Endogenous Correlation and Moral Hazard

Pierre Fleckinger* & René Kirkegaard†

August 2017
PRELIMINARY AND INCOMPLETE

Abstract

We study a contracting problem in which the agent's action is two-dimensional. First, the agent controls the marginal distribution of a performance signal. Second, the agent manipulates the correlation between this performance measure and some exogenous signal like the business cycle. The model allows us to revisit the Informativeness Principle, which originally assumes that the agent's action is one-dimensional and the information structure fixed. In the latter model, the principal is better off the higher the exogenous correlation is between the two signals. However, in the model with endogenous correlation, the principal may be better off incentivizing the agent to lower the correlation between the two signals. The optimal contract then appears less sensitive to exogenous signals than suggested by the standard approach. We examine the difference in the structure of the optimal contract in the two models. Several other applications of the new model are pursued as well.

*MINES ParisTech, PSL Research University & Paris School of Economics, pierre.fleckinger@mines-paristech.fr.

†University of Guelp, rkirkega@uoguelph.ca.

Contents

1 Introduction

Endogenous information structures in incentive problems has been discussed in various contexts.[1] However, most of the literature on the moral hazard problem assumes an exogenous information structure under which the principal devises an optimal incentive formula. We here take first steps toward relaxing this assumption and study the new problem of endogenous dependence structure, and shows its relevance in several applications. For instance, incentive schemes in reality are often plagued by various forms of gaming: typically, in addition to providing a valuable input, the agent is tempted to manipulate ex-post the information flow on which his performance evaluation is based. The information structure used for contracting is hence endogenous. Another example is a situation in which the agent is an advisor to the principal. Then, almost by by definition, the information is endogenous: the agent must be incentivized to provide accurate information, i.e. to offer an information structure to the principal. We study a class of models that encompasses these applications, emphasizing in particular the importance of endogenous correlation of signals used in (optimal) incentive schemes. This allows to revisit Holmström's informativeness principle from a new angle and opens a tractable new toolbox for contract theory.

1.1 Motivation: from exogenous to endogenous correlation

The classic view. Aligning the interest of an agent with that of the principal to overcome moral hazard is the aim of incentive contracting. Optimal contracting relates the observable output of an agent to how much the agent should be rewarded. A broad insight in management essentially asserts that the more closely the output tracks the unobservable input of the agent, the less costly are incentives. Simply put, all information that is relevant to inferring effort ex-post should be used in the incentive formula: this is the Informativeness Principle due to Holmström (1979). The precise statistical sense in which this is true is called the sufficient statistics result. For concreteness, x will denote the observable output and y the additional signal from now on.

[1] Most notably from an ex-ante information perspective, typically in an adverse selection setting with information gathering.

In some instances, the Informativeness Principle is immediately intuitive: if y consists of an additional noisy signal of the effort, then it should be used to increase the confidence that the correct effort is taken in equilibrium. In other instances, this yields a slightly more subtle insight: using an additional signal y in the incentive formula can be useful even if does not provide any direct information on the behavior of the agent. Indeed, through the way this additional signal y *correlates* with the direct signals x, it may still indirectly contain some valuable information.[2] This observation is the key to the incentive theory of benchmarking: introducing a (stochastic) external benchmark (that the agent cannot manipulate) helps reducing incentive costs.

Applications of this benchmarking theory are found in yardstick competition for regulated firms, incentive formula for CEOs that are relative to the industry and bonuses for traders that depend on overall market performance. It is also the same logic that commands using external factors in the incentive formula to reduce the variation of reward that is not attributable to the agent's hidden action.[3]

To be concrete, consider the case of the CEO of a firm which heavily consumes energy.[4] There are a number of contractable indicators available (such as profit, market value, turnovers, operating costs etc.) pertaining directly to the firm. Should external indicators such as the price of energy also be used in the incentive contract? The CEO has no impact on the price of energy, which implies that in isolation this signal does not contain any intrinsic information on what the CEO is actually doing. But if the cost of input increases, the profit of the firm decrease without the CEO being responsible for it. Hence the contract could insure the CEO against this exogenous noise, by using the energy price as an index of an exogenous tough environment, thereby tracking the actual performance more closely. In other cases, a signal akin to the price of energy is not directly available, but for instance (the evolution of) the market value of firms evolving in similar environments is, and can be used as a proxy for exogenous factors that have affected the industry and are beyond the CEO's control.

[2]This seemingly goes against the Controllability Principle asserting that agents should be held accountable only for results they can control.

[3]Hence reconciling the Informativeness and the Controllability Principles.

[4]An interesting account of incentive theory as applied by BP can be found in Roberts (2004). Roberts advocates the filtering out of oil price from the incentive formula.

The alternative view. A partially discordant view has received growing attention in the management literature (Lambert, 2001). The underlying idea can be summarized in the example as: what if the CEO can also influence how dependent on energy the firm is? Then the price of energy is also relevant for incentives, but for a somewhat different reason now: the shareholders want the CEO to take it into account rather than be insulated from its variations. In such a case, neutralizing the price of energy in the incentive formula does not encourage the CEO to pursue a strategy that adapts to external circumstances. Obviously, the consequences for the optimal incentive formula are different when this adaptation aspect is taken into account. In particular, it is less desirable to filter out the effect of y. One way of modeling this has been proposed in the literature: maybe the CEO can obtain early information on the realization on y and change his behavior accordingly (Larmande and Ponssard, 2007; Feriozzi, 2011). Clearly, this is a substantial conceptual step because the incentive problem now features adverse selection on top of moral hazard. We offer a somewhat more direct view on this issue, by modeling the adaptation effort as an explicit choice of correlation in the CEOs strategy.

Such a modeling assumption is motivated by a simple observation: if some factor is part of the incentive formula, then a rational agent will try to do something about it. In the example, even though the CEO can not influence the price of energy, he can still change how sensitive to the price of oil the firm's strategy is. Hence he will try to influence how correlated with the price of energy the profit (or any relevant index used in the incentive formula) is. Taking into account that correlation might be influenced by the agent is hence important to understand how the agent can game the system. Beyond this gaming issue, adaptation is sometimes even an objective of the principal, and implementing correlation is a somewhat different incentive problem than in the standard case where the only task of the agent is to stochastically increase x. Our approach enables a better understanding of how to implement adaptation.

Our modeling strategy results in a multitask problem for the agent, even though he still does not control the *marginal distribution* of y. The dependence between x and y is the second dimension of moral hazard in that case. We demonstrate than in some cases, there is a direct correspondence between choosing correlation and choosing independently strategies that are conditional on y. The corresponding equivalent principal-agent problem features a cost of adaptation. Hence we formalize that endogenous correlation comes at a cost of adaptation that depends on how different the course of actions chosen are depending on the realization of y.

Applications. Accounting for the possibility that the agent manipulates correlation opens a range of new results and predictions that can be useful in explaining pay for luck, asymmetric benchmarking and the absence of negative correlation between risk and incentives in the data.

For instance, we confirm that while benchmarking is extremely useful in many instances, some of the limits that have been pointed out (Traders, herding, sensitivity of the formula, pay for luck) are better understood when the agent can tamper with correlation of the benchmark. Beyond misspecification of the benchmark, and control over the incentive formula (control of the board etc.), it is important to understand how the agent can game the incentive system in this way. In order to properly model the situations just mentioned, taking into account how the agent can influence the dependence to the benchmark was a key missing link in the existing literature.

In other applications, it is even the essential part of the agent's job to control the correlation. Consider the extreme case of a forecasting exercise. The agent is to forecast, say, the price of energy, and x is his report. Then the better the agent works, the more correlated x and y are. We here establish a formal link between information gathering and adverse selection models on the one hand, and the dependence structure in the moral hazard problem on the other hand. This offers a reduced model for moral hazard in forecasting.

1.2 Related literature

[TO BE COMPLETED]

2 Exogenous correlation and the value of information

At the heart of the paper is a parsimonious model that allows us to operationalize correlation between exogenous and endogenous signals in a manner so tractable that it becomes possible to contrast the cases where this correlation is or is not under the agent's control.

In this section, we start the analysis by introducing the core model and giving benchmark results when correlation is exogenous. We then illustrate why the fact that correlation is endogenous matters for the principal and how this affects optimal contracting.

2.1 Information, production and contracts

Formally, the agent controls the marginal distribution of $x \in \{0,1\}$, the performance measure, which is fully characterized by p, the probability of $x = 1$. For instance in a typical model, x is an ex-post profit and effort increases the occurrence of a high profit, and $x = 1$ represent hence a good outcome. We do not however impose any ordered structure, only applications will determine the relevant interpretation. In turn, y is a binary exogenous signal, a measure over which the agent does not have control: $y = 1$ occurs with a given probability q. The signal y could represent the price of oil for a company like BP, the benchmark against which the agent can be compared in a yardstick competition environment, or the state of the business cycle. We will denote a typical realization of (x, y) by $(i, j) \in \{0,1\}^2$, and the corresponding probability by $P_{ij} = Prob[x = i, y = j]$. The two variables x and y can be correlated, through the parameter γ, that will be the focus of our analysis. If $\gamma = 0$, the two random variables are independent. An increase in γ makes the two variables more highly correlated.[5] The following table summarizes the information structure, which each cell containing the corresponding P_{ij}:

$y = 1$	$(1-p)q - \gamma$	$pq + \gamma$
$y = 0$	$(1-p)(1-q) + \gamma$	$p(1-q) - \gamma$
	$x = 0$	$x = 1$

Importantly, note that γ has no impact on the marginal distribution of x, which is entirely determined by p. Note also that the information structure is exhaustively characterized by the vector (p, q, γ): there are only three degrees of freedom since probabilities sum to 1.

[5] An increase in γ consists of what Epstein and Tanny (1980) term a correlation increasing transformation (CIT) in a more general model. Note that the feasibility conditions $0 \leq P_{ij} \leq 1$ must of course hold, which constrains feasible (γ, p) pairs in a way that we make explicit in later sections.

The agent controls the distribution of signals at a cost $c(p, \gamma)$ that for convenience is twice continuously differentiable and convex. However, we emphasize that at this point costs are not necessarily assumed to be monotonic in p and γ. In fact, Section 4 contains applications where it is reasonable to assume that costs are non-monotonic in at least one of the variables. We introduce additional assumptions only in specific applications.

Finally, the wage to the agent will be denoted w, and can be conditional on both realizations of x and y. The vector of conditional wages is hence denoted $\{w_{11}, w_{10}, w_{01}, w_{00}\}$, where the first index pertains to x and the second to y.

2.2 Exogenous Correlation and the Informativeness Principle in the standard model

We here consider first a benchmark case in which the agent cannot influence the correlation, hence it is given at some level γ, and does not enter the agent's costs, i.e. $c(p, \gamma) \equiv k(p)$. In addition, to stay in line with classic moral hazard models, we will for now restrict attention to an economically "standard" model, in which the principal prefers stochastically higher production, which requires effort from the agent. This corresponds to the following set of assumptions:

Definition 1 *A* ***standard model*** *features the following assumptions:*

1. *The principal's gross payoff is increasing in x (and hence in p).*

2. *The principal's gross payoff does not depend on y (and hence neither on q nor on γ).*

3. *The agent's cost is increasing in p.*

Hence in a standard model, the principal directly cares only about x, but neither about y nor about how x and y correlate.[6] If y is used in the contract, this is hence purely for its informational content. Note that we do not require exogenous correlation in this definition.

We will further assume that both players are risk-neutral, that the agent is protected by limited liability (nonnegative wages) and that his participation his ensured.[7] We refer

[6]Typically, the principal can for instance maximize $\mathbb{E}[x] = p$.

[7]For instance if the outside option is lower than the liability, here 0.

to this in short as ”limited liability” in the following. As is well-known, the issue for the principal in such a case is that the agent will optimally receive a limited liability rent, which can be decreased by using extraneous information, provided the additional signal correlates non-trivially with the agent's performance.[8]

In a standard model with limited liability, the problem of the principal can be decomposed into two stages (Grossman and Hart, 1983). Since the principal is risk-neutral, for any given p he wishes to implement, he should optimally do that in the following cost-minimizing way:

$$\min_{w_{ij}} \sum_{i,j} P_{ij} w_{ij}$$

$$s.t. w_{ij} \geq 0 \tag{LL}$$

$$\sum_{i,j} \frac{\partial P_{ij}}{\partial p} w_{ij} = c_p, \tag{IC$_p$}$$

where, since c is convex and the expected wage linear in p, the first-order approach is valid and the incentive constraint has been replaced by its first-order condition, which defines the unique p for any wage scheme. Then, as is well known, the likelihood ratios $\frac{\partial P_{ij}}{\partial p} / P_{ij}$ are the key determinants of the optimal incentive scheme: Holmström (1979) sufficient statistics result asserts that the signal y should be used if and only if these likelihood ratios depend on j. It is straightforward to prove that it is the case here if and only if $\gamma_0 \neq 0$.

As a result, we obtain the following properties of the optimal incentive scheme:

Proposition 1 *Consider a standard model with limited liability. When correlation is exogenous, the optimal incentive scheme for implementing an interior p is such that:*

- *for any γ, $w_{01} = w_{00} = 0$,*
- *if $\gamma < 0$, $w_{11} > 0$ and $w_{10} = 0$,*

[8]In Holmström (1979) setting, the agent is risk-averse and the usual interpretation is that the additional signal helps reducing the risk borne by the agent and hence improves the risk/incentives trade-off. In the limited liability model with a risk agent the parallel is that the additional signal helps reducing the rent to the agent, and hence improves the rent/efficiency trade-off. We show in a later section that the insights obtained in the limited liability case carry over fully to the case of a risk-averse agent.

- *if* $\gamma > 0$, $w_{10} > 0$ *and* $w_{11} = 0$,
- *if* $\gamma = 0$, $w_{11}, w_{10} \geq 0$ *with* $qw_{11} + (1-q)w_{10} = c_p$.

Proof. First, we associate the nonnegative Lagrange multipliers λ_{ij} to the limited liability constraints and μ to the incentive constraint to form the Lagrangian. By the optimality condition for w_{ij}, one has $\frac{\partial P_{ij}}{\partial p}/P_{ij} = 1/\mu - \frac{\lambda_{ij}}{\mu P_{ij}}$, and by the complementary slackness condition of the limited liability constraint $\lambda_{ij} w_{ij} = 0$, it is straightforward to see that w_{ij} is positive only if the corresponding likelihood ratio $\frac{\partial P_{ij}}{\partial p}/P_{ij} = 1/\mu$, and is hence the highest among all pairs (i, j). This implies that w_{00} and w_{01} which corresponding likelihood ratios are always negative have to be zero. Then it remains to compare when w_{11} and w_{10} are associated to the highest likelihood ratio, which comparison leads to $\frac{\partial P_{11}}{\partial p}/P_{11} - \frac{\partial P_{10}}{\partial p}/P_{10} = \frac{q}{pq+\gamma} - \frac{1-q}{p(1-q)-\gamma} \geq 0 \Leftrightarrow \gamma \leq 0$. When $\gamma = 0$, the principal can spread arbitrarily the incentive weight between the realizations of y, which are not informative (clearly, risk-aversion would pin down equal wages here). ■

A few comments are in order regarding this proposition. First, the limited liability model allows a simple characterization of incentive schemes, since generically only one of the wages is positive, the one associated with the highest likelihood ratio. Second, the result has a usual flavor that y should be used as a competitive benchmark ($w_{10} > w_{11}$) when there is positive correlation. To illustrate, suppose that y is the business cycle and correlation is positive. Then a high performance for the firm is more likely when the market is good, hence a high performance should be discounted in such favorable circumstances.

2.3 Comparing Information Systems

The Informativeness Principle says that the exogenous information y should be used to lower implementation costs if it is freely available and informative, i.e. $\gamma \neq 0$ in our setting. However, it does not rank different information systems. The ranking of information systems has been taken up by a subsequent literature, as exemplified by Grossman and Hart (1983) and Kim (1995). This literature assumes that the agent's cost of productive effort is independent of the information system.

In this section we rank information systems within the confines of the standard model with limited liability. The deeper connection to Grossman and Hart (1983) and Kim (1995) is detailed in the extensions in Section 5.

Proposition 1 allows us to explicitly calculate the implementation costs for any given γ when p is interior. Letting $C(p,\gamma)$ denote these costs,

$$C(p,\gamma) = \begin{cases} \left(p - \frac{\gamma}{1-q}\right) c_p & \text{if } \gamma \geq 0 \\ \left(p + \frac{\gamma}{q}\right) c_p & \text{if } \gamma \leq 0 \end{cases}.$$

It is immediately obvious that under the assumption that c_p is independent of γ, implementation costs are strictly decreasing as we move further and further away from $\gamma = 0$ in either direction. The assumption holds if the agent's cost of productive effort does not depend at all on γ or, more generally, if $c(p,\gamma)$ takes the form $c(p,\gamma) = k(p) + \kappa(\gamma)$ such that $c_{p\gamma} = 0$. In either case, the incentive compatibility constraint is unaffected by changes in γ and the limited liability constraint by assumptions renders the participation constraint slack regardless of γ. Then, the principal is able to exploit the improved information implied by the greater (positive or negative) correlation between x and y when γ moves away from $\gamma = 0$. This conclusion is recorded in the following corollary.

Corollary 1 *In a standard model with limited liability, exogenous correlation, and $c_{p\gamma} = 0$, the implementation cost is single-peaked in γ and maximized at $\gamma = 0$.*

This result implies that the principal is better off the more extreme the correlation between x and y is, at least under the assumption that $c_{p\gamma} = 0$. However, going forward, we want the model to be able to accommodate richer interactions between p and γ, allowing for $c_{p\gamma} \neq 0$. A change in γ now affects the incentive compatibility constraint. Nevertheless, the main conclusion of Corollary 1 still holds under a relatively mild regularity assumption on c_p. Specifically we assume that c_p is log-concave in γ. This assumption allows c_p to be convex in γ but rules out that the convexity is too extreme. The assumption also implies that $c_{p\gamma}$ changes sign at most once as γ increases (and, if so, from positive to negative).

In fact, c_p is either increasing, decreasing, or single-peaked in γ. In either case, c_p is minimized at a corner. Stated differently, the incentive compatibility constraint is less demanding at one of the corners. Intuitively, this effect then reinforces the effect from

Corollary 1 and implies that implementation costs must be minimized as γ approaches one of its extreme values. Log-concavity adds enough regularity to the problem that it becomes possible to prove that implementation costs are themselves monotone or single-peaked in γ.

Corollary 2 *In a standard model with limited liability, exogenous correlation, and c_p log-concave in γ, the implementation cost is either monotonically increasing or decreasing in γ or single-peaked in γ.*

Proof. See the Appendix. ■

Thus, the worst level of correlation need no longer be at $\gamma = 0$. This is due to the fact that the incentive compatibility constraint now depends on γ. However, it remains the case that the principal prefers γ to be either as high as possible or as low as possible. Indeed, the latter result may hold even when the regularity condition in Corollary 2 is violated. To illustrate, assume that $p \leq \max\{q, 1-q\}$. Then either $P_{11} \to 0$ as γ converges to the lowest possible value or $P_{10} \to 0$ as γ converges to the highest possible value, or both.[9] Recall that P_{11} and P_{10} are the probability that the positive wage will be paid out if $\gamma < 0$ or $\gamma > 0$, respectively. In other words, the implementation costs can be made arbitrarily small with extreme correlation.[10]

In summary, a main conclusion from the standard model with exogenous correlation is that the principal prefers correlation to be extreme. In fact, Corollary 1 is robust to general risk preferences, as we demonstrate in Section 5. However, we prove in the next subsection that there is a tension between the principal and the agent on what the preferred level of correlation is. In Section 3 we turn to endogenous correlation and prove that the added incentive constraint leads to a conclusion that is directly opposed to the above corollaries: The optimal contract may induce a level of correlation that is not extreme but rather surprisingly small.

[9]These properties follow from an examination of the feasible set. See Section 3 for details. If $p > \max\{q, 1-q\}$ then either P_{01} or P_{00} converge to zero as γ converges to one of its extreme values. However, the wage is zero in the state $(0, 1)$ and the state $(0, 0)$ so implementation costs remain positive.

[10]However, implementation costs cannot be made exactly zero. When γ takes an extreme value, (p, γ) is on the boundary of the feasible set and Proposition 1 does not apply in that case.

2.4 Correlation manipulation by the agent

Suppose that the principal has chosen an optimal incentive scheme according to the previous analysis, and suppose for the sake of exposition that $\gamma > 0$ (the symmetric case features exactly the same logic), so that the agent's expected wage is $(p(1-q)-\gamma)w_{10}$. Then, if possible, the agent would want to decrease the correlation, to increase the occurrence of bonus. Suppose, at the extreme, that correlation is free, i.e. the agent can game the incentive scheme by choosing γ at no cost. Then he would choose the most negative correlation feasible (subject to the constraint that the probabilities P_{ij} are well-defined). While this does not affect the incentive power in terms of p, this does increase the implementation cost for the principal, who has to pay the wage w_{10} more often. More generally, the agent always prefers the opposite correlation to what the principal would prefer, creating an extreme tension. This extreme conclusion relies on risk-neutrality, since for a given p, the principal and the agent play a constant-sum game. But it illustrates why the principal should be concerned by endogenous correlation: the agent is tempted to precisely counter the gains from conditioning on exogenous signal, to such an extent that if correlation is freely manipulable the principal cannot use additional signals in contracting as we shall see. Indeed, if manipulating correlation is costless, the principal faces the following additional incentive constraint:

$$\sum_{i,j} \frac{\partial P_{ij}}{\partial \gamma} w_{ij} = 0, \tag{IC$_\gamma^0$}$$

leading to:

$$w_{11} - w_{01} = w_{10} - w_{00},$$

which says that the incentive wedge in both states j should be equal. It is then easy to see that the only solution for the principal is to give up on using y in the contract,[11] since one must have $w_{i1} = w_{i0}$ for all i. We summarize this observation as:

Proposition 2 *In a standard model with limited liability, if the agent can costlessly choose any feasible correlation level, then the optimal contract does not depend on y.*

This corollary illustrates in a stark way that the intuition from the informativeness principle is weakened when the information structure is endogenous. It is important to note

[11] This result is a particular case of the analysis in the next section, hence we omit the proof.

that the main difference with Holmström (1979) model is that the agent here has a two-dimensional action, and it is the resulting additional incentive constraint that precludes the use of the exogenous signal.[12] Observationally, however, the implication is striking: the contract looks incomplete. Despite being informative and contractible, y is not used by the principal.

3 Endogenous correlation and optimal contracts

In this section we reconsider the general model laid down in section 2.1 under the assumption that γ is endogenous. The second part of the section considers an alternative formulation of the model.

3.1 Endogenous correlation

The agent chooses the pair (p, γ). However, feasibility of (p, γ) evidently requires that any of the four (x, y) outcomes occurs with a probability that is between zero and one. Thus, eight conditions must be satisfied. Four of these are immediately eliminated as being redundant.[13] Depending on the relationship between p and q, two of the remaining four conditions can then be eliminated. This process leads to the following succinct characterization of the feasible set. In particular, (p, γ) is feasible if and only if

$$\gamma \geq \begin{cases} -pq & \text{if } p \leq 1-q \\ -(1-p)(1-q) & \text{if } p \geq 1-q \end{cases} \text{ and } \gamma \leq \begin{cases} (1-p)q & \text{if } p \geq q \\ p(1-q) & \text{if } p \leq q \end{cases}.$$

The feasible set is pictured in Figure 1 when $q < \frac{1}{2}$.

As before, a contract stipulates four wages, one for each possible (x, y) combination. Let w_{ij} denote the wage the agent is paid if $(x, y) = (i, j)$. To begin, assume for concreteness that the agent is risk neutral. His expected utility from action (p, γ) is then

$$\begin{aligned} EU(p, \gamma | w_{00}, w_{01}, w_{10}, w_{11}) &= ((1-p)(1-q) + \gamma) w_{00} + ((1-p)q - \gamma) w_{01} \\ &\quad + (p(1-q) - \gamma) w_{10} + (pq + \gamma) w_{11} - c(p, \gamma). \end{aligned}$$

[12] Fleckinger (2012) shows that if correlation is affected by effort, classic results in multi-agent problems are substantially altered. Still, effort is also one-dimensional in his setting.

[13] Clearly, for one probability to be one, all other three must be zero. This cannot happen whenever $0 < q < 1$.

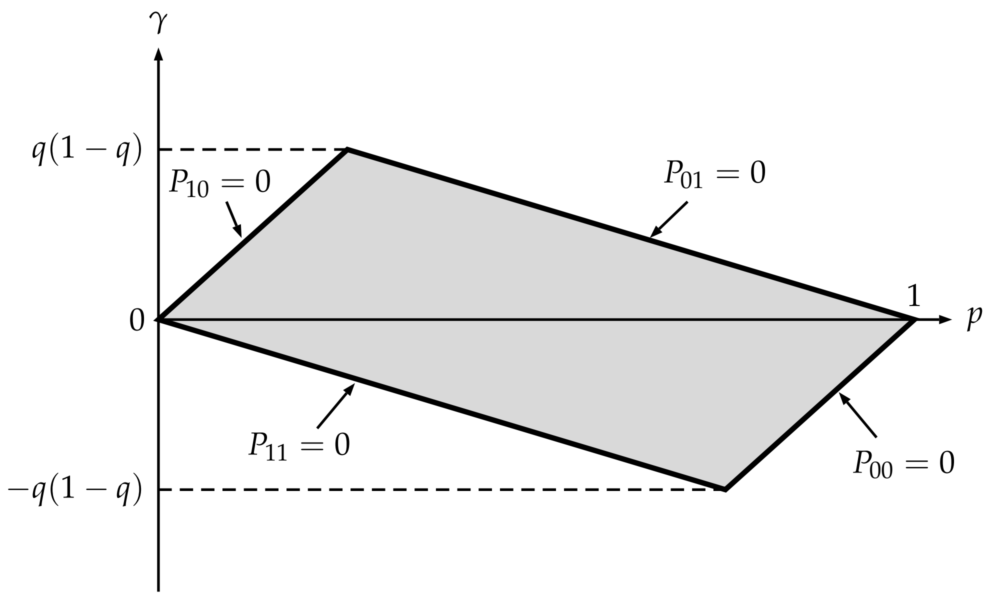

Figure 1: Feasible set.

Note that the first four terms are linear in (p, γ). Hence, expected utility is concave in (p, γ) since $c(p, \gamma)$ is convex in (p, γ). The agent's first-order conditions thus identify the utility maximizing choice of (p, γ), assuming this is interior. Thus, if the principal would like to induce the agent to choose a particular interior (p, γ), then he has to manipulate the contract, $(w_{00}, w_{01}, w_{10}, w_{11})$, to ensure that the agent's first-order conditions are satisfied at that particular (p, γ). In other words, the first-order approach (FOA) is valid. In addition to these conditions, the optimal contract must typically also respect a participation constraint or a limited liability constraint.

It is worth making two remarks at this point. First, note that the above justification of the FOA does not rely on the agent's risk neutrality. The reason is that his expected utility from income remains linear in (p, γ) even if he is risk averse. However, we mainly focus on the risk neutral case in the following. Second, the assumption that $c(p, \gamma)$ is convex is for simplicity only. Kirkegaard (2017, Section 5) considers a slightly more general model in which the cost function is allowed to be non-convex. He proves that the agent's first-order conditions remain sufficient for any implementable (p, γ). The only complication is that not all (p, γ) are implementable. The role of our convexity assumption is thus merely to guarantee that any (p, γ) is implementable.

Assuming the principal wishes to induce an interior action, the agent's first order conditions can be written as

$$\begin{aligned}(1-q)\left[w_{10}-w_{00}\right]+q\left[w_{11}-w_{01}\right] &= c_p(p,\gamma)\\ \left[w_{11}-w_{01}\right]-\left[w_{10}-w_{00}\right] &= c_\gamma(p,\gamma).\end{aligned}$$

The term $w_{10}-w_{00}$ can be thought of as the bonus to a good x outcome given $y=0$. Likewise, $w_{11}-w_{01}$ is the bonus to a good x outcome given $y=1$. Recall that we have made no assumptions about the sign of c_p or c_γ. If, for instance, c_p is negative – such that increasing p marginally reduces costs – then at least one of the "bonuses" must be negative in order to prevent the agent from increasing p. Thus, unsurprisingly, the optimal contract must qualitatively depend on the signs of c_p or c_γ. Indeed, solving the first-order conditions for the bonuses (or penalties) yields

$$\begin{aligned}w_{11}-w_{01} &= c_p(p,\gamma)+(1-q)c_\gamma(p,\gamma)\\ w_{10}-w_{00} &= c_p(p,\gamma)-qc_\gamma(p,\gamma),\end{aligned}$$

where the right hand sides are predetermined by the given (p,γ) that the principal is seeking to induce. Evidently, the cheapest way to achieve a fixed $w_{11}-w_{01}>0$ bonus is to lower w_{11} and w_{01} at the same rate until the $w_{01}\geq 0$ constraint binds. Similarly, the cheapest way to achieve a fixed $w_{11}-w_{01}<0$ penalty is to lower w_{11} until the limited liability constraint binds. Note that these arguments do not rely in any way on the principal's risk preferences. The constraints alone determine the optimal contract. Stated differently, the optimal contract that induces any fixed (p,γ) is independent of the principal's level of risk aversion. The following proposition summarizes the optimal contract for interior actions. Actions on the boundaries are considered in the next section.

Proposition 3 *Consider the limited liability model. The unique optimal contract that induces a given interior (p,γ) is given by*

- $w_{11}=c_p+(1-q)c_\gamma$ *and* $w_{01}=0$ *if* $c_p+(1-q)c_\gamma\geq 0$ *but* $w_{01}=-\left[c_p+(1-q)c_\gamma\right]$ *and* $w_{11}=0$ *if* $c_p+(1-q)c_\gamma\leq 0$*, and*
- $w_{10}=c_p-qc_\gamma$ *and* $w_{00}=0$ *if* $c_p-qc_\gamma\geq 0$ *but* $w_{00}=-\left[c_p-qc_\gamma\right]$ *and* $w_{10}=0$ *if* $c_p-qc_\gamma\leq 0$.

This proposition should be related to proposition 1 where correlation is exogenous. The first observation is that the optimal contract features in general two positive wages, essentially one for each incentive constraints. Note also that it includes proposition 2 where $c_\gamma = 0$ as a special case, which underlines in particular that the content of proposition 2 does not depend on the objective of the principal, but comes primarily from incentive considerations. Also, it is only in cases where effort p is the main driver of costs, i.e. when both $c_p + (1-q)c_\gamma \geq 0$ and $c_p - qc_\gamma \geq 0$, that the shape of the optimal contract is in line with standard results, since then wages are paid only for high x. It is not generically the case, and even in a standard model, when correlation is endogenous, wages can be paid in case of failure ($x = 0$), if correlation is a significant driver of cost. As we will see in other applications, this is a natural property fo optimal incentives for other classes of model.

In light of Corollaries 1 and 2, it is natural to ask what level of correlation the principal optimally induces. We study this in details in the next subsection but offer a motivating example here.

Example 1: Consider the standard model, but now with endogenous correlation. Assume that $c_p + (1-q)c_\gamma > 0$ and $c_p - qc_\gamma > 0$. Then, given any interior (p, γ), it is possible to use Proposition 3 to derive implementation costs. Letting $K(p, \gamma)$ denote these,

$$K(p, \gamma) = pc_p + \gamma c_\gamma.$$

For the purposes of this example, assume that $c(p, \gamma)$ takes the form $c(p, \gamma) = k(p) + \kappa(\gamma)$, as in Corollary 1. Assume moreover that $\kappa(\gamma)$ is strictly convex and minimized at $\gamma = 0$. Then, implementation costs simplify to

$$K(p, \gamma) = pk'(p) + \gamma\kappa'(\gamma).$$

Now fix a value of p that the principal wishes to induce. To do so, he must induce whichever γ minimizes $\gamma\kappa'(\gamma)$. Given the aforementioned assumptions on $\kappa(\gamma)$, we note that $\gamma\kappa'(\gamma) > 0$ if $\gamma \neq 0$. Thus, $\gamma = 0$ is optimal.[14] ▲

[14]Technically, $K(p, \gamma)$ as stated relies on Proposition 3 which assumes (p, γ) is interior. However, in Section 3.3 we prove that $K(p, \gamma)$ is continuous even on the boundary when p is large enough. In this case, $\gamma = 0$ is necessarily globally optimal.

Two remarks to Example 1 are in order. First, it is optimal to induce zero correlation, meaning that x and y turn out to be independent. In contrast, the Informativeness Principle asserts that this is the very worst level of correlation when correlation is exogenous. In fact, Corollary 1 implies that extreme correlation is optimal in the model in Example 1 as long as correlation is exogenous. Comparing Example 1 and Corollary 1 thus illustrates the marked differences in the two models. In particular, we emphasize that it may be in the principal's interest to induce a low level of correlation.

Second, the optimal contract in Example 1 happens to induce the level of correlation that minimizes the agent's costs, $c(p, \gamma)$. This is despite the fact that the participation constraint is always slack, by assumption, in the limited liability model. Thus, $\gamma = 0$ is not chosen to make it cheaper to induce participation, but solely to make it cheaper to ensure incentive compatibility.

In general, however, one should not expect it to be the case that the agent's cost-minimizing level of correlation is induced. The optimal level of correlation is studied in more detail in the next subsection, once again imposing the assumptions of the standard model. It will be shown that the principal will often want to induce a level of correlation that is even smaller than the agent's cost-minimizing level of correlation.

3.2 Correlation Dampening in the Standard Model

We once again consider the standard model with limited liability, but now with endogenous correlation. For simplicity, assume c is strictly convex in γ and let $\gamma_0(p)$ determine the unique value of γ where costs are minimized, given p. We begin by explaining why it is interesting and reasonable to focus on $\gamma = \gamma_0(p)$ as a benchmark. We then move to the main result of the section, which is that the principal under certain conditions finds it optimal to moderate the amount of correlation by implementing a γ value that is smaller than $\gamma_0(p)$. In comparison, recall that Corollaries 1 and 2 imply that if γ is exogenous then the principal is better off with an extreme level of correlation. Finally, we compare optimal contracts under exogenous and endogenous correlation and note that an outsider who thinks γ is exogenous generally overestimates the value of the signal y.

Consider a principal who is unable or unwilling to contract on y. He offers wage w_1 if $x = 1$ and w_0 if $x = 0$. Hence, the agent earns w_1 with probability p and w_0 with probability $1-p$. Note that the agent cannot manipulate his compensation by changing γ. Therefore, the agent always selects the cost-minimizing value of γ, $\gamma_0(p)$. Thus, in this section we ask how γ changes when the principal begins contracting on y.

As before, we fix an interior p that is to be induced, and then ask which γ level the principal should aim for. We confine attention to the limited liability model and assume in addition that $c_p + (1-q)c_\gamma > 0$ and $c_p - qc_\gamma > 0$. The main insight that we want to convey is that the principal often wants to moderate correlation to make it smaller even than $\gamma_0(p)$. We first illustrate this property in a continuation of Example 1.

Example 1 (Continued): As before, assume that $c(p,\gamma)$ takes the form $c(p,\gamma) = k(p) + \kappa(\gamma)$. However, assume now that $\kappa(\gamma)$ is minimized at some $\gamma_0 \neq 0$. Recall that implementation costs for interior (p,γ) are

$$K(p,\gamma) = pk'(p) + \gamma\kappa'(\gamma).$$

By definition, $\kappa'(\gamma_0) = 0$. Hence, $\gamma_0\kappa'(\gamma_0) = 0$. By strict convexity, $\kappa'(\gamma) < 0$ at $\gamma < \gamma_0$ and $\kappa'(\gamma) > 0$ at $\gamma > \gamma_0$. Then, any γ that is strictly between 0 and γ_0 produces $\gamma\kappa'(\gamma) < 0$, whereas any γ outside this interval yields $\gamma\kappa'(\gamma) \geq 0$. Thus, as long as the optimal correlation level to induce, γ^*, is interior, it must be the case that γ^* is strictly between 0 and γ_0. Note that γ^* and γ_0 have the same sign, yet $|\gamma^*| < |\gamma_0|$. Hence, the principal induces a correlation that is smaller than what would be cost-minimizing for the agent. It is only if $\gamma_0 = 0$ (independence is cheapest) that $\gamma^* = \gamma_0$. ▲

Example 1 illustrates that the principal not only wants to move away from extreme correlation but that he may even desire a level of correlation below the agent's cost-minimizing level. To understand the intuition, write the incentive compatibility constraints as

$$(1-q)w_{10} + qw_{11} = c_p \quad (1)$$

$$w_{11} - w_{10} = c_\gamma, \quad (2)$$

utilizing the fact that $w_{00} = w_{01} = 0$. Note that an increase in w_{10} moves the left hand side of the two equations in opposite directions. Increasing w_{10} means that productive effort is rewarded more, but at the same time it punishes effort on γ since increasing γ moves weight from the state $(1,0)$ to the state $(1,1)$. Note also that an important feature of Example 1 is that there is no interaction between p and γ in the agent's cost function, or $c_{p\gamma} = 0$. Thus, a change in γ does not directly impact the incentive compatibility constraint on p since c_p is unaffected. We begin with two preliminary observations. To fix ideas, assume that $\gamma_0 \geq 0$.

First, imagine that $\gamma = 0$. The left hand side of (1) in this case coincides with the expected wage conditional on $x = 1$ (the wage is zero if $x = 0$). This must be constant to give the agent the right incentives for productive effort. Thus, the principal is indifferent between any (w_{10}, w_{11}) that satisfies (1) when $\gamma = 0$. Second, the specific contract that satisfies (1) with $w_{11} = w_{10}$ is equally profitable to the principal for all γ. However, such a contract agrees with (2) if and only if $\gamma = \gamma_0$. Combining the two observations means that the principal is indifferent between the contract that implements $\gamma = 0$ and the contract that implements $\gamma = \gamma_0$.

Implementing a higher γ, $\gamma > \gamma_0$, would mean that w_{11} is paid out increasingly often and w_{10} less often. However, since $c_\gamma > 0$ when $\gamma > \gamma_0$, the only way to induce the agent to increase correlation is to let $w_{11} > w_{10}$. This is unprofitable, because the higher wage would be paid out more often. Thus, implementing $\gamma > \gamma_0$ is inferior to implementing γ_0.

Conversely, imagine implementing negative correlation, $\gamma < 0$. Since $c_\gamma < 0$ when $\gamma < 0$, this requires that $w_{10} > w_{11}$. Note that with negative correlation the higher wage would be paid out more often. Inducing zero correlation is better. Although w_{10} must still exceed w_{11}, the gap must be smaller and the higher wage is paid out less often.

In conclusion, the optimal γ cannot be below zero or above γ_0. Indeed, it must be strictly between the two because γ's in that range require $w_{10} > w_{11}$ for incentive compatibility and, since $\gamma > 0$, the smaller wage is paid out more often.

We next identify easily interpretable sufficient conditions for the correlation dampening effect to hold more generally. In particular, the assumption that $c_{p\gamma} = 0$ is relaxed. Let $\gamma^*(p)$ denote an optimal level of correlation to induce for a fixed p.

Proposition 4 *Assume that $\gamma_0(p) > 0$ is interior and that $c_{p\gamma}(p, \gamma) \geq 0$ for all γ. Then, if $\gamma^*(p)$ is interior, it holds that $\gamma^*(p) < \gamma_0(p)$. Likewise, if $\gamma_0(p) < 0$ is interior and $c_{p\gamma}(p, \gamma) \leq 0$ for all γ then $\gamma^*(p) > \gamma_0(p)$ whenever $\gamma^*(p)$ is interior.*

Proof. As mentioned in the previous section, the cost of implementing any interior

(p, γ) pair is

$$K(p, \gamma) = pc_p + \gamma c_\gamma.$$

The derivative with respect to γ is

$$K_\gamma(p, \gamma) = pc_{p\gamma}(p, \gamma) + \gamma c_{\gamma\gamma}(p, \gamma) + c_\gamma(p, \gamma).$$

Assume first that $\gamma_0(p) > 0$ is interior and that $c_{p\gamma}(p, \gamma) \geq 0$ for all γ. Then, all three terms are positive for any $\gamma \geq \gamma_0(p)$, and the middle term is strictly positive. Hence, $K_\gamma(p, \gamma) > 0$ for all $\gamma \geq \gamma_0(p)$ that are interior. Thus, if $\gamma^*(p)$ is interior it must be the case that $\gamma^*(p) < \gamma_0(p)$. The second part is proven in a similar manner. ∎

Thus, if p and γ are substitutes in the agent's cost function, or $c_{p\gamma} \geq 0$, and $\gamma_0(p) > 0$, then the principal induces a lower level of correlation than $\gamma_0(p)$ (and possibly even negative correlation). By lowering γ, c_p decreases. This makes the incentive compatibility constraint on p less strenuous. Thus, there is a new positive effect of lowering γ below $\gamma_0(p)$ in addition to the effect identified in the discussion of Example 1 above. In short, the principal seeks to "moderate" the amount of correlation compared to the correlation that minimizes the agent's costs.[15] Thus, x and y are less correlated when y is contractible than when it is not. The caveat is that the proposition does not rule out that the principal overdoes it and induces a γ of the opposite sign than $\gamma_0(p)$.

We continue with a comparison of endogenous and exogenous correlation. Assume for concreteness that $\gamma_0(p) > 0$ and, as in the proposition, that $\gamma_0(p) > \gamma^*(p)$. Then, the principal finds it optimal to endogenously induce a level of correlation below a "natural" level. In contrast, we know from Corollary 1 that when γ is exogenous, the principal would be better off if γ increases (assuming $\gamma \geq 0$ to begin with). In this case, the principal would be willing to pay for an exogenous increase in γ above its "natural" level. Of course, the driver of this difference is that γ comes with its own incentive compatibility constraint when it is endogenous.

Recall that under the assumption made thus far, an optimal contract that induces an

[15] It is clearly possible to construct examples where $\gamma^*(p) > \gamma_0(p) > 0$. This could be done by assuming that $c_{\gamma p}$ is negative of a large enough magnitude. However, such examples are arguably less interesting or surprising in light of Corollaries 1 and 2.

interior (p, γ) pair consists of

$$\begin{aligned} w_{11} &= c_p + (1-q)\, c_\gamma > 0 \\ w_{10} &= c_p - q c_\gamma > 0 \\ w_{01} &= w_{00} = 0. \end{aligned}$$

Hence, the optimal contract that induces $\gamma = \gamma_0(p)$ features $w_{11} = w_{10} = c_p$. Thus, the agent's remuneration is independent of y in this case. This is intuitive since he makes no effort at changing the correlation between x and y. Nevertheless, it is interesting to contrast this outcome with the standard intuition based on the Informativeness Principle (see Lemma 1). When γ is exogenous, the agent's compensation depends on y because y is informative about the agent's productive action, p. However when γ is endogenous and the principal seeks to implement $\gamma_0(p)$, his hands are tied by the agent's incentive compatibility constraint. Making pay contingent on y would entice the agent to manipulate the correlation.

To appreciate how the contract changes when the principal optimally manipulates γ, assume that $\gamma_0(p) > 0$ and $\gamma^*(p) < \gamma_0(p)$. Since $c(p, \gamma)$ is convex in γ, it follows that $c_\gamma(p, \gamma^*(p)) < 0$. Hence, the contract now pays $w_{10} > c_p > w_{11} > 0$. Thus, a high x is rewarded more when y is low. This reward structure is necessary in order to induce the agent to lessen the correlation. When $\gamma > 0$ is exogenous, a similar but more extreme reward structure obtains, with $w_{10} > 0$ but $w_{11} = 0$. The reason is that the principal needs to satisfy only a single incentive constraint in the latter case. One interpretation is that the contract is more sensitive to y in the case where γ is exogenous. Thus, our theory may explain why real world contract are often less sensitive to exogenous signals than suggested by the Informativeness Principle.

Finally, consider once again the principal who cannot contract on y. If γ is endogenous, the agent then selects $\gamma = \gamma_0(p)$ as explained earlier. Assume $\gamma_0(p) > 0$. In this case, $w_1 = c_p$ and implementation costs are thus pc_p. Now look at this problem from the point of view of an "outside observer" who thinks that γ is exogenous. If he observes this principal-agent relationship many times, he may estimate that γ is exogenously fixed at $\gamma_0(p)$. Given Corollary 1 and fixing p, he thus believes that implementation costs can be reduced to

$$C(p, \gamma_0(p)) = pc_p(p, \gamma_0(p)) - \frac{\gamma_0(p)}{1-q} c_p(p, \gamma_0(p))$$

whereas implementation costs in reality are

$$\min_{\gamma} K(p,\gamma) = K(p,\gamma^*(p)) = pc_p(p,\gamma^*(p)) + \gamma^*(p)c_\gamma(p,\gamma^*(p)).$$

Proposition 5 *Assume that $\gamma_0(p) \neq 0$ and $\gamma^*(p)$ are interior and that $c(p,\gamma) = k(p) + \kappa(\gamma)$. Then, $K(p,\gamma^*(p)) > C(p,\gamma_0(p))$.*

Proof. Note that

$$\begin{aligned} K(p,\gamma^*) - C(p,\gamma_0) &= pk'(p) + \gamma^*\kappa'(\gamma^*) - pk'(p) + \frac{\gamma_0}{1-q}k'(p) \\ &= \frac{\gamma_0}{1-q}\left[k'(p) + \frac{\gamma^*}{\gamma_0}(1-q)\kappa'(\gamma^*)\right]. \end{aligned}$$

By the argument in Example 1, $\frac{\gamma^*}{\gamma_0} \in (0,1)$. By assumption, $c_p > 0$ and $c_p + (1-q)c_\gamma > 0$. Put together, this proves the result.■

The proposition is fairly intuitive. It is trivially true that $K(p,\gamma_0(p)) > C(p,\gamma_0(p))$ because of the extra incentive constraint when γ is endogenous. It now turns out that the ability to manipulate γ away from $\gamma_0(p)$ is not valuable enough to overcome the cost of the extra constraint. However, this relies in part on the assumption that $c_{p\gamma} = 0$, since this implies that the incentive constraint on p is unaffected by changes in γ. It is conceivable that $c_{p\gamma} \neq 0$ may overturn the result (see the discussion following Proposition 4).

The implication of Proposition 5 is that the outside observer overestimates the cost savings from contracting on y. If is it costly to collect data on y and write a more complicated contract, then this may help explain why outside signals (y) are used less often in practice than what the Informativeness Principle would suggest.

3.3 Contingency planning: Reformulating the model

Consider an agent who is planning ahead and thinking about how to build processes that take into account the possibility that his productivity may depend on future contingencies, as described by y. How much effort the agent devotes to thinking about any given contingency is likely to impact how well he performs in said contingency. Think of the choice variables p_0 and p_1 as how much effort is devoted to preparing for contingency $y = 0$ and $y = 1$, respectively.[16] More concretely, change variables by defining

$$p_0 = \frac{p(1-q)-\gamma}{1-q} = p - \frac{\gamma}{1-q} \tag{3}$$

$$p_1 = \frac{pq+\gamma}{q} = p + \frac{\gamma}{q}. \tag{4}$$

Our model can now be written as:

	$x=0$	$x=1$
$y=1$	$(1-p_1)q$	$p_1 q$
$y=0$	$(1-p_0)(1-q)$	$p_0(1-q)$

Note that p_0 and p_1 denote the conditional probability that $x = 1$ given $y = 0$ and $y = 1$, respectively. In other words, the agent's choice variables determine the conditional distributions of x given y. This formulation of the model is henceforth referred to as the contingency-planning model. Here, (p_0, p_1) is feasible if and only if $(p_0, p_1) \in [0,1] \times [0,1]$. Thus, the feasible set is more easily described than in the (p, γ) formulation of the model.

As before, the marginal distribution of y is completely described by q, and is thus again outside the agent's control. A subtler point is that changes in p_0 and p_1 also change the dependence structure between the two random variables. To see this, note that for any given (p_0, p_1),

$$p = qp_1 + (1-q)\,p_0 \tag{5}$$

$$\gamma = q\,(1-q)\,(p_1 - p_0)\,, \tag{6}$$

meaning in particular that the dependence structure, as captured by γ, depends on p_0 and p_1. Thus, the formulation in terms of (p_0, p_1) pushes the dependence structure into

[16] Hence this interpretation is in line with the state-contingent model of Chambers and Quiggin (1998, 2000): the agent commits to a contingent production plan before the realization of y. A difference however is that y is not observed by the principal in the original version of Chambers and Quiggin (1998), while our focus here is precisely on the use of y in contracting.

the background. Although the two models are mathematically equivalent, the most natural formulation is likely to depend on the application. Hence, we will make use of both formulations and, at times, move from one to the other. However, since our main objective is to shed light on the consequences of endogenizing correlation, we do favor the first model.

The consequence of working with a simpler feasible set in this formulation of the model is that the cost function becomes a little more intricate. Let the transformed cost function be denoted $\widehat{c}(p_0, p_1)$ and note that

$$\widehat{c}(p_0, p_1) = c(qp_1 + (1-q)\, p_0, q\,(1-q)\,(p_1 - p_0)),$$

where the cost function on the right hand side is taken from the (p, γ) formulation of the model. Two observations about moving from one model formulation to the other are now pertinent. First, convexity of the cost function is preserved as we move from one model to the other, since in either case the transformations of variables are linear transformations. It is also clear that the agent's utility from income is linear in (p_0, p_1). Hence, following previous arguments, the FOA is valid in either formulation of the model. Second, note that even if $c(p, \gamma)$ is monotonic in its arguments it is not necessarily the case that $\widehat{c}(p_0, p_1)$ is. For instance, note that increasing p_0 amounts to increasing p but decreasing γ. Thus, care must be taken to keep track of marginal costs when moving between the two models.

In the contingency planning formulation of the limited liability model, the agent's first order conditions take the even simpler form:

$$\begin{aligned} q(w_{11} - w_{01}) &= \widehat{c}_{p_1}(p_0, p_1) \\ (1-q)(w_{10} - w_{00}) &= \widehat{c}_{p_0}(p_0, p_1). \end{aligned}$$

Thus, it is easy to derive the optimal contract that induces any interior (p_0, p_1).

Proposition 6 *Consider the limited liability model. The unique optimal contract that induces a given interior (p_0, p_1) is given by*

- $w_{11} = \frac{\widehat{c}_{p_1}}{q}$ *and* $w_{01} = 0$ *if* $\widehat{c}_{p_1} \geq 0$ *but* $w_{01} = -\frac{\widehat{c}_{p_1}}{q}$ *and* $w_{11} = 0$ *if* $\widehat{c}_{p_1} \leq 0$*, and*
- $w_{10} = \frac{\widehat{c}_{p_0}}{1-q}$ *and* $w_{00} = 0$ *if* $\widehat{c}_{p_0} \geq 0$ *but* $w_{00} = -\frac{\widehat{c}_{p_0}}{1-q}$ *and* $w_{10} = 0$ *if* $\widehat{c}_{p_0} \leq 0$.

Proposition 2 of course agrees with Proposition 1 once it is observed that

$$\widehat{c}_{p_0} = (1-q)\left[c_p - qc_\gamma\right] \tag{7}$$

$$\widehat{c}_{p_1} = q\left[c_p + (1-q)c_\gamma\right]. \tag{8}$$

The contingency planning model also simplifies the characterization of the optimal contract when actions on the boundaries are to be implemented. Consider, for example, implementing some (p_0, p_1) with $p_0 = 0$. In this case, incentive compatibility requires

$$(1-q)(w_{10} - w_{00}) \leq \widehat{c}_{p_0}(p_0, p_1).$$

Note that if $\widehat{c}_{p_0} \geq 0$, then the constraint is satisfied at $w_{10} = w_{00} = 0$. Evidently, the constraint is slack when $\widehat{c}_{p_0} > 0$. On the other hand, if $\widehat{c}_{p_0} < 0$ then the constraint can be written $(1-q)(w_{00} - w_{10}) \geq -\widehat{c}_{p_0} > 0$ and it follows that the optimal contract has $w_{10} = 0$ and $w_{00} = \frac{-\widehat{c}_{p_0}}{1-q}$. The incentive compatibility constraint is binding in this case. Implementing $p_0 = 1$ can be dealt with in a similar manner, as can $p_1 \in \{0, 1\}$. Combined with Proposition 2, the optimal contract can thus be characterized for any (p_0, p_1).

Proposition 7 *Consider the limited liability model. The unique optimal contract that induces a given (p_0, p_1) is characterized in Table 1.*

	$\widehat{c}_{p_0}(p_0, p_1) \geq 0$	$\widehat{c}_{p_0}(p_0, p_1) < 0$
$p_0 = 0$	$w_{10} = 0, w_{00} = 0^*$	$w_{10} = 0, w_{00} = -\frac{\widehat{c}_{p_0}}{1-q}$
$p_0 \in (0,1)$	$w_{10} = \frac{\widehat{c}_{p_0}}{1-q}, w_{00} = 0$	$w_{10} = 0, w_{00} = -\frac{\widehat{c}_{p_0}}{1-q}$
$p_0 = 1$	$w_{10} = \frac{\widehat{c}_{p_0}}{1-q}, w_{00} = 0$	$w_{10} = 0, w_{00} = 0^*$

(i) *w_{00} and w_{10} as functions of p_0.*

	$\widehat{c}_{p_1}(p_0, p_1) \geq 0$	$\widehat{c}_{p_1}(p_0, p_1) < 0$
$p_1 = 0$	$w_{11} = 0, w_{01} = 0^*$	$w_{11} = 0, w_{01} = -\frac{\widehat{c}_{p_1}}{q}$
$p_1 \in (0,1)$	$w_{11} = \frac{\widehat{c}_{p_1}}{q}, w_{01} = 0$	$w_{11} = 0, w_{01} = -\frac{\widehat{c}_{p_1}}{q}$
$p_1 = 1$	$w_{11} = \frac{\widehat{c}_{p_1}}{q}, w_{01} = 0$	$w_{11} = 0, w_{01} = 0^*$

(ii) *w_{01} and w_{11} as functions of p_1.*

Table 1: The optimal contract in the contingency planning formulation of the limited liability model. An asterisk (*) indicates that an incentive compatibility constraint is slack.

Finally, consider once again the original endogenous correlation, or (p, γ), formulation of the model. It may seem more tedious to derive the optimal contract for boundary actions in that set-up, partly because of the shape of the feasible set and partly due to the slightly more complicated form of the incentive compatibility constraints. An easier way to obtain a characterization is to start with Proposition 7 and then use the relationships in (3)-(4) and (7)-(8) to "convert" the optimal contract in the (p_0, p_1) formulation into the (p, γ) formulation. To this end, Figure 2 reproduces Figure 1 but adds information about how the boundaries in the two models relate. Any given (p, γ) can now be readily translated into a form where Proposition 7 can be applied.

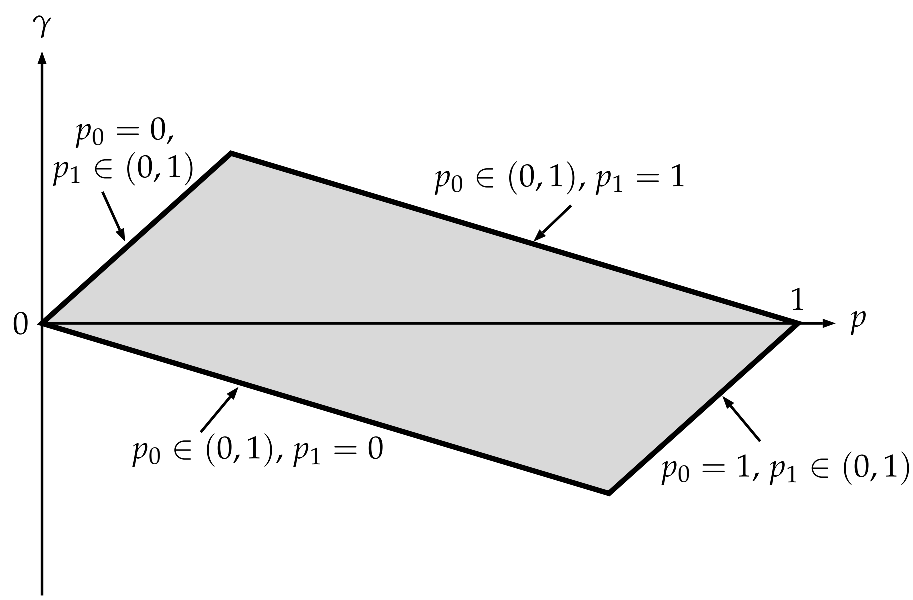

Figure 2: Boundaries in the (p, γ) formulation compared to the (p_0, p_1) formulation.

Now fix p at a large value for which $p > \max\{q, 1-q\}$. Then, as γ moves towards its lowest value we see that $p_0 \to 1$. Likewise, as γ moves towards its highest value, $p_1 \to 1$. Assume moreover that $\widehat{c}_{p_0} > 0$ and $\widehat{c}_{p_1} > 0$. These assumptions are equivalent to those imposed in Section 3.2, specifically that $c_p - qc_\gamma > 0$ and $c_p + (1-q)c_\gamma > 0$. It is now clear from Proposition 7 that costs are continuous in γ, including at the boundaries.

Conversely, implementation costs are discontinuous at one or both boundaries when $p < \max\{q, 1-q\}$.

4 Micro-foundations and applications

The current section is in two parts. The first presents a variety of scenarios, or stories, meant in part to illustrate the applicability of the abstract model. These stories also demonstrate that the agent's costs may not be monotonic in the choice variables. In this sense, they provide micro-foundations for different assumptions on the cost structure. The second part focuses on properties of the optimal contract, including which action it is optimal for the principal to induce. This part requires structure not only about the agent's cost structure but also about the principal's objective function. Both elements generally depend on the application.

4.1 Micro-foundations

4.1.1 A micro-foundation for the forecasting model

In this section, we provide a detailed analysis of the following forecasting scenario, cast in the terms of our model. The goal of the principal is to match the state y, by following the recommendation of the agent x. The effort of the agent hence consists of effort to correlate x to y as much as possible. The typical payoff matrix for the principal is as follows:

$y = 1$	$\pi_{01} \leq 0$	$\pi_{11} > 0$
$y = 0$	$\pi_{00} > 0$	$\pi_{10} \leq 0$
	$x = 0$	$x = 1$

Notice that here it is crucial that the principal's ex-post payoff depends on both realizations of x and y, hence on γ. The timing of the game is:

1. the expert privately invests in the information structure and is committed to it
2. signal x realizes
3. The principal makes a decision based on x, but his payoff depends also on y

4. signal y realizes
5. the expert is paid according to a contract conditional on (x, y) [or any restriction thereof].

The important point is that the cost of the agent should naturally be increasing in γ, which is here the "true" effort, to the opposite of the standard model. Moreover, it is natural to assume that by default, the agent chooses p only based on priors, i.e. a frequency of x's that is directly related to the prior q on y. This reduced-form model can be micro-founded as follows. Consider an agent who does not observe the state directly, but who nevertheless privately receives a non-verifiable and potentially informative signal. The signal is drawn from some distribution $F_0(s)$ in state $y = 0$ and from some distribution $F_1(s|\theta)$ in state $y = 1$. The supports of the two distribution functions coincide. Here, $\theta \geq 0$ is a choice variable of the agent which captures how well he can distinguish the two states from each other. Assume that $F_1(s|0) = F_0(s)$, such that $\theta = 0$ describes no effort – the agent's signal is completely uninformative in this case. Let $f_1(s|\theta)$ denote the density of $F_1(s|\theta)$ and let $f_{1\theta}(s|\theta)$ and $F_{1\theta}(s|\theta)$, respectively, denote their partial derivatives with respect to θ.

The agent updates his beliefs about the true state upon having observed the signal. Assume that $f_1(s|\theta)$ satisfies the (strict) monotone likelihood ratio property (MLRP), or

$$\frac{\partial}{\partial s}\left(\frac{f_{1\theta}(s|\theta)}{f_1(s|\theta)}\right) > 0.$$

It follows that $F_1(s|\theta)$ dominates $F_0(s)$ in terms of the likelihood-ratio for any $\theta > 0$. Thus, for any θ, there exists a threshold signal above which the agent is more likely to believe that $y = 1$ than $y = 0$. It is costly for the agent to improve (increase) θ. For future reference, MLRP also implies that $F_1(s|\theta)$ is strictly decreasing in θ whenever s is interior.

The principal requires the agent to issue a "high" or "low" message, where the former is interpreted as a forecast that $y = 1$ and the latter as a forecast that $y = 0$. The agent is assumed to choose some threshold, z, such that a signal above z will result in a "high" message and a signal below z will result in a "low" message.[17] It is natural to assume that there is no cost of changing z.

[17] There are other ways of partitioning the signal space to determine messages. We assume the threshold structure for simplicity. Intuitively, this is not a restrictive given MLRP, especially not if accurate forecasts are rewarded.

In summary, the agent chooses the pair (z, θ). Given the (z, θ) pair, note that the probability that a "high" message is delivered is

$$\begin{aligned} p &= q\left(1 - F_1(z|\theta)\right) + (1-q)\left(1 - F_0(z)\right) \\ &= 1 - \left[q F_1(z|\theta) + (1-q) F_0(z)\right]. \end{aligned}$$

Likewise, solving for γ yields

$$\gamma = q\,(1-q)\left[F_0(z) - F_1(z|\theta)\right].$$

Note that for any given θ, any $p \in [0,1]$ can be achieved at no cost by simply manipulating the threshold, z. However, doing so also changes γ. Thus, to change p while holding γ constant it is generally necessary to change both z and θ. Consequently, changing p is costly. An exception occurs when $\gamma = 0$, since in this case the $\gamma = 0$ constraint is costless as it is trivially satisfied at $\theta = 0$. Then, z can be used to manipulate p with no side-effects.

The agent's problem is more complicated when x and y are correlated, or $\gamma \neq 0$. As noted above, changing the threshold z does indeed change the marginal distribution of x, but it also changes the correlation unless θ is adjusted at the same time. To summarize, p is free for any given information structure, or θ, but this does not imply that p is free given a fixed target for γ. Indeed, note that for any $\theta > 0$, the MLRP implies that γ has an inverse-u shape in z. It is intuitive that γ is small when z is very small or very large. In these cases, the agent gives the same report for almost all z and so the correlation between x and y naturally tends to be small. To obtain a high γ in these circumstances, it is necessary to invest heavily in θ in order to make the signal sufficiently informative. This observation in turn suggests that high and low p values are costlier to implement than intermediate p values, given a fixed target for γ. This intuition is confirmed and formalized in the following.

To better understand the cost structure, solve for $F_0(z)$ and $F_1(z|\theta)$ given p and γ, yielding

$$\begin{aligned} F_0(z) &= 1 - p + \frac{\gamma}{1-q} \\ F_1(z|\theta) &= 1 - p - \frac{\gamma}{q}. \end{aligned}$$

Given (p, γ), the first equation uniquely determines z. This in turn implies that θ is uniquely determined by the second equation. Note that an increase in γ necessitates an increase in z. To satisfy the second equation, it is then necessary for θ to increase as well. Stated differently, increasing γ is unambiguously costly, given p. This is intuitive since γ measures the informativeness of the agent's message. It is costly to generate a more informative information structure without changing the "average" message (p).

An increase in p necessitates a decrease in z. The simultaneous increase in p and decrease in z leads both sides of the second equation to decrease in value. Thus, it appears ambiguous whether θ is increasing or decreasing in p, given γ. In the following, we prove that costs are u-shaped in p, given $\gamma > 0$. Combining the above two equations yields the condition that

$$F_1(F_0^{-1}(1 - p + \frac{\gamma}{1-q})|\theta) = 1 - p - \frac{\gamma}{q},$$

which describes the relationship between p and θ more directly. Holding γ fixed, let $\theta(p)$ denote the unique value of θ that solves the equation as a function of p. Then, it is easy to show that for any interior p,

$$\theta'(p) = \frac{-1}{F_{1\theta}(z|\theta(p))} \left(1 - \frac{f_1(z|\theta(p))}{f_0(z)}\right),$$

where $z = F_0^{-1}(1 - p + \frac{\gamma}{1-q})$. Given MLRP, $F_{1\theta}(z|\theta) < 0$ for any interior z. Likewise, MLRP implies that $f_1(z|\theta)$ crosses $f_0(z)$ exactly once, from below. Hence, $\theta'(p)$ is first positive and then later negative. In other words, $\theta(p)$ is u-shaped. Since higher θ are costly, it then follows that costs are u-shaped in p, given $\gamma > 0$.

4.1.2 A micro-foundation for costs that are non-monotonic in γ

In the previous scenario, the agent's cost function is u-shaped in p. Here, a complementary scenario is presented in which costs are u-shaped in γ. In a more specialized version of the scenario, $c(p, \gamma)$ is additively separable in p and γ.

We begin with the contingency-planning formulation of the model, where the agent's choice variables are (p_0, p_1). Hence, we think of the agent as preparing for the contingencies $y = 0$ and $y = 1$. The idea is to think about a natural cost structure in this setting, and then translate this back into the (p, γ) formulation to understand what it implies about the properties of $c(p, \gamma)$.

There are several ways of modelling costs under contingency-planning, depending on when the agent incurs costs. For instance, contingency-planning may involve costly investments that are sunk before the state is realized. Alternatively, costs may be incurred only after the state is realized. Here, we examine the latter case in more detail.

Consider an agent whose (commonly known) base skills are adequate for a probability of success of $\underline{p}_0$ in state $y = 0$ and $\underline{p}_1$ in state $y = 1$. The agent can augment these probabilities by purchasing productive machinery. However, he must place the order before the state is realized. The order can be made state-dependent, such that $p_0 - \underline{p}_0 \geq 0$ is the amount of machinery that he takes delivery of in state $y = 0$ and $p_1 - \underline{p}_1 \geq 0$ the corresponding quantity in state $y = 1$. He pays only for the amount that was delivered. Hence, when the order is placed, the agent's expected costs can be written as

$$\widehat{c}(p_0, p_1) = qk(p_1 - \underline{p}_1) + (1-q)\, k(p_0 - \underline{p}_0),$$

where $k(\cdot)$ is some increasing and convex cost function.[18] Note that costs are assumed to be independent of the state. Given (3) and (4),

$$\begin{aligned} c(p,\gamma) &= \widehat{c}\left(p - \frac{\gamma}{1-q}, p + \frac{\gamma}{q}\right) \\ &= qk\left(p + \frac{\gamma}{q} - \underline{p}_1\right) + (1-q)\, k\left(p - \frac{\gamma}{1-q} - \underline{p}_0\right). \end{aligned}$$

Costs are clearly increasing in p. However, note that

$$c_\gamma(p,\gamma) = k'\left(p + \frac{\gamma}{q} - \underline{p}_1\right) - k'\left(p - \frac{\gamma}{1-q} - \underline{p}_0\right).$$

Thus, assuming k is strictly convex, it holds that

$$c_\gamma(p,\gamma) \gtrless 0 \Longleftrightarrow \gamma \gtrless q(1-q)(\underline{p}_1 - \underline{p}_0).$$

Stated differently, $c(p,\gamma)$ is u-shaped in γ and, regardless of p, minimized at $\gamma = q(1-q)(\underline{p}_1 - \underline{p}_0)$.Given (6), note that this level of correlation is achieved if $p_0 = \underline{p}_0$ and $p_1 = \underline{p}_1$,

[18] Recall that we have chosen a parameterization where probabilities are linear in the action but where costs may be non-linear. In the case of buying machinery, it is perhaps initially more logical to think about costs as being linear in quantity and the improvement in productivity as being a concave function of the quantity of machinery. However, since this is just another reparameterization of the problem, there is no substantive difference between the two models.

i.e. if the agent exerts zero effort. In fact, it obtains whenever $p_0 - \underline{p}_0 = p_1 - \underline{p}_1$, i.e. when p_0 and p_1 are increased at the same rate to achieve a higher p. Holding p fixed, a higher level of correlation can be obtained by instead increasing p_1 faster than p_0. Conversely, a lower level of correlation is obtained by increasing p_0 faster than p_1. The last two options both involve an unbalanced increase in p_0 and p_1. This is costly due to the convexity of $k(\cdot)$, given that $p = qp_1 + (1-q)p_0$ is held constant. Note also that the cost minimizing value of γ can be non-zero. Specifically, the cheapest γ is non-zero when the agent's base skills yield a success with different probability in the two states.

Note that if the principal is risk neutral and the value that he attributes to a success or failure, respectively, is the same in both states, then he cares directly only about p, i.e. the probability of success. He takes an interest in γ only for the purposes of minimizing implementation costs. Thus, this recreates the standard model. However, the principal's objective function depends directly on γ if the value of a success or failure depends on the state.

To continue, consider the standard model. Assume that $\underline{p}_1 > \underline{p}_0$, such that the agent's baseline productivity is higher in the good state. We show next that the correlation dampening effect from Section 3 obtains in the current setting as well. This is despite the fact that unlike what we assumed in Section 3, $c_{p\gamma}$ does not have a constant sign here. Specifically, note that $c_{p\gamma}(p,\gamma) = 0$ where $c_\gamma(p,\gamma) = 0$. Thus, $c_{p\gamma}$ generally changes sign in a neighborhood around γ_0. Note that in this model it is trivially true that $c_p + (1-q)c_\gamma > 0$ and $c_p - qc_\gamma > 0$ since the former is proportional to $\widehat{c}_{p_1} = qk'(p_1 - \underline{p}_1)$ and the latter is proportional to $\widehat{c}_{p_0} = (1-q)\,k'(p_0 - \underline{p}_0)$; see (7) and (8).

In the (p_0, p_1) formulation of the model, incentive compatibility in the interior requires that $w_{1j} = k'(p_j - \underline{p}_j)$, $j = 0, 1$. Contingent on state j, expected wage costs are thus $p_j k'(p_j - \underline{p}_j)$. Assuming an interior solution, it is optimal for the principal to equate marginal implementation costs across states, or

$$\frac{\partial p_1 k'(p_1 - \underline{p}_1)}{\partial p_1} = \frac{\partial p_0 k'(p_0 - \underline{p}_0)}{\partial p_0}. \tag{9}$$

Of course, $p_j k'(p_j - \underline{p}_j)$ must be locally convex in order for the first order condition to correctly solve the cost minimization problem. Here, we will impose a stronger assumption that implies convexity. To begin, it is easy to verify that

$$\frac{\partial^2 p_j k'(p_j - \underline{p}_j)}{\partial p_j^2} > -\frac{\partial^2 p_j k'(p_j - \underline{p}_j)}{\partial p_j \partial \underline{p}_j}. \tag{10}$$

Our assumption is that $p_j k'(p_j - \underline{p}_j)$ is submodular in $(p_j, \underline{p}_j)$. Stated differently, the right hand side of the above inequality is non-negative, thereby implying that $p_j k'(p_j - \underline{p}_j)$ is strictly convex in p_j. The assumption is satisfied as long as $k'(p_j - \underline{p}_j)$ is not "too concave" in p_j, or, more accurately, when $k''(p_j - \underline{p}_j) + p_j k'''(p_j - \underline{p}_j) \geq 0$.

The submodularity assumption along with the inequality in (10) makes it possible to utilize (9) to obtain the following comparative statics. First, it is optimal for the principal to induce $p_1 > p_0$. Nevertheless, it holds that $p_1 - \underline{p}_1 < p_0 - \underline{p}_0$. Translating this into the (p, γ) formulation of the model, the first inequality implies that $\gamma^* > 0$ while the second inequality implies that $\gamma^* < \gamma_0$. Hence, as in Example 1, $0 < \gamma^* < \gamma_0$.

I don't know if we want to keep that (or maybe not in its present form)

In the special case where $k(\cdot)$ is quadratic, costs are

$$c(p, \gamma) = q\left(p + \frac{\gamma}{q} - \underline{p}_1\right)^2 + (1-q)\left(p - \frac{\gamma}{1-q} - \underline{p}_0\right)^2,$$

which can be rewritten as

$$c(p, \gamma) = \left[p^2 - 2p\left((1-q)\,\underline{p}_0 + q\underline{p}_1\right)\right] + \left[q\left(\frac{\gamma}{q} - \underline{p}_1\right)^2 + (1-q)\left(\frac{\gamma}{1-q} + \underline{p}_0\right)^2\right].$$

Evidently, $c(p, \gamma)$ is additively separable in p and γ in this case.

4.2 Applications

4.2.1 Pay for luck [preliminary]

Suppose the principal can only observe the sum $x + y$, and assume that his objective is to maximize $\mathbb{E}[x]$. The observability restriction imposes that $w_{10} = w_{01}$, hence if the principal chooses to set a positive transfer when observing $x + y = 1$, he will sometimes reward the agent for a high y and a low x–the essence of pay for luck since the agent does not control y. For a fixed correlation γ, we can apply the analysis in the first section (with exogenous correlation), with the additional restriction that one of the states is just the aggregation of $(1,0)$ and $(0,1)$ and the corresponding likelihood ratio is $\frac{\partial(P_{10}+P_{01})}{\partial p}/(P_{10}+P_{01}) = \frac{1-2q}{p(1-q)+(1-p)q-2\gamma}$. Following the same analysis as in the first lemma, we obtain that the principal prefers to use a pay-for-luck scheme over paying only in state $(1,1)$ if and only if $\gamma \geq q^2$ (note that, aside from feasibility consideration, this is true for any p). Intuitively, it takes correlation to be positive (which is directly in line with lemma 1) and luck to be sufficiently rare (low q) to make pay-for-luck a desirable scheme. Now, of course, if the agent is remunerated only for $x+y=1$, he will have incentives to decrease correlation, which increases the likelihood of the remunerated result. On the contrary, if the agent is paid for $(x,y)=(1,1)$, he will be tempted to increase correlation, which makes in turn pay-for-luck a relatively more attractive scheme.

Note that this is the story with exogenous correlation, and (1) the full implementation problem is not solved, raising interesting questions (2) the goal of the principal is still simple, in that it does not depend directly on y (only through implementation costs).

Here a natural case is when by default the correlation is 0, so that y is just a standard additive noise. When correlation is not manipulable, there is then no pay for luck (only $w_{11}>0$). It would then be interesting to see if the possibility for the agent to manipulate correlation can make pay-for-luck an optimal contract. [To be continued...]

5 Extensions

5.1 Value of Information with risk-aversion

We have so far focused on the case where the agent is risk neutral and a moral hazard problem exists due to limited liability. In this section we revisit the baseline model with exogenous correlation. Both the principal and agent are assumed to be expected utility maximizers, but no restrictions are now imposed on the Bernoulli utility functions. Thus, to mention just a few possibilities, the principal and the agent can be risk averse or risk loving, or have preferences that exhibit loss aversion. Moreover, we no longer impose a limited liability constraint. Hence, the participation constraint comes into play.

The purpose is to prove that even with this minimal structure, a main conclusion from Lemma 1 survives. Specifically, it remains the case that the principal's expected utility is U-shaped in γ and minimized at $\gamma = 0$ (independence). Thus, if $\gamma' > \gamma'' \geq 0$ or $\gamma' < \gamma'' \leq 0$ then the principal weakly prefers γ' to γ''. In comparison, the Informativeness Principle says merely that independence ($\gamma = 0$) is worse than any other level of correlation. Of course, our model is endowed with a more specific signal structure that allows the additional inference.

In line with the standard model, we now assume that costs depend on p only and that γ is exogenous. This assumption implies that the participation constraint is unaffected by an exogenous change in γ.[19]. The cost function $c(p)$ need not be monotonic in p. However, given p, it is assumed that a feasible contract exists (one that satisfies the participation and incentive compatibility constraints) for all feasible values of γ. Also in line with Lemma 1, the principal is assumed to directly care only about the value of x, whereas the only use of y is to potentially lower implementation costs.

Fix some p that the principal wishes to induce. Consider two levels of correlation, γ' and γ'', with $\gamma' > \gamma'' \geq 0$. We want to show that the principal is weakly better off with γ' than with γ''. To do so, we establish that if γ is fixed at γ', then the principal can effectively emulate an environment where γ takes the smaller value γ''. Thus, when $\gamma = \gamma'$, the principal can at the very least guarantee himself the exact same expected payoff as if $\gamma = \gamma''$.

Starting from $\gamma = \gamma'$, the following randomization device can be used by the principal, for any $\varepsilon \in [0, 1]$:

- With probability $1 - \varepsilon$: The agent's pay is based on (x, y) as realized.
- With probability ε: Use x as realized. With probability q this is paired with $y = 1$ and with probability $1 - q$ it is paired with $y = 0$. The agent's pay is then based on the realized value of x and the artificially determined value of y.

Note that the randomization device effectively implements a mixture of the true joint probability distribution and the probability distribution in which x and y are independent. For this reason, the newly created distribution over (x, y) has an intermediate level

[19]The participation constraint never binds in the limited liability model. For this reason, Lemma 1 also holds if costs are $c(p, \gamma) = k(p) + \kappa(\gamma)$, for instance.

of correlation. Specifically, note that the agent is paid according to $(x, y) = (1, 1)$ with probability

$$(1-\varepsilon)\left(pq+\gamma'\right)+\varepsilon pq = pq+(1-\varepsilon)\gamma'.$$

Similarly calculations for the other combinations of (x, y) confirm that the term involving γ' is always replaced by $(1-\varepsilon)\gamma'$. Now, pick $\varepsilon \in (0, 1]$ in such a way that $(1-\varepsilon)\gamma' = \gamma''$. Then, the randomization device recreates a joint probability distribution in which $\gamma = \gamma''$. Next, using the above method to determine (x, y), offer whatever contract would be optimal if γ were to truly take the value γ''. This contract must necessarily satisfy the participation constraint if the true value of γ was $\gamma = \gamma''$. However, since costs by assumption are the same regardless of whether $\gamma = \gamma'$ or $\gamma = \gamma''$, it follows that the participation constraint remains satisfied. At the same time, note that the randomization device does not change the incentive compatibility constraint with respect to p (recall there is no incentive compatibility constraint with respect to γ since this is exogenous). Thus, the agent will still accept the contract and choose the intended p. Thus, the marginal distribution of x, which the principal cares about, is unchanged, and the distribution of payments is the same as would be the case if $\gamma = \gamma''$. In conclusion, the randomization device allows the principal to reproduce the same expected utility as he would get if $\gamma = \gamma''$. It follows that he must weakly prefer γ' to γ'' when $\gamma' > \gamma'' \geq 0$. Of course, an analogous argument applies if $\gamma' < \gamma'' \leq 0$.

The above randomization device establishes that the information system summarized by γ'' is a more "garbled" information system than that summarized by γ'. The reason is that the former adds more noise by mixing in the joint distribution where signals are independent. Thus, the result is closely related to Blackwell's Theorem, although Blackwell's Theorem in its original form applies only to decision problems. See e.g. Blackwell and Girshick (1954) for a review. Grossman and Hart (1983) prove that more garbled information systems are necessarily less valuable to the principal in a standard principal-agent model. The randomization device used above borrows from their argument. Kim (1995) contains a much more thorough examination of information systems, including illuminating discussions of both the Informativeness Principle and Blackwell's Theorem.

We now consider a slightly more general set-up. Specifically, to facilitate comparison with the key results in Section 3.2, we now assume that costs can be written as $c(p, \gamma) = k(p) + \kappa(\gamma)$. Assume that $\kappa(\gamma)$ is strictly convex and minimized at γ_0. Note that this set-up perfectly matches Example 1. Consider some p for which γ_0 is interior. Assume for the sake of argument that $\gamma_0 > 0$. We will show that any $\gamma \in [0, \gamma_0)$ leaves the principal weakly worse off than if $\gamma = \gamma_0$. Thus, the principal's expected utility is maximized either at $\gamma \geq \gamma_0$ or, possibly, at $\gamma < 0$. In other words, an exogenous γ in the range $[0, \gamma_0)$ can never be desirable, compared to γ_0. Recall that Example 1 in contrast shows that when γ is endogenous, it is optimal for the principal to induce some $\gamma \in (0, \gamma_0)$, at least within the confines of the limited liability model.

To prove that $\gamma \in [0, \gamma_0)$ is inferior to $\gamma = \gamma_0$ when correlation is exogenous, let γ_0 take the role of γ' in the above argument. We use the randomization device to replicate the distribution from $\gamma = \gamma'' \in [0, \gamma_0)$. The additive cost structure does not influence the incentive compatibility constraint with respect to p. Thus, the contract that is optimal when $\gamma = \gamma''$ remains incentive compatible when $\gamma = \gamma'$ and the randomization device is used. It remains to show that the contract satisfies the participation constraint. However, since the proposed contract is by definition optimal at $\gamma = \gamma''$, the agent must deem it worthwhile to incur costs $\kappa(\gamma'')$. Hence, the agent must also be willing to incur cost $\kappa(\gamma_0) < \kappa(\gamma'')$. Thus, the participation constraint is satisfied even at $\gamma = \gamma_0$. Consequently, if $\gamma = \gamma_0$ the principal can once again guarantee himself at least the optimal payoff that would arise if $\gamma = \gamma''$.

As a final remark, note that the previous arguments are unaffected by e.g. a minimum wage (or limited liability) or a cap on wages. Thus, the result is robust to a number of exogenous constraints on the contracting environment.

5.2 A double-spanning model with more signal realizations

So far we have assumed that x and y each takes one of two values. In this section we consider more general environments where where x can take one of $n + 1 \geq 2$ values, $x \in \{0, 1, ..., n\}$, and y one of $m + 1 \geq 2$ values, $y \in \{0, 1, ..., m\}$, where n and m are finite.

There are several ways to model such environments. For instance, starting with the contingency planning formulation of the original 2×2 model, one could imagine the agent's action being $m + 1$ dimensional, $\mathbf{p} = (p_0, p_1, ..., p_m)$. The jth element, p_j, determines the conditional distribution of x given $y = j$, $j = 0, 1, ..., m$.

An alternative is to start from the original (p, γ) formulation of the model. The agent's action now has $k+1$ dimensions, with $(p,\gamma) = (p,\gamma_1,\gamma_2,...,\gamma_k)$. As before, p determines the marginal (i.e. unconditional) distribution of x. The γ_i's describe a number of correlation increasing transformations in the sense of Epstein and Tanny (1980). However, there are a plethora of potential correlation increasing transformations when $m,n > 1$.[20] Thus k may be very large.

Note that the contingency planning model and the endogenous correlation model are equivalent if and only if $m = n = 1$. Thus, there is no unambiguous way to generalize the model. Moreover, both avenues outlined above increase the dimensionality of the agent's action. Here, we opt for a simpler generalization. The generalization is inspired by the original (p,γ) formulation of the model as well as the discussion of Blackwell's result in the preceding section.

The agent's action remains two-dimensional. As before, we write it as (p,γ). The idea is to let p determine the marginal distribution of x and γ determine how "garbled" the joint distribution is. We further impose specific functional forms for how p and γ perform these operations. In particular, let π_i denote the marginal probability that $x = i$, $i = 0,1,...n$. Let f and g denote two distinct distributions over x, where f_i and g_i denote the probability that $x = i$ for the two distributions, respectively. Then, we assume that for $p \in [0,1]$,

$$P_i(p) = pf_i + (1-p)g_i. \tag{11}$$

Hence, the marginal distribution of x, $\pi(p)$, is a convex combination of f and g, with the agent determining the weights. For this reason, Grossman and Hart (1983) refer to (11) as a "spanning condition". In the original formulation of our model, where $n = 1$, f and g are degenerate distributions. Specifically, f produces $x = 1$ with probability one whereas g produces $x = 0$ with probability one, or $f_1 = g_0 = 1$, $f_0 = g_1 = 0$.

To continue, define q as the marginal distribution of y, with $q_j > 0$ denoting the probability that $y = j$, $j = 0,1,...,m$. We have now described the marginal distributions of x and y. The role of γ is to determine their joint distribution. To do so, let $\mathbf{A} = \left[a_{ij}\right]_{(n+1)\times(m+1)}$ denote an $(n+1)\times(m+1)$ matrix in which all the rows and columns sum to zero. Then, a joint distribution where $(x,y) = (i,j)$ occurs with probability

$$\pi_i(p)q_j + a_{ij}$$

[20] Let $0 \le i_1 < i_2 \le n$ and $0 \le j_1 < j_2 \le m$. A correlation increasing transformation adds the same probability mass to $(x,y) = (i_1,j_1)$ and $(x,y) = (i_2,j_2)$ while deducting it from (i_1,j_2) and (i_2,j_1).

satisfies the property that the marginal distributions of x and y are $\pi(p)$ and q, respectively. Of course, feasibility still dictates that $\pi_i(p)q_j + a_{ij} \in [0,1]$ for all (i,j). Now, our final assumption is that the joint probability that $(x,y) = (i,j)$ is determined by

$$P_{ij}(p,\gamma) = (1-\gamma)\,\pi_i(p)q_j + \gamma\left(\pi_i(p)q_j + a_{ij}\right). \tag{12}$$

Hence, the joint probability spans $\pi_i(p)q_j$ (independence) and $\pi_i(p)q_j + a_{ij}$ (non-independence), with the weights determined by γ. Thus, in light of conditions (11) and (12), we refer to the model as the *double-spanning model*.

Simplifying yields

$$P_{ij}(p,\gamma) = \pi_i(p)q_j + \gamma a_{ij},$$

which is perhaps more directly comparable to our original formulation of the 2×2 model. In that model, $a_{ij} = 1$ if $i = j$ and $a_{ij} = -1$ if $i \neq j$. Note that γ need not necessarily be positive. Instead, the feasible set of (p,γ) pairs must satisfy $p \in [0,1]$ and $P_{ij}(p,\gamma) \in [0,1]$ for all (i,j). Finally, values of γ close to zero garbles the joint distribution by making it appear more like the noisy distribution in which x and y are independent. Hence, the Blackwell argument from the previous section applies in the benchmark model where γ is considered exogenous. Thus, Corollary 1 extends to any $n \geq 1$ and $m \geq 1$. Note that this holds true even without further assumptions on the primitives f, g, and $\mathbf{A}$.

Given (11), the joint distribution can also be written as

$$P_{ij}(p,\gamma) = \left(pf_i + (1-p)g_i\right)q_j + \gamma a_{ij}.$$

The technical importance of the double-spanning assumption now becomes clear. In particular, the agent's expected utility from income is linear in his action, (p,γ), regardless of his risk preferences. Thus, the first-order approach is valid as long as $c(p,\gamma)$ is convex.

We next turn to a generalization of the correlation dampening results in Section 3.2. Thus, we consider the standard model in which the principal cares directly only about x. To fix ideas, it is natural to assume that the expected value of x is increasing in p, although this assumption plays no formal role in the following. As before, we assume both parties are risk neutral and that the agent is protected by limited liability.

For any interior (p,γ) action, the incentive compatibility constraints with respect to p

and γ are

$$\begin{aligned}\sum_i \sum_j (f_i - g_i)\, q_j w_{ij} &= c_p \\ \sum_i \sum_j a_{ij} w_{ij} &= c_\gamma.\end{aligned}$$

These constraints are useful in calculating implementation costs. Consider any contract $[w_{ij}]_{(n+1)\times(m+1)}$ that is feasible, i.e. satisfies incentive compatibility. The expected wage costs to the principal are

$$\begin{aligned}E[w|p,\gamma] &= \sum_i \sum_j \left[(pf_i + (1-p)g_i)\, q_j + \gamma a_{ij}\right] w_{ij} \\ &= p \sum_i \sum_j (f_i - g_i)\, q_j w_{ij} + \gamma \sum_i \sum_j a_{ij} w_{ij} + \sum_i g_i q_j w_{ij},\end{aligned}$$

or

$$E[w|p,\gamma] = pc_p + \gamma c_y + \sum_i g_i q_j w_{ij}. \tag{13}$$

The first two terms are the same for all feasible contracts. Hence, the principal's problem is equivalent to minimizing the last term, subject to incentive compatibility. Due to the limited liability constraint, the last term is non-negative. At best, it can be made to equal zero.

Imagine now that there is some value of x, i^*, for which $g_{i^*} = 0$ but $f_{i^*} > 0$. In words, the distribution g produces $x = i^*$ with probability zero. In the 2×2 model, $i^* = n = 1$. Imagine moreover that there exists an incentive compatible contract that pays $w_{ij} = 0$ whenever $i \neq i^*$. Only if $x = i^*$ will the agent possibly be paid, depending on j. Under these assumptions it holds that $g_i q_j w_{ij} = 0$ for all (i, j). Thus, the last term in (13) is zero and the contract must be optimal.

Proposition 8 *Fix some interior (p,γ) action to be implemented. Assume that there exists some i^* for which $g_{i^*} = 0$ but $f_{i^*} > 0$. Assume moreover that there is a feasible contract for which $w_{ij} = 0$ whenever $i \neq i^*$. Then, implementations costs are*

$$K(p,\gamma) = pc_p + \gamma c_\gamma. \tag{14}$$

It is instructive to compare Proposition 8 to Proposition 3 in which the optimal contract is characterized in the 2×2 model. In that setting, we reduced the incentive compatibility constraints to

$$\begin{aligned} w_{11} - w_{01} &= c_p + (1-q)c_\gamma \\ w_{10} - w_{00} &= c_p - qc_\gamma. \end{aligned}$$

Thus, there exists a contract that pays strictly positive wages only when $x = i^* = 1$ if and only if $c_p + (1-q)c_\gamma > 0$ and $c_p - qc_\gamma > 0$. We imposed these assumptions in Section 3.2 and used them to establish that implementation costs in that case take the exact same form as in (14). Thus, Proposition 8 describes a generalization of the environment examined in Section 3.2. Since implementation costs are exactly the same, it follows that all the results in Section 3.2 remains valid under the assumptions in Proposition 8.

Corollary 3 *Given the assumptions in Proposition 8, the results in Section 3.2 remain valid for any $n \geq 1$, $m \geq 1$.*

To understand when a contract of the form assumed in Proposition 8 exists, note that the incentive compatibility constraints simplify to

$$\begin{aligned} \sum_j f_{i^*} q_j w_{i^*j} &= c_p \\ \sum_j a_{i^*j} w_{i^*j} &= c_\gamma, \end{aligned}$$

since $g_{i^*} = 0$ and $w_{ij} = 0$ if $i \neq i^*$. If pay is strictly positive only in states (i^*, j') and (i^*, j''), $j' \neq j''$, then the constraints reduce to

$$\begin{aligned} f_{i^*} q_{j'} w_{i^*j'} + f_{i^*} q_{j''} w_{i^*j''} &= c_p \\ a_{i^*j'} w_{i^*j'} + a_{i^*j''} w_{i^*j''} &= c_\gamma. \end{aligned}$$

Under the assumption that $a_{i^*j'} q_{j''} - a_{i^*j''} q_{j'} \neq 0$, these can be solved for

$$\begin{aligned} w_{i^*j'} &= \frac{-c_p a_{i^*j''} + c_\gamma f_{i^*} q_{j''}}{f_{i^*} \left(a_{i^*j'} q_{j''} - a_{i^*j''} q_{j'} \right)} \\ w_{i^*j''} &= \frac{c_p a_{i^*j'} - c_\gamma f_{i^*} q_{j'}}{f_{i^*} \left(a_{i^*j'} q_{j''} - a_{i^*j''} q_{j'} \right)}. \end{aligned}$$

Of course, this is an admissible solution if and only if both are positive. Thus, pick j' and j'' in such a way that $a_{i^*j'} > 0$ and $a_{i^*j''} < 0$. Such j' and j'' exists because the sum of a_{i^*j} over j is zero. Then, the denominators are strictly positive. Both numerators are likewise strictly positive if $c_p > 0$ and $|c_\gamma|$ is small. Recall that these conditions fit with those imposed in Section 3.2. In fact, in the 2×2 model we have $i^* = 1$ and $f_{i^*} = 1$. Let $i' = 1$ and $i'' = 0$, in which case $a_{i^*j'} = a_{11} = 1$, $a_{i^*j''} = a_{10} = -1$, $q_{j'} = q$ and $q_{j''} = 1 - q$. Then, we see that $w_{11} = c_p + (1-q)c_\gamma$ and $w_{10} = c_p - qc_\gamma$, both of which are positive if $c_p > 0$ and $|c_\gamma|$ is small.

Proposition 9 *Assume that there exists some i^* for which $g_{i^*} = 0$ but $f_{i^*} > 0$. Then, for any interior (p, γ) there exists a feasible contract with $w_{ij} = 0$ whenever $i \neq i^*$ if $c_p > 0$ and $|c_\gamma|$ is small.*

6 Conclusion

The informativeness Principle is a very powerful tool in incentive theory. Yet, it could easily be over-interpreted, in particular when the information structure is endogenous. Our approach shows that when the agent can control the correlation of the his performance with external indicators, it becomes harder for the principle to rely on such extraneous information to reduce agency costs. A main insight is that in general the principal is better off attenuating correlation in the optimal contract, or even giving up the use of external indicators when the agent can freely choose correlation on top of effort level. Our approach also offers new reduced forms for several very different moral hazard problems.

A Proofs

A.1 Proof of Corollary 2

Recall that implementation costs are

$$C(p,\gamma) = \begin{cases} \left(p - \frac{\gamma}{1-q}\right) c_p & \text{if } \gamma \geq 0 \\ \left(p + \frac{\gamma}{q}\right) c_p & \text{if } \gamma \leq 0 \end{cases}.$$

Note that $C(p,\gamma)$ is differentiable with respect to γ, except at $\gamma = 0$. The derivative is

$$C_\gamma(p,\gamma) = \begin{cases} \left(p - \frac{\gamma}{1-q}\right) c_{p\gamma} - \frac{c_p}{1-q} & \text{if } \gamma > 0 \\ \left(p + \frac{\gamma}{q}\right) c_{p\gamma} + \frac{c_p}{q} & \text{if } \gamma < 0 \end{cases}.$$

Since $c_p > 0$ by assumption, the derivative can be written

$$C_\gamma(p,\gamma) = \begin{cases} \frac{c_p}{1-q} \left(p(1-q) - \gamma\right) \frac{c_{p\gamma}}{c_p} - 1 & \text{if } \gamma > 0 \\ \frac{c_p}{q} \left(pq + \gamma\right) \frac{c_{p\gamma}}{c_p} + 1 & \text{if } \gamma < 0 \end{cases}.$$

Note that C_γ is larger as γ approaches zero from the left than from the right. Thus, $C(p,\gamma)$ can never have a local minimum at $\gamma = 0$, although it may have a local maximum there.

The second derivative is

$$C_{\gamma\gamma}(p,\gamma) = \begin{cases} \frac{c_p}{1-q} \left(\left(p(1-q) - \gamma\right) \frac{\partial}{\partial\gamma} \frac{c_{p\gamma}}{c_p} - \frac{c_{p\gamma}}{c_p}\right) & \text{if } \gamma > 0 \\ \frac{c_p}{q} \left(\left(pq + \gamma\right) \frac{\partial}{\partial\gamma} \frac{c_{p\gamma}}{c_p} + \frac{c_{p\gamma}}{c_p}\right) & \text{if } \gamma < 0 \end{cases}.$$

If $C(p,\gamma)$ attains a stationary point at $\gamma > 0$ then $\frac{c_{p\gamma}}{c_p} > 0$ is necessary. By log-concavity $\frac{\partial}{\partial\gamma} \frac{c_{p\gamma}}{c_p} \leq 0$. Hence, $C(p,\gamma)$ is locally strictly concave at such a point. Thus, there can be at most one stationary point for $\gamma > 0$. Note also that log-concavity implies that $\frac{c_{p\gamma}}{c_p} > 0$ for all smaller levels of correlation. Thus, there can be no stationary point at $\gamma < 0$. A similar argument show that if there is a stationary point at $\gamma < 0$ then it must be a maximum and there can be no stationary point at $\gamma > 0$. Finally, if $C(p,\gamma)$ has a local maximum at $\gamma = 0$ then there can be no $\gamma \neq 0$ for which there is a stationary point because at least one such hypothetical stationary point much be a minimum. However, this requires convexity, which we ruled out before.

In summary, it has been shown that $C(p, \gamma)$ can have no interior minimum and at most one interior maximum as a function of γ. Thus, $C(p, \gamma)$ is either increasing, decreasing, or single-peaked. This completes the proof.

References

Bebchuk, L.A. and J.M. Fried, 2003. Executive Compensation as an Agency Problem, Journal of Economic Perspectives, 17(3), 71-92.

Bebchuk, L.A. and J.M. Fried, 2004. *Pay without Performance: The Unfulfilled Promise of Executive Compensation*, Harvard University Press, Cambridge:MA.

Bertrand, M. and S. Mullainathan, 2001. Are CEOS Rewarded for Luck? The Ones without Principals Are. Quarterly Journal of Economics, 116(3), 901-932.

Blackwell, D. and Girshick, M.A., 1954, Theory of Games and Statistical Decisions, John Wiley & Sons.

Chambers, R.G. and J. Quiggin, 1998. A state-contingent production approach to principal-agent problems with an application to point-source pollution control, Journal of Public Economics, 70, 441-472.

Chambers, R.G. and J. Quiggin, 2000. *Uncertainty, production, choice, and agency: the state-contingent approach*. Cambridge University Press, Cambridge, UK.

Epstein, L. and S. Tanny, 1980. Increasing Generalized Correlation: A Definition and Some Economic Consequences, Canadian Journal of Economics, 13(1), 16-34.

Fleckinger, P., 2012. Correlation and Relative Performance Evanuation, Journal of Economic Theory 147(1), 93-117.

Grossman, S. J., and O. D. Hart, 1983. An Analysis of the Principal-Agent Problem, Econometrica, 51, 7-45.

Holmström, B., 1979. Moral Hazard and Observability, Bell Journal of Economics, 10(1), 74-91.

Holmström, B. and P. Milgrom, 1991. Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design, Journal of Law, Economics, & Organization, 7, 24-52.

Kim, S.K., 1995. Efficiency of an Information System in an Agency Model, Econometrica, 63 (1), 89-102.

Kirkegaard, R., 2017. Moral hazard and the spanning condition without the first-order approach, Games and Economic Behavior, 102, 373-387.

Lambert, R., 2001. Contracting theory and accounting, Journal of Accounting and Economics, 32, 3-87.

Larmande, F., 2013. Limited liability, the first-order approach, and the ranking of information systems in agencies, Economics Letters, 118, 314-317.

Roberts, J., 2004. *The Modern Firm*, Oxford University Press.

Rogerson, W., 1985. The First-Order Approach to Principal-Agent Problems, Econometrica, 53(6), 1357-1368.