which will be referred to as the separability of $u_N(e(\cdot)\,|\,\hat{e}(\cdot))$ in $e(\cdot)$.

Defining $\hat{E} = \{e(\cdot) \in E^{|S_\theta|} : e(\cdot) \geq \hat{e}(\cdot)\}$ one can apply Lemma A.1 to conclude that there exists an equilibrium $e^*(\cdot) \geq \hat{e}(\cdot)$ if s_θ is not disclosed to the market. For this, one needs to show that there exists a monotone selection from the agent's best response correspondence. But given the separability of $u_N(e(\cdot)\,|\,\hat{e}(\cdot))$ in $e(\cdot)$, the best response correspondence has a product structure so that a maximal best response with respect to the partial order $\geq$ exists. This maximal selection is monotone by the supermodularity of $E(\theta\,|\,q, e(\cdot))$ in q and $e(\cdot)$, as follows from Lemma A.5.

Expected ability is nonincreasing in the ability estimate. Start from a situation where s_θ is not disclosed to the market with equilibrium strategy $\hat{e}(\cdot)$. If now s_θ is disclosed to the market and it believes the agent plays according to $\hat{e}(\cdot)$, then for any s_θ and $e < \hat{e}(s_\theta)$,

$$u_D(\hat{e}^{(s_\theta,e)}(\cdot) \mid \hat{e}(\cdot)) \leq u_N(\hat{e}^{(s_\theta,e)}(\cdot) \mid \hat{e}(\cdot)) \leq u_N(\hat{e}(\cdot) \mid \hat{e}(\cdot)) = u_D(\hat{e}(\cdot) \mid \hat{e}(\cdot)).$$

This follows from an analogous argument as in the previous case. $u_D(e(\cdot) \mid \hat{e}(\cdot))$ is separable in $e(\cdot)$ in that it can be written as $u_D(e(\cdot) \mid \hat{e}(\cdot)) = \sum_{s_\theta} \alpha(e(s_\theta) \mid \hat{e}(s_\theta))$, so the agent wants to deviate upwards to some $\tilde{e}(\cdot) \geq \hat{e}(\cdot)$. By the argument of Lemma A.5, there exists an equilibrium $e^*(\cdot) \geq \hat{e}(\cdot)$ when s_θ is disclosed to the market. ∎

Proof of Proposition 5.7 Suppose the principal can condition the garbling of q on the realization of s_θ, that is she commits to $H(\cdot \mid q, s_\theta)$. If one can show that full revelation of q is optimal for each s_θ, this implies that it remains optimal in a more restricted environment where the garbling of q has to be independent of s_θ, that is she can only commit to $H(\cdot \mid q)$.

Given information structure H and market conjecture $\hat{e}$, the agent's payoff can be written as

$$E_{s_\theta}\left[\sum_q t(q \mid H, s_\theta, \hat{e}) f(q \mid s_\theta, e) - c(e)\right],$$

which follows from s_θ being an exogenous signal. From here, the argument of Lemma A.4 can be reproduced. Given conjecture $\hat{e}$, moving to full revelation after each s_θ induces some effort $e' \geq \hat{e}$ as a straightforward extension of Lemma A.2 shows, and the maximal selection from the best response correspondence is nondecreasing. ∎