

Contributions to Theoretical Economics

Volume 2, Issue 1

2002

Article 5

Risk Averse Supervisors and the Efficiency of Collusion

Antoine Faure-Grimaud*

Jean-Jacques Laffont[†]

David Martimort[‡]

*London School of Economics and Political Science, A.Faure-Grimaud@lse.ac.uk

[†]Universite de Toulouse (IDEI-GREMAQ), laffont@cict.fr

[‡]Universite de Toulouse (IDEI-GREMAQ), martimor@cict.fr

Copyright ©2002 by the authors. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher, bepress, which has been given certain exclusive rights by the author. *Contributions to Theoretical Economics* is one of *The B.E. Journals in Theoretical Economics*, produced by The Berkeley Electronic Press (bepress). <http://www.bepress.com/bejte>.

Risk Averse Supervisors and the Efficiency of Collusion

Antoine Faure-Grimaud, Jean-Jacques Laffont, and David Martimort

Abstract

This paper studies the efficiency of collusion between supervisors and supervisees. Building on Tirole (1986)'s results that deterring collusion with infinitely risk averse supervisors is impossible, while it is costless to do so under risk neutrality, we develop here a theory of collusion based on a trade-off between the risk premia required by (less extreme) risk attitudes and incentives. This allows us to link the efficiency of collusion to the supervisor's risk aversion and to various parameters characterizing the economic environment in which collusion may take place. We are then able to derive implications for the design of organizations, like determining how the number of tasks/agents per supervisor or the level of competition may impact on the cost of collusion, studying the impact of vertical integration on those same costs, or characterizing the role of uncertainty on side-contracting.

KEYWORDS: Supervision, collusion, risk aversion.

1 Introduction

The objective of this paper is to study supervisory structures in organizations and the fundamental trade-offs involved in the design of those structures. We envision supervision as the task of collecting signals about otherwise hidden information on the employees' activities. Central to our analysis is the necessity to deter collusion between supervisors and supervisees and the need to provide incentives to these coalitions to behave in a way which maximizes firms' profits instead of their own objectives. There are two main motivations for this paper.

The first stems from the importance of supervisory activities in organizations. For instance, Chandler (1962) has forcefully argued that changes in supervisory structures constitute the bulk of organizational innovations over the last century. Organizations devote large resources to supervision and do so, having in mind the threat of collusion between supervisors and supervisees. Given the importance of this issue for organizational design, it is crucial to identify the set of parameters likely to affect the efficiency of supervision. Answering this question requires an understanding of how the economic environment of the firm affects the efficiency of collusion inside the firm.

Our second motivation is to fill a gap in the collusion literature between two of its most quoted papers, namely Tirole (1986) and Tirole (1992). These papers offer a stylized model of a firm where the owner (thereafter the principal) has to hire a supervisor to collect information about a productive agent. The supervisor can conceal what he learns and can engage in a collusive side-contract with the agent when doing so favors his own interests. In Tirole (1986), it is assumed that exchanging bribes through this side-contract does not entail any dead-weight loss and only two extreme cases where the supervisor is risk neutral and infinitely risk averse are considered. The results are dramatically different in those two polar cases. Under risk neutrality, collusion is not a problem for the organization. Under infinite risk aversion, supervisory information is useless and collusion is most harmful to the organization.

Alternatively, the methodology followed by Tirole (1992) is to assume that the supervisor and the agent are both risk-neutral although protected by limited liability (an assumption which can be restated as saying that they are infinitely risk averse for negative wealth levels). Some exogenous transaction costs of transferring bribes are introduced to model collusion inefficiencies. The magnitude of these costs appears to crucially affect the performance of the organization. These transaction costs can be thought of as a short-cut for capturing unmodeled frictions in collusion like, for instance, the difficulty of transferring money between the colluding parties in the absence of any enforcement technology. Although this approach has proved to be extremely useful in studying how collusion threats affect economic outcomes, its major limitation lies precisely in its inability to relate the efficiency of collusion to the various parameters characterizing the environment where the firm evolves. Comparative statics exercises in this framework are only valid in so far as the modeler expects the efficiency of collusion not to change with exogenous perturbations of the model. The present paper offers one possible way of solving this difficulty.

As suggested by Tirole (1986), we start from the idea that the efficiency of collusion can be linked to the risk aversion of the supervisor. Surely, risk aversion of a firm's employees is a reasonable empirical assumption, but beyond the two polar cases emphasized by Tirole, little is known of its role in collusion and on the design of collusion-proof organizations. By allowing for some finite degree of risk aversion, we will characterize how the frictions of collusion are affected by the economic environment.

Our analysis shows that the cost for the principal of fighting collusion increases continuously with the supervisor's degree of risk aversion. The principal always prefers to avoid collusion between the supervisor and the agent. However, inducing information revelation by the supervisor requires giving him a reward when he reports an informative signal on the agent and a punishment otherwise. To prevent collusion, the risk averse supervisor is now subject to some risk and inducing him to participate in the grand-contract with the principal before he learns anything about the agent becomes costly. It can only be obtained by giving him a risk premium.¹ Consequently, the principal is worse off with a more risk averse supervisor as this trade-off between inducing information revelation and participation is more acute.

More generally, the trade-off between insurance and incentives will be more or less costly in terms of efficiency depending on various parameters characterizing the environment where collusion takes place. One dimension of the environment is the quality of information sources available to the supervisor. Focusing first on the cost of preventing collusion, we show that this cost is hump-shaped in the precision of information. Increasing the supervisor's information accuracy may subject the supervisor to additional risk and may increase this cost. However, we show that the benefit of increasing control on the agent always outweighs this first effect. Therefore an increase in the precision of supervisory information always increases the principal's welfare.

Second, we also investigate how to design the portfolio of monitoring tasks assigned to a supervisor. An important question is to find out whether it is better to pool different monitoring activities under the control of the same supervisor rather than having several supervisors, one for each activity. With constant risk aversion, incentives to prevent collusion on one task are designed independently of the incentives on other tasks. In such a context, we derive an irrelevance result which can be used as a benchmark: if the supervisor has constant absolute risk aversion, the principal's profit does not depend on the number of tasks allocated to the supervisor.

Third, we discuss the interaction between the external competitive pressure to which a firm in the market place may be subject to and its internal collusion problem. In a simple model with linear demand, we show that the firm's equilibrium output is less sensitive to competitive pressure as the supervisor is more risk averse. Indeed, as the risk premium needed to induce the latter's participation increases, expanding output becomes more difficult as if there were more competitors in the market. Everything happens as if the internal collusion problem exacerbates the competitive pressure and forces each competing firm to reduce its output plan more severely.

¹There is here an analogy with the standard moral hazard problem. Collusion can be viewed as a form of "hidden gaming", and deterring it is akin to inducing the right choice of action from the principal's point of view.

Fourth, we briefly discuss how vertical integration affects the incentives to collude. Vertical integration improves monitoring but it also facilitates collusion. Vertical integration can only be profitable to organizations having supervisors with a sufficiently low degree of risk aversion.

Finally, we study the impact of uncertainty in the firm's environment on collusion within the firm. We show that the timing of communication, i.e., whether the supervisor reports before or after the realization of some uncertainty, can be used by the principal to reduce the risk-premium he pays to ensure the supervisor's participation. By asking for reports before uncertainty is realized, the optimal contract is made less sensitive to the outside environment. Instead, the timing of communication would be irrelevant if there were no collusion or if the collusion technology was fixed as in Tirole (1992).

Few papers have proposed an analysis of the frictions of side-contracting beyond the exogenous transaction costs modeled by Tirole (1992). These contributions can be classified into two main categories: first, those introducing some ad hoc frictions but making them dependent on the environment; second, those giving some deeper foundations to these frictions. In the first class of models, Kofman and Lawarre (1996) analyze a situation where a supervisor can be corruptible or not depending on his preferences. In a political economy model of regulation, LaFont and Tirole (1993, Chapter 11) show that the optimal response to regulatory capture calls for a greater reduction in the power of regulatory incentives as the transaction costs of side-contracting between interest groups and the regulator are lower. Implicitly, the frictions of side-contracting depend on the ability of the group to organize itself and avoid the free-riding problem for collective intervention in the political arena. Still in a regulatory framework, LaFont and Martimort (1999) show that splitting information between two non-cooperating regulators makes collusion with the regulated firm harder. Collusion between a given regulator who is partially informed and the firm is now harder than with a single fully informed regulator since there is asymmetric information in side-contracting. In a model with reciprocal supervision, LaFont and Meleu (1997) argue that reciprocal favors are easier than asymmetric deals and that a norm of reciprocity is easier to enforce than a norm of asymmetric collusion in an organization.² Lastly, Martimort and Verdier (2000) build a Schumpeterian growth model showing that colluding agents are willing to divert resources away from productive activities in order to improve collusive technologies when they have better prospects of remaining in a dominant firm. Frictions there depend on the initial stock of resources available.

Contributions providing foundations for the transaction costs of side-contracting are even scarcer. Martimort (1997, 1999) and Martimort and Verdier (2002) derive conditions that make a collusive agreement self-enforceable. The dead-weight loss of collusion depends on the respective discount rates of the principal and the agents and on the information structure.³ In Faure-Grimaud, LaFont and Martimort (2001), we focus on a static principal-supervisor-agent with soft information where coalition formation is subject to

²In particular, they show that asymmetric supervision may be optimal because it eliminates possible reciprocal favors.

³In a somewhat different vein, Felli (1997) shows how the self-enforceability of a contract can be used by the principal to better fight collusion.

frictions arising from the existence of asymmetric information between the supervisor and the supervisee. The principal can actually play on these frictions to limit the efficiency of side-contracting and to improve the firm's profitability. In Faure-Grimaud, LaFont and Martimort (2000), we study a model of delegation where the principal lets the supervisor directly contract with the agent. In such a model, collusion is by definition no longer an issue since there is no grand-contract ruling the whole organization but instead a sequence of vertical relationships. Nevertheless, the top principal designs the contract of the intermediate supervisor to make him internalize his own objectives. The cost of doing this depends again on the supervisor's risk aversion. Beyond modeling differences⁴, this latter paper does not study the impact of information accuracy nor how the design of organizations affect the frictions of collusion.⁵ In Faure-Grimaud, LaFont and Martimort (1999), we apply the same framework than in the previous paper and study a model of delegated auditing where the probability of audit is chosen endogenously. We show there that the equilibrium probability of audit goes down when the auditor to whom that task is delegated is more risk averse. Finally, Faure-Grimaud and Martimort (2001) study a situation where an intermediary between the principal and the productive agent is always needed to allow the principal to have access to the agent. Because he does not want to bear any risk, this uninformed intermediary may solve the incentive problem vis à vis the agent in a way that the principal finds sub-optimal. This results in some agency costs of intermediation, unless the intermediary is risk neutral.

Section 2 presents our model. Section 3 derives the optimal contract in the case where there is no collusion between the supervisor and the agent. Risk aversion plays no role in this non-cooperative implementation. In Section 4, the optimal collusion-proof grand-contract is derived as well as some comparative statics. This section highlights the role of the supervisor's risk aversion in the design of incentives. Section 5 presents some comparative statics and links the efficiency of collusion to various parameters of the economic environment. Section 6 derives some results about the design of organizations under the threat of collusion. Section 7 concludes. All proofs are gathered in the Appendix.

2 The Model

2.1 Players and Information

We consider a two-tier model of a firm in which productive and supervisory tasks are split. A principal, for instance the firm's owner, contracts with a productive agent and a supervisor. The separation between ownership, production and supervision is motivated by physical constraints. The principal himself is unable to produce or supervise either because the activities of the firm are large in size or because those tasks require some specific skills.

⁴In Faure-Grimaud, LaFont and Martimort (2000), the supervisor's information is soft, the optimal centralized mechanism is not considered and the timing of contracting is different.

⁵Our main concern there was to identify two modelings of collusion and to derive from this identification closed-form formula for the transaction costs of side-contracting.

The agent produces a quantity q of output at a constant marginal cost μ . μ is a piece of private information known only to the agent. It is drawn from a discrete distribution on $\mathcal{L} = \{\mu_1; \mu_2\}$ (we denote $\mu = \mu_2$ if $\mu_1 > 0$) with respective probabilities β and $1 - \beta$.

The supervisor receives a signal α on the agent's marginal cost. This signal α can take either of two possible values. We denote by $\mathcal{T} = \{\alpha; \bar{\alpha}\}$ the set of possible signals. Conditionally on the fact that the agent is efficient, i.e., $\mu = \mu_1$, the supervisor observes α with probability β . Otherwise, the supervisor observes $\bar{\alpha}$. Hence, one can think of α as a piece of revealing evidence on the fact that the agent has type μ_1 . Instead, $\bar{\alpha}$ is a non-revealing signal (but still conveying some information).

The signal is partially verifiable in the sense of Green and LaFont (1986). Only α can be manipulated by the agents who can pretend that $\bar{\alpha}$ has instead been realized. The signal $\bar{\alpha}$ cannot be manipulated at all. One can also think of $\bar{\alpha}$ as a hard information signal which can be hidden (the agents can pretend that a non-revealing signal has been received when it is a revealing one). For instance, supervisory information can be obtained by disclosing documents on the agent's performance.⁶ These pieces of information can be easily hidden if they are revealing; it can be much harder and even impossible to report convincing evidence when there is none.⁷

The joint probabilities p_{ij} on the pairs $(\mu_i; \alpha_j)$ are defined respectively as $p_{11} = \beta^2$, $p_{12} = \beta(1 - \beta)$, $p_{21} = 0$, $p_{22} = (1 - \beta)^2$. The supervisor's signal is not observed by the principal, otherwise a supervisor would not be needed. However, this signal is also learned by the agent. Nature reveals to the agent both his type and the supervisor's information; only the latter is available to the supervisor while the principal observes none of these.⁸

2.2 Preferences.

The supervisor is risk averse and has a CARA utility function⁹ $V = v(s) = \frac{1}{r}(1 - e^{-rs})$; where s is the wage he receives from the principal.¹⁰ The supervisor has no productive role and is only used by the principal to bridge the informational gap with the agent.

The agent is infinitely risk averse below zero wealth and risk neutral above. For

⁶See Bull and Watson (2000) for such a model of evidence disclosure.

⁷The case of supervisory signals which are soft information is analyzed in Baliga (1999) and Faure-Grimaud, LaFont and Martimort (2001). In this latter paper, we show that the possibility of complete manipulation of this information makes it useless in the case where collusion between the supervisor and the agent takes place under symmetric information. Instead, collusion under asymmetric information restores some of the screening ability for the principal at least when the supervisor is risk averse.

⁸This nested information structure is standard in both the literature on collusion and the related literature on delegation in hierarchies (see respectively Tirole (1986, 1992), McAfee and McMillan (1995) and Faure-Grimaud, LaFont and Martimort (2000) among others).

⁹Holmstrom and Milgrom (1990) and Itoh (1993) develop models of collusion between two risk averse agents in a pure moral hazard context. Prendergast and Topel (1996) analyze a model of favoritism in a pure moral hazard context with agents and their supervisor having all CARA utility functions. Relaxing the CARA assumption in our model could be done at the cost of some added complexity without giving many new insights.

¹⁰ $r = 0$ corresponds to the limiting case where the supervisor is risk neutral, $v(s) = s$.

positive payoff, his utility function can thus be written as $U = t \circ \mu q$ where t is the monetary transfer he receives from the principal. The agent produces as long as he gets his reservation utility which is normalized to zero.¹¹

Production of q units of output yields an increasing and concave revenue function $R(q)$ to the principal ($R'(q) > 0$ and $R''(q) \leq 0$). To ensure positive production levels and avoid corner solutions, we assume that the Inada conditions are satisfied, i.e., $R'(0) = +\infty$ and $R'(\infty) = 0$ with $R(0) = 0$. The principal's profit writes as: $\pi = R(q) \circ s \circ t$:

2.3 Contracts

Grand-Contracts: The organization is ruled by the principal through a **grand-contract GC**. From the Revelation Principle,¹² as long as the agent and the supervisor do not collude, there is no loss of generality in restricting the principal to offer truthful direct revelation mechanisms of the kind $GC = \{t(m_a; m_s); s(m_a; m_s); q(m_a; m_s)\}$. m_a is the agent's report to the principal. This report belongs to $\mathcal{M} \subset \mathcal{T}$. m_s is the supervisor's report to the principal which lies instead in \mathcal{T} . To make notations simpler, we denote thereafter t_{ijk} (resp. s_{ijk} and q_{jk}) the agent's transfer (resp. the supervisor's transfer and the agent's output) when the agent reports $(\mu; \alpha)$ and the supervisor reports β instead, for $(i; j; k)$ in $\{1; 2\}^3$. When the agent's and the supervisor's reports coincide, we denote by t_{ij} (resp. s_{ij} and q_j) this transfer (resp. the supervisor's transfer and the agent's output). We also denote by $u_{ij} = t_{ij} \circ \mu q_j$ the agent's ex post information rent when his type is μ and both the supervisor and the agent report α .

The information structure limits somewhat the possible manipulations of the agent's and the supervisor's reports. Indeed, the agent necessarily reports a type μ_i when he reports also that the supervisor's signal is α . Otherwise, the reports would be inconsistent given the common knowledge information structure. Moreover, only α can be manipulated and both the supervisor and the agent can then pretend that α has instead been realized. Again, the reverse is impossible. We denote by \mathcal{I}_j the set of reports on μ compatible with α . From the discussion above, we have thus $\mathcal{I}_1 = \{\mu_1\}$ and $\mathcal{I}_2 = \{\mu_1; \mu_2\}$.

Collusive Side-Contracts: The **side-contract** between the supervisor and the agent consists of **first**, a secret side-transfer \emptyset paid by the agent to the informed supervisor when $\alpha = \beta$ and **second**, a coordination of the supervisor's and the agent's individual reports $(m_a; m_s)$ in this state of nature to the principal. Given that the knowledge of α perfectly reveals the agent's type to the supervisor, collusion takes place under complete information. For simplicity, the supervisor has all the bargaining power at the side-contracting stage and makes a take-it-or-leave-it offer to the agent. The colluding partners

¹¹Our results obtain for any utility function for the agent as long as we maintain this assumption of infinite risk aversion below zero wealth. This assumption is made for tractability as, in the absence of collusion, it leads to a simple trade-off between efficiency and rent extraction. A similar but more complex trade-off would also arise with a positive risk aversion and ex ante contracting as we model here (for pure adverse selection models based on the trade-off between incentives to reveal and insurance see Salanié (1990) and Laffont and Rochet (1999) in the case of a continuum of states).

¹²The Revelation Principle holds in the case of signals which are partially verifiable as in this paper (see Green and Laffont (1986)).

are able to commit to this side-contract.

Finally, it is worth stressing that side-contracting suffers a priori from no exogenous frictions. The collective gains from a joint manipulation of reports can be fully exploited by the collusive partners. One unit of bribe taken from the agent is thus fully pocketed by the supervisor.

2.4 Timing

The timing for the game of contractual offer cum coalition formation is as follows:

- ≤ The principal offers a grand-contract to both the supervisor and the agent.
- ≤ The supervisor and the agent both simultaneously accept or refuse this grand-contract at the ex ante stage, i.e., being still uninformed on the agent's type and the supervisory signal. If either of them refuses, the game ends.
- ≤ The agent learns his productivity parameter μ and the supervisor's signal ω . The supervisor learns only ω .
- ≤ When ω realizes, the supervisor makes a take-it-or-leave-it offer of a collusive side-contract to the agent. If the latter refuses, the grand-contract is played non-cooperatively. If he accepts, the colluding partners commit to a joint manipulation of their reports to the principal and to a bribe.
- ≤ Reports are made, production takes place and transfers within the grand-contract and (possibly) within the side-contract, are paid.

Note that the acceptance of the grand-contract by both the supervisor and the agent takes place before the learning of any information. Hence, the supervisor's and the agent's ex ante participation constraints must be satisfied by this grand-contract. Because of our assumption of infinite risk aversion below zero wealth for the agent, the latter's ex ante participation constraint amounts to a set of ex post participation constraints, one in each state of nature.¹³ With this timing, the principal has the maximal ability to commit by designing the contours of the organization before any learning of information. This seems to be the most relevant assumption in the context of the theory of the firm.

3 Benchmarks

3.1 Costless Supervision

Let us first consider the case where the principal directly receives the signal ω on the agent's private information. This can be viewed as a stylized model of a small firm in

¹³Had the agent been risk neutral, ex ante contracting would allow the principal to extract all the agent's rent. In such a model, the transfers given to the agent can nevertheless be structured to leave no incentives to lie and can destroy the scope for collusion with the supervisor.

which the supervisory task can be performed by the principal himself. Alternatively, if we stick to the interpretation of our model as a picture of a large firm in which supervision is needed, everything happens as if the supervisor would costlessly reveal truthfully his information to the principal before the latter contracts with the agent.

When the principal learns α , he can infer for sure the value of the agent's type and there is no longer any informational gap between him and the agent.

When instead α has been observed, the principal is still uninformed on the agent's type. As it is standard in two-type adverse selection models, the following constraints are of particular importance:¹⁴

≤ The incentive compatibility constraint of an efficient agent when the principal has observed α :

$$u_{12} \geq u_{22} + \phi \mu q_2; \quad (1)$$

≤ The ex ante participation constraint of the infinitely risk averse agent can be decomposed into two relevant ex post participation constraints:

$$u_{22} \geq 0; \quad (2)$$

when the principal has observed α , and

$$u_{11} \geq 0 \quad (3)$$

when the principal has instead observed β .

Accordingly, the optimal contract solves:

$$\max_{\{q_j; u_{ij}\}_{(i,j)}} \sum p_{ij} (R(q_j) - \mu q_j - u_{ij})$$

subject to (1)-(2) and (3).

Solving this problem yields the conditional optimum defined as:

$$R^0(q_{1j}^d) = R^0(q_{2j}^d) = \mu_1 \quad \text{for } j \text{ in } \{1, 2\} \quad (4)$$

$$R^0(q_{22}^d) = \mu_2 + \frac{\int (1 - \phi) \mu}{1 - \int \phi \mu} \quad (5)$$

To reduce the cost of the incentive compatibility constraint (1) and make it less attractive for an efficient agent to mimic an inefficient one, the principal reduces the output produced by an inefficient agent. The efficient agent's output remains equal to its first best value. A positive rent is left to the efficient agent ($u_{12}^d = \phi \mu q_{22}^d$) only when the principal gets a non-revealing supervisory signal. The participation constraints (2) and (3) are both binding. Finally, as the supervisory information becomes less informative, i.e., as ϕ decreases to zero, the output distortion characterized in (5) increases.

¹⁴When the following constraints are binding, as it will be the case at the optimum of the principal's problem, it is easy to show that the remaining incentive and participation constraints are strictly satisfied.

3.2 The Collusion-Free Outcome

Let us now envision the case where the agent and the supervisor do not collude. They report their information non-cooperatively to the principal who is uninformed of the realization of θ . We then look for a truthful Bayesian-Nash equilibrium between the agents.

≤ When $(\mu; \theta)$ has been learned by the agent, the agent's incentive compatibility constraint is:

$$u_{ij} \geq t_{ij} - \mu q_{ij}; \quad \text{for all } \mu \text{ in } j^0 \text{ and } j^1 \geq j: \tag{6}$$

≤ When θ has been learned by the supervisor, the supervisor's incentive compatibility constraint is:¹⁵

$$\sum_i p_{ij} v(s_{ij}) \geq \sum_i p_{ij} v(s_{ij}^0) \quad \text{for all } j^1 \geq j^0: \tag{7}$$

Here, we can use the logic of Nash implementation.¹⁶ The signal θ is a piece of information which is commonly known by the agent and the supervisor. Hence, it can be costlessly extracted by the principal by setting $s_{ij} = t_{ij} = \theta q_{ij} = \theta 1$ when the agent's and supervisor's respective reports on θ differ, i.e., when $j \neq j^0$. In this case, (6) can be reduced to the only relevant incentive constraint (1). Similarly, (7) is necessarily satisfied since its right-hand sides for $j > j^0$ are then infinitely negative.

Finally, the supervisor's ex ante participation constraint writes as:

$$\sum_{(i;j)} p_{ij} v(s_{ij}) \geq 0: \tag{8}$$

The candidate for the optimal contracting outcome with a non-cooperative behavior between the supervisor and the agent is thus the conditionally optimal outcome described in Section 3.1. By always giving to the supervisor a zero wage $s_{ij} = 0$ for all $(i;j)$, the principal eliminates the risk borne by the supervisor and satisfies his ex ante participation constraint (8).

If the principal can perfectly control and forbid communication between the agent and the supervisor, he can achieve the same outcome as with direct supervision. Importantly, this result is independent of the supervisor's degree of risk aversion when the agents do not collude.

In the analysis of this section, we did not insist on unique Nash implementation. It is however easy to ensure uniqueness by offering an arbitrarily small positive payoff to the supervisor if he reports the revealing signal.

¹⁵We multiply this constraint by $\prod_i p_{ij} > 0$ to express the constraint as a function of ex ante probabilities rather than as a function of conditional probabilities. Note that this incentive constraint is Bayesian only when $\theta = \theta$. When $\theta = \theta$, the agent's type is perfectly known by the supervisor.

¹⁶See Maskin (1999).

4 Collusive Behavior

The non-cooperative implementation of the collusion-free outcome above is somewhat unrealistic since it is not immune to a collective manipulation of the agent's and supervisor's reports on ω . Indeed, the supervisor can be bribed by the μ_1 agent when ω has been observed so that they both claim that ω_2 has instead realized. By doing so, the supervisor and the agent can share the rent $\phi \mu_1 q_{22}^d$ which goes to the μ_1 agent when the principal receives a non-revealing signal. Henceforth, we consider the case where the supervisor and the agent collude against the principal through a binding side-contract when $\omega = \omega_2$ has been observed.

4.1 Collusion-Proofness Constraints

Following Tirole (1986) and LaFont and Tirole (1993, Chapter 11), the Collusion-Proofness Principle applies in this environment. There is no loss of generality in restricting the principal to offer collusion-proof grand-contracts. For such a contract, the best side-contract consists of no side-transfer and no collective manipulation of the reports made by both the supervisor and the agent when $\omega = \omega_2$. This last requirement means that the colluding partners must maximize their collective surplus by reporting that μ_1 has realized when ω_2 has been observed. To be collusion-proof, a grand-contract must thus satisfy the following coalition incentive compatibility constraints:

$$s_{11} + u_{11} \geq s_{12} + u_{12} \quad (9)$$

and

$$s_{11} + u_{11} \geq s_{22} + u_{22} + \phi \mu_1 q_{22}^d \quad (10)$$

The right-hand sides above correspond to what the coalition can get by manipulating the agent's and the supervisor's common report on ω_2 and claiming that instead ω_1 has been realized, and that $\mu = \mu_1$ (equation (9)) or that $\mu = \mu_2$ (equation (10)). In this last case, the coalition manipulates the agent's report on his type μ_1 and we get two possible values for these right-hand sides.

Note that (9) will still be binding at the optimal collusion-proof grand-contract because there is no need to give some extra rent to an efficient agent when the supervisor reports truthfully having observed no revealing signal. When this latter constraint is binding, the more stringent constraint between (9) and (10) is that obtained for the highest of the two wages s_{12} and s_{22} . The principal cannot distinguish between these two transfers since the coalition can always pretend to be in the state of nature with the highest supervisory wage. Hence, the principal loses much flexibility in the supervisor's wage and we must necessarily have:

$$s_{12} = s_{22} = s_2 \quad (11)$$

where s_2 is a constant wage received by the supervisor whenever he claims to have observed a non-revealing signal ω_2 .

The relevant coalition incentive compatibility constraint thus writes as:

$$s_{11} + u_{11} \geq s_2 + u_{22} + \phi \mu_1 q_{22}^d \quad (12)$$

Using this simplification in the expression of the supervisor's wages, we can also rewrite the supervisor's ex ante participation constraint as:

$$\int \Delta v(s_{11}) + (1 - \int \Delta v(s_2)) \geq 0; \tag{13}$$

4.2 Characterizing the Optimal Collusion-Proof Contract

The principal maximizes the firm's expected profit subject to coalition incentive, individual incentive and participation constraints. The optimal grand-contract thus solves thus the following problem (denoted thereafter by (P)):

$$\max_{\{q_j; u_{ij}; s_{11}; s_2\}_{(i,j)}} \sum_{(i,j)} p_{ij} (R(q_j) - \mu q_j - s_{ij} - u_{ij})$$

subject to constraints (1)-(2)-(3)-(12) and (13).

Proposition 1 The optimal collusion-proof grand-contract entails:

• Constraints (1)-(2)-(3)-(12) and (13) are all binding. All other omitted constraints are strictly satisfied.

• A decreasing schedule of outputs with no distortion for the most efficient agent

$$q_{i1}^c(r) = q_{i2}^c(r) = q_i^d$$

and a downward distortion for the inefficient agent $q_{b2}^c(r)$ ($q_{b2}^c(r) < q_{b2}^d$) which is implicitly defined by:

$$R'(q_{b2}^c(r)) = \mu_2 + \frac{\int \mu_1 \frac{e^{-r \mu_1 q_{b2}^c(r)}}{\int e^{-r \mu_1 q_{b2}^c(r)}}}{\int e^{-r \mu_1 q_{b2}^c(r)}}; \tag{14}$$

• The supervisor's wages in the different states of nature are respectively given by:

$$s_{11}^c = \mu_1 q_{b2}^c(r) + \frac{1}{r} \ln \int e^{-r \mu_1 q_{b2}^c(r)} > 0; \tag{15}$$

$$s_2^c = \frac{1}{r} \ln \int e^{-r \mu_1 q_{b2}^c(r)} < 0; \tag{16}$$

To better understand the distortion of the collusion-proof contract, let us first start by describing what happens when the supervisor is risk neutral, i.e., $r = 0$. In this case, we can have some supervisory wages such that the coalition incentive constraint (12) and the supervisor's participation constraint (13) are both satisfied and the principal implements the conditionally optimal outcome costlessly. To do so, we find the wages that make those two constraints binding. We observe that s_{11}^c is equal to a strictly positive reward $(\int \mu_1 q_{b2}^c)$; while s_2^c is instead negative $(-\int \mu_1 q_{b2}^c)$.¹⁷ From (14), we

¹⁷Note in fact that the wages defined by (15) and (16) converge towards these values as r goes to zero. However, in the limiting case $r = 0$, other pairs of wages satisfy the coalition incentive compatibility constraint and are such that the participation constraint is binding. Another set of such wages can be obtained by making the supervisor residual claimant for the hierarchy's profit. A scheme $s(q) = R(q) - T$ where T is a fixed-fee designed to extract all the supervisor's rent can also do the job.

also obtain that if $r = 0$ the optimal outputs are those obtained in the absence of collusion: $q_2^c(r = 0) = q_2^d$. Together, this means that collusion has no bite if the supervisor is risk neutral. The intuition is simple. When the supervisor reports that the agent is efficient, the principal must increase the wage he pays him above his wage for the non-revealing report. Indeed, this wage differential must exceed the maximal bribe μq_2^d that the agent is willing to pay. Imposing this risk on the supervisor is not costly for the principal as long as the supervisor is risk neutral. Therefore, there is no need to alter the output levels compared to the collusion-free outcome. Notice however that the collusion threat prevents the principal from giving the supervisor the same wage in all states of nature.

The reader will have recognized an argument often made in pure moral hazard contexts. The choice of whether to hide or not a revealing signal can actually be viewed as a binary moral hazard decision made by the supervisor. It is well-known from the moral hazard literature¹⁸ as well as from the adverse selection literature with ex ante contracting¹⁹ that, under risk neutrality, the principal can achieve the same outcome as in the case where the moral hazard variable could be directly contracted upon by the principal. To achieve this outcome, the principal can simply make the supervisor residual claimant for the firm's profit by selling him the firm for an ex ante fixed-fee.

With risk aversion, implementing group incentives becomes costly for the principal. Risky monetary transfers with the same expected values are worth less for the risk averse supervisor than for the principal. A risk averse supervisor accepts the wage lottery proposed by the principal only if he receives a risk premium. Now, the principal faces a real trade-off between preventing collusion and giving insurance to the supervisor. The total risk borne by the supervisor depends on the production plan since, to deter collusion, the supervisor's wage differential needs to exceed the collusive stake which is worth μq_2^c . The greater the output q_2^c , the more risky should the wage lottery faced by the supervisor be. As a result, allocative distortions now become valuable for the principal. Distorting the output level downwards reduces the risk borne by the supervisor and also the conflict between coalition incentive compatibility and participation constraints.

It is also worth commenting on the form taken by the optimal collusion-proof contract. Everything happens as if the principal relies on a sequential implementation of the second best outcome which seems to be in line with some of the real-world practices observed within the firm. First, the principal calls for a report on θ made by the supervisor only. Conditional on the fact that this report is non-revealing, the principal then asks the agent for a report on his type μ . If the supervisor's report is instead revealing, the principal extracts all the agent's information rent since his type is known to be μ_1 for sure. A process of sequential reports within the organization is thus weakly optimal in face of the threat of collusion.²⁰

¹⁸See Holmstrom (1979), Shavell (1979) and Grossman and Hart (1983) among others.

¹⁹See Salani  (1990) and La ont and Rochet (1999).

²⁰Of course, in practice and as most results in mechanism design, the mechanism above requires strong commitment from the principal. Once the collusive partners have committed themselves not to collude after a collusion-proof contract has been accepted, the principal has an incentive to renegotiate with the supervisor and give him more insurance. This conflict between renegotiation and collusion is analyzed in Felli and Villas-Boas (2000).

5 Comparative Statics

This section discusses how the optimal response of the organization changes with the main parameters of the model.

5.1 The Role of Risk Aversion

We have emphasized that without risk aversion, the coalitional incentive problem can be trivially solved. We now fully characterize how risk aversion affects the optimal production plan, the wage schedule and the principal's welfare.

Proposition 2 The impact of supervisor's risk aversion is the following:

$q_{b2}^c(r)$ is a decreasing function of r with $q_{b2}^c(0) = q_{b2}^d$. When r goes to infinity, $q_{b2}^c(r)$ converges towards q_{b2}^u defined by:

$$R^0(q_{b2}^u) = \mu_2 + \frac{1}{1-\phi} \phi \mu \quad (17)$$

When r goes to zero, the supervisor's wages s_{11} , (resp. s_2), converge to $(1-\phi) \mu q_{b2}^d$, (resp. $\phi \mu q_{b2}^d$). When r goes to infinity, the supervisor's wages s_{11} , (resp. s_2), converge to $\phi \mu q_{b2}^u$, (resp. 0).

The principal's welfare monotonically decreases with r .

The negative impact of r on the ability to use a supervisor is clear.²¹ Preventing collusion requires making the supervisor's wage risky. The cost of doing so is the risk premium that must be given to the supervisor. This cost is increasing in r .

Further intuition can be gained in case of a small uncertainty on cost parameters and thus on collusive stakes (e.g. supposing that $\phi \mu$ is small enough). Using Taylor expansions, we find that:²²

$$s_2^c = \phi \mu q_{b2}^c(r) + \frac{r}{2} (1-\phi) \phi \mu^2 (q_{b2}^c(r))^2 \quad (18)$$

The first term on the right-hand side of (18) represents the supervisor's negative wage received in the case of risk neutrality when the agent is inefficient and the supervisor has reported nothing to the principal. The second term on this right-hand side is actually the risk premium that the principal must pay to the risk averse supervisor to induce his participation. As s_2^c is paid with probability $1-\phi$ and s_{11}^c is paid with probability ϕ ²³

²¹See also Tirole (1986, Proposition 4) and the discussion in Faure-Grimaud, LaFont and Martimort (2001).

²²It is worth stressing that the Taylor expansions below also hold when the agent's utility function is not CARA. The value of r to be used is then the degree of risk aversion at zero wealth.

²³Since $u_{12}^c = u_{22}^c = 0$, (12) yields indeed:

$$s_{11}^c = s_2^c + \phi \mu q_{b2}^c(r):$$

the overall extra agency cost due to risk aversion, i.e., the expected extra wage paid to the supervisor with respect to what he receives in the absence of collusion, then becomes:

$$\int \mathbb{E}^c_{11} + (1 - \int) \mathbb{E}^c_{22} = \frac{r}{2} \int (1 - \int) \mathbb{E} \mu^2 (q_{22}(r))^2: \quad (19)$$

The right-hand side of (19) is precisely the cost of deterring collusion. It is just equal to the risk premium that must be added to the supervisor's expected wage to make him accept the risky lottery of wages necessary to induce collusion-proofness. We already know from the literature on decision-making under uncertainty that, when the so-called Arrow-Pratt approximation holds,²⁴ the risk-premium increases as the square of the size of the risk.

In our context, this property highlights the non-linearity of the cost of deterring collusion (for small collusive stakes) whenever $r > 0$. The expected wage to be paid to the supervisor increases and is convex in the collusive stake $\mathbb{E} \mu_{22}$: Intuitively, this means that collusion is at the margin less costly as collusive stakes are lower. The resulting output distortions can be defined by:

$$R^q(q_{22}(r)) = \mu_2 + \frac{\int}{1 - \int} \mathbb{E} \mu (1 - \int) + r \frac{\int \mathbb{E}}{1 - \int} (1 - \int) \mathbb{E} \mu^2 q_{22}(r): \quad (20)$$

It follows from the previous observations that firms prefer to hire less risk averse supervisors and this preference is more pronounced when collusive stakes are higher. Typically, this is the case when the uncertainty about the agent, as measured by $\mathbb{E} \mu$; is larger.

5.2 The Accuracy of Supervisory Information

Taking risk aversion as given, we now investigate the impact of improving the precision of the supervisory signal. This improvement can be obtained by innovations in monitoring technologies or simply by using external sources of information. In this respect, the performances of firms in related environments subject to correlated shocks provide useful signals to improve control within the hierarchy. Market interactions thus provide information which affects the cost of inside collusion within the firm.

To get an idea of the trade-off involved when information is improved, it is useful to use the Taylor expansion (19). Indeed, this formula holds not only for optimal outputs but also for any other output q_{22} as long as $\mathbb{E} \mu$ is small enough. For a fixed output, we observe that improving the precision of the supervisory information, a priori has an ambiguous impact on the cost of collusion. This cost is proportional to the variance of \mathbb{E} and it has an inverted U-shape because increasing continuously \mathbb{E} first makes the uncertainty borne by the risk averse supervisor increase and then decrease as long as $\int > \frac{1}{2}$.²⁵ On the benefit side, improving supervisory information helps the principal to reduce the agent's information rent.

²⁴See Gollier (2001, p. 22)

²⁵For \mathbb{E} large enough and if $\int > \frac{1}{2}$, there is no trade-off from the principal's point of view: an increase in \mathbb{E} also reduces the expected risk to which the supervisor is subject and thus his expected wage.

Proposition 3 The impact of the accuracy of the supervisory information is the following:

≤ When ≤ increases towards 1 (perfectly informative signal), the inefficient agent's output $q_2^c(r)$ increases towards the output q_2^a defined as:

$$R^Q(q_2^a) = \mu_2 + \frac{\int e^{r\phi} \mu_1}{1 + \int e^{r\phi} \mu_2} \quad (21)$$

≤ The principal's welfare is increasing in ≤

Notice that an increase in the information accuracy cannot hurt the principal: he could always induce the supervisor to report an informative signal with probability less than one, resulting in an outcome equivalent to what is obtained at lower accuracy levels.

Mathematically, the impact of the accuracy of information on the principal's welfare can be best seen by looking at the impact of a change of ≤ on the constrained set and on the principal's objective function (see Figure 1 below).




Figure 1: Optimization in the $(s_2; s_{11})$ plan.

First, note that an increase in \leq leaves unchanged the coalition incentive compatibility constraint. Note also that, for $u_{11} = u_{22} = 0$, this constraint defines a straight line which cuts the north-west quadrant of the $(s_2; s_{11})$ space. At the same time, the set of wages $(s_2; s_{11})$ that can be accepted by the risk averse supervisor lies above a convex curve which turns clockwise around the origin as \leq increases. Therefore, an increase in the accuracy of information results in a smaller constrained set. If the principal's expected welfare was kept the same, this would result in a higher expected cost of preventing collusion with both s_{11} and s_2 being raised. The principal relies more on rewards and less on punishments as the supervisory information becomes more precise. However, at the same time, the principal's objective function changes with \leq and the increase in the expected payment to the supervisor that results from the modification of the constrained set must be compared with the extra saving made from giving less often an information rent to the agent.

This comparison is rather straightforward. Intuitively, the expected reduction of the agent's information rent always exceeds the possible increase in the supervisor's wage cost. Indeed, the supervisor, when he can prove that the agent is efficient, is paid less than what the agent gets in the absence of supervision. Remember that, to satisfy the coalition incentive compatibility constraint, the supervisor must get a reward when he reports an informative signal and a penalty otherwise and that such a risky lottery must be accepted by the supervisor because he does not know the agent's type at the time of joining the firm. The possibility of using a penalty to induce information revelation makes it cheaper for the principal to obtain information from the supervisor than from the infinitely risk averse agent who cannot be punished. The advantage of the supervisor as a source of information comes from his better risk bearing attributes.

This gain of using the supervisor is the driving force behind the Collusion-Proofness Principle in this context. If using the supervisor is less costly than paying the agent directly, it pays to use the former and give him a collusion-proof wage rather than relying solely on the agent's report and giving the latter an information rent.

Notice also that, as supervisory information becomes more precise, the second-best output converges towards the output obtained in a simple principal-agent hierarchy with ex ante contracting and an agent having the same utility function as the supervisor.²⁶ In fact, as the supervisor gets almost perfect information on the agent, this coalition behaves almost as a single agent having a degree of risk aversion which is the minimum between that of the supervisor and that of the agent, i.e., r . Had the principal directly contracted with the agent, the second-best distortion would instead be obtained by replacing the right-hand side of (21) by the usual virtual cost $\bar{c} + \frac{1}{1-\alpha} \phi \mu$ which is greater. Again, this points to the superiority of contracting with a supervisor even if he shares all relevant information with the agent.

Instead, in the neighborhood of $\leq = 0$, the supervisor becomes useless for the principal. The three tier hierarchy reduces to a standard principal-agent pair.

When collusion takes place under complete information, the principal always gains from improving the technology of monitoring. Let us instead suppose that the supervisor

²⁶See LaFont and Martimort (2002, Section 2.11.2) for the derivation of a similar result.

only gets an imperfect signal on the agent's type and is always unsure about the latter's cost parameter. Collusion between the supervisor and the agent then takes place under asymmetric information and some frictions in side-contracting arise from this. This is the setting we have analyzed in Faure-Grimaud, LaFont and Martimort (2001) where, assuming that the supervisor's signal is soft information, we show that the principal may also play on the degree of asymmetric information between the supervisor and the agent to increase those frictions and fight collusion more easily. To do so, he chooses a signal with an interior precision.

We conclude this section with the remark that firms prefer hiring supervisors who are better able to observe the agent's cost. This preference is more pronounced when collusive stakes and the degree of risk aversion of these supervisors are higher. This suggests a substitutability between supervisory accuracy and risk aversion. A more risk averse supervisor can be preferred to a less risk averse one if the former has a more precise signal. It also leads to the prediction that firms evolving in a more risky environment are willing to pay more for accurate supervisory information than firms evolving in stable environments.

6 Collusion Costs and Organizational Design

In this section, we derive from the previous trade-off between insurance and incentives some insights on the design of organizations. Admittedly this section is more exploratory, but our modest goal here is only to show that in a multi-agent context, designing the organization to ease collusion deterrence is a complementary tool to better provision of individual incentives.

6.1 The Possible Irrelevance of Supervisory Structures

The design of supervisory structures for multi-agent firms entails choosing the span of control for a given supervisor on supervised agents. We now give a simple example in which the supervisory structure is actually irrelevant. This example thus offers a benchmark with respect to which one can assess how changes in the information structure may create economies or diseconomies of scope in supervision.

Suppose that the firm is involved in producing two lines of products which are technologically unrelated. There exists no interaction between the product lines either in terms of their demands or their costs. Each of the two different agents produces a different product and the cost parameters of those agents are independently distributed. If one supervisor is chosen for each line of product, the firm can be viewed as twice the replica of the one-product-line firm that we have analyzed so far. In this case, supervisors get signals on the product lines they respectively control which are again independently distributed. The principal's revenue is the sum of the revenues obtained on each product separately.

If a single supervisor is chosen to control both products, he will receive a wage $s_{ij} + s_{kl}$ with probability $p_{ij} p_{kl}$ when information on product line 1 is $f_{\mu_j; \sigma_j}$ and information on product line 2 is $f_{\mu_k; \sigma_k}$.²⁷ The supervisor's ex ante participation constraint in this multi-task environment is thus:

$$\sum_{(i,j;k,l)} p_{ij} p_{kl} v(s_{ij} + s_{kl}) \geq 0 \quad (22)$$

In this context, we consider the collusive side-contracts between the common supervisor and each productive agent that are only bilateral. Coalition incentive constraints then take the same form as previously and we obtain:

Proposition 4 Having one supervisor per product line or only a single supervisor for both product lines yields the same profit to the principal.

Assuming CARA preferences is key for this irrelevance result. With a CARA utility function, adding independent risks has no effect on the supervisor's degree of risk aversion which remains constant. Indeed, when the supervisor is already subject to a first risky wage lottery needed to prevent collusion with the agent producing good 1, preventing further collusion with the agent producing good 2 is no more costly than if the supervisor were not controlling the first activity at all. This would no longer be true with utility functions exhibiting wealth effects.²⁸ More generally, in this framework with risk aversion, the general theory of organizing supervisory structures would closely look like and complement the theory of organizing productive tasks which was developed by Holmstrom and Milgrom (1990) in a pure moral hazard context.

6.2 Collusion and Competition

So far, we have modeled the firm as being insulated from any market interaction. An obvious issue is to investigate how the cost of collusion relates to the competitive pressure on the firm. A first channel is the following: the performances of competing firms in the market place may provide secondary sources of information to the principal by a simple yardstick competition argument. Such sources of information are substitutes to the supervisory information from the principal's point of view. However, market information is cheaper because the principal does not need to reward the supervisor to get it.

Even when yardstick mechanisms cannot be used, maybe because all information obtained from competitors is contained in the market price and this price is a nonverifiable variable, competition has an impact on the way collusion can be fought. Competition affects the size of the firm's output and thus the collusive stakes that the supervisor can

²⁷In full generality, the supervisor's wage should be written as $s_{ij,kl}$. With CARA utility functions, it can be shown that there is no loss of generality in making the additive assumption.

²⁸The risky wage lottery needed to prevent collusion on product line 1 then has a certainty equivalent which may shift the supervisor's degree of risk aversion and affect the cost of preventing collusion on product line 2.

manipulate. On the other hand, fighting collusion is part of the cost borne by the firm and this affects its ability to react to the competitive pressure.

To see this interaction between competition and collusion in more detail, consider the following simple model. Suppose that there are n symmetric firms in a market, competing in quantity. The demand is supposed to be linear $P(Q) = A - Q$ where $Q = \sum_{s=1}^n q_s$: The number of competing firms n can be viewed as measuring the degree of competition on this market.

The internal structure of all these firms is the same as the canonical organization presented before. Moreover, the costs of all firms are perfectly correlated so that, even if there are n firms in the market place, there are only two states of nature: one where all firms have a low cost, or where they all have a high cost.²⁹

To simplify the exposition, we concentrate on the case of a small uncertainty on cost (ϕ small enough). Adapting (20) to the situation, we get that the equilibrium output of a given firm when costs are high satisfies:

$$a - (n + 1)q_{22}^c = \mu + \frac{\int_0^1 \phi \mu(1 - s) + r \int_0^1 (1 - s) \phi \mu^2 q_{22}^c}{1 - \int_0^1 \phi \mu(1 - s) + r \int_0^1 (1 - s) \phi \mu^2}$$

or

$$q_{22}^c = \frac{a - \mu - \frac{\int_0^1 \phi \mu(1 - s) + r \int_0^1 (1 - s) \phi \mu^2}{1 - \int_0^1 \phi \mu(1 - s) + r \int_0^1 (1 - s) \phi \mu^2}}{n + 1}$$

From the above, we get that the elasticity of output with respect to n is given by:

$$\epsilon = \frac{n}{q_{22}^c} \frac{\partial q_{22}^c}{\partial n} = \frac{n}{1 - \int_0^1 \phi \mu(1 - s) + r \int_0^1 (1 - s) \phi \mu^2 + n + 1} < 1:$$

The elasticity of the firm's output with respect to n measures the impact of an increase in the competitive pressure on the firm's output. Of course, as the firm is facing more competitors in the market place, its own output decreases but, at the margin, it decreases less as there are more competitors. We also observe that the elasticity ϵ also decreases with r . The equilibrium output is thus less responsive to an increase in competition as r is large.

In fact, an increase in r plays the same role as an increase in the number of competing firms in determining the equilibrium output of a given firm. With more competing firms, the marginal incentives to expand output decreases because the residual demand faced by a given firm diminishes. With less frictions in internal collusion, those marginal incentives are also lower because the cost of collusion is at the margin higher for large scale production. The degree of risk aversion of the supervisor thus plays the same role as an increase in competitive pressure. As a result, the firm's equilibrium output reacts less to variations in this pressure as r increases.

²⁹This assumption makes the analysis simpler without changing the results. The results also holds when costs are independently distributed. Of course, to justify that yardstick mechanisms are not used in this environment with correlated information we need to assume that a given principal cannot communicate with agents in other competing hierarchies or, as we argued above, that the price which may contain the relevant information for comparison is nonverifiable.

Proposition 5 An increase in the supervisor's risk aversion reduces the sensitivity of the firm's output to competitive pressure.

6.3 Collusion and Vertical Integration

Consider the owner of a buying unit (the principal), who contracts for the provision of a good with a selling unit (the agent). The principal can use a "supervisor", i.e., a member of his own organization to get a signal about the external agent's productivity. We make two assumptions:

A1: Collusion can only take place inside organizations.

This is an extreme assumption but several considerations can motivate the view that collusion between agents in different firms is much more difficult than inside the same firm. The transfer of bribes between agents in separate firms can be easier to detect than inside the same firm. The role of the supervisor as key provider of information may also be reduced in a market relationship as the information can now come from several unrelated sources. Besides, knowing who to bribe can be more of a problem. Also, the provider of information in a market environment may be better able to diversify risk by dealing with several agents. Importantly, the enforceability of collusion is likely to be easier within the same organization. Again, our model is a static one where we assume at the outset that colluding parties can enforce their side-contracts. The usual justification, made formal in Tirole (1992) or Martimort (1999), is to appeal to a repeated game argument. But along these lines, colluding partners in separate firms are less likely to interact in the future so that enforcing colluding agreement is probably more difficult when agents belong to different organizations. All the above arguments lead us to assume that vertical integration eases collusion.

A2: Vertical integration improves information accuracy.

This assumption is often made in the vertical integration literature. For instance, Arrow (1975) has argued that one consequence of vertical integration is to improve information flows between the integrated units. This assumption can in particular be justified when repeated relationships between the supervisor and the agent improve the latter's monitoring.

These two features can easily be captured in our framework. Denote by μ^S the accuracy of supervisory information under separation. Taking the collusion cost in (20) to be zero, the optimal output under separation satisfies:

$$R^0(q_{22}^S(r)) = \mu_2 + \frac{\int_{1-\mu}^1 \phi \mu^{\geq} 1 - \mu^S}{1 - \int_{1-\mu}^1 \phi \mu^{\geq} 1 - \mu^S} : \quad (23)$$

Let us instead suppose that vertical integration takes place. Using again (20), we find that for small $\phi \mu$, the optimal output under integration is given by:

$$R^0(q_{22}^I(r)) = \mu_2 + \frac{\int_{1-\mu}^1 \phi \mu(1 - \mu)}{1 - \int_{1-\mu}^1 \phi \mu(1 - \mu)} + r^I \frac{\int_{1-\mu}^1 (1 - \int_{1-\mu}^1 \phi \mu^2)}{1 - \int_{1-\mu}^1 \phi \mu^2} q_{22}^I(r) : \quad (24)$$

This formula shows that, as r^l becomes small enough, the added cost of collusion within the integrated firm is lower than the extra benefit from better information.

Proposition 6 Vertical integration dominates separation when the supervisor is not too risk averse.

The trade-off is simple: integration facilitates collusion but potentially improves information accuracy. Integration is best when the cost of getting this additional information is not too high, i.e. when collusion is not too difficult to deter. This happens in organizations with supervisors who are not too risk averse. Hence, under our assumptions, we would predict that firms with relatively risk tolerant supervisory structures, characterized by the use of information sensitive incentive schemes, will be the ones willing to take over other firms to vertically integrate.

6.4 Collusion and Uncertainty

The nonlinearity of the collusion costs also has consequences for the behavior of firms under uncertainty. We now investigate how uncertainty in the environment where the firm evolves interacts with the scope for collusion between agents. To model uncertainty in the environment, let us suppose that the principal's revenue is now subject to a multiplicative random shock θ so that $\theta R(q)$ goes to the principal. To our ideas, let us also suppose that θ can take only two values $\underline{\theta}$ and $\bar{\theta}$ with $\text{prob}(\theta = \underline{\theta}) = \lambda$.

Several interpretations can be given to this modeling. For instance, θ can capture some fluctuations in the profitability of the firm's demand: a low θ would correspond to a bust, a high θ to a boom. One would then compare collusion in firms subject to a lot of cyclical variations with what happens in firms in more stable conditions or, for instance, in more mature industries where the state of demand is highly predictable. In a regulatory environment, θ may capture some political fluctuations as different parties in power (the principals) may have different preferences about how large the production of the regulated firm should be.

Under uncertainty, the principal commits to a grand-contract stipulating which production should be made as a function of the supervisor and the agent's reports but also as a function of the realized shock θ that we assume to be verifiable. Agents decide whether or not to accept the contract before the realization of θ . Thus, the timing is identical to the one proposed in Section 2.4 for the first 4 steps. Ex ante contracting will refer to the case where reports are made at step 5, θ is realized at step 6. Ex post contracting refers to the reversed sequence for the last two steps.

Importantly, the scope for collusion depends only on the timing of communication between the principal and his agents. If reports are made before θ is known, agents are forced to collude under uncertainty on the realization of θ . In this case, the risk-averse supervisor asks for a bribe which is the average collusive stake over the different possible realizations of θ , i.e., $E_{\theta}(\phi \mu_{q_2}(\theta)) = \phi [\lambda \mu_{q_2}(\underline{\theta}) + (1 - \lambda) \mu_{q_2}(\bar{\theta})]$.

Using again our approximation (19) valid for small $\phi \mu$, the participation constraint of a supervisor engaged in such an ex ante collusion is:

$$E_{\emptyset}(\int \mathbb{1}_{11}(\emptyset) + (1 - \int \mathbb{1}_{22}(\emptyset)) = \frac{r}{2} \int \mathbb{1}_{11}(\emptyset) (\phi \mu E_{\emptyset}(q_{22}(\emptyset)))^2): \quad (25)$$

Instead, if reports are made after \emptyset is known, agents collude ex post. In this case, the risk-averse supervisor asks for a bribe which is the collusive stake in the realized state of nature. The supervisor's participation constraint becomes:

$$E_{\emptyset}(\int \mathbb{1}_{11}(\emptyset) + (1 - \int \mathbb{1}_{22}(\emptyset)) = \frac{r}{2} \int \mathbb{1}_{11}(\emptyset) E_{\emptyset}((\phi \mu q_{22}(\emptyset))^2): \quad (26)$$

Clearly, ensuring the participation of the risk averse supervisor under ex ante collusion is less costly than with ex post collusion because of the convexity of the risk premium with respect to collusive stakes. This intuition can be made formal as shown in the Appendix and we have:

Proposition 7 In an uncertain environment, the principal prefers to solicit reports of the firm's employees before the realization of uncertainty. This results in outputs levels which fluctuate more than one would obtain with communication taking place after the resolution of uncertainty.

Forcing reports before the realization of the external shock allows the principal to offer wages to the supervisor which are not contingent on that shock. This independence can be seen as an informativeness principle for collusion: the shock on demand is unrelated to the supervisory information and thus, should not be part of an optimal contract for the supervisor. Indeed, offering wages contingent on its realization would only subject the supervisor to additional risk without improving the supervisor's incentives to report the truth. Thus, the principal is better off ensuring the supervisor against this risk.

The principal thus prefers strictly to induce communication before the realization of \emptyset . Of course, the choice of this timing is completely irrelevant when collusion is not an issue. In addition to output distortions and individual incentive schemes, a new tool is now used by the principal to curb coalitional behavior: the timing of communication.

The principal's preference for ex ante reporting also has implications for the optimal production plans. In the case of small $\phi \mu$; we can derive the optimal outputs for each realization of \emptyset under ex ante collusion:

$$Q_{22}^0(q_{22}(\emptyset)) = \mu_2 + \frac{\int}{1 - \int} \phi \mu (1 - \int + r \int \mathbb{1}_{11}(\emptyset) (\phi \mu E_{\emptyset}(q_{22}(\emptyset)))): \quad (27)$$

The output distortions for different values of \emptyset are linked together by the fact that the relevant collusive stake depends now on the average production level. This implies that production plans in different states of the world (i.e. different values of \emptyset) are linked. If for instance $\emptyset < \bar{\emptyset}$, $q_{22}(\emptyset) < E_{\emptyset}(q_{22}(\emptyset))$ and the downward distortion needed in state \emptyset is amplified compared to what would happen with ex post collusion. The organization

implements incentive schemes which are more sensitive to shock realization. If one interprets the sensitivity of output as a measure of the firm's flexibility to adapt to, or to anticipate shocks, we see that the optimal degree of flexibility is affected by the threat of collusion.

7 Conclusion

This paper has proposed an analysis of collusive situations when there is a trade-off between insurance and incentives not to collude. Doing so has enabled us to discuss the frictions of collusive agreements and to link those frictions to some parameters of the external environment where the firm evolves. Our final section has shown that this trade-off has some implications for the design of organizations and some of them were explored here.

It should also be clear that the objective of linking transaction costs of collusion within the organization to the external environment of the firm is a research program with a larger scope than this simple paper. Any theory of frictions in side-contracting is likely to offer in one way or the other some relationships between what happens inside the firm and in its external environment. Instead of being based on the trade-off between risk and incentives considered here, alternative theories of those frictions could build on asymmetric information, repeated self-enforceable relationships, imperfect cultural transmissions, or non-monetary exchanges between colluding partners. But opening the black box of the frictions of collusion seems key to making progress on the understanding of this phenomenon and should give rise to interesting future research.

8 Colophon

We thank seminar participants in Nuffield College Oxford, LSE, UCL, ESSET 98, ECARE in Brussels, Bocconi University, Bologna University, University of Southampton, Cambridge University, Universitat Pompeu Fabra at Barcelona, and EEA2002, as well as Sridhar Arcot, Doh-Shin Jeon, Thomas Mariotti and Jean Tirole for helpful comments. All errors are ours.

Antoine Faure-Grimaud, London School of Economics, FMG, Houghton Street, London WC2A 2AE, GRANDE-BRETAGNE. Email: A.Faure-Grimaud@lse.ac.uk

Jean-Jacques LaFont, Université de Toulouse 1 - Sciences Sociales, Institut d'Economie Industrielle (IDEI), Manufacture des Tabacs (Bât F), 21 allée de Brienne, 31000 Toulouse, FRANCE. Email: lafont@cict.fr

David Martimort, Université de Toulouse 1 - Sciences Sociales, Institut d'Economie Industrielle (IDEI), Manufacture des Tabacs (Bât F), 21 allée de Brienne, 31000 Toulouse, FRANCE. Email: martimor@cict.fr

Appendix

Proof of Proposition 1:

≤ First note that (2) must be binding to reduce the cost of (1). Moreover, to reduce the right-hand side of (12), (1) must be binding.

≤ Hence, we can rewrite (12) as:

$$s_{11} + u_{11} \leq s_2 + \phi \mu q_{22} \quad (28)$$

Taking first outputs as given, the principal's problem becomes:

$$\begin{aligned} \text{Max}_{u_{11}; s_{11}; s_2} & \int \lambda (R(q_{11}) - \mu_1 q_{11} - u_{11} - s_{11}) + \int (1 - \lambda) (R(q_{12}) - \mu_1 q_{12} - s_2 - \phi \mu q_{22}) \\ & + (1 - \lambda) (R(q_{22}) - \mu_2 q_{22} - s_2) \\ & \text{subject to (3)-(28) and (13).} \end{aligned}$$

We denote by λ , μ and π the respective multipliers of these constraints.

≤ Optimizing with respect to u_{11} ; s_{11} and s_2 yields respectively:

$$\lambda \leq \mu - \lambda = 0; \quad (29)$$

$$\lambda \leq \mu + \pi \int \lambda v^0(s_{11}) = 0; \quad (30)$$

$$\lambda (1 - \lambda) \leq \mu + \pi (1 - \lambda) \int \lambda v^0(s_2) = 0; \quad (31)$$

Summing (30) and (31) yields:

$$\pi \int \lambda e^{-rs_{11}} + (1 - \lambda) \int \lambda e^{-rs_2} = 1; \quad (32)$$

Note first that, (32) implies that $\pi > 0$ and thus (13) is binding. Using that $v(\cdot)$ is CARA, we obtain from (13) that $\int \lambda e^{-rs_{11}} + (1 - \lambda) \int \lambda e^{-rs_2} = 1$ and thus $\pi = 1$.

Multiplying (30) by $1 - \lambda$ and (31) by λ and subtracting the second from the first equation, we find that: $\mu = \lambda \int \lambda (v^0(s_2) - v^0(s_{11}))$.

Using (29) and (30) we get $\lambda \int \lambda v^0(s_{11}) > 0$ and thus

$$u_{11} = 0; \quad (33)$$

Hence, we need $s_2 < s_{11}$ to satisfy the coalition incentive constraint (28) and $\mu > 0$. Finally, $\mu > 0$ implies that (28) is binding and thus:

$$s_{11} = s_2 + \phi \mu q_{22}; \quad (34)$$

Inserting this latter condition into (13) yields (16). (34) yields (15).

≤ Optimizing with respect to outputs yields (14).

Proof of Proposition 2: Immediate derivation of (14) yields:

$$\frac{d\alpha_{22}^c(r)}{dr} = \frac{\int \mu^2 \alpha(1 - \int \alpha) e^{r\phi \mu \alpha_{22}^c(r)} + \int \alpha}{(1 - \int) (1 - \int \alpha) e^{r\phi \mu \alpha_{22}^c(r)} + \int \alpha} \frac{d\alpha_{22}^c(r)}{dr} = \frac{\int \mu^2 \alpha(1 - \int \alpha) \alpha_{22}^c(r) e^{r\phi \mu \alpha_{22}^c(r)} + \int \alpha}{(1 - \int) (1 - \int \alpha) e^{r\phi \mu \alpha_{22}^c(r)} + \int \alpha} \tag{35}$$

Hence $\frac{d\alpha_{22}^c(r)}{dr} < 0$.

The limits of the supervisor's wages when r goes to zero have already been derived in the text. When r goes to infinity, the limits are obtained directly from (16) and (15).

The impact on the principal's welfare is computed from:

$$EW = \int \alpha (R(\alpha_{11}^c) - \mu_1 \alpha_{11}^c - s_{11}^c) + \int (1 - \alpha) (R(\alpha_{12}^c) - \mu_1 \alpha_{12}^c - s_2^c - \phi \mu \alpha_{22}^c) + (1 - \int) (R(\alpha_{22}^c) - \mu_2 \alpha_{22}^c - s_2^c)$$

which is a function of r .

Using the Envelop Theorem, we find:

$$\frac{\partial EW}{\partial r} = \int \alpha \frac{\partial s_{11}^c}{\partial r} + (1 - \int \alpha) \frac{\partial s_2^c}{\partial r}$$

where $\frac{\partial s_{11}^c}{\partial r}; \frac{\partial s_2^c}{\partial r}$ denote the partial derivatives of the wages w.r.t. r holding α_{22}^c constant. We obtain that:

$$\frac{\partial EW}{\partial r} = \frac{1}{r} s_2^c + \frac{\int \alpha \phi \mu \alpha_{22}^c e^{r\phi \mu \alpha_{22}^c}}{(1 - \int \alpha) e^{r\phi \mu \alpha_{22}^c} + \int \alpha}$$

To study the sign of this last expression define $f(x)$ as $f(x) = \frac{\int \alpha e^{-x}}{(1 - \int \alpha) e^{-x} + \int \alpha} + \ln(1 - \int \alpha e^{-x})$: In fact, $\frac{\partial EW}{\partial r} < 0$ if and only if for any $x > 0$, $f(x) < 0$: We immediately verify that $f(0) = 0$ and that $f'(x) = \frac{\int \alpha e^{-x} (1 - \int \alpha)}{(1 - \int \alpha) e^{-x} + \int \alpha} < 0$; so the result.

Proof of Proposition 3: The derivation of α_{22}^a is immediate from (21) with α going to one. The principal's welfare is increasing in α as (making use of the Envelop Theorem):

$$\frac{\partial EW}{\partial \alpha} = \frac{\partial s_2^c}{\partial \alpha} = \frac{1}{r} \frac{e^{r\phi \mu \alpha_{22}^c}}{(1 - \int \alpha) e^{r\phi \mu \alpha_{22}^c} + \int \alpha} > 0$$

It is also immediate to observe that $\frac{\partial EW}{\partial \alpha}$ is increasing in r or $\phi \mu \alpha_{22}^c$:

Proof of Proposition 4: The proof follows similar lines to those of Proposition 1. We use symmetry between product lines to simplify notations below.

Note that the supervisor's ex ante participation constraint writes thus as:

$$(\int \alpha^2 v(2s_{11}) + 2 \int \alpha (1 - \int \alpha) v(s_{11} + s_2) + (1 - \int \alpha)^2 v(2s_2)) \geq 0 \tag{36}$$

since the supervisor may have observed either two revealing signals, only one, or none.

Maximizing the principal's expected profit with the standard agent's individual incentive (1) and participation (2) and (3) constraints, the coalitional incentive constraint (12) and the supervisor's participation constraint (36), and first taking again outputs as given, the principal's problem becomes:

$$\begin{aligned} & \text{Max}_{f_{u_{11}}, s_{11}, s_{2g}} \int \lambda^2 (2R(q_{11}) - 2\mu_1 q_{11} - 2u_{11} - 2s_{11}) \\ & + 2 \int \lambda (1 - \lambda) (R(q_{11}) + R(q_{12}) - \mu_1 (q_{12} + q_{11}) - s_2 - s_{11} - \phi \mu q_{22}) \\ & + \int (1 - \lambda)^2 (2R(q_{22}) - 2\mu_2 q_{22} - 2s_2) \\ & \text{subject to (3)-(28) and (36).} \end{aligned}$$

Again, we denote by λ , μ and π the respective multipliers of these constraints.

Optimizing with respect to the supervisor's wages and summing the corresponding first order conditions, we get:

$$\pi \int \lambda e^{-rs_{11}} + (1 - \lambda) \int \lambda e^{-rs_2} = 1; \tag{37}$$

Using that $v(\cdot)$ is CARA, we obtain also from (36) that

$$\int \lambda e^{-rs_{11}} + (1 - \lambda) \int \lambda e^{-rs_2} = 1; \tag{38}$$

Thus $\pi = 1$. From the fact that (28) and (36) are binding, we obtain the same values for s_{11} and s_2 than in the case of Proposition 1.

Proof of Proposition 7:

A first possibility for the principal is to offer contracts after the realization of θ (ex post contracting): For each θ , the principal then solves a program identical to the one solved to obtain Proposition 1, except for the fact that $R(q)$ is replaced by $\bar{R}(q)$: The principal's welfare is thus:

$$\begin{aligned} EW^{\text{ex post}} = E_{\theta} [& \lambda (\bar{R}(q_{11}(\theta)) - \mu_1 q_{11}(\theta)) + \int (1 - \lambda) (\bar{R}(q_{12}(\theta)) - \mu_1 q_{12}(\theta) - \phi \mu q_{22}(\theta)) \\ & + (1 - \lambda) (\bar{R}(q_{22}(\theta)) - \mu_2 q_{22}(\theta) - s_2(\theta))] \end{aligned} \tag{39}$$

where $s_2(\theta) = \frac{1}{r} \ln \left(1 - \lambda \int \lambda e^{-r\phi \mu q_{22}^c(\theta)} \right)$; and with the optimal outputs being equal to their first best levels for a low cost agent and defined by:

$$\bar{R}^0(q_{22}^c(\theta)) = \mu_2 + \frac{\int \lambda e^{-r\phi \mu q_{22}^c(\theta)}}{1 - \lambda \int \lambda e^{-r\phi \mu q_{22}^c(\theta)}}; \tag{40}$$

otherwise. Notice that ensuring $s_{11}(\theta) \geq s_2(\theta) + \phi \mu q_{22}(\theta)$ to deter collusion results in fluctuations of $s_{11}(\theta)$ and $s_2(\theta)$ as the output levels fluctuate according to (40). The supervisor is thus subject to an additional risk with ex post contracting.

Otherwise the principal can force the agent and the supervisor to report their private information before the realization of θ . The principal's program is similar to the previous

one except for the fact that collusive stakes are different. In this case, collusion is deterred if and only if:

$$E_{\emptyset}(s_{11}(\emptyset) + u_{11}(\emptyset)) \geq E_{\emptyset}(s_{22}(\emptyset) + u_{22}(\emptyset)) + c \mu E_{\emptyset}(q_{22}(\emptyset)) :$$

With ex ante contracting, the supervisor's participation constraint becomes:

$$E_{\emptyset}(\int \ln(s_{11}(\emptyset)) + (1 - \alpha) \int \ln(s_2(\emptyset))) \geq 0 :$$

It should be clear that $u_{11}(\emptyset) = u_{22}(\emptyset) = 0$ at the optimum. Moreover, the supervisor's participation constraint is less costly if he receives a constant wage independent of \emptyset and this does not change the coalition incentive constraint. Denote by s_2^A and s_{11}^A these wages.

Suppose that the principal decided to implement the levels of output defined in (40) as part of the optimal ex ante contract (of course, the optimal output levels ex ante are not those defined in (40)). Then the principal welfare would be equal to the same expression as in (39) except that $\alpha s_2(\emptyset) + (1 - \alpha) s_2(\emptyset)$ would then be replaced by $s_2^A = \frac{1}{r} \ln(1 - \alpha) \int \ln(s_2(\emptyset)) + \alpha \int \ln(s_2(\emptyset))$. The convexity of $\frac{1}{r} \ln(1 - \alpha) \int \ln(s_2(\emptyset))$ in x implies that $\alpha s_2(\emptyset) + (1 - \alpha) s_2(\emptyset) > s_2^A$: Hence, ex ante contracting with the same output levels gives to the principal a higher welfare than with ex post contracting.

Then the principal optimizes also w.r.t. outputs and that results in the first best levels for a low cost agent and in outputs satisfying:

$$Q^0(q_{22}^A(\emptyset)) = \mu_2 + \frac{\int \ln(s_2(\emptyset))}{1 - \alpha} c \mu \sqrt{1 - \alpha} \frac{\int \ln(s_2(\emptyset))}{1 - \alpha \int \ln(s_2(\emptyset)) + \alpha \int \ln(s_2(\emptyset))} !$$

where $q_{22}^m = \alpha q_{22}(\emptyset) + (1 - \alpha) q_{22}(\emptyset)$ for a high cost agent.

References

- Arrow, K., 1975, "Vertical Integration and Communication," *Bell Journal of Economics*, 6(1): 173-183.
- Baliga, S., 1999, "Collusion with Soft Information," *Journal of Law, Economics and Organization*, 15: 434-440.
- Baliga, S. and T. Sjostrom, 1998, "Decentralization and Collusion," *Journal of Economic Theory*, 83: 196-232.
- Baron, D. and D. Besanko, 1999, "Informational Alliances," *Review of Economics Studies*, 66: 743-768.
- Bull, J. and J. Watson, 2000, "Evidence Disclosure and Verifiability," Working Paper 2000-16, University of San Diego.
- Chandler, A., 1962, *Strategy and Structure*, MIT Press, Cambridge.
- Faure-Grimaud, A., J.J. LaFont and D. Martimort, 1999, "The Endogenous Transaction Costs of Delegated Auditing," *European Economic Review* (1999), 43: 1039-1048.
- Faure-Grimaud, A., J.J. LaFont and D. Martimort, 2000, "A Theory of Supervision with Endogenous Transaction Costs," *Annals of Economics and Finance*, 1: 231-263.
- Faure-Grimaud, A., J.J. LaFont and D. Martimort, 2001, "Collusion, Delegation and Supervision with Soft Information," forthcoming *Review of Economic Studies*. Mimeo IDEI Toulouse.
- Faure-Grimaud, A. and D. Martimort, 2001, "The Agency Cost of Intermediated Contracting," *Economic Letters*, 71: 75-82.
- Felli, L., 1997, "Preventing Collusion through Delegation," Mimeo STICERD, London.
- Felli, L. and M. Villas-Boas, 2000, "Renegotiation and Collusion in Organizations," *Journal of Economics and Management Strategy*, 9: 453-483.
- Gollier, C., 2001, *The Economics of Risk and Time*, MIT Press, Cambridge.
- Green, J. and J.J. LaFont, 1986, "Partially Verifiable Information and Mechanism Design," *Review of Economic Studies*, 53: 447-456.
- Grossman, S. and O. Hart, 1983, "An Analysis of the Principal-Agent's Problem," *Econometrica*, 51: 7-45.
- Holmstrom, B., 1979, "Moral Hazard and Observability," *Bell Journal of Economics*, 10: 74-91.
- Holmstrom, B. and P. Milgrom, 1990, "Regulating Trade among Agents," *Journal of Institutional and Theoretical Economics*, 146: 85-105.
- Holmstrom, B. and P. Milgrom, 1991, "Multi-task Principal Agent Analysis: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economic and Organization*, 7: 24-52.

- Itoh, H., 1993, "Coalitions, Incentives, and Risk Sharing," *Journal of Economic Theory*, 60: 410-427.
- LaFont, J.J. and D. Martimort, 1997, "The Firm as a Multicontract Organization," *Journal of Economics and Management Strategy*, 6: 201-234.
- LaFont, J.J. and D. Martimort, 1998, "Collusion and Delegation," *Rand Journal of Economics*, 29: 280-305.
- LaFont, J.J. and D. Martimort, 1999, "Separation of Regulators Against Collusive Behavior," *Rand Journal of Economics*, 30: 232-263.
- LaFont, J.J. and D. Martimort, 2000, "Mechanism Design under Collusion and Correlation," *Econometrica*, 68: 309-342.
- LaFont, J.J. and D. Martimort, 2002, *The Theory of Incentives: The Principal-Agent Model*, Princeton University Press, Princeton.
- LaFont, J.J. and M. Meleu, 1997, "Reciprocal Supervision, Collusion and Organizational Design," *Scandinavian Journal of Economics*, 99: 519-540.
- LaFont, J.J. and J.C. Rochet, 1999, "Regulation of a Risk-Averse Firm," in C. d'Aspremont ed. *Social Organizations and Mechanism Design*, De Boeck and Larcier s.a., Bruxelles.
- LaFont, J.J. and J. Tirole, 1993, *A Theory of Incentives in Procurement and Regulation*, MIT Press, Cambridge.
- Martimort, D., 1997, "A Theory of Bureaucratization Based on Reciprocity and Collusive Behavior," *Scandinavian Journal of Economics*, 99: 555-579.
- Martimort, D., 1999, "The Life Cycle of Regulatory Agencies: Dynamic Capture and Transaction Costs," *Review of Economic Studies*, 66: 929-948.
- Martimort D., and T. Verdier, 2000, "The Internal Organization of the Firm, Transaction Costs and Macroeconomic Growth," *Journal of Economic Growth*, 5: 315-340.
- Martimort D., and T. Verdier, 2002, "The Agency Cost of Internal Collusion and Schumpeterian Growth," Mimeo IDEI-Toulouse.
- Maskin, E., 1999, "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66: 23-38.
- McAfee, P. and J. McMillan, 1995, "Organizational Diseconomies of Scope," *Journal of Economics and Management Strategy* 4: 399-426.
- Prendergast, C. and R. Topel, 1996, "Favoritism in Organizations," *Journal of Political Economy*, 104: 959-978.
- Salani , B., 1990, "Selection Adverse et Aversion pour le Risque," *Annales d'Economie et de Statistiques*, 18: 131-150.
- Shavell, S., 1979, "Risk Sharing and Incentives in the Principal and Agent Relationship," *Bell Journal of Economics*, 10: 55-73.
- Tirole, J., 1986, "Hierarchies and Bureaucracies: On the Role of Collusion in Organizations," *Journal of Law, Economic and Organization*, 2: 181-214.

Tirole, J., 1992, "Collusion and the Theory of Organizations," in *Advances in Economic Theory, Sixth World Congress*, vol. 2, ed. J.J. LaFont, 151-206, Cambridge University Press, Cambridge.