Instructions and Incentives in Organizations*

Kieron Meagher[†], Suraj Prasad [‡], and Andrew Wait [§]

April 12, 2022

Abstract

We propose a new perspective on organizations, emphasizing the role of instructions in providing incentives. We examine a conflict of interest between a principal and agent over the course of action in an organization. Instructions on the principal's preferred course, lead to a more informed decision by the agent when executing it, thus steering him towards the course. But less informative instructions are an effective incentive tool only when the agent is sufficiently risk averse. Consequently, hiring a risk averse agent can be optimal. The model yields a rich set of observed organizational configurations based on instructions, risk preferences of agents, and incentive pay.

Key words: instructions, incentives, risk aversion, organization
JEL classifications: D23, D L22, L23.

1 Introduction

Organizations exist to get things done when markets do not work well. Indeed, as Kenneth Arrow emphasizes in his book *Limits of Organization*, they are "a means of achieving the benefits of collective action in situations where the price system fails."[1] The precise nature of organization thus depends on the mechanism that substitutes for markets.

One mechanism, that is well-understood and extensively modeled in the economic literature, involves contracting on a proxy of an agent's performance. This agency theory approach provides a clear set of conditions under which these contracts should be used and their limitations. In particular, this literature highlights the costs of using these contracts when measures of performance are noisy and when agents are risk averse (Holmström (1979)).

A second mechanism, which accords with observed organizations but which has proven less tractable to model, emphasizes the role of instructions – where a worker is told how to do their job. This view is best summarized by Herbert Simon in his book *Administrative Behavior* (Simon (1947)),

*We thank Gabriele Gratton and Hongyi Li for comments.
[†]Research School of Economics, Australian National University. `kieron.meagher@anu.edu.au`
[‡]School of Economics, University of Sydney. `suraj.prasad@sydney.edu.au`
[§]School of Economics, University of Sydney. `andrew.wait@sydney.edu.au`
[1]See Arrow (1974), pp. 33.

where he defines an organization as a "complex pattern of communication" with "advisory and informational services" and "training" constituting some of the central modes to influence an agent's behavior.

The objective of our paper is to develop a simple model of instructions in organizations. Consistent with Herbert Simon's observation above, our model illustrates how instructions – where an agent is given information on how to execute a course of action – can be an important motivational tool. Our framework also allows us to see how both mechanisms – instructions and contracting for performance – interact with one another. We thus provide a richer, and more complete, picture of organizations that we observe in the real world.

In our setting, an agent chooses between two courses of action, and then executes the chosen course by making a decision. Both the principal and agent prefer better execution, with a more "correct" decision on the chosen course. [2] But they have a conflict of interest over which course to choose; the principal prefers the efficient one whereas the agent prefers the other. The selection of the course, and the resulting benefits from its choice and execution (which are realized in the future), are not contractible.

As an example, consider selling a product or service to a customer. Making the sale is itself straightforward. But there are *different ways* in which it can be made: a customer-focused approach tailored to the needs of the customer, or a generic product/service based approach. While the organization may prefer the former (with more reliable sales in the future), the salesperson may find the latter easier. At IBM, for instance, salespeople were encouraged to switch from "pitching newer, faster, more reliable gear" and "hawking hardware" to giving "good advice and software to solve specific business problems." [3] Furthermore, given how different the approaches are, information that improves execution on one of them, may not affect the other at all.

In other examples, a product can be developed to meet the needs of new or existing customers. Service can ensure better long-term adoption of a product or have a short-term focus. Closer to home, an academic can work on research topics that have a high impact or on less important topics. More generally, an organization confronts this conflict in most of its strategic choices for the future: "which products and services it will offer, which customers it will seek to serve, where it will operate" and "what technology it will use" (see Roberts and Saloner (2013)).

Given this setting, we ask the following question: what motivates an agent to choose the efficient course of action?

The traditional agency theory approach to this problem is to make the agent's pay contingent on the course chosen. But in many settings, it may be difficult to verify which course the agent pursues. And even if some performance measures are available, it may lead to dysfunctional behavior by the agent (see Holmström and Milgrom (1991) and Baker (1992), for example). Another option is to tie the agent's pay to some measure of the organization's benefit. But these benefits, as the examples above suggest, may only materialize after sufficient delay, reducing the ability

[2] For example, better execution of a course generates more value to the principal *and* enhances the reputation of the agent.

[3] See "How IBM Teaches Techies to Sell", *Fortune*, Jun. 6, 1989.

to commit to a reward tied to these benefits. Thus, standard “pay for performance” contracts, on their own, may not be an effective tool to address this problem.[4]

We show that instructions – where information on a course leads to a more informed decision by an agent – are a powerful motivational instrument in this setting. Instructions are a pervasive feature in organizations and can be thought of quite broadly: a supervisor guiding a subordinate on aspects of the job, informal and formal training programs where knowledge and skills are taught and imparted to a worker (Becker (1964)), organizational routines which use “rich and relevant information ... to guide action” (Nelson and Winter (1982) and Nelson and Winter (2002)),[5] and detailed instruction manuals and checklists for different contingencies an employee may encounter.

The idea underpinning our paper is a simple one. Executing a course of action requires making a decision. A correct decision – represented by a number on the real line – maximizes the execution benefit of both the principal and the agent. Incorrect decisions – measured by the distance from the correct one – by contrast, dilute benefits for both. The correct decision, however, is uncertain. Information thus plays a central role in these settings.

Instructions – which we model as an agent’s access to a course-specific informative signal – lead to a more informed decision by the agent on a course. This improves the agent’s execution of the course, which in turn induces the agent to choose it, thus realigning interests. Returning to our earlier sales example, IBM instituted an interactive self-study system called InfoWindow to provide information on making customer-focused sales, which required ”more negotiating, coddling, and long-term advising of customers than hawking hardware ever did.” With this system, a trainee could ”practice sales calls with an on-screen actor who portrays a manager” in a specific industry.[6]

In our model, where the correct decision is normally distributed, the signal also reduces the variability of losses associated with poor execution. Thus, incentives to select a course with instructions are stronger, the more risk averse an agent is. As a result, it can be optimal to hire a risk averse agent – they are more responsive to instructions.[7] The more difficult it is for the agent to learn (or the less informative the instructions), the greater the degree of risk aversion needed to induce the desired course.

Our emphasis in this paper is on aspects of work (the way in which it is done and executed) which are difficult to contract on. In many settings, however, there are ”routine” actions which can be verified with noise. For instance, in our sales example, the agent can work hard and contact more marginal customers to boost sales. To paint a more complete picture of motivation within an

[4]Also see Anderson (1985) for a detailed description of these conflicts in selling electronic components and why contracts may not be feasible in this setting.

[5]While we abstract from details in the production process associated with routines, Kvasov and Meagher (2022) model these aspects formally. Also see Chassang (2010) for a model on how routines – where agents use a fixed set of actions – emerge in equilibrium.

[6]The system is programmed to respond differently depending on what the salesperson does. See “How IBM Teaches Techies to Sell”, *Fortune*, Jun. 6, 1989.

[7]Risk aversion, therefore, is one way of formalizing the concept of *docility* put forward by Herbert Simon, which he refers to as being ”tractable, manageable, and above all, teachable.” See Simon (1991), pp. 35.

organization, we allow for a setting where the principal's benefit depends on both of these aspects (courses of action and effort) and examine how instructions and incentive pay interact with one another.[8]

The link between these two instruments – instructions and incentive pay – is a simple one. When instructions are less informative, it increases the need for a risk averse agent to induce the desired course of action. Thus, it is more costly to tie the agent's pay to a noisy performance measure. Said in a different way, solving one incentive problem (inducing the efficient course) through the use of instructions, can make it more difficult to solve another one, which is inducing effort. Overall, enriching the standard agency model, by allowing for instructions, accounts for commonly observed features of organizations: i) the prevalence of risk aversion, ii) low-powered incentive pay, and iii) an ambiguous relationship between uncertainty and the level of incentive pay, where the sign of the relationship depends on the measure of uncertainty considered.

Combining the two mechanisms of instructions and incentive pay also generates a rich set of organizational configurations based on the optimal level of agent risk aversion, the use of instructions, and the use of incentive pay.

When instructions are less informative relative to the agent's bias, the configuration that emerges is a *basic market*. As less informative instructions require very high levels of risk aversion to induce the efficient course of action, the organization settles for the inefficient course, settling for the benefits of a risk neutral agent and high-powered incentive pay (with full residual claimancy) instead.[9] Instructions are less likely to be used in this configuration. One way to think of this configuration is as outsourcing, with little emphasis on the way work is done.

As instructions become more informative, there is a threshold at which the optimal configuration jumps: both in terms of the course of action induced and the agent's risk aversion. So a small change in parameters results in a discontinuous change in the organizational configuration. In this region, agent risk aversion is high (because instructions are still relatively uninformative), incentive pay low-powered, and instructions more likely to be used relative to the ineffective market configuration. We call this configuration a *conventional* organization – conventional within the context of the economic literature which emphasizes a tradeoff between insurance and incentives (Holmström (1979)). The difference with the existing literature is that the terms of the tradeoff, which depends on the degree of risk aversion, are endogenous in our model.

Finally, when instructions are very informative relative to the agent's bias, inducing the efficient course of action with instructions is relatively easy. For this case, instructions are most likely to be used, agents are risk neutral, and incentive pay is high-powered with the agent being a full residual claimant. We call this a *hybrid* configuration. It has features associated with markets like residual claimancy and features associated with organizations like the use of instructions, and leads to the efficient outcome. Franchising, which relies both on detailed instructions

[8]We restrict our attention to linear incentive contracts for this verifiable performance measure.These contracts, which are easy to interpret, are also used by Holmström and Milgrom (1991). Also see Holmström and Milgrom (1987) and Carroll (2015) for settings in which linear contracts are optimal.

[9]This is assuming that the benefits from the course of action and effort to the principal are additive.

and high-powered incentive pay, is an interesting example of this hybrid structure.

These configurations bear a close resemblance to those studied by Anderson and Schmittlein (1984). Looking at sales in the electronic components industry, they document three modes of organization depending on i) the type of sale (customized or not), and ii) the degree of specificity of human assets, which in essence measures the level of difficulty involved in customizing a sale. Passive components – which require little customization – are typically sold through industrial distributors, who have ownership of the goods (and are thus residual claimants), and who are ″less subject to the manufacturer's influence″. This corresponds to the the basic market configuration in our model.

By contrast, active components – which require more customization – are sold through two channels. When customization is difficult – for instance when learning the ″ins and outs″ for a ″salesperson to be effective″, or when a new salesperson with ″experience in a product class″ needs ″extra training″ – manufacturers use ″direct″ sales employees who are paid a salary or salary plus incentives (a conventional organization). [10] When customization is easy, manufacturers are more likely to use sales reps, who operate ″like a sales department in a firm″, and are paid ″commissions″ (a more hybrid configuration).[11][12]

Our framework offers a different perspective on firm boundaries, which are traditionally viewed within the context of incentives through asset ownership (Grossman and Hart (1986) and Hart and Moore (1990)). Building on an observation in Holmström (1999), that employees in firms "rarely own any assets", our emphasis, instead, is on firms "controlling information channels" and selecting agents based on their risk preferences. This allows us to formalize Holmström's insight that effective organization requires a combination of market-like features and firm-like ones: in our paper this takes the form of a combination of residual claimancy and instructions.

The idea that information in the form of a signal from a principal to an agent can affect the agent's behavior is a theme that shows up in Van den Steen (2009) and Che and Kartik (2009) as well. In their setting with heterogeneous priors, information draws divergent beliefs closer to one another, realigning interests.[13] By contrast, in our setting with common priors, the selective release of information for one task bridges differences in preferences across tasks.[14] While instructions lower the agent's cost of execution in our framework, in Rantakari (2021) they reveal payoffs associated with different courses of action.[15] A key difference between our paper and these preceding papers is the role risk aversion plays in the provision of incentives.

[10]Additional factors, that are averaged for the measure of asset specificity, include confidential information, the time taken to get to know accounts, importance of key accounts, and customer loyalty.

[11]Sales reps also sell products for non competing manufacturers, which lies outside the scope of our model.

[12]In addition, Anderson and Schmittlein (1984) also find that a direct sales force is more likely to be used when evaluating performance is difficult. See Holmström and Milgrom (1991) and Holmström and Roberts (1998) for a rationalization of this finding.

[13]Also see the leadership model of Hermalin (1998).

[14]This selective release of information also plays an important role in Jiang and Prasad (2021).

[15]Taking a different approach, Raith (2021) allows for the voluntary transfer of the right to choose an action from an employee to the employer. He also allows for market trade to examine firm boundaries.

2 Model

2.1 Courses of Action and Information

The organization consists of a principal (P) who hires an agent (A). There are two courses of action indexed by a, $a \in \{L, H\}$.

The agent chooses a course of action and then executes the chosen course by making a decision. The agent's decision on course a is represented by the real number x_a. Associated with course a is a "correct" decision μ_a. A's payoff from the execution of course a is $\alpha(b_0 - |x_a - \mu_a|)$ and P's corresponding payoff is $(1-\alpha)(b_0 - |x_a - \mu_a|)$, where $|x_a - \mu_a|$ is the loss from incorrect execution, $\alpha \in (0, 1]$ is the relative share of the value of execution that the worker gets (say because of reputation effects), and where $b_0 > 0$ is a parameter. Thus, both P and A prefer a chosen course to be executed more correctly.

Where preferences diverge, is over the course of action chosen. P gets a benefit of B_a from the course of action a with $B_H = B > 0 = B_L$. A, by contrast, gets a benefit β_a with $\beta_L = \beta > 0 = \beta_H$. We assume that $B > \beta$ so that course H is efficient.

The benefits, from the course chosen, and from its execution, are not contractible.

The correct decision for course a is unknown with μ_a normally distributed with a prior mean of 0 and a variance of $\sigma^2 > 0$. These distributions across courses are independent of one another.

Prior to the agent choosing a course of action, the principal can instruct the agent by providing the agent access to a signal s_a, which is informative about μ_a, for each course of action a, with

$$s_a = \mu_a + \epsilon_a,$$

where ϵ_a is an idiosyncratic error term that is distributed normally with mean 0 and variance σ_ϵ^2. We interpret the agent's access to this signal as the principal *instructing* the agent, with a lower σ_ϵ^2 indicating a more informative signal. To begin with, we assume that instructing the agent on a course of action is costless. Later in our analysis, we impose a cost to the principal of instructing the agent.

For the analysis that follows, it is useful to fix σ^2 and define a new parameter q where $q = 1 - \sqrt{\frac{\sigma_\epsilon^2}{\sigma^2 + \sigma_\epsilon^2}}$. Taking σ^2 as fixed, the parameter q reflects how informative the signal s_a is with $q = 0$ corresponding to an uninformative signal and $q = 1$ corresponding to a fully informative signal.

There are different ways to interpret a larger q. The most natural interpretation is that it is easier for the agent to learn how to execute a course. Another is that the principal is more *able* when it comes to generating information related to execution. Finally, a larger q can also reflect the *managerial or supervisory competence* of the principal to effectively communicate instructions to the agent.

2.2 Preferences

The principal is risk neutral. Her utility is given by

$$U^P = B_a + (1-\alpha)(b_0 - |x_a - \mu_a|) - w.$$

The agent's utility is given by a CARA utility function

$$U^A = -exp^{-r(\beta_a + \alpha(b_0 - |x_a - \mu_a|) + w)},$$

where $r > 0$ is the coefficient of absolute risk aversion. The agent is risk neutral in the limit as r tends to 0.

A's reservation wage and P's reservation utility is 0. We assume that $b_0 > \sqrt{\frac{2}{\pi}}\sigma$, so that the expected payoff to the principal from execution is positive.[16] This ensures that the participation constraint for the principal does not bind.

2.3 Contracts

The course of action, a, the way it is executed, x_a, and the signal, s_a, are all not verifiable. Thus, all that P can do is offer the agent a fixed wage w. In an extension in Section 5, we allow for a different action (effort) that is verifiable with noise, meaning that it can be contracted on.

2.4 Timing and Information

P, who has access to a pool of agents with varying risk preferences, selects an agent A with a coefficient of absolute risk aversion $r \geq 0$ ($r = 0$ corresponds to a risk neutral agent), and offers this agent a take it or leave it contract that specifies a wage w. If A rejects the contract, both parties get their reservation utilities. If A accepts, P decides which courses of action(s) to instruct the agent on (P can, if she wishes, instruct A on both courses). The agent then selects one the course of action and executes it. Finally, payoffs are realized.

The timing and information are summarized in Figure 1.

3 Efficient Benchmark

We first consider a benchmark setting where the course of action is contractible. P chooses the agent's risk preferences r and the agent's course of action a subject to the agent's participation constraint.

Let $\hat{\mu_a}$ denote the correct way to execute course a conditional on the information that the agent has. Since $x_a = E(\hat{\mu}_a)$ minimizes expected losses, A's certainty equivalent at the time of choosing the course of action can be written as[17]

[16] Note that $\mathbb{E}[|x_a - \mu_a|] = \sqrt{\frac{2}{\pi}}(1-q)\sigma$.

[17] See the Appendix for these derivations.

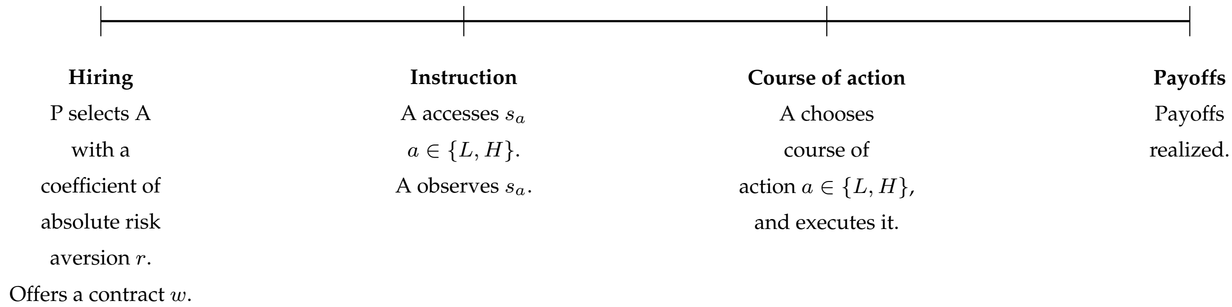

Figure 1: Timeline.

$$CE^A = \beta_a + \alpha b_0 - \frac{r\alpha^2 Var(\hat{\mu_a})}{2} - \frac{1}{r} ln(2\Phi(r\alpha\sqrt{Var(\hat{\mu_a})}) + w$$

where Φ is the distribution function for a standard normal random variable. Notice that since μ_a and s_a are normally distributed, the certainty equivalent above is independent of the realization of the signal and therefore it also corresponds to the agent's certainty equivalent ex-ante.

Suppose P induces the course of action a. Because μ_a and s_a are normally distributed, $Var(\hat{\mu_a})$ is lower when the agent is instructed with the signal on the course a, which in turn increases the agent's certainty equivalent. Thus, it is optimal for P to instruct A on the course she induces. Furthermore, since the participation constraint must bind at the optimum, we can write the participation constraint as[18]

$$w = -\beta_a - \alpha b_0 + \frac{r\alpha^2(1-q)^2\sigma^2}{2} + \frac{1}{r} ln(2\Phi(r\alpha(1-q)\sigma)). \tag{1}$$

Substituting the participation constraint into P's objective function, we can write P's problem as

$$\max_{a,r} \quad B_a + \beta_a + b_0 - (1-\alpha)\sqrt{\frac{2}{\pi}}(1-q)\sigma - (\frac{r\alpha^2(1-q)^2\sigma^2}{2} + \frac{1}{r} ln(2\Phi(r\alpha(1-q)\sigma))).$$

Since the agent's certainty equivalent decreases in the level of the agent's coefficient of absolute risk aversion r, the objective function above is decreasing in r, and is thus maximized by hiring a risk neutral agent.

Finally, because $B > \beta$, P induces H at the optimum. Taking limits in the objective function above (using L'Hopital's rule) as r tends to 0 gives P an expected payoff of

[18] If the participation constraint does not bind at the optimum, P can reduce the wage paid to A and increase her profits, which is a contradiction.

$$B + b_0 - \sqrt{\frac{2}{\pi}}(1-q)\sigma.$$

To summarize, under the first best benchmark, where the course of action is verifiable, P optimally induces the course of action H and hires a risk neutral agent.

4 Optimal Level of Risk Aversion

There are two possibilities to consider: one where P induces the course of action L and the other in which P induces the course of action H. In either case, the participation constraint binds so that equation (1) holds for the induced course. In each of these cases, we solve for the optimal r before turning to the problem of which course the principal chooses to optimally induce.

4.1 Inducing L

Suppose P induces the course of action L. Incentives for A to choose L are stronger when A is selectively instructed only on the course L. Furthermore, instructions on L (because of μ_a being normally distributed) reduce the agent's execution cost and thus increases P's expected utility (through the participation constraint). Thus P optimally instructs A on course L only.

Since the agent optimally chooses $x_a = E(\hat{\mu}_a)$ to minimize expected losses, his certainty equivalent from choosing L is

$$\beta + \alpha b_0 - (\frac{r\alpha^2(1-q)^2\sigma^2}{2} + \frac{1}{r}ln(2\Phi(r\alpha(1-q)\sigma))) + w,$$

and his certainty equivalent from choosing H is

$$\alpha b_0 - (\frac{r\alpha^2\sigma^2}{2} + \frac{1}{r}ln(2\Phi(r\alpha\sigma))) + w.$$

Because $q \geq 0$, $\beta > 0$ and Φ is a strictly increasing function, it follows that the agent prefers L. Thus P's problem (given the binding participation constraint in equation (1)) is

$$\max_r \qquad \beta + b_0 - (1-\alpha)\sqrt{\frac{2}{\pi}}(1-q)\sigma - (\frac{r\alpha^2(1-q)^2\sigma^2}{2} + \frac{1}{r}ln(2\Phi(r\alpha(1-q)\sigma))).$$

As the agent's certainty equivalent is strictly decreasing in r, the optimal solution to the problem is to hire a risk neutral agent. This leads to the following Lemma.

Lemma 1 *Suppose P induces L. Then she optimally hires a risk neutral agent with $r = 0$.*

Her expected payoff in this case is

$$\beta + b_0 - \sqrt{\frac{2}{\pi}}(1-q)\sigma.$$

4.2 Inducing H

Now consider the case where P induces the course of action H. Once again, P optimally instructs A only on the course she induces (course H).

The incentive constraint for the agent to choose the course H is given by

$$-\frac{r\alpha^2(1-q)^2\sigma^2}{2} - \frac{1}{r}ln(2\Phi(r\alpha(1-q)\sigma)) \geq \beta - \frac{r\alpha^2\sigma^2}{2} - \frac{1}{r}ln(2\Phi(r\alpha\sigma)).$$

Define $f(r,q) \equiv \dfrac{r\alpha^2(1-q)^2\sigma^2}{2} + \dfrac{1}{r}ln(2\Phi(r\alpha(1-q)\sigma))$. The incentive constraint can be rearranged as

$$f(r,0) - f(r,q) \quad \geq \quad \beta. \tag{2}$$

There are two cases to consider based on the signal informativeness parameter q (recall that q lies between 0 and 1) when inducing H.

4.2.1 Case 1: $\frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha} \leq q$

Case 1 considers parameters for which the signal is sufficiently informative relative to the agent's bias. For this case, even a risk neutral agent will prefer H to L. To see this, note that taking limits as r tends to 0, allows us to rewrite the incentive constraint in (2) as

$$\sqrt{\frac{2}{\pi}}\alpha q\sigma \geq \beta$$

When the parameters satisfy the condition in Case 1 above, the incentive constraint is satisfied, and thus P can induce the first best outcome for this case.

4.2.2 Case 2: $q < \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}$

Case 2 considers parameters for which the signal is less informative relative to the agent's bias. Note that for r sufficiently small, $f(r,0) - f(r,q) < \beta$, and as r gets arbitrarily large, $f(r,0) - f(r,q)$ tends to infinity. Since $f(r,0) - f(r,q)$ is continuous, there exists a level of risk aversion $r > 0$ such that the agent's incentive compatibility constraint for choosing H holds. Furthermore, since $f(r,q)$ has strictly decreasing differences (see the proof in the Appendix) in r and q, it follows that there is a unique r at which the incentive constraint binds.[19]

[19]Since f has strictly decreasing differences, it follows that $f(r_H,0) - f(r_H,q) > f(r_L,0) - f(r_L,q)$, for $r_H > r_L$.

Definition 1 *Define $\hat{r}$ as the level of r for which the incentive constraint in (2) binds when $q < \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}$, and as zero for all other q.*

Given this definition, the incentive constraint (2) holds if and only if $r \geq \hat{r}$. This variable $\hat{r}$ clearly portrays the main tradeoff that the principal faces. Selecting an agent with a level of risk aversion that is at least as large as $\hat{r}$, allows P to induce the course of action H. However, risk averse agents are costly because they need to be paid a higher risk premium to participate. Thus hiring an agent with a coefficient of absolute risk aversion $\hat{r}$ – the smallest level of risk aversion that helps induce the course H – is the optimal choice for P when she induces course H.

This result is summarized in the lemma below which also looks at how $\hat{r}$ varies with parameters in the model. The proofs of the Lemmas and Propositions that follow are in the Appendix.

Lemma 2 *Suppose P induces the course of action H. Then it is optimal for the principal to hire an agent with $r = \hat{r}$. $\hat{r}$ is weakly decreasing in q and in α, and is weakly increasing in β. The relationship is strict if and only if $q \in [0, \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\})$.*

This lemma says that as instructions get more informative, a lower level of agent risk aversion is necessary to induce the efficient course H. Conversely, less informative instructions render both courses of action roughly similar in terms of their uncertainty for the agent, requiring a higher level of risk aversion to induce H.

4.3 Optimal Course of Action

We now turn to P's choice of which course of action to induce. The following proposition characterizes this choice along with the optimal level of the agent's risk aversion based on the informativeness q of the signal and the agent's bias β.

Proposition 1 *There exist strictly increasing threshold functions $q_0(\beta)$ and $\bar{q}(\beta) = \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}$ with $q_0(\beta) < \min\{1, \bar{q}(\beta)\}$ such that the following hold.*

- *i When instructions are less informative relative to an agent's bias with $q \leq q_0(\beta)$, it is optimal for P to hire a risk neutral agent and induce the course of action L.*
- *ii When instructions have an intermediate level of informativeness relative to an agent's bias with $q > q_0(\beta)$ and $q < \bar{q}(\beta)$, it is optimal for P to hire a risk averse agent with $\hat{r} > 0$ and induce the course of action H. $\hat{r}$ is strictly decreasing in q for this case.*
- *iii When instructions are highly informative relative to an agent's bias with $q \geq \bar{q}(\beta)$, it is optimal for P to hire a risk neutral agent and induce H. P can attain the first best outcome in this case.*

The proposition above says that the principal's optimal choice of which course to choose and which type of agent she should select (based on risk preferences) depends on how informative her instructions are relative to the agent's bias.

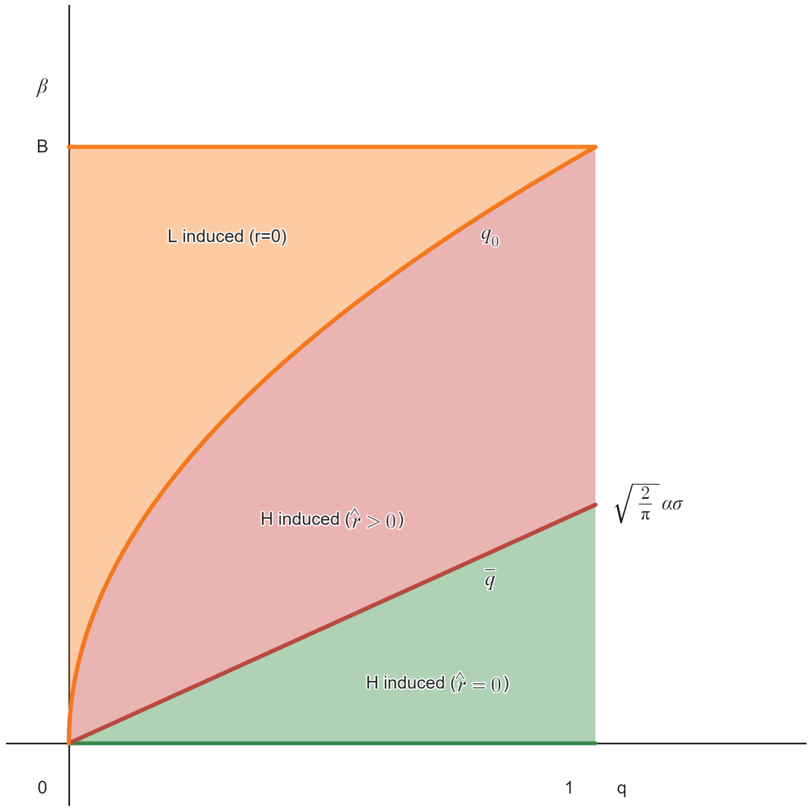

Figure 2: Optimal Course of Action and Agent Risk Aversion.

When the signal is less informative relative to the bias – that is, it is below the threshold, $q_0(\beta)$ – it takes a very high level of risk aversion to induce H. This high level of risk aversion and the associated risk premium make it more profitable for the principal to abandon course H and induce L instead.

As the signal informativeness increases, there is a discontinuous jump in the level of the agent's risk aversion as it becomes optimal to induce H. Finally, for high levels of signal informativeness relative to bias, it becomes feasible to induce H with a risk neutral agent.

Figure 2 depicts these three regions based on the informativeness of the signal q.

The principal's optimal payoff can then be written as

$$EU^{P*} = \begin{cases} \beta + b_0 - \sqrt{\frac{2}{\pi}}(1-q)\sigma & \text{if } q \leq q_0(\beta) \\ B + b_0 - (1-\alpha)\sqrt{\frac{2}{\pi}}(1-q)\sigma - \frac{\hat{r}\alpha^2(1-q)^2\sigma^2}{2} - \frac{1}{\hat{r}}ln(2\Phi(\hat{r}\alpha(1-q)\sigma)) & \text{if } q \in (q_0(\beta), \min\{1, \bar{q}(\beta)\}) \\ B + b_0 - \sqrt{\frac{2}{\pi}}(1-q)\sigma & \text{if } q \geq \bar{q}(\beta) \end{cases}$$

which is strictly increasing in the signal informativeness q.[20]

[20]To see this for the region $q \in (q_0, \bar{q})$, note that the agent's certainty equivalent is strictly decreasing in the level of the coefficient of absolute risk aversion r and $\hat{r}$ is strictly decreasing in q in the interval $(q_0, \bar{q})$.

5 Additional Comparative Statics and Extensions

In this section, we first derive additional comparative static results. We then discuss a few extensions of the model. In the first extension, we allow the agent to exert effort which is contractible. Next, we introduce a cost for the principal to provide instructions.

5.1 Comparative Statics: Execution uncertainty σ^2

So far, we have fixed the execution uncertainty parameter, σ^2, and varied the signal noise parameter, σ^2_ϵ, to see how the quality of instructions affects the selection of an agent based on their risk preferences.

We now reverse the process – we fix the level of noise in the signal and vary the execution uncertainty parameter. When P induces the efficient course H, she once again optimally hires an agent with $r = \hat{r}$. The proposition below characterizes how the optimal level of risk aversion varies with the level of execution uncertainty.

Proposition 2 *Fix σ^2_ϵ. Suppose P induces the course of action H. The optimal level of risk aversion ,$\hat{r}$, is weakly decreasing in σ^2. The relationship is strict if and only if $\sigma < \bar{\sigma}$, where $\bar{\sigma}$ satisfies $(1-\sqrt{\frac{\sigma^2_\epsilon}{\bar{\sigma}^2+\sigma^2_\epsilon}})\bar{\sigma} = \sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}$.*

Proposition 2 says that a lower level of risk aversion is needed to induce the efficient course in a more uncertain environment. That is, uncertainty over how to execute a course plays a similar role to the quality of instructions in providing incentives. In fact, for σ^2 sufficiently large, P can implement the first best by hiring a risk neutral agent.

To gain intuition for the problem, it is useful to write the agent's loss from execution in the certainty equivalent as

$$\frac{r\alpha^2(\Sigma(\sigma,k))^2}{2} + \frac{1}{r}ln(2\Phi(r\alpha\Sigma(\sigma,k)) \tag{3}$$

where

$$\Sigma(\sigma,k) \equiv (1-kq)\sigma \tag{4}$$

and where $k = 1$ corresponds to the case where the agent is instructed and $k = 0$ where the agent is not.

The expressions in (3) and (4) suggest two forces at play. First, with a larger prior variance for executing a course, there is more to learn for the agent – this is captured by the term σ in (4). Information thus has more value for the agent in these settings. Second, the signal is more informative with a higher q when then the prior execution variance is larger.

Both of these effects make it easier for the principal to induce H with a less risk averse agent when the level of execution uncertainty is larger.

5.2 Instructions and Performance Pay for Effort

We now enrich our model. After choosing a course of action and executing it in a certain way (both of which cannot be verified), the agent exerts effort $e \geq 0$ which can be verified with some noise. Effort is costly for the agent but increases the principal's benefit. Returning to our sales example, effort can be interpreted as contacting additional customers.

For this case with effort, the utility functions are:

$$U^{Pe} = \mathcal{B}(a, e) + (1 - \alpha)(b_0 - |x_a - \mu_a|) - w$$

and

$$U^{Ae} = -exp^{-r(\beta_a + \alpha(b_0 - |x_a - \mu_a|) + w - \frac{1}{2}e^2)}$$

where $\mathcal{B}(a, e)$ is the principal's benefit from course a and an effort level e. We assume that $\mathcal{B}(a, e)$ is strictly increasing, twice continuously differentiable, and concave in e, and that it satisfies $\mathcal{B}(H, e) > \mathcal{B}(L, e)$ and $\mathcal{B}_e(H, e) \geq \mathcal{B}_e(L, e)$ for all $e \geq 0$, with $\mathcal{B}(H, 0) = B > \mathcal{B}(L, 0) = 0$.

As mentioned earlier, the key feature of effort that distinguishes it from the course of action chosen is that it can be verified with noise. In particular, P has access to a performance measure of effort $p = e + \epsilon_e$ where ϵ_e is an idiosyncratic error term that is normally distributed with mean 0 and variance σ_e^2. Think of p within the context of our example as current sales. Thus at the start, when P hires A, she can offer A a take it or leave it linear contract of the form $w = t + \gamma p$, where t is a transfer and the slope of the contract γ is the level of incentive pay for the agent.

When P induces the course of action a, her problem (after substituting the participation constraint) is

$$\max_{r,\gamma,e} \quad \mathcal{B}(a, e) + \beta_a + b_0 - (1 - \alpha)\sqrt{\frac{2}{\pi}}(1 - q)\sigma - (\frac{r\alpha^2(1 - q)^2\sigma^2}{2} + \frac{1}{r}ln(2\Phi(r\alpha(1 - q)\sigma))) - \frac{r\gamma^2\sigma_e^2}{2} - \frac{e^2}{2}$$

subject to the effort incentive constraint

$$e = \gamma \tag{5}$$

and subject to (2) when the course H is induced.

At the optimum, the condition $\gamma = \dfrac{\mathcal{B}_e(a, \gamma)}{1 + r\sigma_e^2}$ is satisfied. The following proposition shows how the optimal level of risk aversion and incentive pay vary with the signal informativeness parameter q.

Proposition 3 *Consider a setting where the agent can exert effort. There exists a weakly increasing threshold function $q_{e0}(\beta)$ with $q_{e0}(\beta) \leq 1$ and a strictly increasing threshold function $\bar{q}(\beta) = \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}$ with*

$q_{e0}(\beta) < \bar{q}(\beta)$, such that the following hold.

i When instructions are less informative relative to an agent's bias with $q \leq q_{e0}(\beta)$, it is optimal for P to hire a risk neutral agent and induce the course of action L. The optimal level of incentive pay is constant in q and satisfies $\gamma = \mathcal{B}_e(L, \gamma)$.

ii When instructions have an intermediate level of informativeness relative to an agent's bias with $q > q_{e0}(\beta)$ and $q < \bar{q}(\beta)$, it is optimal for P to hire a risk averse agent with $\hat{r} > 0$ and induce the course of action H. For this case, $\hat{r}$ is strictly decreasing in q, and the optimal level of incentive pay γ, which satisfies $\gamma = \dfrac{\mathcal{B}_e(H, \gamma)}{1 + \hat{r}\sigma_e^2}$, is strictly increasing in q.

iii When instructions are highly informative relative to an agent's bias with $q \geq \bar{q}(\beta)$, it is optimal for P to hire a risk neutral agent and induce H. The optimal level of incentive pay is constant in q and satisfies $\gamma = \mathcal{B}_e(H, \gamma)$. P can attain the first best outcome in this case.

This proposition links instructions in an organization with incentive pay in a very simple way. When instructions are less informative or very informative, hiring a risk neutral agent is optimal. In the former case, this is because P gives up on inducing the efficient course of action H, whereas in the latter case, risk aversion is no longer necessary to induce H, given how effective instructions are. In both cases, the level of incentive pay for the agent does not vary with the informativeness of the signal q.

The level of incentives across these two cases depends on the comparison between the marginal benefit of effort to the principal across both courses of action. When the principal's marginal benefit from the agent's effort is larger for course H relative to course L, then incentives are more high-powered for the high-q case. By contrast, when $\mathcal{B}_e(H, e) = \mathcal{B}_e(L, e)$ for all $e \geq 0$, incentive pay is equally high-powered in both of these extreme cases.

The intermediate case with middling signal informativeness brings us back to the tradeoff between insurance and incentives in Holmström (1979). The key difference is that the terms of this tradeoff, which depends on the degree of risk aversion of the agent, is endogenous in our framework and depends on how informative instructions are.

Figure 3, with the signal informativeness on the horizontal axis and the agent's bias on the vertical axis, depicts the regions where a risk neutral agent is hired and the regions where a risk averse agent is hired, for the case where $\mathcal{B}(a, e) = B_a + e$.

Using (5), the principal's optimal payoff can then be written as

$$EU^{Pe*} = \begin{cases} \mathcal{B}(L, e_L^*) + \beta + b_0 - \sqrt{\frac{2}{\pi}}(1-q)\sigma - \frac{e_L^{*2}}{2} & \text{if } q \leq q_{e0} \\ \mathcal{B}(H, e_H^*) + b_0 - (1-\alpha)\sqrt{\frac{2}{\pi}}(1-q)\sigma & \\ -\frac{\hat{r}\alpha^2(1-q)^2\sigma^2}{2} - \dfrac{1}{\hat{r}} ln(2\Phi(\hat{r}\alpha(1-q)\sigma)) - \frac{e_H^{*2}}{2}(1 + \hat{r}\sigma_e^2) & \text{if } q \in (q_{e0}, \min\{1, \bar{q}\}] \\ \mathcal{B}(H, e_H^*) + b_0 - \sqrt{\frac{2}{\pi}}(1-q)\sigma - \frac{e_H^{*2}}{2} & \text{if } q \in [\bar{q}, 1] \end{cases}$$

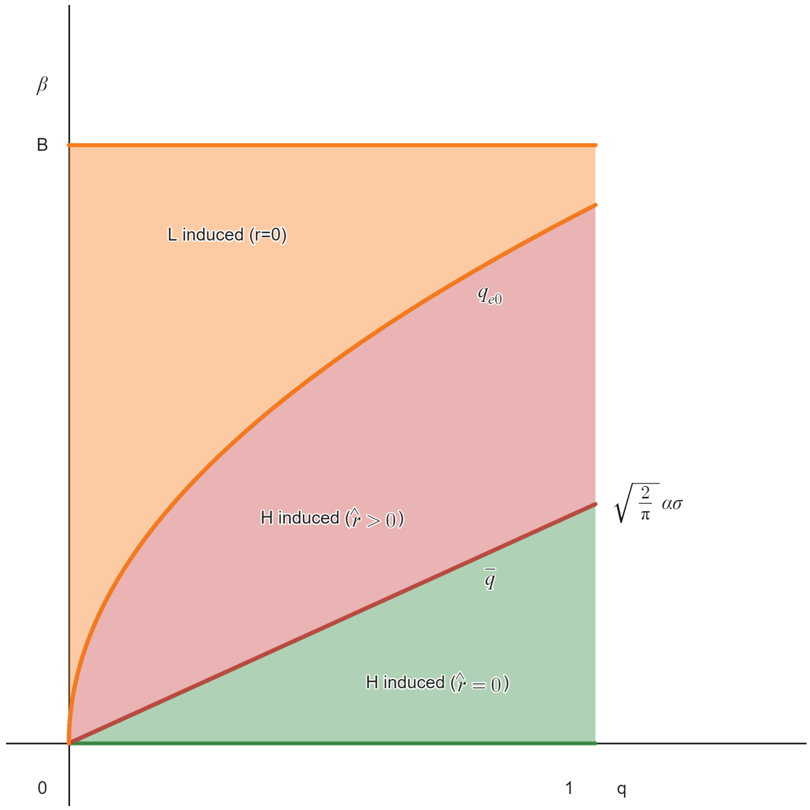

Figure 3: Optimal Course of Action and Agent Risk Aversion: Case with Effort.

where e_L^* satisfies the equation $e = \mathcal{B}_e(L, e)$ and e_H^* satisfies the equation $e = \dfrac{\mathcal{B}_e(H, e)}{1 + \hat{r}\sigma_e^2}$. Once again, the principal's optimal utility EU^{Pe*} is strictly increasing in q.[21]

5.3 Costly Instruction

In this subsection, we add a cost of providing instructions on a course of action. Let $C > 0$ denote this cost.

When P incurs this cost, her optimal expected utility is $EU^{Pe*} - C$.[22] When she does not incur the cost, all she can do is induce L, giving her an expected utility of $\mathcal{B}(L, e_L^*) + \beta + b_0 - \sqrt{\frac{2}{\pi}}\sigma - \frac{e_L^{*2}}{2}$.

Let C^* denote the threshold cost of instructions beyond which instructing the agent is no longer optimal. Thus

$$C^* = EU^{Pe*} - (\mathcal{B}(L, e_L^*) + \beta + b_0 - \sqrt{\frac{2}{\pi}}\sigma - \frac{e_L^{*2}}{2}). \tag{6}$$

Since EU^{Pe*} is strictly increasing in q, the following proposition holds.

[21]For the regions where a risk neutral agent is hired, this follows from the fact that the expected execution cost is strictly decreasing in q. For the intermediate region where $q \in (q_{e0}, \min\{1, \bar{q}\})$, this follows from the fact that e_H^* is strictly decreasing in r, $\hat{r}$ is strictly decreasing in q, and that the agent's certainty equivalent in strictly decreasing in r.

[22]Note that since instruction is costly, it is optimal for P to instruct the agent only for the course of action she induces.

Proposition 4 *The threshold cost of instructions C^* at which point it is no longer optimal to instruct the agent is strictly increasing in q.*

This proposition says that instructions are more likely, the more informative instructions are.

6 Discussion

We now discuss our results. We first assume a specific functional form for the principal's benefit function $\mathcal{B}$ and highlight different organizational configurations that emerge from our model for this form. We then discuss some empirical implications from our model.

6.1 Organizational Configurations

Consider a benefit function for the principal with $\mathcal{B} = B_a + e$. For this case, where the benefit from the course of action and effort are additive, the optimal level of incentives takes the simple form of $\dfrac{1}{1 + r\sigma_e^2}$ which is easy to interpret. In particular, when the agent is risk neutral, the optimal incentive level of $\gamma = 1$ can be interpreted as the agent being a residual claimant. Given our results in Propositions 3 and 4, three main organizational configurations emerge.

6.1.1 $q \in [0, q_{e0}]$: Basic Market

In this region, the inability to contract on a course of action and the low informative content of instructions prevent the socially efficient course of action from being induced. The structure however has market-like features with the agent being a full residual claimant. It is worth pointing out, that these high powered incentives for this case depend on the specific functional form for the principal's benefit assumed above.

6.1.2 $q \in (q_{e0}, \bar{q})$: Conventional organization with a tradeoff between insurance and incentives

In this region, at the threshold of q_{e0}, the organizational structure abruptly jumps to hiring risk averse agents because instructions are effective enough to induce the high course of action. This risk aversion, however, imposes a cost on the organization in terms of risk sharing, as in Holmström (1979). The risk aversion of the agent also leads to low-powered incentive pay. Furthermore, instructions are more likely to be used relative to the ineffective market configuration.

As instructions become more informative, the degree of agent risk aversion drops and incentive pay becomes more high-powered.

To summarize, this region resembles features in Holmström (1979) but there are two main differences. The first is that the degree of agent risk aversion is endogenous. Second, instructions play an important role in the organization.

6.1.3 $q \in [\bar{q}, 1]$: Hybrid organization

This is a hybrid structure with instructions highly likely to be used and with high-powered incentive pay. It thus has features traditionally associated with organizations (instructions) and features associated with markets (residual claimancy). An example of this configuration is franchising, where franchisees receive significant training and instruction on how to run various aspects of the business and have high-powered incentive pay.

6.2 Empirical Implications

There are three main implications that arise from the model. All of these are *conditional* on the efficient course of action being induced.

First, more informative instructions should lead to higher powered incentive pay. This implication is consistent with the literature on complementary human resource practices. In particular, as Ichniowski, Shaw, and Prennushi (1997) show within the context of steel mills, high performing workplaces that are characterized by high levels of training and communication between line managers and operators over "issues related to performance and quality", are also more likely to use incentive pay.

Second, more informative instructions should lead to the selection of more risk neutral workers. While evidence from Ichniowski, Shaw, and Prennushi (1997) does point to a positive relationship between the quality of training and communication and an "extensive selection procedure used to hire new workers", it is unclear what role risk preferences play in the hiring process.

Finally, the relationship between uncertainty and incentive pay depends on the measure of uncertainty used. Incentive pay and uncertainty over the performance measure for effort (σ_e^2) have the standard negative relationship predicted by agency theory. But incentive pay and uncertainty over executing a course (σ^2) have a positive relationship. Allowing for instructions and endogeneity of risk aversion and highlighting the role of instructions is thus perhaps a contributing factor to the ambiguous estimated relationship between uncertainty and the use of incentive pay; see Prendergast (2002).

7 Conclusion

Economists have long emphasized communication patterns between a superior and a subordinate – where the subordinate is told what to do – as a defining feature of a firm (Coase (1937) and Simon (1947)). Yet, formalizing these views has proven to be very difficult.

We develop a model to explore a closely related notion where a superior instructs a subordinate by specifying how the subordinate can carry out the work. We show that instructions on how to execute a course of action can be a useful way to steer the subordinate in a direction that the organization finds more valuable. We also show that these instructions work more effectively when the subordinate is risk averse, so that hiring a risk averse worker can be optimal.

A central insight from our analysis is that risk aversion can add value by realigning interests between a principal and an agent. Highlighting the economic benefits of risk aversion thus complements classical agency theory, which only emphasizes the costs. As in this literature, we assume that an agent's risk preferences are observable. But we go a step further in thinking about how agents are selected based on these preferences. In practice, firms often rely on a variety of instruments – interviews, references, and psychometric tests – to assess preferences at the hiring stage.

Our analysis also makes a formal connection between risk aversion and compliance, a view put forward by scholars of organizations – for example, by William Whyte in his well-known book, "The Organization Man" (see Whyte (1957)).

Finally, the model generates a rich set of organizational configurations that combines features of markets (residual claimancy) with features of firms (instructions) offering a different perspective on firm boundaries.

Appendix

Derivation of the agent's optimal execution choice.

Let $g(\hat{\mu_a})$ be the density function of μ_a conditional on the information that the agent has. Then A's expected utility can be written as

$$\begin{aligned} EU^A &= -\int_{-\infty}^{\infty} exp(-r(\beta_a + \alpha(b_0 - |\hat{\mu_a} - x|) + w))g(\hat{\mu_a})d\hat{\mu_a} \\ &= -exp(-r(\beta_a + \alpha b_0 + w))(\int_x^{\infty} exp(r\alpha(\hat{\mu_a} - x)g(\hat{\mu_a})d\hat{\mu_a} + \int_{-\infty}^{x} exp(r\alpha(x - \hat{\mu_a})g(\hat{\mu_a})d\hat{\mu_a}) \end{aligned}$$

The first order conditions for an interior optimum yield

$$\int_x^{\infty} exp(r\alpha(\hat{\mu_a} - x)g(\hat{\mu_a})d\hat{\mu_a} - \int_{-\infty}^{x} exp(r\alpha(x - \hat{\mu_a})g(\hat{\mu_a})d\hat{\mu_a} = 0,$$

which are satisfied at $x = E(\hat{\mu_a})$ given the symmetry of the density function g. The second order conditions depend on the sign of

$$-\int_x^{\infty} r^2\alpha^2 exp(r\alpha(\hat{\mu_a} - x)g(\hat{\mu_a})d\hat{\mu_a} - \int_{-\infty}^{x} r^2\alpha^2 exp(r\alpha(x - \hat{\mu_a})g(\hat{\mu_a})d\hat{\mu_a} - 2r\alpha g(x),$$

which is negative. Thus $x = E(\hat{\mu_a})$ is the agent's optimal way to execute the course of action a.■

Derivation of the agent's certainty equivalent.

Let $Y = |\hat{\mu_a} - E(\hat{\mu_a})|$. Then Y has a folded normal distribution (see Leone, Nelson, and Nottingham (1961)) with a density function equal to 0 for $y < 0$ and $\frac{\sqrt{2}}{\sqrt{\pi Var(\hat{\mu_a})}} exp(\frac{-y^2}{2Var(\hat{\mu_a})})$ for $y > 0$. Thus, A's expected utility can be written as

$$\begin{aligned} EU^A &= -\mathbb{E}[exp(-r(\beta_a + \alpha b_0 - \alpha y + w))] \\ &= -exp(-r(\beta_a + \alpha b_0 + w))\mathbb{E}[exp(r\alpha y)] \\ &= -exp(-r(\beta_a + \alpha b_0 + w))exp(\frac{r^2\alpha^2 Var(\hat{\mu_a})}{2})2\Phi(r\alpha\sqrt{Var(\hat{\mu_a})}) \\ &= -exp(-r(\beta_a + \alpha b_0 - r\alpha^2\frac{Var(\hat{\mu_a})}{2} + w))exp(r\frac{ln(2\Phi(r\alpha\sqrt{Var(\hat{\mu_a})}))}{r}) \\ &= -exp(-r(\beta_a + \alpha b_0 - r\alpha^2\frac{Var(\hat{\mu_a})}{2} - \frac{1}{r}ln(2\Phi(r\alpha\sqrt{Var(\hat{\mu_a})}) + w)), \end{aligned}$$

where the third line follows from the moment generating function of the folded normal distribution (see Tsagris, Beneki, and Hassani (2014), pp.16) and where the fourth line uses the fact that the log and exponential functions are inverses of each other.■

Cross Partials

i **Proof that $f(r,q)$ has strictly decreasing differences in r and q.**

The cross partial of f with respect to r and q is,

$$\begin{aligned}\frac{\partial^2 f}{\partial r \partial q} &= \alpha^2(1-q)\sigma^2(-1-\frac{\Phi''(r\alpha(1-q)\sigma)}{\Phi(r\alpha(1-q)\sigma)}+\frac{(\Phi'(r\alpha(1-q)\sigma))^2}{(\Phi(r\alpha(1-q)\sigma))^2}) &\text{(A.1)}\\ &= \alpha^2(1-q)\sigma^2(-1+\frac{\Phi'(r\alpha(1-q)\sigma)}{\Phi(r\alpha(1-q)\sigma)}(r\alpha(1-q)\sigma+\frac{\Phi'(r\alpha(1-q)\sigma)}{\Phi(r\alpha(1-q)\sigma)})). &\text{(A.2)}\end{aligned}$$

Define $h(z) \equiv \frac{\Phi'(z)}{\Phi(z)}(z+\frac{\Phi'(z)}{\Phi(z)})$. Then

$$\begin{aligned}h(z)-1 &= \frac{\Phi'(z)}{\Phi(z)}(z+\frac{\Phi'(z)}{\Phi(z)})-1 &\text{(A.3)}\\ &= \frac{\Phi'(z)}{(\Phi(z))^2}(\Phi(z)z+\Phi'(z)-\frac{(\Phi(z))^2}{\Phi'(z)}). &\text{(A.4)}\end{aligned}$$

Notice that $h(z) < 1$ if and only if $\Phi(z)z+\Phi'(z)-\frac{(\Phi(z))^2}{\Phi'(z)} < 0$. When $z=0$, this inequality holds. Furthermore, the derivative of the expression on the left hand side of the inequality above is $-\Phi(z)(1+\frac{\Phi(z)z}{\Phi'(z)})$ which is negative when $z \geq 0$. Thus the cross partial in A.2 is negative so that f has strictly decreasing differences in r and q.■

ii **Proof that f has strictly decreasing differences in α and q.**

The cross partial of f with respect to α and q is,

$$\frac{\partial^2 f}{\partial \alpha \partial q} = -\alpha(1-q)r\sigma^2+\alpha(1-q)r\sigma^2(-1-\frac{\Phi''(r\alpha(1-q)\sigma)}{\Phi(r\alpha(1-q)\sigma)}+\frac{\Phi'(r\alpha(1-q)\sigma))^2}{(\Phi(r\alpha(1-q)\sigma))^2})-\frac{\Phi'(r\alpha(1-q)\sigma)}{\Phi(r\alpha(1-q)\sigma)}\sigma.$$

which is negative from part (i) above.

Proof of Lemma 2.

Consider two possible cases. When $\frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha} < q$, the first best can be implemented so that the optimal level of risk aversion r is set to 0.

Next suppose $q < \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}$. Since the incentive constraint in (2) holds with equality, it follows that

$$\frac{\partial \hat{r}}{\partial \beta} = \frac{1}{f_r(r,0) - f_r(r,q)} \tag{A.6}$$

and

$$\frac{\partial \hat{r}}{\partial q} = \frac{f_q(r,q)}{f_r(r,0) - f_r(r,q)} \tag{A.7}$$

and

$$\frac{\partial \hat{r}}{\partial \alpha} = \frac{f_\alpha(r,q;\alpha) - f_\alpha(r,0;\alpha)}{f_r(r,0;\alpha) - f_r(r,q;\alpha)} \tag{A.8}$$

Since f has strictly decreasing differences in r and q, and strictly decreasing differences in α and q, and since f is strictly decreasing in q, it follows that $\frac{d\hat{r}}{d\beta} > 0$, $\frac{d\hat{r}}{dq} < 0$ and $\frac{d\hat{r}}{d\alpha} < 0$. ■

Proof of Proposition 1.

Part (iii) follows from the fact that the first best is implementable when $\frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha} \leq q$.

To see why parts (i) and (ii) hold, fix σ, β and α, and let $q < \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\}$ It is optimal for P to induce the course L if and only if

$$B - \beta \leq f(\hat{r}, q) - f(0, q) \tag{A.9}$$

where $f(0,q)$ is the limit of $f(r,q)$ as r tends to 0.

Consider the right hand side of (A.9): $f(\hat{r},q) - f(0,q)$. As q tends to 0, $\hat{r}$ tends to infinity, and thus the right hand side of (A.9) tends to infinity. Also, the right hand side of (A.9) is equal to 0 when $q = \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\}$. Since f is continuous, there exists a $q \in (0, \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\})$ such that (A.9) holds with equality.

Furthermore, the right hand side of (A.9) is strictly decreasing in q. To see this, consider two values of q: q' and q'' with $q'' > q'$. Then

$$\begin{aligned} f(\hat{r}(q''), q'') - f(0, q'') &< f(\hat{r}(q'), q'') - f(0, q'') && \text{(A.10)} \\ &< f(\hat{r}(q'), q') - f(0, q'), && \text{(A.11)} \end{aligned}$$

where the first inequality follows from the fact that $\hat{r}$ is strictly decreasing in q and the fact that A's certainty equivalent is strictly decreasing in r. And the second inequality holds because f has strictly decreasing differences in r and q. Thus there is a unique $q_0 \in (0, \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\beta\})$ such that (A.9) holds with equality.

To see how q_0 varies with β, notice that

$$\frac{dq_0}{d\beta} = -\frac{f_r\frac{\partial \hat{r}}{\partial \beta} + 1}{f_q(\hat{r}, q) - f_q(0, q) + f_r\frac{\partial \hat{r}}{\partial q}}.$$

Since f is strictly increasing in r and strictly decreasing in q, since $\frac{\partial \hat{r}}{\partial \beta} > 0$ and $\frac{\partial \hat{r}}{\partial q} < 0$ (from (A.12) and (A.8)), and since f has strictly decreasing differences in r and q, it follows that the derivative above is positive.■

Proof of Proposition 2.

Consider two possible cases. When $\sigma \geq \bar{\sigma}$, it follows that $\frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha} < q$. Thus, the first best can be implemented so that the optimal level of risk aversion r is set to 0.

Next suppose $\sigma < \bar{\sigma}$. Note that $f(r, q)$ can be written as

$$F(r, \Sigma(\sigma, k)) \equiv \frac{r\alpha^2(\Sigma(\sigma, k))^2}{2} + \frac{1}{r}ln(2\Phi(r\alpha\Sigma(\sigma, k)),$$

where $\Sigma(\sigma, k) \equiv (1 - kq)\sigma$. Then from (2),

$$\frac{\partial \hat{r}}{\partial \sigma} = \frac{-((F_\Sigma\Sigma_\sigma)|_{k=0} - (F_\Sigma\Sigma_\sigma)|_{k=1})}{F_r|_{k=0} - F_r|_{k=1}}. \tag{A.12}$$

Consider the denominator. The cross partial of F with respect to r and Σ is $\alpha^2\Sigma(1 - h(r\alpha\Sigma))$ which is positive from (A.4). Since Σ is strictly decreasing in k, the denominator is positive.

Next, consider the numerator $-((F_\Sigma\Sigma_\sigma)|_{k=0} - (F_\Sigma\Sigma_\sigma)|_{k=1})$. First note that the second order partial derivative $F_{\Sigma\Sigma} = \alpha^2 r(1-h(r\alpha\Sigma))$, which is positive from (A.4). Thus, $F_\Sigma|_{k=0} > F_\Sigma|_{k=1}$. Furthermore, $\Sigma_\sigma|_{k=0} = 1 > 1 - \sigma q'(\sigma) = \Sigma_\sigma|_{k=1}$. So the sign of the numerator is negative. Hence, $\frac{\partial \hat{r}}{\partial \sigma}$ is negative.■

Proof of Proposition 3.

Part (iii), once again, follows from the fact that the first best is implementable when $\frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha} \leq q$.

To see why parts (i) and (ii) hold, fix σ, β, and α, and let $q < \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\}$. Let e_a^* denote the optimal level of effort induced by the principal when she induces the course of action a.

Given (5), it is optimal for P to induce the course L if and only if

$$\mathcal{B}(H, e_H^*) - \frac{e_H^{*2}}{2}(1 + \hat{r}\sigma_e^2) - (\mathcal{B}(L, e_L^*) - \frac{e_L^{*2}}{2}) - \beta \leq f(\hat{r}, q) - f(0, q) \tag{A.13}$$

where $f(0, q)$ is the limit of $f(r, q)$ as r tends to 0.

From the proof of Proposition 1, we know that the right hand side of (A.13) is strictly decreasing in q, tends to infinity as q tends to 0, and is equal to 0 when $q = \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\}$.

Consider the left hand side of (A.13). As q tends to 0, $\hat{r}$ tends to infinity, and thus from (5), e_H^* tends to 0 in the limit. Thus the left hand side of (A.13) tends to $B - (\mathcal{B}(L, e_L^*) - \frac{e_L^{*2}}{2}) - \beta$ as q tends to 0. Also, because $\hat{r}$ is strictly decreasing in q and the principal's expected utility is strictly decreasing in r, the left side of (A.13) is strictly increasing in q.

Thus there are two possible cases depending on the value of the left hand side of (A.13) at $q = \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\}$. First, when this value is positive, there exists a unique $q_{e0} \in (0, \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\})$ such that (A.13) holds, and above which course H is induced. As $\hat{r}$ is strictly decreasing in q, the optimal level of incentives is strictly increasing in q for this case. Second, when the left hand side of (A.13) is negative, the principal optimally induces course L for all $q \in [0, \min\{1, \frac{1}{\sigma}\sqrt{\frac{\pi}{2}}\frac{\beta}{\alpha}\}]$.

To see how q_{e0} varies with β, notice that when (A.13) holds with equality, the derivative $\frac{dq_{e0}}{d\beta}$ can be simplified and written as

$$\frac{dq_{e0}}{d\beta} = \frac{1 + (\frac{1}{2}e_H^{*2}\sigma_e^2 + f_r)\frac{\partial \hat{r}}{\partial \beta}}{-\frac{1}{2}(e_H^{*2}\sigma_e^2 + f_r)\frac{\partial \hat{r}}{\partial q} + f_q(0, q) - f_q(\hat{r}, q)}.$$

Since $\frac{\partial e_H^*}{\partial r} < 0$, since $\frac{\partial \hat{r}}{\partial \beta} > 0$ and $\frac{\partial \hat{r}}{\partial q} < 0$ (from (A.12) and (A.8)), and since f has strictly decreasing differences in q and r, the sign of the derivative above is positive.■

References

ANDERSON, E. (1985): "The Salesperson as outside Agent or Employee: A Transaction Cost Analysis," *Marketing Science*, 4(3), 234–254.

ANDERSON, E., AND D. C. SCHMITTLEIN (1984): "Integration of the Sales Force: An Empirical Examination," *RAND Journal of Economics*, 15(3), 385–395.

ARROW, K. (1974): *The Limits of Organization*. W.W. North & Co., New York.

BAKER, G. (1992): "Incentive Contracts and Performance Measurement," *Journal of Political Economy*, 100, 598–614.

BECKER, G. (1964): *Human Capital: a Theoretical and Empirical Analysis, with Special Reference to Schooling*. University of Chicago Press, Chicago.

CARROLL, G. (2015): "Robustness and Linear Contracts," *American Economic Review*, 105(2), 536–563.

CHASSANG, S. (2010): "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts," *American Economic Review*, 100(1), 448–465.

CHE, Y.-K., AND N. KARTIK (2009): "Opinions as Incentives," *Journal of Political Economy*, 117(5), 815–860.

COASE, R. (1937): "The Nature of the Firm," *Economica*, 4(16), 386–405.

GROSSMAN, S., AND O. HART (1986): "The Costs and Benefits of Integration: A Theory of Lateral and Vertical Integration," *Journal of Political Economy*, 94, 691–719.

HART, O., AND J. MOORE (1990): "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98, 1119–1158.

HERMALIN, B. (1998): "Toward an economic theory of leadership: Leading by example," *American Economic Review*, 88(5), 1188–1206.

HOLMSTRÖM, B. (1979): "Moral Hazard and Observability," *Bell Journal of Economics*, 10, 74–91.

——— (1999): "The Firm as a Subeconomy," *Journal of Law, Economics and Organization*, 15(1), 74–101.

HOLMSTRÖM, B., AND P. MILGROM (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55, 303–328.

——— (1991): "Multitask Principal-Agent Analysis: Incentive Contracts, Job Ownership and Asset Design," *Journal of Law, Economics and Organization*, 7, 24–52.

HOLMSTRÖM, B., AND J. ROBERTS (1998): "The Boundaries of the Firm Revisited," *Journal of Economic Perspectives*, 12(4), 73–94.

ICHNIOWSKI, C., K. SHAW, AND G. PRENNUSHI (1997): "The Effects of Human Resource Management Practices on Productivity," *American Economic Review*, 86, 291–313.

JIANG, Y., AND S. PRASAD (2021): "Of Opinionated Bosses and Yes Men," Working paper.

KVASOV, D., AND K. MEAGHER (2022): "A Knowledge Based theory of Intrafirm Production with Applications to Mergers, R&D and New Products.," Working paper.

LEONE, F. C., L. S. NELSON, AND R. B. NOTTINGHAM (1961): "The folded normal distribution," *Technometrics*, 3(4), 543–550.

NELSON, R., AND S. WINTER (1982): *An Evolutionary Theory of Economic Change.* Belknap, Cambridge.

NELSON, R., AND S. WINTER (2002): "Evolutionary Theorizing in Economics," *Journal of Economic Perspectives*, 16(2), 23–46.

PRENDERGAST, C. (2002): "The Tenuous Trade-off between Risk and Incentives," *Journal of Political Economy*, 110(5), 1071–1102.

RAITH, M. (2021): "Employment as a Relational Obligation to Work.," Working paper.

RANTAKARI, H. (2021): "Simon Says? Equilibrium Obedience and the Limits of Authority.," Working paper.

ROBERTS, J., AND G. SALONER (2013): *In R. Gibbons, and J. Roberts, eds., The Handbook of Organizational Economics*.chap. Strategy and Organization. Princeton University Press, Princeton.

SIMON, H. (1947): *Administrative Behavior*. The Free Press, New York.

——— (1991): "Organizations and Markets," *The Journal of Economic Perspectives*, 5(2), 25–44.

TSAGRIS, M., C. BENEKI, AND H. HASSANI (2014): "On the Folded Normal Distribution," *Mathematics*, 2(1), 12–28.

VAN DEN STEEN, E. J. (2009): "Authority versus Persuasion," *American Economic Review Papers and Proceedings*, 99(2), 448–453.

WHYTE, W. H. J. (1957): *The Organization Man*. Garden City, NY: Doubleday.