Collusion and Selective Supervision*

Alberto Motta†

First Version: January 2007

This Version: November 2010

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 specifying 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, Vianney Dequiedt, Gregory Pavlov, Bob Gibbons, Isabella Dobrescu, Alfredo Burlando, Chris Bidner, Barton Lipman, Michael Manove and Stefan Napel for helpful discussions. This paper has benefited with comments from seminar participants at University Toulouse 1, Boston University, University of New South Wales, University of Sydney, University of Melbourne, 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 actually 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 that 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. 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. Departing from this restriction is useful because of one important assumption: the agent and the supervisor cannot collude on participation decisions (hereafter referred to as *no-collusion in participation decision*).[5] 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. 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. 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, which are

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

[5]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. We are thankful to David Martimort for his comment. 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.

related to the structure of the supervisor's information and 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,[6] 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, the present contribution seems to 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. It follows that allowing for collusion on participation decisions may be a more interesting way of thinking about the problem in this particular framework.

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

[6] 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 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.

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 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.[7] This paper tackles the issue from a different angle, suggesting a new and 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.[8] 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 though their setting is different from the one considered here, there is one aspect which is common to both contributions: participation decisions convey information about the types of the players. Dequiedt

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

[8] 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 our 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.

(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,[9] but none of them have addressed this question yet. Among them, Mookherjee and Tsumagari (2004) analyze this problem in a supervision setup.[10] 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 of generality, A is assumed to be risk neutral.[11] A utility function is given by

$$t - \theta q,$$

[9] 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.

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

[11] 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.

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. 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 $\theta_i' \geq \theta_i$. One aspect is worth noting: FLM (2003) assume that $p(\theta_i, \tau_j) > 0$. 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 the sake of 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 redundant, i.e., they convey the same information.[12] 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 $p(\theta_i, \tau_j) > 0$ 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.[13]

S salary is s, which represents her monetary transfer from P. Her utility function is given by $U_S(.)$, with $U_S'(.) > 0$ and $U_S''(.) \leq 0$. It follows that S could be either risk neutral or risk averse. P payoff for a given output q, transfer level t and wage s is

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

[12] 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.

[13] If $N = 3$, and the conditional probabilities $p(\theta_i|\tau_j)$ are

	τ_1	τ_2	τ_3
θ_1	B_1	B_1	0
θ_2	B_2	B_2	0
θ_3	0	0	1

then the information structure for S reduces to Celik's (2009) connected partition case $\{\{\theta_1, \theta_2\}, \theta_3\}$.

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.

2.1 Direct Supervision

To begin with, 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 the sake of 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. Let us also denote by $u_{ij} = t_{ij} - \theta_i q_{ij}$ A information rent in state (i, j). 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 participation constraints,

$$u_{ij} \geq 0 \quad \text{for all } i, \tag{3}$$

and the binding upward adjacent incentive compatibility constraints,[14]

$$u_{ij} \geq u_{i'j} + (\theta_{i'} - \theta_i)q_{i'j} \quad \text{for all } (i', i) \text{ such that } i' > i. \tag{4}$$

[14] At the optimum of P problem the following constraints are binding. It is easy to show that the remaining 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[15] 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.[16] For the sake of simplicity, we will assume that in this particular case $q_{ij}^{sb} = q_{ij-1}^{sb}$. Having this schedule in place, it is possible to show that:

Proposition 1 *The output for any type i is at least as distorted after the observation of j than after the observation of j', where $j' \geq j$:*

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

[15]Note that the optimal output for the most efficient type when $p(\theta_1|\tau_i) = 0$ could be set to zero beacuse type $\theta_1$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 $p(\theta_1|\tau_i) = 0$.

[16]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.

This proposition 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. Proposition 1 entails that the information rents for all types are larger after the observation of τ'_j than after τ_j:

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

It follows that A strictly prefers higher signals to be realized.

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 implementable. 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: FML's (2003) 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 reports type θ_i and P knows τ_j are,[17]

Type	P knows τ_1		P knows τ_2		P knows τ_3	
	Output	Agent U	Output	Agent U	Output	Agent U
$\theta_1 = 0.25$	4	0.4	4	0.7	4	0.8
$\theta_2 = 0.5$	0.8	0.2	1.6	0.3	1.9	0.4
$\theta_3 = 1$	0.4	0	0.6	0	0.8	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: Celik's (2009) 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 S information 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 redundant. When P receives the information directly (*direct supervision*) the second-best outcome is implementable. The optimal quantities (refereed to as "Output" in the table) and utility levels of A (referred to as U in the table) when A reports type θ_i and P knows τ_j are,

Type	P knows τ_1		P knows τ_2		P knows τ_3	
	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

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

As before, the results are rounded to the second decimal place. 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 the extra information rent with A.

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.

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 this literature on collusion, this side-contract is assumed to be enforceable.[18] If A or S refuse the side-contract, the game is played non-cooperatively. 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

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

non-cooperatively play the truthful equilibrium of an individually incentive compatible mechanism. By definition, $u_{ij} = t_{ij} - \theta_i q_{ij}$. 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. Acceptance of the side-contract by all A types imposes the following *ex post* participation constraints:

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

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

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

The optimal side-contract solves the following problem:

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

subject to (10) and (9).

When the informative signal allows S to exactly observe A type, (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 for all (i, j). Notice that any $\phi^*_{\tau_j}(\theta_i)$ that solves (12) must also solve (11) but not vice versa. Therefore, collusion is more difficult to prevent when there is no residual asymmetric information between A and S.

3 Selective Supervision Mechanism

3.1 Timing

The precise implementation of our mechanism depends crucially on the timing of S information. Let us denote by *Timing 1* the framework in which S receives her information before P has the opportunity to offer her the mechanism. This is the setting adopted by Celik (2009) and FLM (2003):

- At date -1, S learns τ and A learns θ and τ.
- At date 0, P offers a mechanism to S and A.
- At date 1, A and S decide whether to accept or refuse the mechanism.
- 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.
- At date 3, production and transfers take place.

This framework (*Timing 1*) rests on the assumption that S receives her informative signal in the first stage of the game. This assumption is unsatisfactory because it outlines a type of supervision that fits a limited range of cases. 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* the framework in which S receives the signal τ at date 1.5. The next section proposes a numerical example to highlight the basic features of a Selective Supervision Mechanism.

3.2 Numerical Example (continued)

3.2.1 Example 1: FML's (2003) Information Structure

This numerical example is identical to the one proposed in the previous Section, except that S and A are now allowed to collude. In order to prevent collusion, P can design a Selective Supervision Mechanism. By doing so, the second-best outcome can be implemented even when A and S collude. The Selective Supervision Mechanism is designed in the following fashion: P offers a menu of contracts (or organizational structures) to A. 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). 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 reports type θ_i and S reports τ_j. As before, the results are rounded to the second decimal place (where possible). The original results are presented in Appendix 1.

Contract 1	No-supervision		
Type	Output	U	s
$\theta_1 = 0.25$	4	0.4	0
$\theta_2 = 0.5$	0.8	0.2	0
$\theta_3 = 1$	0.4	0	0

Contract 2	S reports τ_1			S reports τ_2		
Type	Output	U	s	Output	U	s
$\theta_1 = 0.25$	4	0.4	0.3	4	0.7	0
$\theta_2 = 0.5$	0.8	0.2	0.1	1.6	0.3	0
$\theta_3 = 1$	0.4	0	-0.2	0.6	0	0

Contract 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	0.4	0.4	4	0.7	0.1	4	0.8	0
$\theta_2 = 0.5$	0.8	0.2	0.2	1.6	0.3	0.1	1.9	0.4	0
$\theta_3 = 1$	0.4	0	-0.2	0.6	0	-0.03	0.8	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 (-0.06). If S has observed the signal τ_2, she is better off reporting the true signal τ_2 (0.01) 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). The same applies to Contract 2 and (trivially) to Contract 1. Moreover, 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 0.8 when the type is θ_1 and 0.4 when the type is θ_2. The potential for collusion arises only when A has type θ_3. In this case, S could receive a negative payoff. Consider a simple example where 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 (-0.03). 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 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 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.

This kind of Selective Supervision Mechanism achieves the second-best outcome by endogenizing the scope of supervision. When there is high probability that the type is 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 likely to be 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 no supervision is involved.

3.2.2 Selective Supervision Mechanism: Celik's (2009) Information Structure

The Selective Supervision Mechanism that achieves the second-best outcome in Example 2 is,[19]

Contract 1 and 2	No-supervision		
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

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

A fast inspection reveals that this mechanism is collusion-proof and A optimally selects Contract j when the signal is τ_j. This implements the second-best outcome. The formal proof is similar to the one for Example 1 and is omitted. As before, this kind of Selective Supervision Mechanism 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 the type is inefficient, S has the possibility of managing multiple signals.

In what follows, we will show that the Selective Supervision Mechanism always achieves the second-best outcome in Celik's (2009) and FLM's (2003) frameworks. On the contrary, it may not achieve the second-best outcome when the timing of the game is *Timing 2*, the information structure is the one proposed by FLM (2003), and $N > 2$.

3.3 Timing 1

Under *Timing 1,* the implementation of the Selective Supervision Mechanism is somewhat trivial. If S learns τ at date -1, P can costlessy 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

[19] As before, the results are rounded to the second decimal place.

and agree to respond to P collusively.

Let us denote by s_{ijk} (respectively t_{ijk} and q_{ijk}) S's wage (respectively A's 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 the Selective Supervision Mechanism in such a way 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 3 (Non-Cooperative Implementation). As before, write $s_{ijj} = s_{ij}$ and $t_{ijj} = t_{ij}$. Using the logic of Nash implementation, P can costlessly elicit τ by inducing the agent and the supervisor 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.

Despite being trivial in general, the implementation of the Selective Supervision Mechanism is still interesting when S signal is binary, i.e., $T = \{\tau_1, \tau_2\}$. This is the case in both Celik (2009) and FLM(2003). When S signal is binary, P does need to offer S a menu of contracts. It is sufficient to offer an unique contract that is refused by S when τ_2 is realized. S participation decision conveys important information to P and allows him to costlessy elicit τ. The structure of such a Selective Supervision Mechanism is simple: S gets zero wage $s_{ij}^{sb} = 0$, and A is offered an individually incentive-compatible mechanism that includes q_{Nj}^{sb} only when S refuses the contract. By not offering the output schedule for the most inefficient type q_{Nj}^{sb}, P can effectively reduce the information rent for the other type(s) when S accepts the contract. Given that S and A cannot collude in participation decisions, the logic of the Nash implementation applies again. P can design a Selective Supervision Mechanism that induces S to refuse the contract when θ_N is realized .[20] This Selective Supervision Mechanism is feasible only if S is not indispensable for production, which is a departure from Celik's (2009) and FML's (2003) assumption. Clearly, if S is indispensable for production, P needs to design a Selective Supervision Mechanism that entails a menu of contracts (as described in the first part of this paragraph).

Admittedly, it feels a bit 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 a menu of contracts. All considered, our result seems to suggest that a departure

[20]Notes available from the author.

from the assumption of *no-collusion in participation decision* is required to develop a sound theory of collusive contracting under *Timing 1*. On the contrary, under *Timing 2* the implementation of the Selective Supervision Mechanism is more challenging and the second-best outcome can be achieved only under certain parametric conditions. Therefore, the assumption of *no-collusion in participation decision* seems to be more natural under the latter timing.

3.4 Timing 2

In this section we discuss the implementation of the Selective Supervision Mechanism under *Timing 2*. P offers a menu of N contracts

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

Let us denote by Γ_n the generic contract where $n \in [1, 2, ..., N]$ 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, Γ_1 allows S to send only one message τ_1, Γ_2 allows S to send two messages τ_1 and τ_2 and so on. Each contract Γ_n consists in a triple

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

where q_{ij}^{sb}, t_{ij}^{sb} are the second-best outputs and transfers. The optimal wages depend on the information structure. Consider first FML's (2003) 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}^{*} are scheduled in the following fashion: s_{nj}^{*} solves

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

By construction, s^*_{nj} must be negative. These wages are designed to guarantee that S is willing to report her signal truthfully.

Consider now Celik's (2009) information structure, where there exists at least on type i such that $p(\theta_i, \tau_j) = 0$. S wage is given by

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

It is possible to show that in both cases

Proposition 2 *the wage schedule induces S to report her signal truthfully:*

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

This proposition helps to clarify the outcome of the game when A and S behave non-cooperatively. The proposition 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 our main result.

Proposition 3 *In the case where*

(i) the information structure is the one proposed by FLM and

(ii) there are more than 2 types of agents

Selective Supervision Mechanism (SC) implements the second-best outcome if both the distribution of the agent's production cost and the distribution of the supervisor's signal are not too dispersed. In all other cases, SC always implements the second-best outcome.

As we mentioned before, the Selective Supervision Mechanism may not achieve the second-best outcome when the timing of the game is *Timing 2*, the information structure is the one proposed by FLM (2003), and $N > 2$. This is due to the fact that s^*_{nj} could be a negative payoff large enough to induce S to prefer to misreport θ_n. This is more likely to occur when both the distribution of the production cost and the distribution of the informative signal are dispersed. In order to avoid this misreport, P must increase S wage. This causes a distortion that pushes SC away from second-best. One response to these findings is that there might exist an alternative Selective Supervision Mechanism (with a different wage schedule s_{ij}) that implements the second-best outcome even when the distribution of cost and signal are very dispersed. But a fast inspection reveals that this cannot be the case. In fact, it is possible to show that the Selective Supervision Mechanism proposed in this section highlights a general problem: inducing S to report her signal truthfully requires *at least* one wage s_{ij} to be negative for a given signal j. Increases in the dispersion of cost and signal induces this negative payoff to further decrease: a threshold level is eventually reached where the negative wage is small enough and the second-best outcome is no longer implementable.

To conclude, our result suggests that collusion may still be a problem even under the assumption of *no-collusion in participation decision*. Under *Timing 1* it is always costless to prevent collusion, but this might not be the case when *Timing 2* is considered.

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 Celik's (2009) information structure our result is still robust to these aspects. Indeed, our proof is based on the strong notion of collusion-proofness expressed in (12). On the contrary, under *Timing 2* and FLM's (2003) information structure, 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 Celik's (2009) information structure where $p(\theta_i|\tau_i) = 1$ for all i.

4.2 Implication for Decentralization

The recent literature evaluating delegation when agents collude offers an intriguing puzzle. FLM (2003) and Celik (2009) represent two influential papers in this literature. Despite the very similar setting they consider, 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.[21] The crucial assumption of no collusion in participation decisions drives this result. 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 the 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.[22] 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 rec-

[21] 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.

[22] 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.

ommendation with different degrees of approbation, restricted visa permits, tax amnesties for specific types of evasions, or work contracts subject to different degrees of discretion.

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.

6 Appendix

6.1 Appendix 1: Numerical Example

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 reports type θ_i and S reports τ_j. The following table refers to the same numerical example presented in the article, but the payoffs are not rounded:

Contr. 1	No-supervision		
Type	Output	U	s
θ_1	4	0.4105	0
θ_2	0.8421	0.2	0
θ_3	0.4	0	0

Contr. 2	S reports τ_1			S reports τ_2		
Type	Output	U	s	Output	U	s
θ_1	4	0.4105	0.3346	4	0.7451	0
θ_2	0.8421	0.2	0.1333	1.6471	0.3333	0
θ_3	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
θ_1	4	0.4105	0.431 2	4	0.7451	0.096 6	4	0.8417	0
θ_2	0.8421	0.2	0.175	1.6471	0.3333	0.041 7	1.8667	0.3750	0
θ_3	0.4	0	-0.2519	0.6667	0	-0.03237	0.7500	0	0

If S and A behave non-cooperatively, S reports her signal truthfully. To see this point take Contract 3 and suppose that S has observed τ_3. Then S expected utility when she reports τ_3 is $\sum_1^3 p(\theta_i|\tau_3)0 = 0$. This expected utility is (weakly) higher than the one S would obtain by reporting τ_2 (0) or τ_1 (-0.0528). When S learns the signal τ_2, she is better off reporting the true signal τ_2 (0.0129) rather than reporting τ_3 (0) or τ_1 (0). Similarly, when S learns the signal τ_1 she prefers to report truthfully (0.1622) rather than reporting τ_2 (0.05338) or τ_1 (0). Moreover, each contract is also collusion-proof, i.e., A and S do not collude. To see 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 sum of A and S payoffs is 0.8417 when the type is θ_1 and 0.3750 when the type is θ_2. The potential for collusion arises when A has type θ_3. In this case, S could receive a negative payoff. Suppose that 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 Proposition 1

We have already shown that the second-best always implements the first-best output $q_{1j}^{sb} = q_1^{fb}$ for the most efficient agent. This entails $q_{1j}^{sb} \leq q_{1j'}^{sb}$ for $j' \geq j$.

Consider 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) $p_{i'j} > 0$ and $p_{i'j'} = 0$ implies that $p(\theta_i'|\tau_j) > 0$ 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 Proposition 2

The proof has two parts. In the first part we prove that (17) holds when at least one output schedule is equal to zero, i.e., Celik's (2009) information setup. In the second part we prove that (17) holds when all output schedule are strictly positive, i.e., FML's (2003) information setup. Before proceeding, note that each contract

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

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

To start with, note that from $q_{ij}^{sb} = 0$, follows $p(\theta_i, \tau_j) = 0$. In this case our information structure reduces to a connected partition structure. 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) $U(s_{ij'}^{*}) = U(s_{ij}^{*})$. This is the case when j' and j are redundant 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 follows directly from our information structure. Recall that in this part of the proof we are considering Celik's (2009) connected partition structure. In this case, for any signal $j > 1$ we have the following condition: if $p(\theta_i, \tau_j) > 0$, then for all other non-redundant signals $j' \epsilon [j-1, j-2, ..., 1]$ we must have that $p(\theta_i, \tau_{j'}) = 0$. Thus the *conditionally optimal* second-best outputs are $q_{ij'}^{sb} = 0$. From (16) follows that $s_{ij'}^{*} = 0$.

(c) $U(s_{ij'}^{*}) \geq U(s_{ij}^{*})$ for all i such that $p(\theta_i, \tau_j) = 0$. This follows directly from (16).

Naturally, S would benefit from the misreport only in case (c): but those wages $s_{ij'}^{*}$ 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 redundant 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 θ_i such that $p(\theta_i, \tau_j) > 0$. This follows directly from (16).

(c) $U(s_{ij'}^{*}) \geq U(s_{ij}^{*})$ for all θ_i such that $p(\theta_i, \tau_j) = 0$. This follows again from (16) and from the information structure.

As before, S would benefit from the misreport only in case (c): but those wages $s^*_{ij'}$ are never realized when S learns j. This establishes that (17) holds when at least one output schedule is equal to zero, i.e., $q^{sb}_{ij} = 0$.

6.3.2 Part 2

In the second part we prove that (17) holds when all output schedule are strictly positive, i.e., $q^{sb}_{ij} > 0$. 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 any 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 any j and j' such that $j'=j+1$. We now prove that this result holds for any 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,

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

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 $j'=j+3$, $j'=1+4$ 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. In the first part we consider the case where at least one output schedule is equal to zero, i.e., Celik's (2009) information setup. In the second part we consider the case where all output schedules are strictly positive, i.e., FML's (2003) information setup. In both parts we will first prove that a generic contract Γ_n is collusion-proof. Secondly, we will prove that SC implements the second-best outcome.

6.4.1 Part 1

If there exists at least one output schedule q_{ij}^{sb} such that $q_{ij}^{sb} = 0$ then $p(\theta_i, \tau_j) = 0$. By assumption, our information structure reduces to a connected partition structure: this structure includes the case where S exactly observes A type. To prove that Γ_n is collusion-proof (even when S exactly observes A type) it is sufficient 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

$$U_S(s_{ij}^*) \geq U_S(s_{i'j'}^* + t_{i'j'} - \theta_i q_{i'j'}^{sb} - U_{ij}) \tag{28}$$

for all (i, i', j, j'). Notice that (16) implies that

$$s_{ij}^* = u_{in}^{sb} - u_{ij}^{sb} \tag{29}$$

for all i such that $q_{ij}^{sb} \neq 0$. Recall that by definition $u_{i'j'}^{sb} = t_{i'j'}^{sb} - \theta_{i'} q_{i'j'}^{sb}$ and

$$t_{i'j'} - \theta_i q_{i'j'}^{sb} = u_{i'j'}^{sb} + (\theta_{i'} - \theta_i) q_{i'j'}^{sb}. \tag{30}$$

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

$$U_S(u_{in}^{sb} - u_{ij}^{sb}) \geq U_S(u_{i'n}^{sb} + (\theta_{i'} - \theta_i) q_{i'j'}^{sb} - U_{ij})$$

This condition always holds: to see this point note that (4) and (6) implies that $u_{in}^{sb} \geq u_{i'n}^{sb} + (\theta_{i'} - \theta_i) q_{i'j'}^{sb}$ and from (9) follows that $U_{ij} \geq u_{ij}^{sb}$. This establishes that 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_n. By assumption, at this stage A and S

cannot collude. Therefore, A optimally select 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 (a) S receives a wage equal to zero regardless of the realization of A type (b) A produces the *conditionally-optimal* second best output level q_{ij}^{sb} and receives transfer t_{ij}^{sb}. This proves that SC is collusion-proof and implements the *conditionally-optimal* second best.

6.4.2 Part 2

If $q_{ij}^{sb} > 0$ for all (i,j) then $p(\theta_i, \tau_j) > 0$ for all (i,j). Thus, to prove that Γ_n is collusion-proof 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). In order for S not to be able to find a profitable misreport the following condition must hold:

$$\sum_{i=1}^{n} p(\theta_i|\tau_j) U_S\left(s_{ij}^*\right) \geq \sum_{i=1}^{n} p(\theta_i|\tau_j) U_S\left(s_{i'j'}^* + t_{i'j'} - \theta_i q_{i'j'} - U_{ij}\right)$$

for all (i, i', j, j') and subject to (10) and (9). Using (30), rearrange this expression and obtain,

$$\sum_{i=1}^{n} p(\theta_i|\tau_j) U_S\left(s_{ij}^*\right) \geq \sum_{i=1}^{n} p(\theta_i|\tau_j) U_S\left(s_{i'j'}^* + u_{i'j'}^{sb} + (\theta_{i'} - \theta_i) q_{i'j'}^{sb} - U_{ij}\right) \tag{31}$$

It is straightforward to prove that S, having observed signal j and learned that A has type $i \neq n$ 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 $i' \neq n$. Using (29) twice, (31) reduces to

$$U_S\left(u_{in}^{sb} - u_{ij}^{sb}\right) \geq U_S\left(u_{i'n}^{sb} + (\theta_{i'} - \theta_i) q_{i'j'}^{sb} - U_{ij}\right). \tag{32}$$

To prove that this condition is satisfied, note that (4) and (6) imply that $u_{in}^{sb} \geq u_{i'n}^{sb} + (\theta_{i'} - \theta_i) q_{i'j'}^{sb}$ and from (9) follows $U_{ij} \geq u_{ij}^{sb}$.

When $i' = n$, (15) implies that $s_{nj'}^* < 0$. Recall that (4) implies $u_{ij}^{sb} > u_{nj'}^{sb} + (\theta_n - \theta_i) q_{nj'}^{sb}$ and from (9) follows that $U_{ij} \geq u_{ij}^{sb}$. Thus, the following condition must hold

$$U_S\left(u_{in}^{sb} - u_{ij}^{sb}\right) > U_S\left(s_{nj'} + u_{nj'}^{sb} + (\theta_n - \theta_i) q_{nj'}^{sb} - U_{ij}\right).$$

We are left to consider one last case: namely, we need to show the conditions such that S, having observed signal j and learned that A has type $i = n$ through side-contracting, prefers to tell the truth rather than to report that the state of nature is (i', j'). To begin with, consider the case where $i' \neq n$. The benefit of this misreport lies in the fact that S avoids the negative wage s^*_{nj}. Nonetheless, given that u^{sb}_{ij} satisfies (4), it is costly for S to induce type n to misreport his type as any other type $i' \neq n$. It follows that S has to provide type n with the following positive transfer to misreport his type as $i' \neq n$,

$$b_{nj} = (\theta_n - \theta_{i'})q^{sb}_{i'j'} - u^{sb}_{i'j'}.$$

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 (n, j) also affects the information rent paid to all other A types. Namely, the need to provide type n with the transfer b_{nj} increases the information rent for the other types. From (10) derives that all types $i > n$ must be provided with extra information rent. This coupled with (9) implies that $U_{ij} > u^{sb}_{ij}$. Let us denote with $\widehat{U}_{ij} = U_{ij} - u^{sb}_{ij}$ this extra information rent. Therefore, S having observed signal j and learned that A has type $i = n$ through side-contracting, prefers to tell the truth rather than to report that the state of nature is (i', j') if

$$p(\theta_n|\tau_j)U_S\left(s^*_{nj}\right) + \sum_{i=1}^{n-1} p(\theta_i|\tau_j)U_S\left(s^*_{ij}\right) \geq p(\theta_n|\tau_j)\left(-b_{nj}\right) + \sum_{i=1}^{n-1} p(\theta_i|\tau_j)U_S\left(s^*_{ij} - \widehat{U}_{ij}\right) \tag{33}$$

for all (i', j') where $i' \neq n$. This condition is satisfied when s^*_{nj} is large enough, i.e., close enough to zero. Let us denote by $\widetilde{s}_{nj}$ the threshold value of s^*_{nj} such that (33) is satisfied.

S can also try to avoid the negative payoff s^*_{nj} by reporting that the state of nature is (n, j') when the true state is (n, j). This misrepresentation is never beneficial for all $j' < j$ because in this case $s^*_{nj} > s^*_{nj'}$. Nonetheless it could be beneficial for $j' > j$. In the case where the benefit of the misreport lies in the fact that S improves her payoff by $s^*_{nj'} - s^*_{nj}$ when type n is realized. But again 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 (n, j) also affects the information rent paid to all other A types. Given that $q_{nj'} \geq q_{nj}$ for all $j' > j$, S has to provide all types $i > n$ with an extra information rent equal to $(\theta_n - \theta_{n-1})\left(q_{nj'} - q_{nj}\right)$.

In conclusion, S, having observed signal j and learned that A has type n through side-contracting, prefers to tell the truth rather than to report that the state of nature is (n, j') if

$$p(\theta_n|\tau_j)U_S\left(s^*_{nj}\right) + \sum_{i=1}^{n-1} p(\theta_i|\tau_j)U_S\left(s^*_{ij}\right) \geq \tag{34}$$

$$p(\theta_n|\tau_j)U_S\left(s^*_{nj'} - s^*_{nj}\right) + \sum_{i=1}^{n-1} p(\theta_i|\tau_j)U_S\left(s^*_{ij} - (\theta_n - \theta_{n-1})\left(q_{nj'} - q_{nj}\right)\right)$$

for all $j' > j$. As before, this condition is satisfied when s^*_{nj} is large enough, i.e., close enough to zero. Denote by $\widehat{s}_{nj}$ the value of s^*_{nj} such that (34) is satisfied. Summarizing, the generic contract Γ_N is collusion-proof if

$$s^*_{nj} \geq \max\left[\widehat{s}_{nj}, \widetilde{s}_{nj}\right]. \tag{35}$$

If condition (35) holds then Γ_n 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 Γ_n. 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 (a) S receives a wage equal to zero regardless of the realization of A cost type, and (b) 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 SC is collusion-proof. One final aspect is worth noticing: it is easy to see that condition (35) is always satisfied when $N = 2$. Therefore, SC always implements the second-best outcome in the setting proposed by FML(2003). Moreover, from the construction of s^*_{nj}, it follows that (35) is more likely to hold when the distribution of the agent's production cost and the distribution of the supervisor's signal are not too dispersed.

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.

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

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

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