Collusion and Selective Supervision*

Alberto Motta†

First Version: January 2007

This Version: July 2011

Abstract

This paper studies a mechanism-design problem involving a principal-supervisor-agent in which collusion between supervisor and agent can only occur after they have decided to participate in the mechanism. We show how collusion can be eliminated at no cost via the use of a mechanism in which the principal endogenously determines the scope of supervision. A simple example of such a mechanism is one in which the agent bypasses the supervisor and directly contracts with the principal in some states of the world. The result that collusion can be eliminated at no cost in this environment highlights the important assumptions required for collusion to be a salient issue in the existing literature. The result is robust to alternative information structures, collusive behaviours and specification of agent's types. Applications include work contracts with different degrees of supervision, self-reporting of crimes, tax amnesties, immigration amnesties and mechanisms based on recommendation letters.

Key Words: Collusion, supervision, selective supervision, delegation, mechanism design, revelation principle.

JEL Classification: D82, C72, L51.

*This paper was written during my visit at Boston University. I wish to express my thanks to the Department of Economics for its hospitality and support. I am especially grateful to my advisors Antonio Nicolò and Dilip Mookherjee for their support, comments and encouragement. I also thank David Martimort, Gorkem Celik, Yeon-Koo Che, Jinwoo Kim, Vianney Dequiedt, Gregory Pavlov, Alessandro Pavan, Giacomo Calzolari, Bob Gibbons, Isabella Dobrescu, Alfredo Burlando, Chris Bidner, and Barton Lipman for helpful discussions. This paper has benefited with comments from seminar participants at Columbia University, University Toulouse 1, Boston University, University of New South Wales, University of Sydney, University of Melbourne, Australian National University, Monash University and University of Padua.

†University of New South Wales, School of Economics, Australian School of Business, Sydney. E-mail: motta@unsw.edu.au

1 Introduction

Third-party supervision is commonly observed within economic organizations.[1] Usually the need for supervisory activity originates in an information asymmetry between the residual claimant of a productive activity (the *principal*) and the party that carries out the productive activity (the *agent*). The role of the *supervisor*[2] is to provide the principal with information concerning actions or characteristics of the agent. This creates a potential for collusion between supervisor and agent, wherein the agent bribes the supervisor to conceal information from the principal. Most studies conclude that collusion is a problem, and that eliminating it is costly for the principal.[3] In this context, the role of collusion in limiting the scope for incentives, the value of hiring supervisors, and the delegation to supervisors, have been examined by many authors in a standard framework. This includes asymmetric information between the colluding parties, and the inability to collude prior to making a decision to participate in the mechanism.

This paper focuses on a tool for combating collusion that has been previously overlooked. This tool is based on the idea of *selective supervision*, where the supervisor may not be engaged by the principal in certain states of the world. Take, for example, a simple mechanism where the agent selects between a regime with supervision and a regime without it. The choice between being supervised or directly contracting with the principal reveals useful information to the principal and reduces the scope of collusion. In the standard framework, we show that it costlessly eliminates collusion.

Thus, if collusion can be easily eliminated in the standard framework, what are the real sources of a collusion problem? To answer this questions we explore several variations of the standard setup in terms of its underlying assumptions. One crucial assumption is the timing of the supervisor's information, i.e., whether the supervisor receives her information before or after being employed by the principal. Based on this assumption, we distinguish between situations where the principal consults an expert (someone who has already received the information) and situations where he hires an auditor (someone who will investigate the agent after being employed).

The previous literature (Celik, 2009, and Faure-Grimaud, Laffont and Martimort, 2003 - FLM,

[1] Owners of a firm usually delegate the responsibility for supervising production to top managers; stockholders rely on auditors to acquire information about management conduct; managers ask employees to report on the performance of coworkers; and Governments make use of agencies to regulate firms, auditors to examine tax returns, and inspectors to detect illegal immigration.

[2] We refer to the supervisor and the agent respectively as *she* and *he*.

[3] See for example Tirole (1986), Laffont and Tirole (1991,1993), Lambert-Mogilianksy (1998), Faure-Grimaud and Martimort (2001), Faure-Grimaud, Laffont and Martimort (2003), and Celik (2009).

hereafter) focus on the first case, where the supervisor is outlined as an expert. Among the assumptions adopted in this literature, one is particularly relevant: the agent and the supervisor can collude only after they have accepted the mechanism offered by the principal (hereafter referred to as *no-collusion in participation decision*). Unlike the mechanisms proposed by Celik (2009) and FLM (2003), *selective supervision* can costlessly eliminate collusion.[4] This is due to the fact that we *do not* restrict attention to direct revelation mechanisms with full participation. Furthermore, we allow the players to report all the information available to them.[5] Departing from these restrictions is useful because it allows us to further exploit the assumption of *no-collusion in participation decision.* Given that participation decisions are collusion-free, the principal can design a rich mechanism where the supervisor's participation decision is used to capture some information on the agent's characteristics.[6] Interestingly, the implementation of *selective supervision* does not rest on special assumptions about the accuracy of the supervisor's information: the principal can eliminate collusion at no cost even when there is no residual asymmetric information between the supervisor and the agent. Second, we do not require any restrictions on the allocation of bargaining power inside the coalition. Third, the result does not depend on the identity of the coalition member who offers and initiates the collusive agreement. Fourth, the mechanism holds for a quite general specification of the agent's production costs and does not rely on special assumptions about players' utility functions.

In the second part of the paper, we depart from the previous literature by considering the second information timing: namely, the supervisor is an auditor who receives her information after being employed by the principal. This is a realistic case because the supervisor's information is often acquired through an inspection or lengthy investigation, which takes place following the acceptance of the contract. Under this latter timing, the implementation of *selective supervision* is more challenging. The reason is that the supervisor has no information on the agent's characteristics when she makes her participation decision. Therefore, the principal cannot use her participation decision to extract information. In fact, the principal can only use the agent's participation decision to achieve this goal.

[4]FLM's (2003) mechanism eliminates collusion at zero cost when the supervisor is risk neutral and there are two possible production costs.

[5]Following Celik (2009) and FLM (2003), we assume that the agent learns both his production cost and the supervisor's information. We allow the agent to report all his information and not only his production cost.

[6]If the principal can offer only a single mechanism (i.e., menu of mechanisms are not allowed), participation decisions are limited to two possibilities: "accept" or "refuse" the mechanism. In this case, if the supervisor is assumed to be indispensable for production, Celik's (2009) and FLM's (2003) results are general and *selective supervision* cannot improve their mechanisms. Under these circumstances, the Revelation Principle can be invoked to support the idea that restricting attention to direct revelation and full participation mechanisms is without loss of generality.

But extracting information from the agent is more complex: unlike the supervisor, the agent has a productive role and the principal must ensure that his incentive compatibility constraints are met. Extracting information from the agent's participation decision might conflict with these constraints. As a result, *selective supervision* can costlessly eliminate collusion only under certain conditions, related to the supervisor's information structure and to the specification of the agent's production costs.

Admittedly, the collusion-proof implementation presented in this paper heavily relies on the assumption of *no collusion in participation decision*. Far from strenuously trying to make a case in favor of this assumption, which is nevertheless plausible in many realistic situations,[7] this paper intends to shed light on those factors that make collusion truly problematic by identifying the factors which are less so.

From this perspective, our results suggest that the assumption of *no collusion in participation decision* is more or less plausible depending on the timing of the supervisor's information. If the supervisor receives her information before being employed, the assumption of *no collusion in participation decision* unravels the collusion problem: the latter can be eliminated at no cost. Clearly, the fact that collusion can be easily overcome is in contrast to its persistence in the real-world and feels a bit artificial. Allowing for collusion on participation decisions may be a more interesting way of thinking about the problem.

But this conclusion does not hold when we consider the second information timing (i.e., the supervisor is outlined as an auditor.) In this case, collusion might be harmful to the principal and increasingly so with the number and dispersion of the possible production costs. The salience of collusion also depends on the structure of the supervisor's information, where the seemingly technical distinction between a structure based on signals (FML, 2003) and one based on connected partitions (Celik, 2009) turns out to play a crucial role.

The remainder of the paper is organized as follows. In the next section we discuss some applications and present the related literature. Section 2 proposes the general model. Section 3 presents the selective

[7] It is plausible to assume that there might be several supervisors and agents that can be employed by the principal. In this case, the agent and the supervisor may be matched together after they have decided to participate. This situation is particularly plausible when the supervisor is an auditor. Under these circumstances, their failure to coordinate participation decisions is due to the impossibility of signing a preemptive side-contract with all eligible supervisors. Consider now the case where the supervisors are experts (they receive their information before being employed). The principal could decide to hire the supervisor after the agent has made his participation decision. In some cases, the use of job rotation for supervisors achieves the same result. In some other cases, the principal can avoid the disclosure of the agent's identity at the participation stage. This precaution makes it difficult for the supervisor to collude since she faces a potentially vast population of eligible agents.

supervision mechanism. Section 4 provides some additional comments. Section 5 concludes. All proofs are given in the Appendix.

1.1 Applications and Related Literature

Consider a *selective supervision* mechanism where the agent can decide whether to be supervised or not. Depending on the agent's characteristics, he may prefer one regime to the other. There are many real-world examples of this kind of mechanism. Within firms, the scope and intensity of supervision usually varies depending on the agent's characteristics. *Selective supervision* sheds light on the formation of such hierarchy structures and explains them as a result of the threat of collusion. Apart from firms, other real-world examples include self-reporting of illegal acts, wherein offenders can choose to report their illegal acts directly to the principal by choosing a mechanism that bypasses the supervisor. The literature on law enforcement has long highlighted that self-reporting allows the government to save money by reducing enforcement costs.[8] This paper tackles the issue from a different angle, suggesting a different advantage to the use of self-reporting: namely, the reduction of the costs associated with the threat of collusion. Another example is that of tax amnesties where the agent is induced to report his type directly to the principal, bypassing the supervisor's inspection. The same applies to immigration amnesties.[9] These applications are further discussed in the last part of Section 4.

The concept of *selective supervision* shares some similarities with the mechanism proposed by Dequiedt (2006) and Celik and Peters (2010). The latter studies an example of a mechanism-design problem where the players can coordinate their actions in a default game. They show that some allocation rules are implementable only with mechanisms that will be rejected on the equilibrium path. It may be useful to re-label aspects of their framework to highlight the similarities with the environment considered here. The two players in their model can be thought of as the supervisor and the agent in our framework. The default game corresponds to the principal's mechanism in this paper, whereas the coordination-mechanism corresponds to the collusive side-contract between the agent and the supervisor. Consequently, the scope of the present contribution goes beyond the one proposed by Celik and Peters (2010) in that it considers endogenously determined "default" games. Even

[8] See Kaplow and Shavell (1994), Innes (1999, 2001).

[9] Some of these issues are explored in a separate paper by Burlando and Motta (2008a). They analyze the impact of self-reporting on law enforcement when officers are corruptible. They show that a budget-constrained government may prefer an enforcement system based on corruption rather than one based on legal fines. They conclude that the government can use self-reporting as a way to clean up corrupt enforcement agencies. Unlike this paper, they do not adopt a mechanism design approach. Moreover our contribution considers a larger class of mechanisms, wherein self-reporting with a binary information structure for the supervisor is only one simple application.

though their setting is different from the one considered here, there is one common aspect to both contributions: participation decisions convey information about the types of the players. Dequiedt (2006) considers a similar point in the mechanism design literature that assumes that each agent has a veto power.

In a related paper, Che and Kim (2006) study a general collusion setup where agents cannot collude prior to making their decision to participate in the mechanism. They conclude that the second-best payoff is implementable when players are risk neutral. Given the restrictions they impose on the correlation of information of the colluding parties, their result does not apply to the setup considered in this paper. On the contrary, it remains an open and intriguing question whether or not the strategy proposed in our framework applies to no-supervision setups such as the one they proposed.

We suspect that *selective supervision* is useful in a setting where collusion occurs prior to participation, though in that context it is unlikely to costlessly eliminate collusion. This issue remains to be explored in future research. A few interesting papers have already studied the implementation of collusion-proof mechanisms when agents can collude on their participation decisions,[10] but none of them have addressed this question yet. Among them, Mookherjee and Tsumagari (2004) analyze this problem in a supervision setup.[11] However they focus on a different question with respect to the one analyzed in our contribution. Namely, they consider two productive agents and explore the possibility that collusion may rationalize delegation to intermediaries uninvolved in production. In particular, they do not focus on the identification of the optimal mechanism in the presence of collusion.

2 The General Model

This section proposes a setting that accommodates as special cases both Celik's (2009) and FLM's (2003) frameworks. There is a productive agent (A) who bears the cost of production. Without loss

[10] Pavlov (2008), Che and Kim (2009) and Dequiedt (2007) consider auctions where bidders collude prior to participating. Che and Kim (2009) study an optimal collusion-proof auction in an environment where subsets of bidders may collude not just on their bids but also on their participation. They find that informational asymmetry facing the potential colluders can be significantly exploited to reduce their possibility to collude. Dequiedt (2007) considers two bidders with binary types. He finds that the seller can, at most, collect her reserve price when a bidder's valuation exceeds that price, if and only if a cartel can commit to certain punishment. Pavlov (2008) independently studies a problem similar to Che and Kim (2009) and reaches similar conclusions. Quesada (2004) studies collusion initiated by an informed party under asymmetric information.

[11] Mookherjee (2006) provides an excellent survey of this strand of the literature.

of generality, A is assumed to be risk neutral.[12] A utility function is given by

$$t - \theta q,$$

where t denotes the transfer he receives from the principal (P), q is the output level and θ represents the unitary cost of production, which takes n possible values from the set $\Theta = \{\theta_1, \theta_2, ..., \theta_N\}$, where $0 < \theta_1 \leq ... \leq \theta_{N-1} \leq \theta_N$. The distribution of the cost, $f(\theta)$, is common knowledge while A knows the realization of θ. The supervisor (S) receives a signal τ on A cost. This signal is also received by A. τ is drawn from a discrete distribution on $T = \{\tau_1, \tau_2, ..., \tau_N\}$. The joint probabilities on (θ_i, τ_j) are defined as $p_{ij} = Prob(\theta = \theta_i, \tau = \tau_j)$ with $\sum_{j=1}^{n} p_{ij} > 0$ for all i and $\sum_{i=1}^{n} p_{ij} > 0$ for all j. From the joint distribution above, one can derive the conditional probabilities $p(\theta_i|\tau_j)$. There is a positive correlation between signals and types when the monotone likelihood ratio property is satisfied,

$$p(\theta_i'|\tau_j')p(\theta_i|\tau_j) - p(\theta_i|\tau_j')p(\theta_i'|\tau_j) \geq 0 \tag{1}$$

for all $(\tau_j, \tau_j', \theta_i, \theta_i')$ such that $\tau_j' \geq \tau_j$ and $\theta_i' \geq \theta_i$.

S salary is s, which represents her monetary transfer from P. Her utility function is given by $U_S(.)$, with $U_S'(.) > 0$ and $U_S''(.) < 0$. It follows that S is risk averse.[13] P payoff for a given output q, transfer level t and wage s is

$$W(q) - t - s,$$

where $W'(q) > 0$, $W''(q) < 0$, for all q, and $\lim_{q\to 0} W'(0) = \infty$, $\lim_{q\to\infty} W'(q) = 0$. These conditions ensure positive production regardless of A cost type θ. P can commit to a contract, consisting in a triple

$$\Gamma = \{q(m_s, m_a), t(m_s, m_a), s(m_s, m_a)\}.$$

This contract defines the outcome and the monetary transfer respectively for A and S as a function of S and A messages, which are denoted as m_s and m_a and belong respectively to the message spaces M_s and M_a. If the contract is rejected, the game ends with zero production and no monetary transfer to the players. In other words, the outside option is normalized to zero for both A and S.

[12]The results would not change with a risk averse agent. Indeed, his ex-post participation and incentive constraints would be identical, and only those constraints are relevant for the analysis.

[13]The results would not change with a risk neutral supervisor.

One aspect is worth noting: FLM (2003) assume that $p(\theta_i, \tau_j) > 0$ for all (i, j). Unlike them, we only require $p(\theta_i, \tau_j) \geq 0$. This allows us to include Celik's (2009) information setup within the same framework. For simplicity, we further assume that if $p(\theta_i, \tau_j) = 0$, then all signals τ_j' such that $p(\theta_i, \tau_j') > 0$ are informationally equivalent (i.e., they convey the same information.)[14] This assumption guarantees that S information structure is equivalent to Celik's (2009) when $p(\theta_i, \tau_j) = 0$. Notice that our framework reduces to FLM (2003) when $p(\theta_i, \tau_j) > 0$, for all (i, j), and $N = 2$. Moreover, it also includes Celik (2009) as a special case when $N = 3$, and the informative signal is designed as a connected partition information structure.[15] In what follows, we refer to FLM (2003) and Celik (2009) information structures respectively as *signal-based* and *partition-based*.

2.1 Direct Supervision

Consider the case where the principal directly receives the signal on A private information. The Revelation Principle ensures that there is no loss of generality in looking for the optimal contract within the class of direct truthful revelation mechanisms of the form $\{t_\tau(\theta), q_\tau(\theta)\}$ where θ is A report on his unit-cost of production to P. For simplicity, denote by t_{ij} (q_{ij}) A transfer (output schedule) when A reports that he has type θ_i and P knows τ_j. A state denoted by (i, j) is a realization of a cost type θ_i and a signal τ_j. Let us denote A information rent in state (i, j) by $u_{ij} = t_{ij} - \theta_i q_{ij}$. When P observes a signal τ_j, the standard treatment of this problem suggests that an output profile $\{q_{ij}\}_{i \in [1,2,\ldots,N]}$ is implementable through a contract if and only if it is weakly decreasing. This condition is satisfied when the following *monotonicity constraint* holds:

$$q_{ij} \geq q_{i'j} \quad \text{for all } i' > i. \tag{2}$$

Furthermore, the agent's lowest utility levels that are compatible with this implementation are revealed by the binding participation constraint for the highest cost type,

$$u_{Nj} = 0, \tag{3}$$

[14]If there are two signals τ_j' and τ_j'' such that $p(\theta_i, \tau_j') > 0$ and $p(\theta_i, \tau_j'') > 0$, they must convey the same information, i.e., $p(\theta_i, \tau_j') = p(\theta_i, \tau_j'')$ for all θ_i.

[15]If $N = 3$, and the conditional probabilities are $p(\theta_1|\tau_1) = p(\theta_1|\tau_2) = B_1$, $p(\theta_2|\tau_1) = p(\theta_2|\tau_2) = B_2$, $p(\theta_3|\tau_1) = p(\theta_3|\tau_2) = 0$, $p(\theta_1|\tau_3) = p(\theta_2|\tau_3) = 0$, and $p(\theta_3|\tau_3) = 1$, then the information structure for S reduces to Celik's (2009) connected partition case $\{\{\theta_1, \theta_2\}, \theta_3\}$.

and the upward adjacent incentive compatibility constraints for the other types,

$$u_{ij} \geq u_{i+1j} + (\theta_{i+1} - \theta_i)q_{i+1j} \qquad \text{for all } i \neq N. \tag{4}$$

At the optimum of P problem, these constraints are binding; It is easy to show that the remaining incentive compatibility constraints are strictly satisfied. When P observes a signal τ_j, he updates his beliefs on A type. The conditional probabilities are $p(\theta_1|\tau_j) = p_{1j}/\sum_{i=1}^{n} p_{ij}$ for $j = 1, 2..., N$. The optimal contract solves

$$\max_{\{q_{ij},u_{ij}\}_{i\in\{1,2,...,N\}}} \sum_{i=1}^{n} p(\theta_i|\tau_j)\left[W(q_{ij}) - \theta_i q_{ij} - u_{ij}\right],$$

subject to (2), (3), and (4).

The assumption $\lim_{q\to 0} W'(0) = \infty$ ensures positive production regardless of A cost type θ. The solution to this problem yields the *conditionally-optimal* second-best (hereafter referred to as second-best), which implements the first-best outputs $q_{1j}^{sb} = q_1^{fb}$ for the most efficient agent[16] and outputs q_{ij}^{sb} for the other ones. To begin with, consider the optimal outputs for those types θ_i such that $p(\theta_i|\tau_j) > 0$. The optimal q_{ij} solves,

$$\begin{aligned} &W'(q_{1j}^{fb}) = \theta_1, \\ &W'(q_{ij}^{sb}) = \theta_i + \sum_{z=1}^{i-1} \tfrac{p_{zj}}{p_{ij}}(\theta_i - \theta_{i-1}) \qquad \text{for } i \neq 1. \end{aligned} \tag{5}$$

Consider now the optimal outputs for those types that have zero probability to be realized when P receives the signal j, i.e., $p(\theta_i|\tau_j) = 0$. In this case the quantity assigned to type i has no direct impact on P utility in that it does not affect $W(.)$. But q_{ij} could affect P utility indirectly by increasing the information rents that P has to forgo to those types that are more efficient than i. In this case, it is optimal to assign type i a zero-output schedule $q_{ij}^{sb} = 0$. On the other hand, if the choice of q_{ij} does not affect condition (4), then any quantity q_{ij} large-at-will can be optimally assigned, provided that the monotonicity constraint (2) is satisfied.[17] For the sake of exposition, we will assume that in this

[16]Note that the optimal output for the most efficient type when $p(\theta_1|\tau_i) = 0$ could be set to zero beacuse type θ_1 has zero probability to be realized following the signal τ_i. Nevertheless, the quantity produced by type θ_1 has no material effect on the incentive compatibility contraints. Therefore, for the sake of simplicity, we will assume that the optimal output for the most efficient type solves $W'(q_{i1}) = \theta_1$ even if $p(\theta_1|\tau_i) = 0$.

[17]For example, q_{ij} does not affect condition (4) when all the types i' that are more efficient than i have zero probability of being realized when the signal is j.

particular case $q_{ij}^{sb} = q_{ij-1}^{sb}$. Having this schedule in place, it is possible to show that:

Lemma 1 *For all i and for each pair of signals j, j' with $j' > j$,*

$$q_{ij}^{sb} \leq q_{ij'}^{sb}. \tag{6}$$

This Lemma clarifies how the informative signal affects the optimal outputs. When τ_j decreases, A is more likely to be efficient. Reducing the information rents calls then for a greater reduction of the outputs for the less efficient types. Lemma 1 entails that A information rent is increasing in the realization of the signal:

$$u_{ij}^{sb} \leq u_{ij'}^{sb} \text{ , for all } (i, j, j') \text{ such that } j' > j. \tag{7}$$

2.1.1 Numerical Example

Consider a simple numerical example where $W = \ln(q)$ and $\Theta = \{\theta_1, \theta_2, \theta_3\} = \{0.25, 0.5, 1\}$. Under perfect information, P observes θ, and the first-best quantities $q_1^{fb} = 4$, $q_2^{fb} = 2$ and $q_3^{fb} = 1$ are implemented. A receives a transfer t equal to his production cost: regardless of what type is realized, A receives zero information rent, i.e., $U_i = 0$. On the contrary, under asymmetric information, P has to provide A with some information rent in order to induce him to report his type. It follows from the standard treatment of this problem that reducing the information rent of the efficient types calls for a reduction in the output schedules of the inefficient ones. The extent of this distortion depends on the information available to P. The informative signal τ serves exactly this purpose: following a certain realization of τ, P can optimally re-adjust the output schedule, alleviating the asymmetric information problem.

Example 1: Signal-Based Information Structure. We refer to *direct supervision* to indicate the case where P directly receives the signal τ. Under *direct supervision*, P can implement the second-best outcome. For example, consider the following conditional probabilities $p(\theta_i|\tau_j)$

	τ_1	τ_2	τ_3
θ_1	55%	15%	5%
θ_2	20%	35%	35%
θ_3	25%	50%	60%

The optimal quantities (referred to as "Output" in the table) and utility levels of A (referred to as U in the table) when A truthfully reports type θ_i and P knows τ_j are,

	P knows τ_1		P knows τ_2		P knows τ_3	
Type	Output	U	Output	U	Output	U
$\theta_1 = 0.25$	4	0.4105	4	0.7451	4	0.8417
$\theta_2 = 0.5$	0.8421	0.2	1.6471	0.3333	1.8667	0.3750
$\theta_3 = 1$	0.4	0	0.6667	0	0.7500	0

When P receives the signal τ_1, he infers that A is more likely to be efficient. Reducing A information rent calls then for a large reduction of the output schedule of types θ_2 and θ_3. On the contrary, when P observes the signal τ_3 he knows that A is less likely to be efficient. Being that the information rents are less of a concern, P increments the output schedule for θ_2 and θ_3.

Example 2: Partition-Based Information Structure. At this stage, it might be useful to propose a second numerical example that uses Celik's (2009) information structure. In this example the signal is a connected partition of the type space,

	τ_1	τ_2	τ_3
θ_1	50%	50%	0%
θ_2	50%	50%	0%
θ_3	0%	0%	100%

where the signals τ_1 and τ_2 are informationally equivalent. When P receives the information directly (*direct supervision*), the second-best outcome is implementable. The optimal quantities (referred to as "Output" in the table) and utility levels of A (referred to as U in the table) when A truthfully reports type θ_i and P knows τ_j are,[18]

	P knows τ_1		P knows τ_2		P knows τ_3	
Type	Output	Agent U	Output	Agent U	Output	Agent U
$\theta_1 = 0.25$	4	0.3	4	0.3	4	0.8
$\theta_2 = 0.5$	1.3	0	1.3	0	1.3	0.5
$\theta_3 = 1$	0	0	0	0	1	0

When P learns τ_3, he knows that A has type θ_3. It follows that quantities $q_{13}^{sb} = 4$ and $q_{23}^{sb} = 1.3$ can be selected arbitrarily: they will never be chosen in equilibrium and they do not affect the incentive compatibility constraints for type θ_3 (provided that the monotonicity constraint (2) is satisfied).

If P does not receive the signal τ, he has to elicit this information from S. Whenever S and A collude, the mechanism presented above is no longer implementable. Regardless of the real realization of τ, S prefers to report τ_3 and then share A extra information rent.

2.2 Non-cooperative implementation

When S and A do not collude, FLM (2003) show that there is no loss of generality in restricting P to use direct truthful revelation mechanisms. Let us denote by s_{ijk} (respectively t_{ijk} and q_{ijk}) S wage (respectively A transfer and the output target) when A reports that he has type i and that S signal is k, and when S reports she has observed j. For the sake of simplicity, write $s_{ijj} = s_{ij}$ and $t_{ijj} = t_{ij}$. Using the logic of Nash implementation, FLM (2003) show that P can costlessly elicit τ by inducing A and S to reveal their signal. A incentive constraints can be reduced to the following relevant incentive constraints:

$$u_{ij} \geq u_{i'j} + (\theta_{i'} - \theta_i) q_{i'j} \quad \text{for all } (i, i', j)\,, \tag{8}$$

and S gets zero wage $s_{ij}^{sb} = 0$ for all (i, j). Therefore, P can achieve the same outcome as with direct supervision. Moreover, FLM (2003) show that the out-of-equilibrium wages for S can be designed to ensure a unique Nash implementation. This outcome is feasible if S and A do not cooperate or P is capable of preventing them from communicating.

[18] The results are rounded to the second decimal place.

2.3 Cooperative implementation

The process of collusion is formalized by assuming that S makes a take-it-or-leave-it offer to A, after the acceptance of the grand-contract by both parties. The side-contract is a pair $SC = \{\phi(.), b(.)\}$ where $\phi(.)$ is a collective manipulation of the messages (m_s, m_a) sent to P, while $b(.)$ is A transfer received from S. As standard in the literature on collusion, this side-contract is assumed to be enforceable.[19] If A or S refuse the side-contract, the game is played non-cooperatively. Following FLM (2003) and Celik (2009), we assume that S does not change his beliefs on A following the latter's refusal of collusion, i.e., S has passive beliefs. Let us denote by u_{ij} the status quo payoff that A receives when his type is i and S has received signal j, and they non-cooperatively play the truthful equilibrium of an individually incentive compatible mechanism. We define $u_{ij} \equiv t_{ij} - \theta_i q_{ij}$. A information rent obtained from truthfully playing the side-contract is instead $U_{ij} = b_{ij} + t(\phi_{\tau_j}(\theta_i)) - \theta_i q(\phi_{\tau_j}(\theta_i))$, where $\phi_{\tau_j}(\theta_i)$ denotes the manipulation of reports induced by the collusive side-contract when A reports having type i to S and the latter has observed j. For simplicity, we assume that the manipulation function $\phi_{\tau_j}(\theta_i)$ is deterministic. Our main results would not change with non-deterministic manipulation functions.[20] Acceptance of the side-contract by all A types imposes the following *ex post* participation constraints:

$$U_{ij} \geq u_{ij} \quad \text{for all } (i, j)\,. \tag{9}$$

Note that there is no loss of generality in assuming that the side-contract is accepted by all types.[21] We are now left to identify the manipulations that are available to S at the side-contracting stage. First, S cannot always distinguish the different A types. To circumvent this problem S must provide A the incentive not to imitate the other types. Accordingly, for $SC = \{\phi(.), b(.)\}$ to be an available side-contract for S, the following constraints must be satisfied

$$U_{ij} \geq U_{i'j} + (\theta_{i'} - \theta_i) q(\phi_{\tau_j}(\theta_{i'})) \quad \text{for all } (i', i)\,. \tag{10}$$

[19] Relaxation of the enforceability assumption is considered by Martimort (1999), Abdulkadiroglu and Chung (2003), and Khalil and Lawarree (2006).

[20] Notes available from the author.

[21] If S wants to exclude one type from the manipulation she can set $\phi^*_{\tau_j}(\theta_i)$ equal to (θ_i, τ_j), $b^*(i, j)$ equal to zero, and $U_{ij} = u_{ij}$ for that type. Given that we assumed that S has passive beliefs, this is equivalent to excluding the type from the side-contract.

The optimal side-contract solves the following problem:

$$\max_{\left\{\phi_{\tau_j}(\theta_i),U_{ij}\right\}_{i\in\{1,2,\ldots,N\}}} \sum_{i=1}^{n} p(\theta_i|\tau_j)U_S(s(\phi_{\tau_j}(\theta_i)) + t(\phi_{\tau_j}(\theta_i)) - \theta_i q(\phi_{\tau_j}(\theta_i)) - U_{ij}), \tag{11}$$

subject to (10) and (9).

Note that the partition-based information structure includes the special case where $p(\theta_i|\tau_i) = 1$, for all i, i.e., S exactly observes A type. In this case, (11) becomes

$$\max_{\left\{\phi_{\tau_j}(\theta_i),U_{ij}\right\}} U_S(s(\phi_{\tau_j}(\theta_i)) + t(\phi_{\tau_j}(\theta_i)) - \theta_i q(\phi_{\tau_j}(\theta_i)) - U_{ij}) \tag{12}$$

subject to (9).

The mechanism offered by P is *collusion-proof* if the optimal side-contract proposed by S to A and accepted by all A types entails no manipulation of reports and zero side-transfers. P mechanism is thus *collusion-proof* when the optimal manipulation of reports $\phi^*_{\tau_j}(\theta_i)$ is equal to (θ_i, τ_j), for all (i, j), and the optimal transfer $b^*(i, j)$ is equal to zero, and $U_{ij} = u_{ij}$ for all (i, j).

3 Selective Supervision Mechanism

Unlike FLM (2003) and Celik (2009), we allow P to offer a menu of contracts (or organizational structures). Each contract entails a different scope for supervision, ranging from no-supervision (S does not communicate with P) to full-supervision (the message space for S is the set of all possible signals τ). Furthermore, we allow A to report all his information and not only his cost type.[22] Part of the information is conveyed through the selection of P contract.

Formally, let $\mathcal{P}(T)$ denote the partition set of the message set $T = \{\tau_1, \tau_2, ..., \tau_N\}$ and let $\rho \in \mathcal{P}(T)$ denote an element of it, that is a subset of messages available to S. In a Selective Supervision Mechanism (SSM, hereafter) P offers a menu of contracts. A contract Γ is a quadruple $\{\rho, q_{ij}, t_{ij}, s_{ij}\}$, where q_{ij}, t_{ij} and s_{ij} denote respectively the outcome and the monetary transfers when A reports that he has type $\theta_i \in \Theta$ and S reports that the signal is $\tau_j \in \rho$. A menu of contracts, which is offered by P to either A or S, is a subset of the universal set of contracts.

The precise implementation of a SSM depends crucially on the timing of S information. Let us

[22] Following Celik (2009) and FLM (2003), we assume that the agent learns both his production cost and the supervisor's information.

denote by *Timing 1* the framework in which S receives her information before P has the opportunity to offer the mechanism. This is the setting adopted by Celik (2009) and FLM (2003). Under this assumption, S can be thought as an "informed third party" or a "witness" who happened to learn some information about A even before P had shown any interest in contracting with her. Oftentimes, S information is instead acquired after an inspection or lengthy investigation, which takes place following the acceptance of P mechanism. Let us denote by *Timing 2* this latter framework in which S receives the signal τ after accepting P mechanism.

In what follows, we show that, under *Timing 1,* Selective Supervision always achieves the *conditionally-optimal* second best. On the contrary, under *Timing 2*, the implementation of the SSM is more challenging and the second-best outcome can be achieved only under certain conditions. In order to clarify these conditions, we present a simple SSM. This mechanism always achieves the second best outcome if (i) the information structure is partition-based or (ii) the information structure is signal-based and there are only two types of agents, i.e., $N = 2$. For the remaining case (the information structure is signal-based and $N > 2$), our final proposition contains a negative result: there exists a distribution of production costs and supervisory signals that prevents *any* SSM from achieving the second best outcome.

The fact that collusion can be easily overcome under *Timing 1*, but not under *Timing 2,* seems to suggest that the assumption of *no-collusion in participation decision* is more natural under the latter timing. This conclusion is also consistent with the following observation. Under *Timing 1*, S is an "informed third party" who happens to learn some information about A before accepting P contract. If S knows A *before* the acceptance of the contract, it is only plausible to presume that they could collude on their participation decisions. On the other hand, under *Timing 2*, S needs to investigate A *after* the acceptance of the contract; this setup admits the possibility that S and A are oblivious of their respective identities before accepting the contract. Under these circumstances, *no-collusion in participation decision* could be a plausible assumption.

3.1 Timing 1

Under *Timing 1,* the implementation of the SSM is as follows: (i) At date -1, S learns τ and A learns θ and τ, (ii) At date 0, P offers a menu of contracts (or organizational structures) to S, who can either select one contract or refuse all contracts. If S refuses the game ends, (iii) At date 1, A decides whether to accept or refuse the contract selected by S at date 0. If A refuses the game ends, (iv) At date 2, S and A can stipulate a side-contract. If they do not stipulate a side-contract, the mechanism is played non-cooperatively by A and S, (v) At date 3, production and transfers take place.

If S learns τ at date -1, P can costlessly elicit τ by offering a menu of contracts to S. By assumption (*no collusion in participation decisions*), S selects the contract in a non-cooperative fashion because she cannot collude with A at this stage. Once the contract is selected, A and S can behave cooperatively and agree to respond to P collusively.

Let us denote by s_{ijk} (respectively t_{ijk} and q_{ijk}) S wage (respectively A transfer and the output target) when A reports that he has type i and that S signal is k and when S selects the contract j. P can design a SSM such that selecting contract j is equivalent to reporting the signal τ_j. Having this schedule in place, it is easy to see that this framework is equivalent to what discussed in Section 2.2 (Non-Cooperative Implementation). As before, write $s_{ijj} = s_{ij}$ and $t_{ijj} = t_{ij}$. Because τ is a piece of information which is commonly known by S and A at date -1, using the logic of Nash implementation, P can costlessly elicit τ by inducing A and S to reveal their signal. A incentive constraints can be reduced to the following relevant incentive constraints:

$$u_{ij} \geq u_{i'j} + (\theta_{i'} - \theta_i)q_{i'j} \quad \text{for all } (i, i', j)\,, \tag{13}$$

and S gets zero wage $s_{ij}^{sb} = 0$ for all (i, j). Therefore, P can achieve the same outcome as with direct supervision.

One response to these findings is that they are based on non-standard participation decisions: The participation decision in this paper is different in that S at the participation stage have to decide what contract to choose out of a menu of contracts presented by P. On the contrary, standard participation decision would involve a binary decision: accept or refuse the mechanism. Our response to this observation is twofold.

First, we acknowledge that it could feel artificial to bypass the problem of collusion by allowing P to offer a menu of contracts, whose selection is, by assumption, collusion-free. Nonetheless, it is also natural to expect P to be able to offer such a menu.

Second, it is possible to show that SSM can achieve the second-best outcome even under standard participation decision. This is always the case when S signal is binary. The signal is binary in both FLM (2003), where $T = \{\tau_1, \tau_2\}$, and Celik (2009), where $T = \{\tau_1, \tau_2, \tau_3\}$ and the signals τ_1 and τ_2 are informationally equivalent. The next section offers an example.

3.1.1 An Illustrative Example With Binary Information

When S signal is binary, the implementation of the SSM is interesting because it does not require P to offer a menu of contracts. Still, the second-best outcome is implemented. P can achieve this outcome by offering a single contract that is refused by S when a certain signal is realized. S participation decision conveys important information to P and allows him to costlessly elicit τ. The following numerical example is based on the partition-based information structure and is identical to the one proposed in Section 2.1.1 (Example 2), provided that S and A are now allowed to collude. The tables below shows the optimal quantities q^{sb} (referred to as "output" in the table), utility levels of A (referred to as U in the table) and wages for S (referred to as s in the table) when A truthfully reports his type θ_i,

	S accepts the mechanism		
Type	Output	U	s
$\theta_1 = 0.25$	$4(= q_{11}^{sb})$	0.3	0
$\theta_2 = 0.5$	$1.3(= q_{21}^{sb})$	0	0
$\theta_3 = 1$	0	0	-0.65

	S refuses the mechanism		
Type	Output	U	s
$\theta_1 = 0.25$	0	0	-
$\theta_2 = 0.5$	0	0	-
$\theta_3 = 1$	$1(= q_{33}^{sb})$	0	-

The structure of such a SSM is simple. A is offered an incentive-compatible mechanism that includes q_{33}^{sb} only when S refuses the contract. By not offering the output schedule for the most inefficient type (q_{33}^{sb}), P can effectively reduce the information rent for the other type(s) when S accepts the contract.[23] Given that S and A cannot collude in participation decisions, the logic of the Nash implementation applies. P can design a SSM that induces S to refuse the contract when τ_3 is realized. A simple sketch of the intuition is offered. To begin with, recall that collusion takes place after the acceptance of the grand-contract by both parties. Therefore, the threat of coalition formation arises only in the case where S accepts the contract. When S accepts, A and S cannot find any profitable collective manipulation. Suppose that S accepts the contract; the collusive coalition clearly has no stake in misreporting types θ_2 and θ_1. On the other hand, S would like to misreport type θ_3 in the attempt to avoid the negative transfer (-0.65). A simple inspection reveals that S is indifferent between paying -0.65 to P or offering a bribe $(\theta_3 - \theta_2)q_{21}^{sb}(= -0.65)$ to convince type θ_3 to misreport his type as

[23] Note that S does not send any message. She just decides whether to participate or not.

θ_2. Also, S strictly prefers paying -0.65 than offering a bribe $(\theta_2 - \theta_1)q_{21}^{sb} - (\theta_3 - \theta_1)\,q_{11}^{sb}(= -2.7)$ to induce type θ_3 to mimic type θ_1. Accordingly, if A and S accept the contract they shall respond to it in a non-cooperative fashion.[24] A crucial aspect of this mechanism is related to S participation decision. Notice that the participation constraint for S holds when she observes τ_1 or τ_2 (i.e., the partition $\{\theta_1, \theta_2\}$), whereas it doesn't hold when the signal τ_3 (i.e., partition $\{\theta_3\}$) is observed. Given that S refusal of the contract indicates that the realized type is θ_3, the only relevant constraints when S refuses the contract are the participation and the incentive compatibility constraints for type θ_3. This extra information, conveyed by S participation decision, is acquired by P without forgoing any rents to S. Similarly, only the participation and the incentive compatibility constraints for types θ_2 and θ_1 must hold when S accepts the contract.

Clearly, this SSM could be relabeled in the following fashion: S selects from a menu consisting of two contracts. These contracts are identical to the ones described in previous the table, provided that S wages are replaced with $s = 0$ in the left-hand-side contract. Nonetheless, the objective of this section is to show that SSM can achieve the second-best outcome, even when P is not allowed to use menus of contracts. When S signal is not binary, the mechanism presented in this section does not achieve the second-best outcome but it still improves over a standard direct revelation mechanisms with full participation.[25]

3.2 Timing 2

In the previous section we established that the second-best outcome can always be implemented under *Timing 1*. Does this result hold for *Timing 2* as well? In what follows, we prove that the SSM can implement the second-best outcome in a large number of cases, but not all. To show this point, we discuss the implementation of a simple SSM, which is compatible with *Timing 2*. The simple SSM implements the second best when (i) the information structure is partition-based, or (ii) the information structure is signal-based and there are two 2 types of agents. The remaining case (the information structure is signal-based and $N > 2$) is discussed in the next section.

3.2.1 The simple SSM

The simple SSM is designed in the following fashion: (i) At date -1, A learns θ and τ, (ii) At date 0, P offers a menu of contracts (or organizational structures) to A, who can either select one contract or

[24] In other words, A sends his message non-cooperatively whenever the side-contract fails to be established.

[25] If further clarifications are required, detailed notes are available from the author.

refuse all contracts. If A refuses the game ends, (iii) At date 1, S decides whether to accept or refuse the contract selected by A at date 0. If S refuses the game ends, (iv) At date 2, S learns τ, (v) At date 3, S and A can stipulate a side-contract. If they do not stipulate a side-contract, the mechanism is played non-cooperatively by A and S, (vi) At date 4, production and transfers take place.[26] The menu of contracts is

$$PC = \{\Gamma_1, \Gamma_2, ..., \Gamma_{N-1}, \Gamma_N\}.$$

where Γ_n denotes the generic contract and $n \in [1, 2, ..., N]$ indicates the number of messages available to S. Under Γ_n, S message space contains n message(s) τ_j such that $j \leq n$. For example, in Γ_1 the only message available to S is τ_1. In Γ_2 there are two messages available, τ_1 and τ_2, and so on. Each contract Γ_n consists in a quadruple

$$\Gamma_n = \left\{\rho_n, q_{ij}^{sb}, t_{ij}^{sb}, s_{ij}^*\right\}, \qquad \text{for all } (i,j) \text{ such that } j \leq n \text{ and } i \leq N,$$

where q_{ij}^{sb}, t_{ij}^{sb} are the second-best outputs and transfers. The optimal wages depend on the information structure. Consider first the signal-based information setup, where $p(\theta_i, \tau_j) > 0$ for all (i,j). S wage is given by

$$\begin{aligned} s_{in}^* &= 0, && \text{for all } i, \\ s_{ij}^* &= u_{in}^{sb} - u_{ij}^{sb}, && \text{for all } (i,j) \text{ such that } i \neq N \text{ and } j \neq n. \end{aligned} \tag{14}$$

Note that these wages are either positive or equal to zero. They are specifically designed to ensure that the sum of the payoffs of S and A is equal to A information rent when τ_n is realized. The remaining optimal wages s_{Nj}^* solve

$$\begin{aligned} &\sum_{i=1}^{N-1} p(\theta_i, \tau_{j+1}) U_S\left(s_{ij}^*\right) + p(\theta_N, \tau_{j+1}) U_S\left(s_{Nj}\right) = \\ &\sum_{i=1}^{N} p(\theta_i, \tau_{j+1}) U_S\left(s_{ij+1}^*\right). \end{aligned} \tag{15}$$

By construction, s_{Nj}^* must be negative. These wages are designed to guarantee that S is willing to report her signal truthfully, when the game is played non-cooperatively.

Consider now the partition-based information structure, where there exists at least on type i such

[26] Note that the simple SSM is also compatible with *Timing 2*. S information can be acquired either at date -1 or at date 2. In both cases, A knows θ and τ when he selects a contract from the menu.

that $p(\theta_i, \tau_j) = 0$. S wage is given by

$$\begin{aligned} s^*_{ij} &= 0 && \text{for all } (i,j) \text{ such that } q^{sb}_{ij} = 0 \\ s^*_{ij} &= u^{sb}_{in} - u^{sb}_{ij} && \text{for all } (i,j) \text{ such that } q^{sb}_{ij} \neq 0. \end{aligned} \tag{16}$$

By definition,[27] if $p(\theta_i|\tau_j) = 0$, the optimal outputs q^{sb}_{ij} are equal to zero only if they affect condition (4). Otherwise, they are $q^{sb}_{ij} = q^{sb}_{ij-1}$. It is possible to show that, in both information structures,

Lemma 2 *If S and A behave non-cooperatively, the wage schedule of the simple Selective Supervision Mechanism induces S to report her signal truthfully:*

$$\sum_{i=1}^{N} p(\theta_i, \tau_j) U_S \left(s^*_{ij}\right) \geq \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S \left(s^*_{ij'}\right), \quad \textit{for all } (j, j'). \tag{17}$$

This Lemma helps clarify the outcome of the game when A and S behave non-cooperatively. It implies that S reports her signal truthfully and from (4) follows that A reports his true type. This outcome is important because it constitutes A outside option from colluding with S. In fact, if A refuses S side-contract, they end up playing P mechanism non-cooperatively. Having this schedule in place, we go on to present the next result.

Proposition 3 *The simple Selective Supervision Mechanism implements the second-best outcome if*

(i) the information structure is partition-based, or

(ii) the information structure is signal-based and there are two 2 types of agents,

The following section provides a numerical example that features a partition-based structure. In Section 3.3 we offer an example of the simple SSM when the information structure is signal-based.

One last remark on the simple SSM. Note that A knows θ and τ when he selects a contract. Therefore, his choice could influence S beliefs about θ. However, A choice at date 0 is designed to extract a very specific information: by choosing the contract at date 0, A reveals τ. Note that S already knows τ at date 3. Therefore, along the equilibrium path, A choice at date 0 does not alter S beliefs about θ at date 3. Finally, consider this example: when τ_1 is realized, suppose that A is expected to pick Γ_1, but he picks Γ_2 instead. Now, at date 3, S knows that she is out-of-equilibrium.

[27] See Section 2.1

In this case, we assume that S does not use A choice at date 0 to update his out-of-equilibrium beliefs about θ, i.e., S has passive believes.

3.2.2 Partition-Based Information Structure

We present the simple SSM that achieves the second-best outcome in Example 2 (Section 2.1.1). For simplicity, suppose that S is risk neutral. The table below shows the optimal quantities q^{sb} (referred to as "output" in the table), utility levels of A (referred to as U in the table) and wages for S (referred to as s in the table) when A truthfully reports his type θ_i,

Contr. 1	S reports τ_1		
Type	Output	U	s
$\theta_1 = 0.25$	4	0.3	0
$\theta_2 = 0.5$	1.3	0	0
$\theta_3 = 1$	0	0	0

Contr. 2	S reports τ_1			S reports τ_2		
Type	Output	U	s	Output	U	s
$\theta_1 = 0.25$	4	0.3	0	4	0.3	0
$\theta_2 = 0.5$	1.3	0	0	1.3	0	0
$\theta_3 = 1$	0	0	0	0	0	0

Contr. 3	S reports τ_1			S reports τ_2			S reports τ_3		
Type	Output	U	s	Output	U	s	Output	U	s
$\theta_1 = 0.25$	4	0.3	0.5	4	0.3	0.5	4	0.8	0
$\theta_2 = 0.5$	1.3	0	0.5	1.3	0	0.5	1.3	0.5	0
$\theta_3 = 1$	0	0	0	0	0	0	1	0	0

From Proposition 3 follows that this simple SSM implements the second-best outcome. The proof of Proposition 3 is in Appendix. We provide a sketch of the intuition. To begin with, suppose that S and A behave non-cooperatively. In this case S reports her signal truthfully. To see this point, consider Contract 3. Note that for each signal available to S, the contract is incentive compatible and A reports her type truthfully. Therefore, if S has observed τ_3 (i.e., the type is θ_3), her expected utility when she reports τ_3 is 0. This utility is (weakly) larger than the one S would obtain by reporting τ_2 (0) or τ_1 (0). If S has observed the signal τ_2, she (weakly) prefers to report the true signal τ_2 ($\sum_1^2 p(\theta_i|\tau_2) * 0.5 = 0.5$) rather than reporting τ_1 (0.5) or τ_3 (0). Similarly, when S learns the signal τ_1, she prefers to report truthfully (0.5) rather than reporting τ_2 (0.5) or τ_3 (0). The same applies trivially for Contract 1 and 2.

Each contract is also collusion-proof. To see this point, note that in each contract the sum of A and S payoffs is exactly the same, regardless of the signal reported by S. Moreover, the output schedule for type θ_3 and θ_2 is also the same. This entails that, even if A-S would act as single decision-making unity, they cannot find a profitable manipulation. Thus, the mechanism is played non-cooperatively and A optimally selects Contract j when the signal is τ_j.

This kind of SSM achieves the second-best outcome by endogenizing the scope of supervision. When the type is the most inefficient (following signal τ_3) Contract 3 is selected: the organizational structure is then complex, with a large scope for supervision. S has the possibility of managing a large number of signals, including the altogether bad news (τ_3) that the type is highly inefficient. On the contrary, when there is high probability that A is efficient (following signal τ_1), Contract 1 is selected and the organization collapses into an informal contract between P and A where supervision is minimal.[28]

3.3 Timing 2: Impossibility Result

Proposition 3 leaves us with a last case to analyze: the information structure is signal-based, and $N > 2$. For this remaining case we have a negative result. We look at all possible SSM and we show that

Proposition 4 *There is no SSM that implements the second-best outcome for each distribution of the cost θ and the signal τ, when (i) $N > 2$, (ii) the information structure is signal-based, (iii) and the timing of the game is Timing 2.*

The intuition behind the proof is simple. First, we derive some necessary conditions for the SSM to implement the second-best outcome. Second, we show that there is a distribution of the cost θ and the signal τ that violates these conditions. The crucial necessary conditions are (i) there must be at

[28] Note that in this particular example the SSM would work even if Contract 1 is not offered. This is due to the fact that the information structure is partition-based and some signals are informationally equivalent.

least one contract where S has at least two messages available, and (ii) inducing S to report her signal truthfully requires *at least* one wage to be negative in a certain state of the world. In the proof, we show that it is possible to find a distribution of costs and signals such that this negative wage prevents the SSM from implementing the second best.

Interestingly, the partition-based structure does not feature the same pathology. From (16) follows that S wage can be set (weakly) greater than zero in all states of the world. This is due to the fact that S learns whether or not a certain type has zero probability to be realized for a given signal. Moreover, each signal specifies a different partition, with no overlap between partitions. Thus, there is no need to set negative wages to induce S to report her signal truthfully (see Proof of Lemma 2, Part 1).

Although Proposition 4 shows that there is no guarantee that the second-best outcome can be implemented in the signal-based structure with $N > 2$, it is still possible to show that a simple SSM *can* do so under certain parametric conditions. An example is provided in the following section.

3.3.1 Signal-Based Information Structure

This numerical example is identical to the one proposed in Section 2.1.1 (Example 1), except that S and A are now allowed to collude. For simplicity, suppose that S is risk neutral. The structure of the mechanism is based on the simple SSM presented in Section 3.2. The table below shows the optimal quantities q^{sb} (referred to as "output" in the table), utility levels of A (referred to as U in the table) and wages for S (referred to as s in the table) when A truthfully reports his type θ_i and S reports τ_j.

Contr. 1	S reports τ_1		
Type	Output	U	s
$\theta_1 = 0.25$	4	0.4105	0
$\theta_2 = 0.5$	0.8421	0.2	0
$\theta_3 = 1$	0.4	0	0

Contr. 2	S reports τ_1			S reports τ_2		
Type	Output	U	s	Output	U	s
$\theta_1 = 0.25$	4	0.4105	0.3346	4	0.7451	0
$\theta_2 = 0.5$	0.8421	0.2	0.1333	1.6471	0.3333	0
$\theta_3 = 1$	0.4	0	-0.1937	0.6667	0	0

Contr. 3	S reports τ_1			S reports τ_2			S reports τ_3		
Type	Output	U	s	Output	U	s	Output	U	s
$\theta_1 = 0.25$	4	0.4105	0.4312	4	0.7451	0.0966	4	0.8417	0
$\theta_2 = 0.5$	0.8421	0.2	0.175	1.6471	0.3333	0.0417	1.8667	0.3750	0
$\theta_3 = 1$	0.4	0	-0.2519	0.6667	0	-0.03237	0.7500	0	0

In what follows, we provide a simple sketch of the basic features of the mechanism. The complete analysis in offered in Appendix 1. To begin with, suppose that S and A behave non-cooperatively. In this case S reports her signal truthfully. To see this point, consider Contract 3. If S has observed τ_3, her expected utility when she reports τ_3 is $\sum_1^3 p(\theta_i|\tau_3) * 0 = 0$. This expected utility is (weakly) larger than the one S would obtain by reporting τ_2 (0) or τ_1 (-0.068). If S has observed the signal τ_2, she is better off reporting the true signal τ_2 (0.012) rather than reporting τ_1 (0) or τ_3 (0). Similarly, when S learns the signal τ_1, she prefers to report truthfully (0.2) rather than reporting τ_2 (0.05) or τ_3 (0). Now consider Contract 2: If S has observed τ_3 or τ_2, her expected utility when she reports τ_2 is (weakly) larger than the one S would obtain by reporting τ_1. Finally, it is easy to check that S truthfully reports τ_1.

Each contract is also collusion-proof. To prove this point, note that when A has type θ_1 or θ_2 the coalition A-S has no stake in misreporting the signal. Indeed, the sum of A and S payoffs is exactly the same, regardless of the signal reported by S: the coalition payoff is 0.8417 when the type is θ_1, and 0.3750 when the type is θ_2. This observation alone does not guarantee that S and A do not want to manipulate both τ and θ. Nevertheless, one can easily check that the *monotonicity constraints* (2) hold *across* all contracts, i.e., $q_{ij} \geq q_{i'j}$ for all (i, i', j) such that $i' > i$. This entails that $U_S(s^*_{ij}) \geq U_S(s^*_{i'j'} + t^{sb}_{i'j'} - \theta_i q^{sb}_{i'j'} - U_{ij})$ for all (i, i', j, j') with $i \neq 3$ and $U_{ij} \geq u^{sb}_{ij}$ (the latter condition follows from (9)). In other words, there is no direct benefit from misreporting type θ_1 or θ_2. Nonetheless, the side-contract has to be offered before A has revealed his type to S, and S thus takes into account that changing what she commits to announce in a certain state affects the information rent paid to the other types (i.e., it could affect U_{ij} for the other types). Given that (9) must hold, by misreporting type θ_1 (or θ_2), S cannot reduce the information rent for the other types below their outside options, i.e., u^{sb}_{ij}. Therefore, the manipulation where only type θ_1 is misreported (i.e., $\phi_{\tau_j}(\theta_1) = (\theta_{i'}, \tau_{j'})$, $\phi_{\tau_j}(\theta_2) = (\theta_2, \tau_j)$ and $\phi_{\tau_j}(\theta_3) = (\theta_3, \tau_j)$ for all (i', j, j')) cannot be beneficial. Indeed, under this manipulation, the information rents for type θ_2 and θ_3 are already at their minimum possible level, i.e., $U_{2j} = u^{sb}_{2j}$ and $U_{3j} = u^{sb}_{3j}$ for all j.[29] It is straightforward to see that the same argument applies to the manipulation where only type θ_2 is misreported. We are left to prove that the manipulation of type θ_3 only is also not beneficial. This will be sufficient to show that the are no profitable manipulations available in the simple SSM. To see this point note the following argument: If (i) the misreport of each A type is not beneficial per se, i.e., $U_S(s^*_{ij}) \geq U_S(s^*_{i'j'} + t^{sb}_{i'j'} - \theta_i q^{sb}_{i'j'} - U_{ij})$ for all (i, i', j, j'), and (ii) the information rents for the other types cannot be reduced below u^{sb}_{ij} for all (i, j), then the optimal manipulation of reports $\phi^*_{\tau_j}(\theta_i)$ is equal to (θ_i, τ_j), for all (i, j), the optimal transfer $b^*(i, j)$ is equal to zero, and $U^*_{ij} = u^{sb}_{ij}$ for all (i, j).

We are now left to prove that the manipulation of type θ_3 only is not beneficial. In the case where A has type θ_3, S could receive a negative payoff. Consider a simple example where A has selected Contract 3 and S has observed τ_2. S would then like to report τ_3 when A has type θ_3 and get zero instead of a negative payoff (-0.03237). But this side-contract has to be offered before A has revealed his type to S and S takes thus into account that changing what she commits to announce in state (θ_3, τ_2) also increases the information rent paid to the other types. In Appendix 1 we show that S is never willing to offer a side-contract to A. Thus, A optimally selects Contract j when the signal is τ_j. This implements the second-best outcome.

Notice that this outcome is feasible because of one crucial assumption: namely, *no-collusion in*

[29] Given that (9) must hold, the misreport of type θ_1 cannot further reduce the information rent for type θ_2 or θ_3.

participation decisions. This assumption entails that A and S cannot collude on the selection of Contract j. A and S can collude only after A decides what contract to accept. In other words, they can collude only "within" each Contract j.

As before, this kind of SSM endogenizes the scope of supervision. When there is high probability that A is efficient, the organizational structure is simple, with no supervisory activity. On the other hand, when there is high probability that the type is inefficient, S has the possibility of managing multiple signals.

4 Remarks on Collusion-Proof Implementation

4.1 Collusive Behavior and Supervisory Information

A couple of issues concerning bargaining power and collusive behaviors are worth noting. Under *Timing 1*, results do not depend on the distribution of the bargaining power allocation inside the coalition nor do they rest on the identity of the coalition member who offers and initiate the collusive agreement. Under *Timing 2* and partition-based information structure, our result is still robust to these aspects. Indeed, the proof is based on the strong notion of collusion-proofness expressed in (12). On the contrary, under *Timing 2* and signal-based information structure, the proof is based on the weak notion of collusion-proofness in (11). This entails that we still require S to make a take-it-or-leave-it offer to A.

Finally, collusion-proof implementation does not depend on special assumptions about the accuracy of S information. For example, suppose S learns A cost, i.e., there is no residual asymmetric information between A and S. The mechanism simply reduces to a special case of partition-based information structure where $p(\theta_i|\tau_i) = 1$, for all i.

4.2 Implications for Decentralization

The recent literature evaluating delegation when agents collude offers an intriguing puzzle. Despite the very similar setting considered in FLM (2003) and Celik (2009), the results of these papers are strikingly different: FLM (2003) find that delegation is always equivalent to centralization, whereas Celik (2009) finds that centralization is superior in general. The results of this paper seems to confirm that centralization performs better than delegation.[30] The crucial assumption that drives this result

[30] On the contrary, Baliga and Sjostrom (1998), and Laffont and Martimort (1998) consider a setup that does not involve supervision, showing that under certain conditions delegation is the optimal organizational response to collusion.

is no collusion in participation decisions. P can improve his payoff by contracting directly with S and A. This is due to the fact that participation decisions can be exploited to extract supplementary information. This is not possible under decentralized contracting.

4.3 Applications

The precise implementation of *selective supervision* depends on both the nature and the timing of supervisor's information. First, consider the case where the information is binary. If the supervisor receives this information after the acceptance of the principal's offer, the example of the agent choosing between a regime with or without supervision applies.[31] On the other hand, if the supervisor receives her information before the acceptance decision there is an alternative mechanism available. In this case, the principal could offer a mechanism in which the supervisor can opt out in some states of the world. An example fitting this case is that of an advisor who is asked to write a recommendation letter for a student. The advisor may refuse to write the letter, revealing some information about the agent's type. By the same token, foreign embassies have the discretion to refuse immigration permits to applicants whom they do not consider suitable for admission. Similarly, hiring committees may refuse to offer interviews to certain candidates. Failure to receive interviews signal a portion of the private information available to the committees. In all these cases, the supervisor's decision to opt out conveys information about the applicants' characteristics.

When the nature of the supervisor's information is not binary, the implementation of *selective supervision* becomes more nuanced. In this case, the principal proposes a menu of mechanisms. Each contract specifies a different scope for supervisory activity, where the scope of supervision refers to the dimension of the message space available to the supervisor. Applications include work contracts subject to different degrees of discretion, self-reporting schemes limited to certain crimes, letters of recommendation with different degrees of approbation, restricted visa permits, tax amnesties for specific types of evasions, or work contracts subject to different degrees of discretion.

[31] Burlando and Motta (2008b) consider a framework with two types of agents and limited liability for the supervisor. Unlike this paper, they consider hard-information supervision where, with some probability, the supervisor learns the true agent's type, otherwise she learns nothing. They show that there exists a mechanism that eliminates agency costs by providing the productive agent with the possibility of avoiding inspection. When the productive agent is risk averse, the mechanism also provides him with an insurance coverage: as a consequence, this mechanism would be worthwhile even abstracting from collusion.

5 Conclusions

This paper analyzes the role of supervision in organizations involving both supervisory and productive tasks when these two tasks are performed by parties prone to collusion. The main contribution of the paper is to show the role of endogenous selection of supervisory activity by the principal. If collusion between supervisor and agent can occur only after they have decided to participate in the mechanism, endogenous selection of supervisory activity can costlessly eliminate collusion. This conclusion is robust to alternative information structures, collusive behaviors and specification of agent's types. Surprisingly, the cost related to collusion can be fully eliminated even when there is no residual asymmetric information between the agent and the supervisor. This paper, therefore, presents results in contrast to the important insight gained from Laffont and Martimort (1997, 2000) that agents' asymmetric information constitutes an obstacle to collusive arrangements. Rather, this paper highlights the fact that the inability to collude prior to making a decision to participate in the mechanism represents an "Achilles' heel" of collusive coalition. The result that collusion can be eliminated at no cost in this environment allows us to highlight the important assumptions required for collusion to be a salient issue in the existing literature.

We suspect that *selective supervision* is useful in a setting where collusion occurs prior to participation, although in that context, it is unlikely to costlessly eliminate collusion. This issue remains to be explored in future research. Another limitation of our analysis is that we consider a one-shot mechanism with one principal and one agent. Analysis of dynamic mechanisms with multiple principals is left for future research (see, for example, Pavan and Calzolari, 2009.)

6 Appendix

6.1 Appendix 1: Numerical Example (Section 3.3.1)

As we showed in Section 3.3.1, the potential for collusion arises when A has type θ_3. In this case, S could receive a negative payoff. Suppose that A has selected Contract 3 and S has observed τ_2. S would then like to report τ_3 when A has type θ_3, and get zero instead of a negative payoff (-0.0324). But this side-contract has to be offered before A has revealed his type to S and S takes thus into account that changing what she commits to announce in state (θ_3, τ_2) also affects the information rent paid to all other A types. During the side-contracting stage, type θ_2 is now willing to report to S that he has type θ_3 and earn some extra information rent (0.0417). To prevent this, S has to provide type

θ_2 with this extra information rent (0.0417). Moreover, in order for type θ_1 not to be willing to report to S that he has type θ_2 he also must see his information rent increased by the same amount (0.0417). Therefore, the costs of this side-contract $(0.0417 + 0.0417)$ outweigh the benefits (0.0328). Clearly S would not like to report τ_1 when A has type θ_3 and get (-0.1937) instead of (-0.0324). Another possible strategy would entail reporting that the state is (θ_2, τ_2) when the true state is (θ_3, τ_2). In this case to convince type θ_3 to accept the misreport S has to forgo (0.4902). It is sufficient to see that the benefits when type θ_3 is realized (0.0324) are not enough to cover the costs (0.4902). Another possible strategy would entail reporting that the state is (θ_1, τ_2) when the true state is (θ_3, τ_2). In this case to convince type θ_3 to accept the misreport S has to forgo (2.2549). It is sufficient to check that the benefits when type θ_3 is realized (0.0324) are not enough to cover the costs (2.2549). Another possible strategy would entail reporting that the state is (θ_2, τ_3) when the true state is (θ_3, τ_2). In this case to convince type θ_3 to accept the misreport S has to forgo (0.5583). It is sufficient to check that the benefits when type θ_3 is realized (0.0324) are not enough to cover the costs (0.5583). Another possible strategy would entail reporting that the state is (θ_1, τ_3) when the true state is (θ_3, τ_2). In this case to convince type θ_3 to accept the misreport S has to forgo (2.158 3). It is sufficient to check that the benefits when type θ_3 is realized (0.0324) are not enough to cover the costs (2. 158 3). Another possible strategy would entail reporting that the state is (θ_2, τ_1) when the true state is (θ_3, τ_2). In this case to convince type θ_3 to accept the misreport S has to forgo (0.2210). It is sufficient to check that the benefits when type θ_3 is realized (0.0324) are not enough to cover the costs (0.2210). The last possible strategy would entail reporting that the state is (θ_1, τ_1) when the true state is (θ_3, τ_2). In this case to convince type θ_3 to accept the misreport S has to forgo (2. 5895). It is sufficient to check that the benefits when type θ_3 is realized (0.03237) are not enough to cover the costs (2. 589 5). Using the same method, it is easy to check that S is not willing to misreport θ_3 when she has observed τ_1. This proves that Contract 3 is collusion-proof. Finally, Contract 2 is also collusion-proof. To see this point, it is sufficient to note that S salary in state (θ_3, τ_1) is larger in Contract 2 than in Contract 3. Therefore, if S is not willing to misreport θ_3 in Contract 3, she is also not willing to misreport type θ_3 in Contract 2. This establishes that the mechanism is collusion-proof.

6.2 Proof of Lemma 1

We have already shown that the second-best outcome always implements the first-best output $q_{1j}^{sb} = q_1^{fb}$ for the most efficient agent for all j. Consider now the case where $i \neq 1$. Rearrange (1) and obtain

$$\frac{p(\theta_i|\tau_j)}{p(\theta_i'|\tau_j)} = \frac{p_{ij}}{p_{i'j}} \geq \frac{p(\theta_i|\tau_j')}{p(\theta_i'|\tau_j')} = \frac{p_{ij'}}{p_{i'j'}} \tag{18}$$

for all (i, i', j, j') such that $j' \geq j$, $i' \geq i > 1$, $p_{i'j} \neq 0$ and $p_{i'j'} \neq 0$. Sum up (18) for all i smaller than i' and obtain

$$\sum_{z=1}^{i'-1} \frac{p_{zj}}{p_{i'j}} \geq \sum_{z=1}^{i'-1} \frac{p_{zj'}}{p_{i'j'}} \tag{19}$$

Recall that optimality requires (5) to be satisfied. Condition (19) ensures that

$$\theta_{i'} + \sum_{z=1}^{i'-1} \frac{p_{zj}}{p_{i'j}}(\theta_{i'} - \theta_{i'-1}) \geq \theta_{i'} + \sum_{z=1}^{i'-1} \frac{p_{zj'}}{p_{i'j'}}(\theta_{i'} - \theta_{i'-1}).$$

This condition coupled with the fact that $W''(.) < 0$ establishes that $q_{i'j}^{sb} \leq q_{i'j'}^{sb}$ for all (i', j, j') such that $j' \geq j$, $p_{i'j} \neq 0$ and $p_{i'j'} \neq 0$. Notice that $q_{i'j}^{sb} \leq q_{i'j'}^{sb}$ holds for all $i' \in [2, ..., N]$. We are now left to prove that this result holds when $p_{i'j} = 0$ and/or $p_{i'j'} = 0$.

There are three cases: (a) $p_{i'j} = 0$ and $p_{i'j'} = 0$, (b) $p_{i'j} = 0$ and $p_{i'j'} > 0$, (c) $p_{i'j} > 0$ and $p_{i'j'} = 0$. Under (a) it is straightforward to see that $q_{i'j}^{sb} = 0 \leq q_{i'j'}^{sb} = 0$. Under (b) optimality requires that P assigns type $\theta_{i'}$ a (strictly) positive production schedule $q_{i'j'}^{sb} > 0$ when the signal is τ_j' and zero when the signal is τ_j. Naturally, $q_{i'j}^{sb} = 0 < q_{i'j'}^{sb}$. Under (c) note that $p_{i'j} > 0$ and $p_{i'j'} = 0$ can be denoted as $p(\theta_i'|\tau_j) > 0$ and $p(\theta_i'|\tau_j') = 0$. In this case condition (1) is met only if $p(\theta_i|\tau_j') = 0$ for all $\theta_i < \theta_i'$. But this implies that $q_{i'j'}$ (the quantity produced by type θ_i' when τ_j' is realized) affects neither $W(.)$ nor the information rents that P has to forgo. To see this point note that all types more efficient than θ_i' (i.e. all $\theta_i < \theta_i'$) have also zero probability to be realized when P receives τ_j'. Therefore, any quantity $q_{i'j'}$ large-at-will can be assigned, provided that the monotonicity constraint (2) is satisfied. More specifically P can optimally select $q_{i'j'}$ such that $q_{i'j}^{sb} \leq q_{i'j'}^{sb}$. This entails that $q_{i'j}^{sb} \leq q_{i'j'}^{sb}$ for all $j' \geq j$ and all $i' \in [2, ..., N]$.

Relable the last result and obtain, $q_{ij}^{sb} \leq q_{ij'}^{sb}$ for all $j' \geq j$ and all $i \in [2, ..., N]$. In the first part of the proof we showed that $q_{1j}^{sb} \leq q_{1j'}^{sb}$ for $j' \geq j$. These two results establish that $q_{ij}^{sb} \leq q_{ij'}^{sb}$ for all $j' \geq j$ and all i.

6.3 Proof of Lemma 2

The proof has two parts. In the first part we prove that (17) holds for the partition-based information structure. In the second part we prove that (17) holds for the signal-based information setup. Before proceeding, note that each contract

$$\Gamma_n = \left\{\rho_n, q_{ij}^{sb}, t_{ij}^{sb}, s_{ij}^*\right\}$$

is an incentive compatible mechanism that satisfies (4). By construction, A is induced to report his true type when he plays the mechanism non-cooperatively.

6.3.1 Part 1

We first consider the partition-based information structure. Here, we assume that A and S behave non-cooperatively. When S receives the signal j he has no stake in misreporting her signal as $j' < j$ because $U(s_{ij}^*) \geq U(s_{ij'}^*)$ for all types i that have a positive probability to be realized when S receives the signal j. To see this point, note that there are only three possibilities:

(a) $U(s_{ij'}^*) = U(s_{ij}^*)$ for all i. This is the case where j' and j are informationally equivalent signals with $p(\theta_i, \tau_j) = p(\theta_i, \tau_j')$ for all i.

(b) $U(s_{ij'}^*) = 0$ for all i such that $p(\theta_i, \tau_j) > 0$. This is the case where j' and j are not informationally equivalent. It follows directly from the information structure. For any signal $j > 1$ we have the following condition: if $p(\theta_i, \tau_j) > 0$, then $p(\theta_i, \tau_{j'}) = 0$ for all other non-informationally equivalent signals $j' \epsilon [j-1, ..., 1]$. Thus the *conditionally optimal* second-best outputs are $q_{ij'}^{sb} = q_{ij'-1}^{sb} = ... = q_{i1}^{sb} = 0$. From (16) follows that $s_{ij'}^* = 0$.

(c) The remaining case (all i such that $p(\theta_i, \tau_j) = 0$) is irrelevant because these wages are never realized when S learns j.

This proves that (17) holds for all $j' < j$. We are left to prove that (17) holds also for all $j' > j$. Note that if S receives the signal j he has no stake in misreporting her signal as $j' > j$ because there are only three possibilities:

(a) $U(s_{ij'}^*) = U(s_{ij}^*)$ when j' and j are informationally equivalent signals with $p(\theta_i, \tau_j) = p(\theta_i, \tau_j')$ for all i.

(b) $U(s_{ij'}^*) \leq U(s_{ij}^*)$ for all i such that $p(\theta_i, \tau_j) > 0$. This is the case where j' and j are not informationally equivalent. $U(s_{ij'}^*) \leq U(s_{ij}^*)$ follows directly from (16).

(c) The remaining case (all i such that $p(\theta_i, \tau_j) = 0$) is irrelevant because these wages are never realized when S learns j.

This establishes that (17) holds when at least one output schedule is equal to zero, i.e., $q_{ij}^{sb} = 0$.

6.3.2 Part 2

In the second part we prove that (17) holds when all output schedule are strictly positive, i.e., $q_{ij}^{sb} > 0$. This part is divided into two sections.

Section (a): We prove that (17) holds for all (j, j') such that $j' > j$. Using (15), compute the optimal value of $U_S(s_{Nj+1})$ and $U_S(s_{Nj})$

$$U_S\left(s_{Nj+1}^*\right) = \frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_{j+2})\left[U_S\left(s_{ij+2}^*\right) - U_S\left(s_{ij+1}^*\right)\right] + p(\theta_N, \tau_{j+2})U_S\left(s_{Nj+2}^*\right)}{p(\theta_N, \tau_{j+2})} \tag{20}$$

$$U_S\left(s_{Nj}^*\right) = \frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_{j+1})\left[U_S\left(s_{ij+1}^*\right) - U_S\left(s_{ij}^*\right)\right] + p(\theta_N, \tau_{j+1})U_S\left(s_{Nj+1}^*\right)}{p(\theta_N, \tau_{j+1})} \tag{21}$$

We are now ready to prove that (17) is satisfied for each pair j and $j' = j + 1$. Rearrange (17) and obtain

$$\sum_{i=1}^{N-1} p(\theta_i, \tau_j)U_S\left(s_{ij}^*\right) + p(\theta_N, \tau_j)U_S\left(s_{Nj}^*\right) \geq \sum_{i=1}^{N-1} p(\theta_i, \tau_j)U_S\left(s_{ij+1}^*\right) + p(\theta_N, \tau_j)U_S\left(s_{Nj+1}^*\right)$$

Substituting (20) and (21) into the former expression yields

$$\sum_{i=1}^{N-1} p(\theta_i, \tau_j)U_S\left(s_{ij}^*\right) + p(\theta_N, \tau_j)\frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_{j+1})\left[U_S\left(s_{ij+1}^*\right) - U_S\left(s_{ij}^*\right)\right] + p(\theta_N, \tau_{j+1})U_S\left(s_{Nj+1}^*\right)}{p(\theta_N, \tau_{j+1})} \geq$$
$$\sum_{i=1}^{N-1} p(\theta_i, \tau_j)U_S\left(s_{ij+1}^*\right) + p(\theta_N, \tau_j)\frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_{j+2})\left[U_S\left(s_{ij+2}^*\right) - U_S\left(s_{ij+1}^*\right)\right] + p(\theta_N, \tau_{j+2})U_S\left(s_{Nj+2}^*\right)}{p(\theta_N, \tau_{j+2})}$$

This expression can be written as

$$\frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_j) \left[U_S\left(s_{ij}^*\right) - U_S\left(s_{ij+1}^*\right)\right]}{p(\theta_N, \tau_j)} - \frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_{j+1}) \left[U_S\left(s_{ij}^*\right) - U_S\left(s_{ij+1}^*\right)\right]}{p(\theta_N, \tau_{j+1})} \geq \tag{22}$$

$$\frac{\sum_{i=1}^{N} p(\theta_i, \tau_{j+2}) \left[U_S\left(s_{ij+2}^*\right) - U_S\left(s_{ij+1}^*\right)\right]}{p(\theta_N, \tau_{j+2})}$$

From (20) follows that $\sum_{i=1}^{N} p(\theta_i, \tau_{j+2}) \left[U_S\left(s_{ij+2}^*\right) - U_S\left(s_{ij+1}^*\right)\right] = 0$ and (14) ensures that $\sum_{i=1}^{N-1} \left[U_S\left(s_{ij}^*\right) - U_S\left(s_{ij+1}^*\right)\right] \geq 0$, therefore (22) can be rewritten as

$$\frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_j)}{p(\theta_N, \tau_j)} \geq \frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_{j+1})}{p(\theta_N, \tau_{j+1})}$$

which is always satisfied when the monotone likelihood ratio property applies. This proves that (17) is satisfied for each pair j and j' such that $j' = j + 1$. We now prove that this result holds for each pair j and $j' = j + 2$,

$$\sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s_{ij}^*\right) \geq \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s_{ij+2}^*\right) \tag{23}$$

We adopt the same solution concept as before. Using (15), compute the optimal value of $U_S\left(s_{Nj+2}\right)$ and $U_S\left(s_{Nj}\right)$ and then substitute them into (23)

$$\sum_{i=1}^{N-1} p(\theta_i, \tau_j) U_S\left(s_{ij}^*\right) + p(\theta_N, \tau_j) \frac{\sum_{i=1}^{N-1} p(\theta_i, \tau_{j+1}) \left[U_S\left(s_{ij+1}^*\right) - U_S\left(s_{ij}^*\right)\right] + p(\theta_N, \tau_{j+1}) U_S\left(s_{nj+1}^*\right)}{p(\theta_N, \tau_{j+1})} \geq$$

$$\sum_{i=1}^{N-1} p(\theta_i, \tau_j) U_S\left(s_{ij+2}^*\right) + p(\theta_N, \tau_j) \frac{\sum_{i=1}^{n-1} p(\theta_i, \tau_{j+3}) \left[U_S\left(s_{ij+3}^*\right) - U_S\left(s_{ij+2}^*\right)\right] + p(\theta_N, \tau_{j+3}) U_S\left(s_{Nj+3}^*\right)}{p(\theta_N, \tau_{j+3})}$$

After several rearrangements, this expression can be rewritten as

$$\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_j)\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_j)}-\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_{j+1})\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+1}\right)\right]}{p(\theta_N,\tau_{j+1})}\geq$$
$$\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_{j+3})\left[U_S\left(s^*_{ij+3}\right)-U_S\left(s^*_{ij+2}\right)\right]+p(\theta_N,\tau_{j+3})U_S\left(s^*_{ij+3}\right)}{p(\theta_N,\tau_{j+3})}-U_S\left(s^*_{Nj+1}\right)$$

Adding and subtracting $\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_{j+1})\left[U_S\left(s^*_{ij+1}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_{j+1})}$ (in the LHS) and $U_S\left(s^*_{ij+2}\right)$ (in the RHS) yields

$$\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_j)\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_j)}-\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_{j+1})\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_{j+1})}+$$
$$\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_{j+1})\left[U_S\left(s^*_{ij+1}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_{j+1})}\geq$$
$$\frac{\sum_{i=1}^{N} p(\theta_i,\tau_{j+3})\left[U_S\left(s^*_{ij+3}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_{j+3})}-U_S\left(s^*_{Nj+1}\right)+U_S\left(s^*_{Nj+2}\right)$$

Rearrange and obtain,

$$\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_j)\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_j)}-\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_{j+1})\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_{j+1})}\geq$$
$$\frac{\sum_{i=1}^{N} p(\theta_i,\tau_{j+3})\left[U_S\left(s^*_{ij+3}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_{j+3})}-\frac{\sum_{i=1}^{N} p(\theta_i,\tau_{j+1})\left[U_S\left(s^*_{ij+1}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_{j+1})} \tag{24}$$

From (15) we know that $\sum_{i=1}^{N} p(\theta_i,\tau_{j+3})\left[U_S\left(s^*_{ij+3}\right)-U_S\left(s^*_{ij+2}\right)\right]=0$, and above we have just proved that $\sum_{i=1}^{N} p(\theta_i,\tau_{j+1})\left[U_S\left(s^*_{ij+1}\right)-U_S\left(s^*_{ij+2}\right)\right]\geq 0$. Therefore the right-hand-side of (24) must be negative. Recall that (14) implies that $\sum_{i=1}^{N-1}\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+2}\right)\right]\geq 0$, therefore the monotone likelihood ratio property ensures that $\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_j)\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_j)}-\frac{\sum_{i=1}^{N-1} p(\theta_i,\tau_{j+1})\left[U_S\left(s^*_{ij}\right)-U_S\left(s^*_{ij+2}\right)\right]}{p(\theta_N,\tau_{j+1})}\geq 0.$

This establishes that (17) is satisfied for any j and $j'=j+2$. A fast inspection reveals that the same strategy can be used recursively to prove that (17) is also satisfied for $j'=j+3$, $j'=j+4$ and so on. Clearly for $j = n$ (17) is trivially satisfied. This establishes the first section of part 2.

Section (b): We prove that (17) holds for all (j, j') such that $j' < j$. From (15) follows that

$$\sum_{i=1}^{N} p(\theta_i, \tau_{j-1}) U_S\left(s^*_{ij-1}\right) = \sum_{i=1}^{N} p(\theta_i, \tau_{j-1}) U_S\left(s^*_{ij-2}\right) \tag{25}$$

Note that the monotone likelihood ratio property also implies first-order stochastic dominance. The following conditions must hold

$$\sum_{i=1}^{N} p(\theta_i, \tau_{j-1}) U_S\left(s^*_{ij-2}\right) \geq \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij-2}\right) \tag{26}$$

$$\sum_{i=1}^{N} p(\theta_i, \tau_{j-1}) U_S\left(s^*_{ij-1}\right) \geq \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij-1}\right). \tag{27}$$

Substitute (25) into (27), divide (26) by (27) and obtain

$$\sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij-1}\right) \geq \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij-2}\right)$$

Using (15) we have

$$\sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij-1}\right) = \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij}\right) \geq \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij-2}\right)$$

Clearly, this is a recursive argument: For example consider $j-2$ and $j-3$. From (15) follows that

$$\sum_{i=1}^{N} p(\theta_i, \tau_{j-2}) U_S\left(s^*_{ij-2}\right) = \sum_{i=1}^{N} p(\theta_i, \tau_{j-2}) U_S\left(s^*_{ij-3}\right).$$

Proceeding as before and using first-order stochastic dominance, we obtain

$$\sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij-2}\right) \geq \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij-3}\right)$$

Applying the same solution concept for all $j' < j$ yields

$$\sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij}\right) \geq \sum_{i=1}^{N} p(\theta_i, \tau_j) U_S\left(s^*_{ij'}\right) \quad \text{for all } (j, j') \text{ such that } j' < j.$$

This establishes the proof.

6.4 Proof of Proposition 3

The proof has two parts. One considers the partition-based information setup, whereas the other one considers the signal-based information setup with $N = 2$. In both parts we will first prove that a generic contract Γ_n is collusion-proof. Secondly, we will prove that PC implements the second-best outcome.

6.4.1 Part 1

In this part of the proof we consider the partition-based structure. Note that this structure includes the special case where $p(\theta_i|\tau_i) = 1$, for all i, i.e., S exactly observes A type. Therefore, in order to prove that Γ_n is collusion-proof (even when S exactly observes A type), we need to show that $\phi^*_{\tau_j}(\theta_i) = (i, j)$ solves (12) and the optimal transfer $b^*(i, j)$ is equal to zero for all (i, j). For a given U_{ij}, we need to prove that

$$U_S(s^*_{ij}) \geq U_S(s^*_{i'j'} + t^{sb}_{i'j'} - \theta_i q^{sb}_{i'j'} - U_{ij}), \tag{28}$$

for all (i, i', j, j') such that $p(\theta_i, \tau_j) > 0$. To begin with, let us discuss the trivial case where S, having observed signal j and learned that A has type i through side-contracting, reports that the state of nature is (i', j') where $q^{sb}_{i'j'} = 0$, $t^{sb}_{i'j'} = 0$, and $s^*_{i'j'} = 0$. From (16) follows that $U_S(s^*_{ij}) \geq 0$ for all (i, j). Therefore,

$$U_S(s^*_{ij}) \geq U_S(s^*_{i'j'} + t^{sb}_{i'j'} - \theta_i q^{sb}_{i'j'} - U_{ij}) = U_S(-U_{ij}).$$

Consider now the remaining manipulations (i', j') with $q^{sb}_{i'j'} \neq 0$. Notice that from (16) follows that

$$s^*_{ij} = u^{sb}_{in} - u^{sb}_{ij}. \tag{29}$$

Recall that by definition $u^{sb}_{i'j'} = t^{sb}_{i'j'} - \theta_{i'} q^{sb}_{i'j'}$. Subtract $\theta_i q^{sb}_{i'j'}$ from both sides of this expression and rearrange to obtain

$$t^{sb}_{i'j'} - \theta_i q^{sb}_{i'j'} = u^{sb}_{i'j'} + (\theta_{i'} - \theta_i) q^{sb}_{i'j'}. \tag{30}$$

Substituting the latter expression into (28) and using (29) twice we have

$$U_S(u^{sb}_{in} - u^{sb}_{ij}) \geq U_S(u^{sb}_{i'n} + (\theta_{i'} - \theta_i) q^{sb}_{i'j'} - U_{ij}). \tag{31}$$

This condition holds for all (i, i', j, j'). To see this, recall the properties of the connected partition information structure: If there are two signals τ'_j and τ''_j such that $p(\theta_i, \tau'_j) > 0$ and $p(\theta_i, \tau''_j) > 0$, they must convey the same information, i.e., $p(\theta_i, \tau'_j) = p(\theta_i, \tau''_j)$ for all θ_i. This entails that $q^{sb}_{i'n} = q^{sb}_{i'j'}$, for all (i', j') such that $p(\theta'_i, \tau'_j) > 0$. In the case where $p(\theta'_i, \tau'_j) = 0$ and $q^{sb}_{i'j'} \neq 0$, by definition $q^{sb}_{i'j'} = q^{sb}_{i'j'-1}$.[32] Therefore, $q^{sb}_{i'n} = q^{sb}_{i'j'}$, for all (i', j') such that $q^{sb}_{i'j'} \neq 0$. Note that, at the optimum of the Direct Supervision problem, the incentive-compatibility constraints are satisfied. For $j = n$ this entails that $u^{sb}_{in} \geq u^{sb}_{i'n} + (\theta_{i'} - \theta_i)q^{sb}_{i'n}$. It follows that

$$u^{sb}_{in} \geq u^{sb}_{i'n} + (\theta_{i'} - \theta_i)q^{sb}_{i'n} = u^{sb}_{i'n} + (\theta_{i'} - \theta_i)q^{sb}_{i'j'}.$$

From (9) follows that $U_{ij} \geq u^{sb}_{ij}$ for all (i, j). This establishes that condition (31) holds for all (i, i', j, j'). Thus, also (28) holds for all (i, i', j, j'). In other words, there is no stake in misreporting any A type.

If there is any residual asymmetric information within the coalition (i.e., S does not observe A type), we need a last step for our proof to be complete. Recall that the side-contract has to be offered before A has revealed his type to S, and S thus takes into account that changing what she commits to announce in a certain state affects the information rent paid to the other types (i.e., it could affect U_{ij} for the other types). Given that (9) must hold, by misreporting any type, S cannot reduce the information rent for the other types below their outside options, i.e., u^{sb}_{ij}. It follows that the optimal manipulation of reports $\phi^*_{\tau_j}(\theta_i)$ is equal to (θ_i, τ_j), for all (i, j), the optimal transfer $b^*(i, j)$ is equal to zero, for all (i, j), and the optimal information rent is $U^*_{ij} = u^{sb}_{ij}$ for all (i, j). Notice that, under this manipulation, the information rents for all types are already at their minimum levels, i.e., $U_{ij} = u^{sb}_{ij}$ for all (i, j). Thus, the generic contract Γ_n is collusion-proof: A and S truthfully report (i, j) when S and A observe the signal j and A reports that he has type i. In the first stage of the game A selects one contract Γ. By assumption, at this stage A and S cannot collude. Therefore, A optimally selects the contract Γ_j correspondent to the signal τ_j that he has observed. A fast inspection reveals that A would not benefit from choosing any other contract $\Gamma_{j'}$. When A observes τ_j and select Γ_j we have that (i) S receives a wage equal to zero regardless of the realization of A type, (ii) A produces the *conditionally-optimal* second best output level q^{sb}_{ij} and receives transfer t^{sb}_{ij}. This proves that PC is collusion-proof and implements the *conditionally-optimal* second best.

[32]See Section 2.1. If $p(\theta_i|\tau_j) = 0$ the optimal outputs q^{sb}_{ij} are equal to zero only if they affect condition (4). Otherwise, they are equal to $q^{sb}_{ij} = q^{sb}_{ij-1}$.

6.4.2 Part 2

In this part of the proof we consider the signal-based structure when $N = 2$. By construction, Γ_1 is collusion-proof: S can send only one message. Thus, we only need to prove that Γ_2 is collusion-proof. To this purpose, we need to show that $\phi_{\tau_j}(\theta_i) = (i, j)$ solves (11) and the optimal transfer $b^*(i, j)$ is equal to zero for all (i, j).

Consider the manipulations where S, having observed signal j and learned that A has type $i = 1$ through side-contracting, prefers to tell the truth rather than to report that the state of nature is (i', j'). First, consider the case where $i' = N = 2$: (15) implies that $s^*_{Nj'} < 0$. Recall that (4) implies $u^{sb}_{in} > u^{sb}_{Nj'} + (\theta_N - \theta_i)q^{sb}_{Nj'}$ and from (9) follows that $U_{ij} \geq u^{sb}_{ij}$. Thus, the following condition must hold:

$$U_S\left(u^{sb}_{in} - u^{sb}_{ij}\right) > U_S\left(s^*_{Nj'} + u^{sb}_{Nj'} + (\theta_N - \theta_i)q^{sb}_{Nj'} - U_{ij}\right).$$

Using (30) and (29), this expression can be rewritten as $U_S\left(s^*_{ij}\right) \geq U_S\left(s^*_{i'j'} + t^{sb}_{i'j'} - \theta_i q^{sb}_{i'j'} - U_{ij}\right)$ for $i' = 2$, $i = 1$ and all (j, j').

Consider now the case where $i' = 1$. From (5) follows that $q^{sb}_{1n} = q^{sb}_{1j'}$ for all j'. Thus,

$$u^{sb}_{1n} \geq u^{sb}_{1n} + (\theta_1 - \theta_1)q^{sb}_{1n} = u^{sb}_{1n} + (\theta_1 - \theta_1)q^{sb}_{1j'} \quad \text{for all } j'.$$

Given that $N = 2$, this is sufficient to prove that $u^{sb}_{in} \geq u^{sb}_{i'n} + (\theta_{i'} - \theta_i)q^{sb}_{i'j'}$ for $i = i' = 1$ and for all j'. Moreover, from (9) follows $U_{ij} \geq u^{sb}_{ij}$. Thus, $U_S\left(u^{sb}_{in} - u^{sb}_{ij}\right) \geq U_S\left(u^{sb}_{i'n} + (\theta_{i'} - \theta_i)q^{sb}_{i'j'} - U_{ij}\right)$ for $i = i' = 1$ and all (j, j'). Using (30) and (29) twice, this can be rewritten as $U_S\left(s^*_{ij}\right) \geq U_S\left(s^*_{i'j'} + t^{sb}_{i'j'} - \theta_i q^{sb}_{i'j'} - U_{ij}\right)$ for $i = i' = 1$ and all (j, j').

So far, we proved that $U_S\left(s^*_{ij}\right) \geq U_S\left(s^*_{i'j'} + t^{sb}_{i'j'} - \theta_i q^{sb}_{i'j'} - U_{ij}\right)$ for $i = 1$ and for all (i', j, j'). In other words, the misreport of type 1 is never *directly* beneficial to S. Could it be *indirectly* beneficial? Note that the manipulation of type 1 has to be offered before A has revealed his type to S, and S thus takes into account that changing what she commits to announce in state $(1, j)$ also affects the information rent paid to type 2, i.e., it could affect U_{2j}. If the information rent paid to type 2 decreases as a result of the manipulation of type 1, then the latter could be beneficial, at least indirectly. Whenever type 2 is not misreported, i.e., $\phi_{\tau_j}(\theta_2) = (\theta_2, \tau_j)$, then the optimal information rent for type 2 is $U^*_{2j} = u^{sb}_{2j}$. Note that (9) entails $U_{2j} \geq u^{sb}_{2j}$. Therefore, by misreporting type 1, S cannot further reduce the information rent for type 2 (if type 2 is not misreported, i.e., $\phi_{\tau_j}(\theta_2) =$

$(\theta_2, \tau_j))$. In what follows, we show that, indeed, it is always optimal for S not to misreport type 2.

To prove this, we need to consider the manipulations where S, having observed signal j and learned that A has type $i = N = 2$ through side-contracting, reports that the state of nature is (i', j'). In what follows, we consider all the possible manipulations and show that none of them is beneficial for S. Before considering these manipulations, note that (15) entails $p(\theta_1, \tau_2)U_S\left(s^*_{11}\right) + p(\theta_2, \tau_2)U_S\left(s^*_{21}\right) = 0$, and from (16) follows $s^*_{11} = u^{sb}_{12} - u^{sb}_{11} = (\theta_2 - \theta_1)\left(q^{sb}_{22} - q^{sb}_{21}\right)$, and $s^*_{12} = s^*_{22} = 0$.

1) Consider the manipulation where S, having observed signal $j = 1$ and learned that A has type $i = 2$ through side-contracting, reports that the state of nature is $(i', j) = (1, 1)$. Given that u^{sb}_{ij} satisfies (4), it is costly for S to induce type 2 to misreport his type as type 1. Recall that $u^{sb}_{2j} = 0$. Thus, in order for type 2 to accept the side contract, S has to leave him with at least the same level of utility (see (9)). It follows that S has to provide type 2 with the following positive transfer to misreport his type as 1:

$$b = (\theta_2 - \theta_1)q^{sb}_{11} - u^{sb}_{11} = (\theta_2 - \theta_1)\left(q^{sb}_{11} - q^{sb}_{21}\right) > 0.$$

This side-contract has to be offered before A has revealed his type to S, and S thus takes into account that changing what she commits to announce in state $(2, 1)$ also affects the information rent paid to type 1. From (10) follows that offering such a side-contract has a cost for S since type 1 rent increases from $(\theta_2 - \theta_1)q^{sb}_{21}$ to $(\theta_2 - \theta_1)q^{sb}_{11}$ when $j = 1$ realizes. The *indirect* expected cost of this manipulation borne by S is then $p(\theta_1, \tau_1)U_S((\theta_2 - \theta_1)\left(q^{sb}_{21} - q^{sb}_{11}\right)) = p(\theta_1, \tau_1)U_S(-b)$. Therefore, S utility from this manipulation is

$$p(\theta_1, \tau_1)U_S(-b) + p(\theta_2, \tau_1)U_S(s^*_{11} - b).$$

Given that $s^*_{11} - b = (\theta_2 - \theta_1)\left(q^{sb}_{22} - q^{sb}_{11}\right) < 0$, S is not willing to undertake this manipulation.[33] Note that this result holds irrespective of what the side-contract instructs type 1 to report, i.e., for all $\phi_{\tau_1}(\theta_1) = (\theta_i, \tau_j)$ with $i = \{1, 2\}$ and $j = \{1, 2\}$.

2) Consider the manipulation where S, having observed signal $j = 1$ and learned that A has type $i = 2$ through side-contracting, reports that the state of nature is $(i', j') = (1, 2)$. As before, S has to

[33] Note that this trivially holds for $N = 2$ but not for $N > 2$.

provide type 2 with the following positive transfer to misreport his type as 1:

$$b = (\theta_2 - \theta_1)q_{12}^{sb} - u_{12}^{sb} = (\theta_2 - \theta_1)\left(q_{12}^{sb} - q_{22}^{sb}\right) > 0.$$

This side-contract has to be offered before A has revealed his type to S. As before, from (10) follows that offering such a side-contract has a cost for S since type 1 rent increases from $(\theta_2 - \theta_1)q_{21}^{sb}$ to $(\theta_2 - \theta_1)q_{12}^{sb}$ when $j = 1$ realizes. The *indirect* expected cost of this manipulation borne by S is then $p(\theta_1, \tau_1)U_S((\theta_2 - \theta_1)\left(q_{21}^{sb} - q_{12}^{sb}\right))$, where $(\theta_2 - \theta_1)\left(q_{21}^{sb} - q_{12}^{sb}\right) < 0$. Thus, S utility from this manipulation is

$$p(\theta_1, \tau_1)U_S((\theta_2 - \theta_1)\left(q_{21}^{sb} - q_{12}^{sb}\right)) + p(\theta_2, \tau_1)U_S(s_{12}^* - b).$$

Given that $s_{12}^* - b = -b < 0$, S is not willing to undertake this manipulation. Note that this result holds irrespective of what the side-contract instructs type 1 to report, i.e., for all $\phi_{\tau_1}(\theta_1) = (\theta_i, \tau_j)$ with $i = \{1, 2\}$ and $j = \{1, 2\}$.

3) Consider the manipulation where S, having observed signal $j = 2$ and learned that A has type $i = 2$ through side-contracting, reports that the state of nature is $(i', j) = (1, 2)$. S has to provide type 2 with the following positive transfer to misreport his type as 1:

$$b = (\theta_2 - \theta_1)q_{12}^{sb} - u_{12}^{sb} = (\theta_2 - \theta_1)\left(q_{12}^{sb} - q_{22}^{sb}\right) > 0.$$

This side-contract has to be offered before A has revealed his type to S. Again, from (10) follows that offering such a side-contract has a cost for S since type 1 rent increases from $(\theta_2 - \theta_1)q_{22}^{sb}$ to $(\theta_2 - \theta_1)q_{12}^{sb}$ when $j = 2$ realizes. The expected cost of this manipulation borne by S is then $p(\theta_1, \tau_2)U_S((\theta_2 - \theta_1)\left(q_{22}^{sb} - q_{12}^{sb}\right)) = p(\theta_1, \tau_2)U_S(-b)$. Therefore, S utility from this manipulation is

$$p(\theta_1, \tau_2)U_S(-b) + p(\theta_2, \tau_2)U_S(s_{12}^* - b).$$

Given that $s_{12}^* - b = -b < 0$, S is not willing to undertake this manipulation. Note that this result holds irrespective of what the side-contract instructs type 1 to report, i.e., for all $\phi_{\tau_2}(\theta_1) = (\theta_i, \tau_j)$ with $i = \{1, 2\}$ and $j = \{1, 2\}$.

4) Consider the manipulation where S, having observed signal $j = 2$ and learned that A has type 2 through side-contracting, reports that the state of nature is $(i', j') = (1, 1)$. S has to provide type 2 with the following positive transfer to misreport his type as 1:

$$b = (\theta_2 - \theta_1)q_{11}^{sb} - u_{11}^{sb} = (\theta_2 - \theta_1)\left(q_{11}^{sb} - q_{21}^{sb}\right).$$

This side-contract has to be offered before A has revealed his type to S. Again, from (10) follows that offering such a side-contract has a cost for S since type 1 rent increases $(\theta_2 - \theta_1)q_{22}^{sb}$ to $(\theta_2 - \theta_1)q_{11}^{sb}$ when $j = 2$ realizes. The expected cost of this manipulation borne by S is then $p(\theta_1, \tau_2)U_S((\theta_2 - \theta_1)\left(q_{22}^{sb} - q_{11}^{sb}\right))$, where $(\theta_2 - \theta_1)\left(q_{22}^{sb} - q_{11}^{sb}\right) < 0$. Therefore, S utility from this manipulation is

$$p(\theta_1, \tau_2)U_S((\theta_2 - \theta_1)\left(q_{22}^{sb} - q_{11}^{sb}\right)) + p(\theta_2, \tau_2)U_S(s_{11}^* - b).$$

Given that $s_{11}^* - b = (\theta_2 - \theta_1)\left(q_{22}^{sb} - q_{11}^{sb}\right) < 0$, S is not willing to undertake this manipulation. Note that this result holds irrespective of what the side-contract instructs type 1 to report, i.e., for all $\phi_{\tau_2}(\theta_1) = (\theta_i, \tau_j)$ with $i = \{1, 2\}$ and $j = \{1, 2\}$.

5) Consider the manipulation where S, having observed signal $j = 1$ and learned that A has type $i = 2$ through side-contracting, reports that the state of nature is $(i, j') = (2, 2)$. Such a manipulation increases S wage in this state of the world by $s_{22}^* - s_{21}^* = -s_{21}^* > 0$. In this case, S does not need to provide type 2 with any transfer: type 2 obtains zero utility both in state $(2, 2)$ and $(2, 1)$. Nonetheless, this side-contract has to be offered before A has revealed his type to S, and S thus takes into account that changing what she commits to announce in state $(2, 1)$ also affects the information rent paid to type 1. Offering such a side-contract has a cost for S since type 1 rent increases from $(\theta_2 - \theta_1)q_{21}^{sb}$ to $(\theta_2 - \theta_1)q_{22}^{sb}$ when $j = 1$ realizes. The expected cost of this manipulation borne by S is then $p(\theta_1, \tau_1)(\theta_2 - \theta_1)\left(q_{22}^{sb} - q_{21}^{sb}\right)$ which is stricitly larger than its possible expected benefit, i.e., $p(\theta_2, \tau_1)(-s_{21}^*)$. To see this, note that (15) entails $p(\theta_2, \tau_2)U_S\left(s_{21}^*\right) + p(\theta_1, \tau_2)U_S\left(s_{11}^*\right) = 0$. Rearrange to obtain,

$$U_S\left(s_{21}^*\right) = -\frac{p(\theta_1, \tau_2)}{p(\theta_2, \tau_2)}U_S\left(s_{11}^*\right).$$

If S were to be risk neutral, $-s_{21}^* = \frac{p(\theta_1,\tau_2)}{p(\theta_2,\tau_2)}s_{11}^* = \frac{p(\theta_1,\tau_2)}{p(\theta_2,\tau_2)}(\theta_2 - \theta_1)\left(q_{22}^{sb} - q_{21}^{sb}\right)$. In this case, the expected cost of the manipulation borne by S is $p(\theta_1, \tau_1)(\theta_2 - \theta_1)\left(q_{22}^{sb} - q_{21}^{sb}\right)$, which larger than its possible expected benefit $p(\theta_2,\tau_1)(-s^*_{21}) = p(\theta_2,\tau_1)\frac{p(\theta_1,\tau_2)}{p(\theta_2,\tau_2)}(\theta_2-\theta_1)\left(q^{sb}_{22}-q^{sb}_{21}\right)$.[34] Clearly, with a risk averse S, $-s^*_{21}$ would be even smaller and the expected benefit further reduced. This proves that this manipulation is not selected by S. Note that this result holds irrespective of what the side-contract instructs type 1 to report, i.e., for all $\phi_{\tau_1}(\theta_1) = (\theta_i,\tau_j)$ with $i=\{1,2\}$ and $j=\{1,2\}$.

6) Finally, consider the manipulation where S, having observed signal $j=2$ and learned that A has type $i=2$ through side-contracting, reports that the state of nature is $(i,j') = (2,1)$. Such a manipulation would decrease S wage in this state of the world by $s^*_{21}-s^*_{22} = s^*_{21} < 0$. This side-contract has to be offered before A has revealed his type to S, and S thus takes into account that changing what she commits to announce in state $(2,2)$ also affects the information rent paid to type 1. Offering such a side-contract has a potential benefit for S since type 1 information rent could decrease from $(\theta_2-\theta_1)q^{sb}_{22}$ to $(\theta_2-\theta_1)q^{sb}_{21}$ when $j=2$ realizes. Nonetheless, (9) implies that $U_{12} \geq (\theta_2-\theta_1)q^{sb}_{22}$. Thus the above manipulation cannot reduce type 1 information rent. It follows that S is not willing to pursue this manipulation. Note that this result holds irrespective of what the side-contract instructs type 1 to report, i.e., for all $\phi_{\tau_2}(\theta_1) = (\theta_i,\tau_j)$ with $i=\{1,2\}$ and $j=\{1,2\}$.

This proves that Γ_2 is collusion-proof: A and S truthfully report (i,j) when S and A observe the signal j and A has type i. In the first stage of the game, A selects one contract. By assumption, at this stage A and S cannot collude. Therefore, A optimally selects the contract Γ_j correspondent to the signal τ_j that he has observed. A fast inspection reveals that A would not benefit from choosing any other contract $\Gamma_{j'}$. When A observes τ_j and selects Γ_j we have that (i) S receives a wage equal to zero regardless of the realization of A cost type, and (ii) A produces the conditionally-optimal second-best output level q^{sb}_{ij} and receives transfer t^{sb}_{ij}. This implements the conditionally-optimal second best-outcome and proves that PC is collusion-proof.

6.5 Proof of Proposition 4

In this proof we derive some necessary conditions for a generic SSM to implement the second-best outcome. Then, we show that there is a distribution of the cost θ and the signal τ that violate these conditions. The generic SSM is defined in Section 3. The timing is (i) At date -1, A learns θ and

[34]Note that from the monotone likelihood ratio property follows $\frac{p(\theta_1|\tau_1)}{p(\theta_2|\tau_1)} \geq \frac{p(\theta_1|\tau_2)}{p(\theta_2|\tau_2)}$ and so,

$$p(\theta_1,\tau_1)(\theta_2-\theta_1)\left(q^{sb}_{22}-q^{sb}_{21}\right) \geq p(\theta_2,\tau_1)\frac{p(\theta_1,\tau_2)}{p(\theta_2,\tau_2)}(\theta_2-\theta_1)\left(q^{sb}_{22}-q^{sb}_{21}\right)$$

τ, (ii) At date 0, P offers a menu of contracts to A, who can either select one contract or refuse all contracts. If A refuses the game ends, (iii) At date 1, S decides whether to accept or refuse the contract selected by A at date 0. If S refuses the game ends, (iv) At date 2, S learns τ, (v) At date 3, S and A can stipulate a side-contract. If they do not stipulate a side-contract, the mechanism is played non-cooperatively by A and S, (vi) At date 4, production and transfers take place.

To begin, consider the case where S can send only one signal, i.e., $\rho = \{\tau_1\}$ in each contract. Having this schedule in place, if S decides to participate in P mechanism, her salary depends exclusively on the contract selected by A and the message sent by A to P. A necessary condition for the SSM to implement the second-best outcome is the following: when τ_j is realized, an A who has type i must (i) select a contract Γ where

$$\sum_{i=1}^{n} p(\theta_i, \tau_j) s_i = 0,$$

and (ii) select an output schedule equal to q_{ij}^{sb}, with an associated information rent equal to u_{ij}^{sb}. Point (i) ensures that P extracts S information without forgoing any rent to her. Given that S is risk averse, the only way to achieve this is to offer a salary schedule $s_i^* = 0$, for all i and for all contracts Γ selected by A along the equilibrium path[35]; any other wage schedule would require P to forgo a strictly positive risk premium because S would have to be compensated for bearing the cost of uncertainty. Note that S learns τ only at date 2. Therefore, the acceptance of the contract by S takes place before learning τ. Hence, S interim participation constraints must be satisfied.[36] Given that $s_i^* = 0$ for all i, S interim participation constraints are trivially satisfied along the equilibrium path. Having this schedule in place, it is easy to see that this SSM is equivalent to no-supervision. A would simply select the contract that yields the larger information rent. Denote by $\Gamma_{N(i)}$ the contract where an A who has type i obtains an information rent equal to u_{iN}^{sb}. This information rent is larger than the information rent available in any other contract because from (7) follows that $u_{iN}^{sb} \geq u_{ij}^{sb}$ for all j. Therefore, offering a menu of contracts is not beneficial to P. Thus, the problem reduces to the one described in the Direct Supervision Section, provided that P does not observe τ. This proves that the SSM menu of contracts must feature at least one contract Γ in which S can send at least two messages. Moreover, it is also necessary to prevent an A who has type i from selecting $\Gamma_{N(i)}$ in all states of the world. To this purpose, an A who has type i and has observed $\tau_j \neq \tau_N$ must not obtain u_{iN}^{sb} by selecting $\Gamma_{N(i)}$.

[35] Note that s_i (S salary) depends exclusively on A message

[36] Instead, since A is informed on his own type and on the S signal at the time of accepting the side-contract, A ex-post participation constraints must be satisfied.

Take a SSM that strictly satisfies these two conditions. In such a SSM there is a contract, $\Gamma_{N(i)}$, that allows S to send two signals, $\rho_{N(i)} = \{\tau_N, \tau_{N'}\}$, with $N > N'$. In what follows, let us focus on this contract $\Gamma_{N(i)}$. We further restrict attention to the case where the realized signal τ observed by A at date 0 is either τ_N or $\tau_{N'}$. By construction, contract $\Gamma_{N(i)}$ must ensure that S truthfully report her signal when the game is played non-cooperatively; this is a necessary (but not sufficient) condition to prevent an A who has type i and has observed $\tau_{N'}$ from obtaining u^{sb}_{iN} by (i) selecting $\Gamma_{N(i)}$ and (ii) rejecting S side-contract. It follows that the wage schedule in contract $\Gamma_{N(i)}$ must satisfy:

$$\begin{aligned} \sum_{i=1}^{n} p(\theta_i, \tau_N) U_S\left(s^*_{iN}\right) = 0 &\geq \sum_{i=1}^{n} p(\theta_i, \tau_N) U_S\left(s_{iN'}\right), \\ \sum_{i=1}^{n} p(\theta_i, \tau_{N'}) U_S\left(s_{iN'}\right) \geq 0 &= \sum_{i=1}^{n} p(\theta_i, \tau_{N'}) U_S\left(s^*_{iN}\right). \end{aligned} \tag{32}$$

First, let us analyze the trivial case where the contract $\Gamma_{N(i)}$ specifies a wage schedule $s_{iN'} = 0$ for all i. Note that an A who has type i, having observed signal $\tau_{N'}$, must obtain $u^{sb}_{iN'}$ along the equilibrium path (as we mentioned above in point (ii), this is a necessary condition for the SSM to implement the second-best outcome.) On the contrary, by selecting $\Gamma_{N(i)}$ he would obtain (say) $u_{iN'}$ (if the game is played non-cooperatively and (32) holds). Clearly, if $u_{iN'} > u^{sb}_{iN'}$, A prefers to select $\Gamma_{N(i)}$ when the realized signal is $\tau_{N'}$. To avoid this, the following condition must hold:

$$u^{sb}_{iN'} \geq u_{iN'} \tag{33}$$

From (7) follows that $u^{sb}_{iN} \geq u^{sb}_{iN'}$ for all (i, N') with $N > N'$. This coupled with (33) entails

$$u^{sb}_{iN} \geq u^{sb}_{iN'} \geq u_{iN'}, \tag{34}$$

for all (i, N') with $N > N'$. In particular, for $i \neq N$, the last inequalities becomes $u^{sb}_{iN} > u^{sb}_{iN'} \geq u_{iN'}$. Having this schedule in place, there is a simple misreport that is beneficial to S: after having observed $\tau_{N'}$ and learned that A has type i through side-contracting, S prefers to report that the state of nature is (i, N) rather than tell the truth, i.e. (i, N'). To prove this, it is sufficient to check that S can make the following take-it-or-leave-it offer to A. First, S and A report that the state of nature is (i, N), obtaining a collective utility from the misreport equal to u^{sb}_{iN}. Second, the take-it-or-leave-it offer entitles A to keep $u_{iN'}$ (the same utility he would get by playing the game non-cooperatively), while S pockets the difference, i.e., $u^{sb}_{iN} - u_{iN'} > 0$. This misreport leaves A indifferent and S strictly better off (i.e., she obtains $u^{sb}_{iN} - u_{iN'}$ when the state (i, N') is realized instead of $s_{iN'} = 0$ in all states of the world.)

This proves that SSM cannot implement the second best outcome if $s_{iN'} = 0$, for all i. It follows that we must have $s_{iN'} \neq 0$ for at least one value of i. Note that (32) entails $\sum_{i=1}^{n} p(\theta_i, \tau_N) U_S(s_{iN'}) \leq 0$. It follows that there is at least one value of i (say, $i = \widetilde{i}$) such that $s_{\widetilde{i}N'}$ is strictly negative. From (32) also follows that $\sum_{i=1}^{n} p(\theta_i, \tau_{N'}) U_S(s_{iN'}) \geq 0$. Therefore, there is at least one value of i (say, $i = \widehat{i}$) such that $s_{\widehat{i}j'}$ is strictly positive.

To proceed with the rest of the proof, we need to show the conditions such that S, having observed signal N' and learned that A has type $i = \widetilde{i}$ through side-contracting, prefers to tell the truth rather than to report that the state of nature is $(\widetilde{i}, N)$. The benefit of this misreport lies in the fact that S obtains $s^*_{\widetilde{i}N} = 0$, instead of $s_{\widetilde{i}N'} < 0$. Thus, the benefit is $B = -s_{\widetilde{i}N'}$. But this side-contract has to be offered before A has revealed his type to S, and S thus takes into account that changing what she commits to announce in state $(\widetilde{i}, N)$ also affects the information rent paid to other A types. If $q^{sb}_{\widetilde{i}N} > q_{\widetilde{i}N'}$, from (10) and (9) follow that S might have to provide types $i < \widetilde{i}$ with an extra information rent equal to $(\theta_{\widetilde{i}} - \theta_{\widetilde{i}-1})(q_{\widetilde{i}N} - q_{\widetilde{i}N'})$. One aspect is worth noting: the cost of misreporting type $i = 1$ is equal to zero for S. This is due to the fact that the side-contract involving type $i = 1$ does not affect the information rent paid to the other A types. Thus, $s_{1N'}$ cannot be set to be negative. Then, $(\theta_{\widetilde{i}} - \theta_{\widetilde{i}-1})(q_{\widetilde{i}N} - q_{\widetilde{i}N'})$ is the potential cost of S misreport.[37]

The rest of the proof is straightforward: let us denote as K the difference between $\theta_{\widetilde{i}}$ and $\theta_{\widetilde{i}-1}$ such that $K = \theta_{\widetilde{i}} - \theta_{\widetilde{i}-1}$. For all i, keep $p(\theta_i, \tau_N)$ and $p(\theta_i, \tau_{N'})$ constant and reduce K. From (5) follows that $\lim_{K \to 0} q^{sb}_{\widetilde{i}N} = q^{sb}_{\widetilde{i}-1N}$. Thus, $\lim_{K \to 0} (\theta_{\widetilde{i}} - \theta_{\widetilde{i}-1}) \left(q^{sb}_{\widetilde{i}N} - q_{\widetilde{i}N'} \right) = 0$. This entails that the total cost of S misreport decreases with K down to a lower-bound equal to zero. On the other hand,

[37] Consider the case where S reports $\tau_{N'}$. It follows from the standard treatment of this problem that an output profile $q_{iN'}$ for $i = \{1, 2, ..., N\}$ is implementable through a contract if and only if it is weakly decreasing, i.e., (2) holds. Then, at the side-contracting stage, the agent's lowest utility levels that are compatible with this implementation are revealed by the upward adjacent incentive compatibility constraints for all types but the most inefficient one, i.e, (10) becomes $U_{iN'} = U_{i+1N'} + (\theta_{i+1} - \theta_i) q_{i+1N'}$ for all $i \neq N$. It is easy to show that the remaining incentive compatibility constraints are strictly satisfied. This proves that, if $q^{sb}_{\widetilde{i}N} > q_{\widetilde{i}N'}$, S *might* have to provide types $i < \widetilde{i}$ with an extra information rent equal to $(\theta_{\widetilde{i}} - \theta_{\widetilde{i}-1})(q_{\widetilde{i}N} - q_{\widetilde{i}N'})$. If $q^{sb}_{\widetilde{i}N} < q_{\widetilde{i}N'}$, the misreport actually reduces the information rent by $(\theta_{\widetilde{i}} - \theta_{\widetilde{i}-1})(q_{\widetilde{i}N} - q_{\widetilde{i}N'})$. In the case where $q^{sb}_{\widetilde{i}N}$ is so small as to violate (2), i.e., $q^{sb}_{\widetilde{i}N} < q^{sb}_{\widetilde{i}+1N'}$, S cannot prevent types $i > \widetilde{i}$ from misreporting as type $\widetilde{i}$ at the side-contracting stage. Thus, S gets a payoff equal to zero when types $i = \left\{\widetilde{i}, ..., N\right\}$ are realized. Given that $s_{1N'}$ must be larger or equal to zero, this misreport could still be beneficial for any $\widetilde{i} = \{2, ..., N\}$.

given that $p(\theta_i, \tau_N)$ and $p(\theta_i, \tau_{N'})$ are kept constant, there is at least a wage $s_{\widetilde{i}N'}$ that still needs to be strictly negative. Thus the benefit of the misreport must also be strictly positive, i.e., $B = -s_{\widetilde{i}N'} > 0$. Therefore, using a simple continuity argument, it is possible to prove there exists a value of K small enough such that the benefit of the misreport is larger than the total cost.

In this case, an A who has type $\widetilde{i}$ and has observed $\tau_{N'} \neq \tau_N$ can receive $u^{sb}_{\widetilde{i}N}$ and produce $q^{sb}_{\widetilde{i}N}$ by (i) selecting $\Gamma_{N(\widetilde{i})}$ and (ii) accepting S side-contract. Recall that one of the necessary conditions for the SSM to implement the second-best was the following: when $\tau_{N'}$ is realized an A who has type $\widetilde{i}$ must select an output schedule equal to $q^{sb}_{\widetilde{i}N'}$, with an associated information rent equal to $u^{sb}_{\widetilde{i}N'}$. We have just proved that there is a value of the productive costs $\theta_{\widetilde{i}}$ and $\theta_{\widetilde{i}-1}$ such that A produces $q^{sb}_{\widetilde{i}N}$ instead. By properly manipulating $p(\theta_i, \tau_N)$ and $p(\theta_i, \tau_{N'})$ for all i, the difference $q^{sb}_{\widetilde{i}N} - q^{sb}_{\widetilde{i}N'}$ can be made substantially large even if $K \to 0$ and $\theta_{\widetilde{i}} = \theta_{\widetilde{i}-1}$. If $q^{sb}_{\widetilde{i}N} - q^{sb}_{\widetilde{i}N'}$ is large, so is the associated inefficiency due to the misreport. Thus, the SSM does not achieve the second best outcome. This proves that there is a distribution of costs and signals that violates one of the necessary conditions for SSM to implement the second-best outcome. The proof does go through only if $N > 2$.[38] QED.

References

[1] Abdulkadiroglu, A. and Chung, K. (2003). "Auction design with tacit collusion," Columbia University working paper.

[2] Baliga, S. and Sjostrom, T. (1998). "Decentralization and Collusion," Journal of Economic Theory, 83: 196-232.

[3] Burlando, A. and Motta, A. (2008a). "Can Self Reporting Reduce Corruption in Law Enforcement?," University of Padua working paper.

[4] Burlando, A. and Motta, A. (2008b). "Conditional Delegation and Optimal Supervision," Boston University working paper.

[5] Celik, G. (2009). "Mechanism Design with Collusive Supervision," Journal of Economic Theory, 144 (1): 69-95.

[6] Celik, G., Peters, M. (2010). "Equilibrium Rejection of a Mechanism," UBC Department of Economics working paper.

[7] Che, Y. and Kim, J. (2006). "Robustly collusion-proof implementation," Econometrica, 74: 1063-1107.

[8] Che, Y. and Kim, J. (2009). "Optimal collusion-proof auctions," Journal of Economic Theory, 144(2): 565-603.

[9] Dequiedt, V. (2006). "Ratification and Veto Constraints in Mechanism Design," working paper.

[10] Dequiedt, V. (2007). "Efficient collusion in optimal auctions," Journal of Economic Theory, vol. 136(1): 302-323.

[11] Faure-Grimaud A., Laffont J. J. and Martimort, D. (2003). "Collusion, Delegation and Supervision with Soft Information," Review of Economic Studies, 70: 253-275.

[38] Consider the case where $N = 2$. If $K \to 0$ the problem reduces to a mechanism with a single A, i.e., $\theta_2 = \theta_1$. Then, there is effectively no need for supervision. Moreover, (32) does not need to hold (i.e., there is no need for $s_{\widetilde{i}N'}$ to be negative.) Our proof does not go through if $N = 2$. On the other hand, when $N > 2$ and $K \to 0$ (i.e, $\theta_{\widetilde{i}} \to \theta_{\widetilde{i}-1}$) the problem does not reduce to a single A type mechanism. Then, there is still need for at least one $s_{\widetilde{i}N'}$ to be strictly negative, (32) must hold and our proof holds.

[12] Faure-Grimaud A. and Martimort, D. (2001). "On some agency costs of intermediated contracting", Economics Letters, 71: 75-82.

[13] Innes, R. (1999). "Remediation and self-reporting in optimal law enforcement," Journal of Public Economics, Elsevier, 72(3): 379-393.

[14] Innes, R. (2001). "Violator Avoidance Activities and Self-Reporting in Optimal Law Enforcement," Journal of Law, Economics, and Organization 17(1): 239-56.

[15] Kaplow, L., and Shavell, S. (1994). "Optimal Law Enforcement with Self Reporting of Behavior," Journal of Political Economy, 102(3): 583-606.

[16] Khalil, F. and Lawarree, J. (2006). "Incentives for corruptible auditors in the absence of commitment," Journal of Industrial Economics, 15(2): 457-478.

[17] Laffont, J. J. and Martimort D. (1997). "Collusion under asymmetric information," Econometrica, 65: 875-911.

[18] Laffont, J. J. and Martimort, D. (1998). "Collusion and delegation," Rand Journal of Economics, 29: 280-305.

[19] Laffont, J. J. and Martimort, D. (2000). "Mechanism design with collusion and correlation," Econometrica, 68: 309-342.

[20] Laffont, J. J. and Tirole, J. (1991). "The Politics of Government Decision-Making: A Theory of Regulatory Capture," The Quarterly Journal of Economics, 106(4): 1089-1127.

[21] Laffont, J. J. and Tirole, J. (1993). "A Theory of Incentives in Procurement and Delegation," Cambridge, MA: MIT Press.

[22] Lambert-Mogilianksy, A. (1998). "On Optimality of Illegal Collusion in Contracts," Review of Economic Design, 3: 303-348.

[23] Martimort, D. (1999). "The life cycle of regulatory agencies: Dynamic capture and transaction costs," The Review of Economic Studies, 66: 929-947.

[24] Melumad, N., and Reichelstein, S. (1984). "Value of Communication in Agencies," The Journal of Economic Theory, 47: 334-368.

[25] Mookherjee, D. (2006). "Decentralization, hierarchies, and incentives: A mechanism design perspective," Journal of Economic Literature, 44(2): 367-390.

[26] Mookherjee, D. and Reichelstein, S. (1990). "Implementation via Augmented Revelation Mechanisms," Review of Economic Studies 57: 453-475.

[27] Mookherjee, D. and Tsumagari, M. (2004). "The Organization of Supplier Networks: Effects of Delegation and Intermediation," Econometrica, 72(4): 1179-1219.

[28] Pavan, A. and Calzolari, G., 2009. "Sequential contracting with multiple principals," Journal of Economic Theory, Elsevier, vol. 144(2), pages 503-531, March.

[29] Pavlov, G. (2008). "Auction Design in the Presence of Collusion," Theoretical Economics 3: 383-429.

[30] Quesada, L. (2004). "Collusion as an informed principal problem," University of Wisconsin working paper.

[31] Tirole, J. (1986). "Hierarchies and Bureaucracies," Journal of Law, Economics, and Organization, 2: 181-214.