



Consulting collusive experts

Dilip Mookherjee^a, Alberto Motta^b, Masatoshi Tsumagari^{c,*}

^a Department of Economics, Boston University, 270 Bay State Road, Boston, MA, 02215, United States of America

^b School of Economics, University of New South Wales, Australian School of Business, Sydney 2052, Australia

^c Department of Economics, Keio University, 2-15-45 Mita Minato-ku, Tokyo 108-8345, Japan



ARTICLE INFO

Article history:

Received 29 August 2018

Available online 11 May 2020

JEL classification:

D82

D86

L23

Keywords:

Mechanism design

Collusion

Delegation

Expert

ABSTRACT

In designing a contract with an agent privately informed about its cost, should a principal consult an expert who has already received a partially informative signal of the agent's cost? The expert has a prior relationship with the agent, facilitating (weak) ex ante collusion which coordinates their participation and reporting decisions with accompanying side-payments. While delegating contracting with the agent to the expert is never profitable, we show that consulting the expert is typically valuable. Changes in bargaining power within the coalition have no effect, while altruism of the expert towards the agent makes the principal worse off.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Consider a Principal (P) (e.g., lender, customer, employer or government) seeking to procure a service from an agent (A) (e.g., borrower, contractor, worker or firm) that is privately informed about its cost. In order to limit the agent's information rent, P seeks the advice of an expert (M) (e.g., rating agency, consultant, manager or regulator) endowed with information concerning technology and supply conditions in the relevant industry. The expert has a prior relationship with the agent, which creates the potential for ex ante collusion: besides coordinating their reports to P, they can also coordinate their decision whether to participate in the mechanism. Prior literature on hierarchical mechanism design has focused mainly on contexts of interim collusion, restricted to coordination of reporting decisions, after coalition partners have independently agreed to participate (Faure-Grimaud, Laffont and Martimort (FLM hereafter, 2003), Celik (2009)). Interim collusion seems more appropriate in auditing or supervision contexts where P assigns an auditor with no prior relationship with the agent, and the auditor and agent come into contact with one another after agreeing to participate. When M and A know each other prior to contracting with P, they can coordinate their participation decisions, thereby enlarging the scope of collusion. In particular, the principal can no longer extract collusion rents of M at the participation stage. This raises a number of new questions concerning effects of different organizational designs on P's payoff. Under what conditions can P still benefit from the expertise of M? Is it optimal for P to contract with M alone, and delegate contracting with the agent to M? Do the costs of collusion vary with allocation of bargaining power between M and A, or altruism of M towards A?

* Corresponding author.

E-mail addresses: dilipm@bu.edu (D. Mookherjee), motta@unsw.edu.au (A. Motta), tsuma@econ.keio.ac.jp (M. Tsumagari).

The existence of prior connections between experts and agents is common in many real world contexts: e.g., credit rating agencies and borrowers, contractors and subcontractors, company Directors and CEOs, managers and workers, or regulators and private utilities. However, its consequences for the design of hierarchical contracts have not received much attention.¹ Apart from incorporating collusion in participation, our model is similar in the following respects to existing literature on collusion. M's signal of A's cost is partially informative. Both M and A observe the realization of this signal, resulting in one-sided asymmetric information within the coalition, which represents the sole friction in collusion. If they fail to agree on a side contract they play non-cooperatively thereafter (referred to as *weak collusion*). Hence M and A enter into a deviating side contract only if it results in an interim Pareto improving allocation for the coalition. In some other technical respects, we consider a more general model: e.g., we allow offered side-contracts to be rejected on the equilibrium path, and updating of beliefs off the equilibrium path following rejection.² We also consider a continuum type space for A's cost, and show how classical Myersonian mechanism design methods based on 'virtual' types can be extended (using techniques in Jullien (2000)) to incorporate ex ante collusion constraints, while allowing for general information structures for M. This is in contrast to previous analyses which have focused on discrete (two or three) type cases and specific information structures.

The equilibrium concept we employ is a refinement (PBE(c)) of Perfect Bayesian Equilibrium (PBE), which requires non-existence of any other (interim) Pareto superior PBE in the game played by M and A (following offer of a contract by P, and any given realization of M and A's signals). The motivation for this refinement is that in its absence, collusion ceases to be costly, in the sense that P's second-best profit can always be implemented by a PBE involving a particular lack of coordination between M and A. Hence the refinement is essential to capture the problem posed by collusion. We prove a version of the Collusion-Proofness Principle (Tirole (1995)): any allocation can be implemented as the outcome of a PBE(c) of some mechanism if and only if it is ex ante collusion proof (i.e., can also be achieved as the outcome of a corresponding direct revelation mechanism, where the coalition has no incentive to enter into a non-null collusive side-contract) and satisfies interim participation constraints for M and A. This result does not require arbitrary restrictions on off-equilibrium path beliefs within the coalition.³ It is subsequently used to show that P's problem reduces to contracting with the coalition treated as a single composite agent with a shadow (unit) cost function for delivering the good, which incorporates incentive and ex post participation constraints for the coalition. P exercises some control over the nature of this shadow cost function, as it depends on outside option payoffs A would earn from noncooperative play of P's mechanism.

Using the preceding results, we derive some qualitative properties of optimal mechanisms in the presence of ex ante collusion. First, it is **never** optimal for P to delegate authority to M to (sub-)contract with A, i.e., where P does not personally enter into any non-null contract with A. In other words, centralized contracting is essential for P to be able to benefit from consulting M. The intuitive reason is that delegation entails P giving up altogether on attempting to control the shadow cost of the coalition, which then ends up being higher than A's own cost. Consequently, P would then do better to bypass M entirely and contract directly with A without receiving any advice from M. A similar result is obtained by Mookherjee and Tsumagari (2004) in a team production setting involving ex ante collusion between two agents privately informed about their respective costs. This result demonstrates a stark difference between ex ante and interim collusion, as in the latter delegation has been shown to be optimal for some specific settings (Baliga and Sjöström (1998), FLM (2003)). Collusion in participation decisions is akin to a limited liability constraint for the coalition, which prevents P from being able to extract M's collusion rents upfront at the participation stage, resulting in double marginalization of rents (DMR).

This raises the question whether consulting M can nevertheless be valuable with centralized contracting. We show the answer is generally yes.⁴ The proof is based on showing that small variations can be constructed around the optimal contract with A where M is not consulted, which raise P's expected payoff while preserving ex ante collusion-proofness and M's interim participation constraints. Following any realized signal of M, the variation entails raising the output procured over a range of cost types (where the information rents of A assessed according to M's interim beliefs are lower relative to those assessed using P's prior beliefs), and lowering it over another range (where the opposite is true). These variations are achieved with a suitable choice of A's outside option payoffs corresponding to these cost types. The informativeness of M's signal implies the existence of such pairs of cost ranges, where M and P have different marginal rates of substitution with respect to the induced output variations. Hence 'mutual gains from trade' are generated by such a variation, which generate higher ex ante profits for P while preserving M's interim participation constraints.

The following additional results are subsequently obtained:

- We provide sufficient conditions for second-best profits to be unattainable by P under ex ante collusion. This contrasts with results obtained by Che and Kim (2006) or Motta (2009) for interim collusion settings who show that second-best profits can be achieved quite generally.

¹ Implications of ex ante collusion in the context of auction design (where bidders collude on participation and bids) have been studied by Pavlov (2008), Che and Kim (2009) and Che et al. (2018).

² These are elaborated later.

³ This pertains to beliefs over A's cost formed by M following rejection by A of an offered side-contract, which occurs off the equilibrium path.

⁴ Our base model assumes the good to be procured is perfectly divisible. With an indivisible good, it turns out it may sometimes not be valuable to consult M, owing to lack of scope for adjusting the quantity procured based on M's advice. We discuss this in the section on extensions of the model; further details of this case are provided in the online Appendix.

- Changes in bargaining power between M and A over the side-contract do not alter optimal allocations or payoffs achieved. Hence efforts to manipulate ‘control rights’ between M and A with the intent of altering the allocation of bargaining power within the coalition, would not lower collusion costs. A modified version of the ‘Coase Theorem’ therefore holds, despite the presence of asymmetric information within the coalition. This turns out to be a consequence of the standard assumption that failure to collude results in noncooperative play in P’s mechanism.⁵
- Collusion costs would increase if M were altruistic towards A, implying the need for P to consult experts without ‘personal’ connections with A.⁶ The result is not a priori obvious, owing to two offsetting effects. Increased altruism facilitates collusion by lowering frictions within the coalition, but also reduces the severity of the DMR problem by limiting the extent to which M seeks to gain personally from lowering the price offered to A. The former effect outweighs the latter when contracting is centralized, while the opposite turns out to be true under delegated contracting.
- In the specific case of an indivisible good and two possible signal realizations for M, we provide a more detailed characterization of the optimal contract under ex ante collusion, and how it differs from optimal contracts under interim collusion. These are illustrated via numerical solutions for a parametric example. When P’s valuation of the good is low, the second-best is achievable in both settings. Over an intermediate range of valuations, the second best is not achievable, optimal contracts in the two contexts coincide and incentives are ‘low-powered’. For higher valuations, optimal contracts diverge: under interim collusion it entails delegation and ‘high powered’ contracts, resembling a franchise arrangement which is infeasible in the presence of ex ante collusion. Finally for sufficiently high valuations, consulting the expert is not valuable in ex ante collusion, but is always valuable in interim collusion.

The paper is organized as follows. Section 2 discusses relation to existing literature in more detail. Section 3 presents the model featuring a perfectly divisible good and a general information structure for M involving a finite number of possible signals. We first explain the solution concept used, and then provide a characterization of feasible allocations under ex ante collusion. Using this, Section 4 presents the main results: suboptimality of delegation, value of the expert with centralized contracting, and conditions under which the second-best cannot be implemented. Section 5 discusses extensions incorporating alternative allocations of bargaining power within the coalition, altruistic experts, and the case of an indivisible good. Finally, Section 6 discusses implications of our results, extensions and shortcomings. Proofs are presented in the Appendix, while supplementary results are contained in an online Appendix.⁷

2. Related literature

The most directly related papers are those of Faure-Grimaud, Laffont and Martimort (FLM, 2003) and Celik (2009) on consequences of interim collusion between a supervisor and an agent who collude only on reports but not on participation decisions. Such a context limits the scope of collusion, by allowing the principal to extract part of the collusion rents of the supervisor at the participation stage. As FLM show, delegation can be optimal under some circumstances, e.g., when types and signals are binary, with full support for conditional distributions. Celik (2009) on the other hand shows delegation performs worse than no supervision if there are three types and two signals, and the support of conditional distributions depends on the signal. Our results show that Celik’s results (which apply to a specific information structure in the context of interim collusion) extends to general information structures under ex ante collusion.

Other contributions of our paper are of a more technical nature. Since Laffont and Martimort (1997), the literature on collusion has typically restricted attention to equilibria in which collusive side contract offers are never rejected on the equilibrium path, and in which beliefs are passive (i.e., off-equilibrium-path rejections do not result in any updating of prior beliefs). Celik and Peters (2011) have argued that this entails a loss of generality, by constructing an example where the set of achievable payoffs by colluding parties can be enlarged by allowing equilibrium path rejection. Our theory does not impose any such restrictions: we allow equilibrium path rejections of side contract offers, and beliefs to be updated following such rejection both on and off the equilibrium path. Moreover, we study a continuum type space, while our results happen to apply to a discrete type space as well.

In Mookherjee et al. (2018) we apply the model of this paper to the setting of procurement of an indivisible good, and use it to study trade-offs between ‘outsourcing’ and ‘direct foreign investment’ as two alternative ways for a multinational company to procure an input from another country. Outsourcing corresponds to delegation which is subject to DMR, while centralized contracts entail an additional fixed cost (associated with costs of direct communication or contracting between P and A). Finally, Mookherjee and Tsumagari (2017) extend the analysis to incorporate strong collusion, where M and A can commit to ‘extortion’ threats if the other party refuses to participate in the collusion. In that context the allocation of bargaining power between M and A turns out to matter: P benefits from consulting M if and only if M has higher bargaining power than A.

⁵ Mookherjee and Tsumagari (2017) show this result no longer holds in settings of ‘strong’ collusion, where side contracts include commitments to ‘extortion’ threats concerning strategies they will employ by either party should the other party refuse to participate in the collusion.

⁶ By altruism we mean an expert whose payoff is a convex combination of his own compensation and the agent’s payoff, with a lower welfare weight on the latter than on own-compensation.

⁷ The link to the online Appendix is <http://people.bu.edu/dilipm/publications/OnlineApp-gebrevJan2020v1.pdf>.

We now discuss the relation to the wider literature on collusion. Early contributions (e.g., Tirole (1986), Laffont and Tirole (1993)) focused on contexts of ‘hard’ information where supervisors could only hide information but could not report untruthfully. The subsequent literature examines the case of ‘soft’ information where no constraints on allowable reports are imposed. They can be classified by the context (auctions, team production or supervision), the nature of collusion (ex ante or interim, weak or strong collusion),⁸ and whether type spaces are discrete or continuous.

A large part of existing literature deals with auctions and team production, where there are multiple privately informed agents and no supervisor. For auctions, Dequiedt (2007) considers strong ex ante collusion with binary agent types and shows that efficient collusion is possible, implying that the second-best may not be achieved. In contrast, Pavlov (2008) considers a model with continuous types where the second-best can be achieved with weak ex ante collusion, and Che and Kim (2009) find the same result with either weak or strong ex ante collusion with continuous types. Team production with binary types is studied by Laffont and Martimort (1997), who show the second best can be achieved with weak interim collusion; this analysis is extended to a public goods context in Laffont and Martimort (2000) to explore the role of correlated types. Baliga and Sjöström (1998) consider a team setting with two productive agents that collude, involving moral hazard and limited liability rather than adverse selection. They show that delegation to one of the agents is an optimal response to collusion for a wide set of parameter values. Che and Kim (2006) show how second-best allocations can be achieved in a general mechanism design context (including team production) under weak interim collusion. Quesada (2004) on the other hand shows strong ex ante collusion is costly in a team production model with binary types. Mookherjee and Tsumagari (2004) show delegation to one of the agents is worse than centralized contracting in the presence of weak ex ante collusion. The logic of this result is similar to that underlying our result that delegation to the expert is worse than not consulting the expert at all. Their paper also considers delegation to an expert who is perfectly informed about the costs of each agent, and show that its value relative to centralized contracting depends on complementarity or substitutability between inputs supplied by different agents. The current paper differs insofar as there is only one agent, and there is asymmetric information within the expert-agent coalition owing to the expert receiving a noisy signal of the agent’s cost.

3. The model

3.1. Environment

A delivers a verifiable output q to P at a personal cost of θq . Output is perfectly divisible: the range of feasible outputs is \mathfrak{R}_+ . P’s return from q is $V(q)$ where $V(q)$ is twice continuously differentiable, increasing and strictly concave satisfying $\lim_{q \rightarrow 0} V'(q) = +\infty$, $\lim_{q \rightarrow +\infty} V'(q) = 0$ and $V(0) = 0$. These conditions imply that $q^*(\theta) \equiv \operatorname{argmax}_q [V(q) - \theta q]$ is continuously differentiable, positive on $\theta \in [0, \infty)$ and strictly decreasing.

A is privately informed about the realization of θ ; P and M share a common prior $F(\theta)$ over θ on the interval $\Theta \equiv [\underline{\theta}, \bar{\theta}] \subset \mathfrak{R}_+$. F has a density function $f(\theta)$ which is continuously differentiable and everywhere positive on its support. The ‘virtual cost’ $H(\theta) \equiv \theta + \frac{F(\theta)}{f(\theta)}$ is strictly increasing in θ ; this assumption simplifies the analysis but is inessential to the results.

M plays no role in production, and costlessly acquires an informative signal η about θ . The underlying assumption is that the relevant knowledge concerning A’s cost realization has already been acquired by M prior to contracting. The set of possible realizations of η is \mathcal{N} , a finite set with $\#\mathcal{N} \geq 2$. The finiteness of this set is assumed for technical convenience, and is inessential as long as M’s information regarding θ is not perfect. It is common knowledge that the realization of η is observed by both M and A. $a(\eta | \theta) \in [0, 1]$ denotes the probability of η conditional on θ . The function $a(\eta | \theta)$ is continuously differentiable and positive on $\Theta(\eta)$, where $\Theta(\eta)$ denotes the set of values of θ for which signal η can arise with positive probability. We assume $\Theta(\eta)$ is an interval for every $\eta \in \mathcal{N}$. Define $\underline{\theta}(\eta) \equiv \inf \Theta(\eta)$ and $\bar{\theta}(\eta) \equiv \sup \Theta(\eta)$. We assume that for any $\eta \in \mathcal{N}$, $a(\eta | \theta)$ is not a constant function on Θ , and there are some subsets of θ with positive measure such that $a(\eta | \theta) \neq a(\eta' | \theta)$ for any $\eta, \eta' \in \mathcal{N}$. In this sense each possible signal realization conveys information about the agent’s cost. The information conveyed is partial, since \mathcal{N} is finite. This formulation includes both cases of full support and partition information structures.

The conditional density function and the conditional distribution function are respectively denoted by $f(\theta | \eta) \equiv f(\theta)a(\eta | \theta)/p(\eta)$ (where $p(\eta) \equiv \int_{\underline{\theta}(\eta)}^{\bar{\theta}(\eta)} f(\tilde{\theta})a(\eta | \tilde{\theta})d\tilde{\theta}$) and $F(\theta | \eta) \equiv \int_{\underline{\theta}(\eta)}^{\theta} f(\tilde{\theta} | \eta)d\tilde{\theta}$. The ‘virtual’ cost conditional on the signal η is $h(\theta | \eta) \equiv \theta + \frac{F(\theta|\eta)}{f(\theta|\eta)}$. We do not impose any monotonicity assumption for $h(\theta | \eta)$. Let $\hat{h}(\theta | \eta)$ be constructed from $h(\theta | \eta)$ and $F(\theta | \eta)$ by the ironing procedure introduced by Myerson (1981) (see the online Appendix for details regarding this procedure).

All players are risk neutral. P’s objective is to maximize the expected value of $V(q)$, less expected payment to A and M, represented by X_A and X_M respectively. M’s objective is to maximize expected transfers $X_M - t$ where t is a transfer from

⁸ Ex ante collusion permits collusion over both reporting and participation decisions, while interim collusion pertains only to collusion over reporting. Weak collusion refers to collusion for mutual gain, where failure to agree to collude is followed by noncooperative play. Under strong collusion, each partner commits to a threat pertaining to how it would play in P’s mechanism, should the other partner refuse to collude, thereby permitting extortion as well as mutual gain.

M to A. A seeks to maximize expected transfers received, less expected production costs, $X_A + t - \theta q$. Both A and M have outside options equal to 0.

In this environment, a (*deterministic*) allocation is represented by

$$(u_A, u_M, q) = \{(u_A(\theta, \eta), u_M(\theta, \eta), q(\theta, \eta)) \in \mathfrak{R}^2 \times \mathfrak{R}_+ \mid (\theta, \eta) \in K\}$$

where $K \equiv \{(\theta, \eta) \mid \eta \in \mathcal{N}, \theta \in \Theta(\eta)\}$, u_M, u_A denote M and A's payoff respectively, and q represents the production level. P's payoff equals $u_P = V(q) - u_M - u_A - \theta q$. These payoffs relate to transfers and productions as follows: $u_A \equiv X_A + t - \theta q$; $u_M \equiv X_M - t$; $u_P \equiv V(q) - X_M - X_A$.

For technical convenience we will also consider randomized allocations (in which P selects a probability distribution over (u_A, u_M, q) in any given state), though it will turn out that P does not gain from such allocations.⁹ In a randomized allocation, $(u_A(\theta, \eta), u_M(\theta, \eta), q(\theta, \eta))$ denotes the expected payoffs of A, M and the expected output, conditional on the state (θ, η) .

In the absence of collusion where P costlessly learns the realization of η , it is well known (e.g., adapting arguments of Baron and Myerson (1982)) that the resulting optimal or *second-best allocation* $(u_A^{SB}, u_M^{SB}, q^{SB})$ is as follows:

$$u_A^{SB}(\theta, \eta) = \int_{\theta}^{\bar{\theta}(\eta)} q^{SB}(y, \eta) dy,$$

$$E[u_M^{SB}(\theta, \eta) \mid \eta] = 0$$

and

$$q^{SB}(\theta, \eta) \equiv q^*(\hat{h}(\theta \mid \eta)) = \arg \max_q [V(q) - \hat{h}(\theta \mid \eta)q].$$

3.2. The Ex Ante collusion game

Owing to risk-neutrality of all parties, concavity of V and linearity of A's payoff in q , it is easy to check that P can restrict attention to a deterministic grand contract:

$$GC = (X_A(m_A, m_M), X_M(m_A, m_M), q(m_A, m_M); \mathcal{M}_A, \mathcal{M}_M)$$

where \mathcal{M}_A (resp. \mathcal{M}_M) is a message set for A (resp. M).¹⁰ This mechanism assigns a deterministic allocation, i.e. transfers X_M, X_A and output q , for any message $(m_A, m_M) \in \mathcal{M}_A \times \mathcal{M}_M$. \mathcal{M}_A includes A's exit option $e_A \in \mathcal{M}_A$, with the property that $m_A = e_A$ implies $X_A = q = 0$ for any $m_M \in \mathcal{M}_M$. Similarly \mathcal{M}_M includes M's exit option $e_M \in \mathcal{M}_M$, where $m_M = e_M$ implies $X_M = 0$ for any $m_A \in \mathcal{M}_A$.

The timing of events is as follows.

- (C1) A observes θ and η , M observes η .
- (C2) P offers a grand contract GC .
- (C3) M and A play the side contract game described in more detail below.

As in existing literature, we assume the side-contract is costlessly enforceable. Moreover we assume M can make a take-it-or-leave-it offer of a side-contract. This assumption turns out to be inessential: Section 5.1 explains how the same results obtain with side contracts offered by an uninformed third party that assigns arbitrary welfare weights to M and A.

Our focus is on the subgame C3 played by M and A, following stages C1 and C2. Conditional on their signals and a grand contract offered by P, C3 consists of the following three stages.

- (i) M offers a side-contract SC which determines for any $\tilde{\theta} \in \Theta(\eta)$ to be privately reported by A to M, a probability distribution over joint messages $(m_A, m_M) \in \mathcal{M}_A \times \mathcal{M}_M$, and a side payment from M to A.¹¹ Formally, it is a pair of functions $\{\tilde{m}(\tilde{\theta}, \eta), t(\tilde{\theta}, \eta)\}$ where $\tilde{m}(\tilde{\theta}, \eta) : \Theta(\eta) \times \{\eta\} \rightarrow \Delta(\mathcal{M}_A \times \mathcal{M}_M)$, and $t : \Theta(\eta) \times \{\eta\} \rightarrow \mathfrak{R}$. $\Delta(\mathcal{M}_A \times \mathcal{M}_M)$

⁹ This owes to the linearity of A's payoff in the output produced, while $V(q)$ is strictly concave. Given any randomized allocation which is feasible, its deterministic equivalent allocation (in which the probability distribution over (u_A, u_M, q) in any state is replaced by its corresponding expected value) is also feasible, and generates higher payoff to P.

¹⁰ Randomized contracts are optimal in Ortner and Chassang (2018) owing to their assumption that the contract offered to M by P is not observed by A. In our context, contracts are observed by both M and A. Given risk-neutrality and linearity of A's payoff in q , randomization does not affect either payoffs or information, and thus has no impact on the equilibrium outcome of side-contracting game. On the other hand, P prefers a deterministic output owing to strict concavity of $V(q)$. So there are no benefits from randomized mechanisms. A similar argument is provided in Strausz (2006, footnote 4).

¹¹ The option of randomizing over possible messages is useful for technical reasons. Owing to quasilinearity of payoffs, there is no need to randomize over side transfers.

denotes the set of probability measures over $\mathcal{M}_A \times \mathcal{M}_M$. The case where M does not offer a side contract is represented by a null side-contract (NSC) with zero side payments ($t(\theta, \eta) \equiv 0$), and messages $(m_A(\theta, \eta), m_M(\eta))$ which constitutes a non-cooperative equilibrium of the grand contract under prior beliefs.

(ii) A either accepts or rejects the SC offered, and the game continues as follows.

(iii) If A accepts the offered SC, he sends a private report $\theta' \in \Theta(\eta)$ to M, following which the SC is executed.¹² If A rejects SC, A and M then play the grand contract noncooperatively.

The subgame differs from the case of *interim collusion* since messages include both participation decisions and reports. In interim collusion, (C3) is replaced by (IC3) which consists of the following stages: (i') A and M independently decide whether or not to participate; (ii') if either or both decide not to participate, the game ends with everyone getting zero payoffs. Otherwise the game continues, with M offering a side contract which maps A's (internal) type report θ and M's signal η into (a probability distribution over) joint message sent to P, and t' a side payment between M and A. (iii') A accepts or rejects the offered side contract. If A accepts, the side contract is implemented, otherwise they play the grand contract noncooperatively.

3.3. Equilibrium concept

Returning to the context of ex ante collusion, we explain the equilibrium concept to be used for the game defined by (C3) that follows any choice of a grand contract by P. It is natural to confine attention to Perfect Bayesian Equilibria (PBE) of this game.¹³

If the mechanism design problem is stated as selection of an allocation by P subject to the constraint that it can be achieved as the outcome of some PBE following a choice of a grand contract, it is implicitly presumed that P is free to select any PBE if the game happens to have multiple PBEs. However, it can be shown that in such a formulation the problem of collusion can be completely overcome by the Principal. Specifically, any second best allocation can be achieved as the outcome of a PBE of game (C3) following a suitable choice of a grand contract (besides realization of players' types). A formal proof of this is provided in the online Appendix.

A heuristic description of how this can be achieved is as follows. P selects a grand contract and recommends a non-cooperative equilibrium of this contract in which (i) conditional on participation by M, noncooperative play results in the second-best allocation; (ii) M is paid nothing; and (iii) if M does not participate, P offers A a 'gilded' contract providing the latter a high payoff in all states. On the equilibrium path M always offers a null side contract. If A rejects any offer of a non-null side-contract, they mutually believe that subsequently M will not participate in the grand contract, and A will receive the gilded contract. This forms a PBE as rejection of any non-null side contract is sequentially rational for A given A's belief that M will exit following any rejection. Exiting is sequentially rational for M given his belief that A will reject the side contract and they will subsequently play the grand contract noncooperatively, whence M will be paid nothing. Collusion is thus overcome by P exploiting a lack of coordination among A and M over continuation beliefs and play of the side contracting game. This denies the essence of collusive activity, which involves coordination by the colluding parties 'behind the Principal's back'. The following concept of collusion-proofness incorporates this by allowing M and A to collectively coordinate on the choice of side-contracting equilibria that are Pareto-undominated (for the coalition) relative to the given status quo.

Definition 1. A PBE(c) is a Perfect Bayesian Equilibrium (PBE) of C3 with the following property. There does not exist some signal realization η for which there is some other Perfect Bayesian Equilibrium (PBE) of C3 in which (conditional on η) M's payoff is strictly higher and A's payoff not lower for any type $\theta \in \Theta(\eta)$.

Definition 2. An allocation (u_A, u_M, q) is EAC feasible if there exists a grand contract and a PBE(c) of the subsequent game which results in this allocation.

3.4. Collusion-proofness principle

We now show a version of the Collusion-Proofness Principle (Tirole (1995)) holds in this model. Stated informally, formulating collusion constraints by the PBE(c) refinement implies that P loses nothing by confining attention to revelation mechanisms which satisfy both individual and coalition incentive and participation constraints – i.e., where (a) M does not benefit by offering A any non-null side contract, so M and A play the game defined by the grand contract non-cooperatively; and (b) in the noncooperative play of the grand contract, both report truthfully and agree to participate. In this setting, the revelation mechanism requires each party to report their willingness to participate, and (conditional on agreeing to participate) their information (i.e., (θ, η) for A, and η for M). Refusal to participate by either party constrains P to offer that

¹² Standard arguments show that the restriction to direct revelation mechanisms for the side contract entails no loss of generality.

¹³ For definition of PBE, see Fudenberg and Tirole (1991).

party a null contract. If both participate, there is cross-checking of reports by M and A, providing scope for P to punish inconsistent cross-reports.

To state this formally, we need to introduce the following definitions. Given any (conditional expected) allocation (u_A, u_M, q) , the corresponding *coalitional contract* is a pair of functions $(X(m), Q(m))$ on domain $m \in \hat{\mathcal{M}} \equiv K \cup \{e\}$ (where $K \equiv \{(\theta, \eta) \mid \theta \in \Theta(\eta), \eta \in \mathcal{N}\}$) obtained as follows: if $m = (\theta, \eta) \in K$,

$$(X(m), Q(m)) = (u_A(\theta, \eta) + \theta q(\theta, \eta) + u_M(\theta, \eta), q(\theta, \eta))$$

and $(X(e), Q(e)) = (0, 0)$. Here the message e represents a coordinated decision for both M and A to exit from P's mechanism, and message (θ, η) represents a coordinated decision for M and A to agree to participate and send the common report (θ, η) to P.

The coalitional contract can be interpreted as a contract that P enters into with the coalition treated as a single composite agent, which consists of the coalition $\{M, A\}$ whose members exchange side payments and coordinate their messages to P. The stipulated payoff for the agent represents her outside option while bargaining with M within the coalition (which corresponds to the payoff that A expects to attain by rejecting the side contract offered by M and playing the grand contract noncooperatively). Observe that the allocation (u_A, u_M, q) determines the coalitional contract as well as payoff function u_A of A. It is easy to see that the converse is also true: any coalitional contract combined with a payoff function u_A for A determines an allocation.¹⁴ Hence an allocation is equivalently represented by a coalitional contract and a payoff function for A.

Let $\Delta(\hat{\mathcal{M}})$ denote the set of the probability measures on $\hat{\mathcal{M}}$, and use $\tilde{m} \in \Delta(\hat{\mathcal{M}})$ to denote a randomized message submitted by the coalition to P. With a slight abuse of notation, we shall denote the corresponding coalitional contract by $(X(\tilde{m}), Q(\tilde{m}))$, which is defined on $\Delta(\hat{\mathcal{M}})$. $\tilde{m} = (\theta, \eta)$ or e will be used to denote the probability measure concentrated at (θ, η) or e respectively.

Given a coalitional contract $(X(\tilde{m}), Q(\tilde{m}))$ and payoff function u_A for A, the *side contracting problem* $P(\eta)$ for any $\eta \in \mathcal{N}$ is the following. Select $(\tilde{m}(\theta \mid \eta), \tilde{u}_A(\theta, \eta))$ to maximize M's expected payoff

$$E[X(\tilde{m}(\theta \mid \eta)) - \theta Q(\tilde{m}(\theta \mid \eta)) - \tilde{u}_A(\theta, \eta) \mid \eta] \tag{1}$$

subject to $\tilde{m}(\theta \mid \eta) \in \Delta(\hat{\mathcal{M}})$,

$$\tilde{u}_A(\theta, \eta) \geq \tilde{u}_A(\theta', \eta) + (\theta' - \theta) Q(\tilde{m}(\theta' \mid \eta)) \tag{2}$$

for any $\theta, \theta' \in \Theta(\eta)$, and

$$\tilde{u}_A(\theta, \eta) \geq u_A(\theta, \eta) \tag{3}$$

for all $\theta \in \Theta(\eta)$.

The side contract is interpreted as specifying a coordinated report $\tilde{m}(\theta \mid \eta)$ to P, accompanied by a side-payment which ends up offering a payoff $\tilde{u}_A(\theta, \eta)$, if A reports θ to M. Constraint (2) requires truthful reporting by A to be incentive compatible, while constraint (3) requires A to be offered a payoff at least as the outside option payoff $u_A(\theta, \eta)$.

An allocation is defined to be *ex ante collusion proof* if the optimal side contract designed by M in response is null, in the sense that a truthful joint report of the state (θ, η) combined with the state-contingent payoffs stipulated by the allocation (i.e., with no side transfers between M and A) solves the side contracting problem described above.

Definition 3. The (conditional expected) allocation $(u_A(\theta, \eta), u_M(\theta, \eta), q(\theta, \eta)) : K \rightarrow \mathbb{R}^2 \times \mathbb{R}_+$ is *ex ante collusion proof (EACP)* if for every $\eta \in \mathcal{N}$: $(\tilde{m}(\theta \mid \eta), \tilde{u}_A(\theta, \eta)) = ((\theta, \eta), u_A(\theta, \eta))$ solves problem $P(\eta)$.

We are now in a position to state the main result of this section.

Proposition 1. An allocation (u_A, u_M, q) is EAC feasible if and only if it is an EACP allocation satisfying interim participation constraints

$$E[u_M(\theta, \eta) \mid \eta] \geq 0 \text{ for all } \eta \tag{4}$$

$$u_A(\theta, \eta) \geq 0 \text{ for all } (\theta, \eta) \tag{5}$$

Note that the PBE(c) notion allows for collusion to occur (i.e., a non-null side contract to be offered and accepted by some types of A), and also for side-contract offers to be rejected by some types of A, both on and off the equilibrium path. Hence the focus on EACP allocations does not require arbitrary restrictions on side contract outcomes, e.g., which rule out the possibility of equilibrium-path rejections by A of the side contract offered by M. The problem discussed by Celik and

¹⁴ Specifically, $q(\theta, \eta) = Q(\theta, \eta)$, and $u_M(\theta, \eta) = X(\theta, \eta) - u_A(\theta, \eta) - \theta Q(\theta, \eta)$.

Peters (2011) therefore does not apply to this setting.¹⁵ Moreover, we do not impose any assumption on beliefs apart from those incorporated in the PBE(c) notion; in particular we do not require any additional assumption of passive beliefs.¹⁶

4. Main results

4.1. Suboptimality of delegated contracting

First consider the special case of *Delegation to M (DM)* where P delegates authority to M over contracting with A. Here the GC designed by P involves a null contract for A: the latter submits no report to P directly, and receives no production instructions or payments from P. P contracts only with M, requiring the latter to send a message m_M to P which determines the output $Q(m_M)$ and aggregate payment $X(m_M)$ to the (M, A) coalition. Following receipt of this offer, M designs a side contract for A which selects an output $q(m_A)$ and payment $X_A(m_A)$ to the latter as a function of a message m_A sent by A to M, provided A accepts the side contract. After receiving A's message (and conditional on A agreeing to participate), M submits a participation decision and message m_M to P. In contrast to the interim collusion setting, M can postpone submission of the participation decision *after* receiving a report from A.

Our first main result is that delegation is *never* optimal in ex ante collusion, as it is strictly dominated by the case where M is not consulted at all, which we refer to as *No Monitor (NM)*.

Proposition 2. *Delegation to M generates lower expected profit for the Principal compared to the optimal NM mechanism with no monitor.*

The FLM result concerning optimality of delegation in an interim collusion setting with two cost types therefore does not extend to ex ante collusion.¹⁷ The underlying argument (extending Proposition 1 in Mookherjee and Tsumagari (2004) for a setting with two agents and no supervisor or intermediary) is simple and very general (e.g., it can be shown to extend to a discrete type setting also). P contracts for delivery of the good with M, so the problem reduces to contracting with a single agent M. In order to deliver the good to P, M needs to procure it in turn from A. The cost that M expects to incur equals A's virtual cost function $h(\theta|\eta)$ corresponding to the signal observed by M. This is unambiguously higher than the delivery cost θ of A if P were to contract directly with A. This is the well-known problem of double marginalization of rents (DMR), arising due to exercise of monopsony power by M in side-contracting with A. Unlike the context of interim collusion, M can postpone her own participation decision *after* receiving A's report. This effectively translates into a kind of 'limited liability' constraint for M, which prevents P from taxing away upfront the rents earned by M.¹⁸

Given this result, we hereafter focus on centralized contracting, where P offers a non-null contract to both M and A in GC. The side-contracting problem $P(\eta)$ then involves type-dependent outside option payoffs for A which correspond to what A can attain by rejecting the offered side contract and playing P's mechanism noncooperatively. By designing contracts for A and offering positive rather than zero outside option payoffs $u_A(\theta, \eta)$, P can manipulate the side contracting problem by making it harder for M to design a profitable side contract (as constraint (3) becomes tighter).

¹⁵ They show in the context of a model of a two-firm cartel that such restrictions can entail a loss of generality. Rejection of a side contract by some types of A can communicate information to M about A's type, affecting subsequent play and resulting payoffs in the noncooperative play of the grand contract. Celik and Peters (2011) show that there can be collusive allocations amongst cartel members which can only be supported by side-contract offers which are rejected with positive probability on the equilibrium path. This problem does not arise in our setting as the side contract is offered by M, and the side contract does not have to satisfy interim participation constraints for M. However, the Celik-Peters problem could conceivably arise in a setting where the side contract is offered by a third party. Even in that context, it turns out that the problem can be overcome with a suitable modification of the side contracting game. Specifically, we augment the agent's message space at the participation stage beyond the binary exit-participation decision to directly communicate some information about her cost. It is then unnecessary for the agent to choose the exit option in order to send a cost signal. With such an augmented message space, we show that any allocation achieved by the coalition with rejection on the equilibrium path is also achievable in an equilibrium where the agent always participates. For further details see the online Appendix.

¹⁶ This assumption states that M's beliefs regarding A's type θ are the same as her prior beliefs, if an offered side contract is rejected by some types of A off the equilibrium path, or if a null side contract is offered by M. In the latter case the PBE restriction implies that M's beliefs will remain unchanged. In the case of rejection of a non-null side contract, we allow M's beliefs to change. For further details see the proof of Proposition 1.

¹⁷ It can be shown, however, that the optimal allocation under ex ante collusion can be achieved by a modified form of delegation, where P communicates and transacts only with M on the equilibrium path. In this arrangement, M is 'normally' expected to contract on behalf of the coalition $\{M, A\}$ with P, sending a joint participation decision and report of the state (θ, η) to P after having entered into a side contract with A. However A has the option of circumventing this 'normal' procedure and asking P to activate a grand contract in which A and M will send independent reports and participation decisions to P. The presence of this option ensures that A has sufficient bargaining power within the coalition; it does not have to be 'actually' used on the equilibrium path.

¹⁸ While it is relatively easy to show that DM cannot dominate NM, the proof establishes the stronger result that DM is **strictly** dominated by NM. The proof of strict domination is straightforward in the case that $h(\theta|\eta)$ is continuous and nondecreasing in θ over a common support $[\underline{\theta}, \bar{\theta}]$ for every η . In that case an argument based on Proposition 1 in Mookherjee and Tsumagari (2004) can be applied. In the general case there are a number of additional technical complications, but we show that the result still goes through.

4.2. Characterization of EACP allocations

In order to understand the implications of the EACP property, we start by focusing on the side contract design problem $P(\eta)$ faced by M following signal realization η . The following discussion essentially follows the analysis of Jullien (2000) of an agency model with type dependent outside options. We highlight the key steps that help illuminate subsequent definitions and concepts that are used below.

Given the coalitional contract $X(m)$, $Q(m)$ and A's payoff function $u_A(\theta, \eta)$, and the signal η , M selects the following controls while deciding which side contract to offer: the message $\tilde{m}(\theta | \eta)$ sent to P, and the side payment from M to A defining A's rents $\tilde{u}_A(\theta, \eta)$. Standard arguments imply that A's incentive constraint (2) is equivalent to

$$\tilde{u}_A(\theta, \eta) = \tilde{u}_A(\bar{\theta}(\eta), \eta) + \int_{\theta}^{\bar{\theta}(\eta)} Q(\tilde{m}(y | \eta)) dy \tag{6}$$

and $Q(\tilde{m}(\theta | \eta))$ is non-increasing in θ . Hence the problem can be rewritten as

$$\max E[X(\tilde{m}(\theta | \eta)) - \theta Q(\tilde{m}(\theta | \eta)) - \tilde{u}_A(\theta, \eta) | \eta]$$

subject to $\tilde{m}(\theta | \eta) \in \Delta(\hat{\mathcal{M}})$ where $\hat{\mathcal{M}} \equiv K \cup \{e\}$, the incentive constraint (6), the participation constraint (3), and $Q(\tilde{m}(\theta | \eta))$ non-increasing in θ . So the solution maximizes (subject to the constraint $Q(\tilde{m}(\theta | \eta))$ is non-increasing in θ) the following Lagrangian expression corresponding to some non-decreasing function $\lambda(\theta | \eta)$:¹⁹

$$\begin{aligned} \mathcal{L} &\equiv E[X(\tilde{m}(\theta | \eta)) - \theta Q(\tilde{m}(\theta | \eta)) - \tilde{u}_A(\theta, \eta) | \eta] \\ &+ \int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} [\tilde{u}_A(\theta, \eta) - u_A(\theta, \eta)] d\lambda(\theta | \eta) \end{aligned}$$

where $\tilde{u}_A(\theta, \eta)$ is given by (6), and $X(\tilde{m})$, $Q(\tilde{m})$ denote expected values of $X(m)$, $Q(m)$ taken with respect to probability measure \tilde{m} . The slope of the function λ , when it exists, corresponds to Kuhn-Tucker multipliers for the agent's participation constraint (3).

Using the agent's incentive constraint (6), standard arguments imply that the first term in the Lagrangian $E[X(\tilde{m}(\theta | \eta)) - \theta Q(\tilde{m}(\theta | \eta)) - \tilde{u}_A(\theta, \eta) | \eta]$ equals

$$E[X(\tilde{m}(\theta | \eta)) - \{\theta + \frac{F(\theta | \eta)}{f(\theta | \eta)}\} Q(\tilde{m}(\theta | \eta)) | \eta] - \tilde{u}_A(\bar{\theta}(\eta), \eta) \tag{7}$$

Turning next to the second term, integration by parts combined with (6) yields:

$$\begin{aligned} &\int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} \tilde{u}_A(\theta, \eta) d\lambda(\theta | \eta) = \lambda(\bar{\theta}(\eta) | \eta) \tilde{u}_A(\bar{\theta}(\eta), \eta) - \lambda(\underline{\theta}(\eta) | \eta) \tilde{u}_A(\underline{\theta}(\eta), \eta) \\ &+ \int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} \lambda(\theta | \eta) Q(\tilde{m}(\theta | \eta)) d\theta \\ &= [\lambda(\bar{\theta}(\eta) | \eta) - \lambda(\underline{\theta}(\eta) | \eta)] \tilde{u}_A(\bar{\theta}(\eta), \eta) + \int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} [\lambda(\theta | \eta) - \lambda(\underline{\theta}(\eta) | \eta)] Q(\tilde{m}(\theta | \eta)) d\theta \\ &= \Lambda(\bar{\theta}(\eta) | \eta) \tilde{u}_A(\bar{\theta}(\eta), \eta) + \int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} \Lambda(\theta | \eta) Q(\tilde{m}(\theta | \eta)) d\theta \\ &= \Lambda(\bar{\theta}(\eta) | \eta) \tilde{u}_A(\bar{\theta}(\eta), \eta) + E[\frac{\Lambda(\theta | \eta)}{f(\theta | \eta)} Q(\tilde{m}(\theta | \eta)) | \eta] \end{aligned}$$

where $\Lambda(\theta | \eta)$ denotes $[\lambda(\theta | \eta) - \lambda(\underline{\theta}(\eta) | \eta)]$. Clearly Λ is nonnegative and nondecreasing in θ , with $\Lambda(\underline{\theta}(\eta) | \eta) = 0$, and its slope with respect to θ (when it exists) corresponds to the Kuhn Tucker multiplier on the participation constraint.

Moreover, $\Lambda(\bar{\theta}(\eta) | \eta) = 1$, owing to the following argument. Consider the effect of raising uniformly A's outside option function from $u_A(\theta, \eta)$ to $u_A(\theta, \eta) + \Delta$ where Δ is an arbitrary positive scalar. It is evident that the solution is unchanged, except that $\tilde{u}_A(\theta, \eta)$ is raised uniformly by Δ . Hence the maximized payoff of M must fall by Δ , implying that

¹⁹ Since randomized side contracts can be chosen, the objective function is concave, the feasible set is convex and has non-empty interior; hence the Kuhn-Tucker theorem can be applied. It is also easily checked that M can confine attention to a deterministic function $\tilde{u}_A(\theta, \eta)$ without any loss of profit.

$$\int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} \Delta d\lambda(\theta | \eta) = [\lambda(\bar{\theta}(\eta) | \eta) - \lambda(\underline{\theta}(\eta) | \eta)]\Delta = \Delta,$$

implying $\Lambda(\bar{\theta}(\eta) | \eta) = \lambda(\bar{\theta}(\eta) | \eta) - \lambda(\underline{\theta}(\eta) | \eta) = 1$ in the optimal solution.

Combining the two terms of the Lagrangian, we therefore obtain

$$\begin{aligned} \mathcal{L} \equiv & \int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} [X(\tilde{m}(\theta | \eta)) - \pi(\theta | \eta)Q(\tilde{m}(\theta | \eta))]dF(\theta | \eta) \\ & - \int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} u_A(\theta, \eta)d\Lambda(\theta | \eta) \end{aligned} \tag{8}$$

where $\pi(\theta | \eta) \equiv \theta + \frac{F(\theta|\eta) - \Lambda(\theta|\eta)}{f(\theta|\eta)}$.

Hence M will select the joint message $\tilde{m}(\theta | \eta)$ to point-wise maximize $X(m) - \pi(\theta | \eta)Q(m)$, if the monotonicity constraint on $Q(\tilde{m}(\theta | \eta))$ is ignored. Ignoring this problem for the time being (we explain below how to address it), this result – that M selects the joint message $\tilde{m}(\theta | \eta)$ to point-wise maximize $X(m) - \pi(\theta | \eta)Q(m)$, where $\pi(\theta | \eta) = \theta + \frac{F(\theta|\eta) - \Lambda(\theta|\eta)}{f(\theta|\eta)}$ – has a simple intuitive explanation. When the agent’s outside option payoff is constant, as in the case of delegation, the agent’s participation constraint does not bind for any intra-marginal type θ in M’s side contract problem. Hence $\Lambda(\theta | \eta)$ equals zero (for all intra-marginal types), and $\pi(\theta | \eta)$ reduces to the familiar expression for the ‘virtual’ unit cost $\theta + \frac{F(\theta|\eta)}{f(\theta|\eta)}$ incurred by M in procuring the good from A. Under delegation, this is effectively the unit cost incurred by the coalition, i.e., M, in delivering the good to P, which M is privately informed about (since P does not communicate or contract directly with A). This results in the DMR problem in delegation, i.e., P could do better procuring the item directly from A who incurs the cost θ rather than the higher virtual cost incurred by M.

In centralized contracting, the DMR problem can be ameliorated by P contracting directly with A as well as M, which results in a type dependent participation constraint in the side contracting problem faced by M. Then the Lagrange multiplier $\Lambda(\theta | \eta)$ on the participation constraint can be strictly positive, resulting in a lower ‘shadow’ virtual cost $\pi(\theta | \eta) = \theta + \frac{F(\theta|\eta)}{f(\theta|\eta)} - \frac{\Lambda(\theta|\eta)}{f(\theta|\eta)}$ incurred by M in delivering a unit of the good. Being forced to pay a higher rent to A then curbs M’s inclination to ‘under-procure’ from A in delegation resulting from M’s monopsony power, thereby limiting M’s rent and raising the output delivered to P.

Observe also that if the optimal side contract is degenerate and concentrated at (θ, η) , it must be the case that

$$X(\theta, \eta) - \pi(\theta | \eta)Q(\theta, \eta) \geq X(\tilde{m}') - \pi(\theta | \eta)Q(\tilde{m}')$$

for any $\tilde{m}' \in \Delta(\mathcal{M})$. In other words, the coalition behaves as if it were a single composite agent with a payoff function $X - \pi(\theta | \eta)Q$ in state (θ, η) . This enables us to derive corresponding coalitional incentive and participation constraints:

$$X(\theta, \eta) - \pi(\theta | \eta)Q(\theta, \eta) \geq X(\theta', \eta') - \pi(\theta | \eta)Q(\theta', \eta')$$

$$X(\theta, \eta) - \pi(\theta | \eta)Q(\theta, \eta) \geq 0$$

for any $(\theta, \eta), (\theta', \eta') \in K$.

This discussion motivates the following definition of the coalition shadow cost function.

Definition 4. For any $\eta \in \mathcal{N}$, $Y(\eta)$ is a collection of **coalition shadow cost (CSC)** functions $\pi(\cdot | \eta) : \Theta(\eta) \rightarrow \Re$ which satisfy the following property. For any function in this collection, there exists a real-valued function $\Lambda(\theta|\eta)$ which is non-decreasing in $\theta \in \Theta(\eta)$ with $\Lambda(\underline{\theta}(\eta) | \eta) = 0$ and $\Lambda(\bar{\theta}(\eta) | \eta) = 1$, such that

$$\pi(\theta|\eta) \equiv \theta + \frac{F(\theta | \eta) - \Lambda(\theta | \eta)}{f(\theta | \eta)} \tag{9}$$

We now explain how the analysis is modified if the output monotonicity constraint turns out to bind (which will happen if $\pi(\theta | \eta)$ is decreasing in θ over some interval). In this case we follow the standard ironing procedure of Baron and Myerson (1982): the coalitional shadow cost function $\pi(\theta | \eta)$ is replaced by its ironed version, and we add the requirement that $q(\theta, \eta)$ is flat on any interval where π differs from z .²⁰ The coalitional payoff function is $X - z(\theta | \eta)Q$, and $z(\theta | \eta)$ replaces $\pi(\theta | \eta)$ in the coalition incentive and participation constraints.

²⁰ We provide relevant details of the ironing procedure in the online Appendix.

Definition 5. For any $\eta \in \mathcal{N}$, the set of **coalition virtual cost (CVC)** functions is the set

$$Z(\eta) \equiv \{z(\cdot | \eta) \text{ is the ironed version of some } \pi(\cdot | \eta) \in Y(\eta)\}$$

of functions obtained by applying the ironing procedure to the set $Y(\eta)$ of CSC functions.²¹ Denote by $\Theta(\pi(\cdot | \eta), \eta)$ the corresponding pooling region of θ where $\pi(\cdot | \eta)$ is flattened by the ironing procedure.

The coalitional constraints derived above (combined with A's incentive constraint (2) with $\tilde{m}(\theta | \eta) = (\theta, \eta)$) are therefore necessary for ex ante collusion proofness. They also turn out to be sufficient, as shown in the Appendix. Hence we can represent the EACP property by these constraints. We summarize these results below.

Proposition 3. *The allocation (u_A, u_M, q) is EACP if and only if the following conditions hold for every η . Let $(X(m), Q(m))$ denote the corresponding coalitional contract.*

There exists a CVC function $z(\cdot | \eta) \in Z(\eta)$ such that

(i) *For every $(\theta, \eta), (\theta', \eta') \in K \equiv \{(\theta, \eta) | \theta \in \Theta(\eta), \eta \in \mathcal{N}\}$,*

$$X(\theta, \eta) - z(\theta | \eta)Q(\theta, \eta) \geq X(\theta', \eta') - z(\theta' | \eta)Q(\theta', \eta')$$

$$X(\theta, \eta) - z(\theta | \eta)Q(\theta, \eta) \geq 0.$$

(ii) *For any $\theta, \theta' \in \Theta(\eta)$, $u_A(\theta, \eta) \geq u_A(\theta', \eta) + (\theta' - \theta)Q(\theta', \eta)$.*

(iii) *$Q(\theta, \eta)$ does not vary with θ on any interval over which the CVC function z is flattened by the ironing procedure.²²*

This result shows that the allocation is EACP if there exists a coalitional (shadow) cost function $z(\cdot | \eta)$ for which the coalitional contract satisfies the coalitional incentive and participation constraint (represented by (i)); the payoff function u_A is incentive compatible for A (represented by (ii)); and the output function Q is consistent with the cost function $z(\cdot | \eta)$ in the way required by (iii).

Combined with Proposition 1, this result implies P's problem of designing an optimal contract can be reduced to choosing (a) the Kuhn-Tucker multiplier function $\Lambda(\theta | \eta)$ on the support $[\underline{\theta}(\eta), \bar{\theta}(\eta)]$ which is nondecreasing, taking the value zero at $\underline{\theta}(\eta)$ and 1 at $\bar{\theta}(\eta)$, which determines the shadow cost function $\pi(\theta | \eta)$ (and its ironed version $z(\theta | \eta)$) for the coalition; (b) a coalitional contract $X(\theta, \eta)$, $Q(\theta, \eta)$ satisfying (i) and (iii) above, and (c) the payoff function $u_A(\theta, \eta)$ satisfying (ii). The choice of $X(\cdot)$, $Q(\cdot)$ and $u_A(\cdot)$ functions will determine M's payoff in each state. Finally, in order to be feasible, interim participation constraints must hold for both M and A.

Characterizing features of an optimal EACP contract is difficult in general, except in very special cases.²³ We focus therefore on some qualitative features of an optimal organizational design, such as whether consulting M is profitable, whether collusion imposes costs on P, and some determinants of these costs.

4.3. Value of consulting the expert

The first question that naturally arises is the value to P of consulting the expert M. Delegation to M has already been shown to be suboptimal, relative to not consulting M at all. Does there generally exist a centralized contract involving M which generates a higher profit to P than is possible in the absence of M? We can show this is indeed the case.

Proposition 4. *Consulting M is always valuable, i.e., there exists a centralized contract which enables P to attain a higher expected profit than would be possible in the absence of M.*

We sketch the argument here informally; the formal proof is in the Appendix. Start with the optimal contract for A in the absence of M, which can be solved using standard techniques (e.g., as in Baron and Myerson (1982)): A is paid $X^*(\theta)$ and asked to produce $Q^*(\theta)$ following a cost report of θ , where $Q^*(\theta)$ maximizes $V(Q) - \{\theta + \frac{F(\theta)}{f(\theta)}\}Q$, so

²¹ The ironing procedure ensures these functions are continuous and non-decreasing. For further details, see the online Appendix.

²² Condition (iii) is needed for the following reason. Even though the coalition incentive constraint (i) implies the monotonicity of $Q(\theta, \eta)$ on $\Theta(\eta)$, it does not exclude the possibility that $z(\theta' | \eta) = z(\theta | \eta)$ and $Q(\theta', \eta) < Q(\theta, \eta)$ for some $\theta > \theta'$, associated with choice of $X(\theta, \eta)$ and $X(\theta', \eta)$ such that

$$X(\theta, \eta) - z(\theta | \eta)Q(\theta, \eta) = X(\theta', \eta) - z(\theta' | \eta)Q(\theta', \eta).$$

However if $z(\theta' | \eta) = z(\theta | \eta)$ arises as a result of ironing procedure, θ and θ' must be included in the same bunching interval, which would require that $Q(\theta', \eta) = Q(\theta, \eta)$. In other words, the coalition incentive constraint (i) does not imply (iii).

²³ For instance, in the case where the good is indivisible, it is possible to obtain detailed properties of optimal contracts and compute solutions numerically for particular distributions. See Section 5.3 below.

$$V'(Q^*(\theta)) = \theta + \frac{F(\theta)}{f(\theta)} \equiv H(\theta) \tag{10}$$

while

$$X^*(\theta) = \theta Q^*(\theta) + \int_{\theta}^{\bar{\theta}} Q^*(y) dy. \tag{11}$$

This is a particular EACP contract, with a shadow cost function $\pi(\theta | \eta) = z(\theta | \eta) = \theta$, i.e., where the Kuhn-Tucker multiplier function is $\Lambda(\theta | \eta) = F(\theta | \eta)$, and where $u_A(\theta) = X^*(\theta) - \theta Q^*(\theta)$, $u_M(\theta) = 0$ for all θ .

Given any possible signal η that may be observed by M, Proposition 3 implies that P while consulting M can select an EACP contract corresponding to any non-decreasing shadow/virtual cost function $\pi(\theta | \eta) = z(\theta | \eta)$ for the coalition which deviates from θ in the interior of the support $[\underline{\theta}(\eta), \bar{\theta}(\eta)]$. This can be combined with payments and outputs according to $X(\theta, \eta) = X^*(z(\theta | \eta))$, $Q(\theta, \eta) = Q^*(z(\theta | \eta))$. In other words, it is as if P contracts with a single composite agent whose cost type is $z(\theta | \eta)$ instead of θ – i.e., P is dealing with an agent with a different cost distribution from A when P contracts without consulting M. Corresponding to outputs $Q^*(z(\theta | \eta))$, A’s rent is selected so as to satisfy A’s incentive constraint:

$$u_A(\theta, \eta) = \int_{\theta}^{\bar{\theta}(\eta)} Q^*(z(y | \eta)) dy + \int_{\bar{\theta}(\eta)}^{\bar{\theta}} Q^*(y) dy.$$

If the resulting M’s rent $u_M(\theta, \eta)$ satisfies M’s interim participation constraint

$$E[u_M(\theta, \eta) | \eta] = E[X^*(z(\theta | \eta)) - \theta Q^*(z(\theta | \eta)) - u_A(\theta, \eta) | \eta] \geq 0,$$

$(X(\theta, \eta), Q(\theta, \eta), u_A(\theta, \eta))$ will be feasible under ex ante collusion. We explain now how such a contract can be constructed to generate a higher profit for P.

To simplify the notation, use $\tilde{\theta}(\theta)$ to denote the random variable $z(\theta | \eta)$ in signal state η . This random variable which represents the altered cost type of the supplying agent, has support $[\underline{\theta}(\eta), \bar{\theta}(\eta)]$ and can be expressed as a function of θ . We shall select small random perturbations around A’s cost, where $\tilde{\theta}$ will deviate slightly from θ over some chosen subintervals of $[\underline{\theta}(\eta), \bar{\theta}(\eta)]$. The resulting payments and outputs then can be expressed as a function of θ :

$$x(\theta) \equiv X^*(\tilde{\theta}(\theta)) = \tilde{\theta}(\theta) Q^*(\tilde{\theta}(\theta)) + \int_{\tilde{\theta}(\theta)}^{\bar{\theta}} Q^*(y) dy \tag{12}$$

$$q(\theta) \equiv Q^*(\tilde{\theta}(\theta)) \tag{13}$$

resulting in expected profit for P in state η of

$$\begin{aligned} \Pi(\eta) &\equiv E[V(q(\theta)) - x(\theta) | \eta] \\ &= E[V(Q^*(\tilde{\theta}(\theta))) - \tilde{\theta}(\theta) Q^*(\tilde{\theta}(\theta)) - \int_{\tilde{\theta}(\theta)}^{\bar{\theta}} Q^*(y) dy | \eta] \end{aligned} \tag{14}$$

Consider a small variation $d\tilde{\theta}(\theta)$ in the function $\tilde{\theta}(\theta)$. This induces the following variation in P’s expected profit (14):

$$d\Pi(\eta) = E[\{V'(Q^*(\tilde{\theta}(\theta))) - \tilde{\theta}(\theta)\} dQ^*(\tilde{\theta}(\theta)) | \eta] \tag{15}$$

where $dQ^*(\tilde{\theta}(\theta)) \equiv \frac{\partial Q^*(\tilde{\theta}(\theta))}{\partial \tilde{\theta}} d\tilde{\theta}(\theta)$ is the associated variation in the output function. Using (10), this reduces to

$$d\Pi(\eta) = E[\{H(\tilde{\theta}(\theta)) - \tilde{\theta}(\theta)\} dQ^*(\tilde{\theta}(\theta)) | \eta] = E[\frac{F(\tilde{\theta}(\theta))}{f(\tilde{\theta}(\theta))} dQ^*(\tilde{\theta}(\theta)) | \eta] \tag{16}$$

In words, the rate at which P’s expected profit changes in response to the induced variation in output equals the informational rent $\frac{F(\tilde{\theta}(\theta))}{f(\tilde{\theta}(\theta))}$ of the agent (the inverse hazard rate of P’s prior belief distribution). Intuitively, this reflects the extent of ‘under-production’ or ‘under-procurement’ resulting from the informational rents accruing to the agent. An increase in the output is worth more to P ($V'(Q^*(\tilde{\theta})) = H(\tilde{\theta})$) than what it costs the agent ($\tilde{\theta}$) to deliver.

Now consider the associated variation in M's interim rent in state η . This rent can be expressed as

$$\begin{aligned} \Pi_M(\eta) &\equiv E[u_M(\theta, \eta) \mid \eta] = E[x(\theta) - \theta q(\theta) - u_A(\theta, \eta) \mid \eta] \\ &= E[X^*(\tilde{\theta}(\theta)) - h(\theta \mid \eta) Q^*(\tilde{\theta}(\theta)) \mid \eta] - \int_{\tilde{\theta}(\eta)}^{\tilde{\theta}} Q^*(y) dy \\ &= E[\{\tilde{\theta}(\theta) - h(\theta \mid \eta)\} Q^*(\tilde{\theta}(\theta)) + \int_{\tilde{\theta}(\theta)}^{\tilde{\theta}(\eta)} Q^*(y) dy \mid \eta] \end{aligned} \tag{17}$$

where $h(\theta \mid \eta) \equiv \theta + \frac{F(\theta|\eta)}{f(\theta|\eta)}$, the virtual cost incurred by M in procuring from A 'within the coalition'. Hence the variation induced in M's interim rent is

$$d\Pi_M(\eta) = E[\{\tilde{\theta}(\theta) - h(\theta \mid \eta)\} dQ^*(\tilde{\theta}(\theta)) \mid \eta] = E[\{\tilde{\theta}(\theta) - \theta - \frac{F(\theta \mid \eta)}{f(\theta \mid \eta)}\} dQ^*(\tilde{\theta}(\theta)) \mid \eta] \tag{18}$$

If $\tilde{\theta}(\theta)$ is 'close' to θ , M's interim rent changes with the induced output variation at a rate $\frac{F(\theta|\eta)}{f(\theta|\eta)}$, the informational rent paid by M to A 'within the coalition', which equals the inverse hazard rate of M's posterior beliefs over A's cost θ . Hence a rise in the output delivered induced by a fall in $\tilde{\theta}$ causes M's rent to fall at a rate equal to the inverse hazard rate of $F(\theta \mid \eta)$.

The variation thus causes P's profit and M's rent to move in opposite directions, and at rates which differ owing the difference in their beliefs. This difference creates room for 'gains from trade' between P and M. Since M's signal is informative, there must exist some state η^{**} and costs θ_L, θ_H for which

$$\left[\frac{F(\theta_L \mid \eta^{**})}{f(\theta_L \mid \eta^{**})} \right] / \left[\frac{F(\theta_L)}{f(\theta_L)} \right] \neq \left[\frac{F(\theta_H \mid \eta^{**})}{f(\theta_H \mid \eta^{**})} \right] / \left[\frac{F(\theta_H)}{f(\theta_H)} \right] \tag{19}$$

This implies the existence of output variations dQ_L, dQ_H such that

$$\begin{aligned} \frac{F(\theta_L \mid \eta^{**})}{f(\theta_L \mid \eta^{**})} dQ_L + \frac{F(\theta_H \mid \eta^{**})}{f(\theta_H \mid \eta^{**})} dQ_H &= 0 \\ \frac{F(\theta_L)}{f(\theta_L)} dQ_L + \frac{F(\theta_H)}{f(\theta_H)} dQ_H &> 0 \end{aligned} \tag{20}$$

i.e., P's expected profits rise while M's interim rent is unchanged. Hence it is possible for P to consult M and ask her to report if the state is η^{**} , in which case a variation can be constructed for $\tilde{\theta}(\theta)$ around the identity function $I(\theta) = \theta$ in small neighborhoods of θ_L, θ_H in opposite directions, while restricting the variation to non-decreasing functions $\tilde{\theta}(\theta)$.

4.4. Costs of Ex Ante collusion

Now that we have shown that consulting M does allow P to earn higher profit, the next question is by how much. An upper bound is provided by the second-best profit, which if achievable would mean that ex ante collusion ends up imposes no cost to P. The next result provides sufficient conditions for this to not be possible.

Proposition 5. *With ex ante collusion, the second-best payoff cannot be attained by P if:*

- (i) *The support of θ does not vary with the signal: $\Theta(\eta) = \Theta$ for any $\eta \in \mathcal{N}$;*
- (ii) *There exists $\eta^* \in \mathcal{N}$ such that $f(\theta|\eta^*)$ and $\frac{f(\theta|\eta^*)}{f(\theta|\eta)}$ are both strictly decreasing in θ for any $\eta \neq \eta^*$;*
- (iii) $\underline{\theta} f(\underline{\theta} \mid \eta^*) > 1$;
- (iv) $V'''(q) \leq \frac{(V''(q))^2}{V'(q)}$ for any $q \in \{\tilde{q} \mid \tilde{q} = q^{SB}(\theta, \eta) \text{ for some } (\theta, \eta) \in K\}$.

Condition (ii) includes a weaker version of the monotone likelihood ratio property: there is a signal realization η^* which is 'better' news about θ than any other realization, in the sense of shifting weight in favor of low realizations of θ . It additionally requires that the conditional density $f(\theta|\eta^*)$ is strictly decreasing in θ , i.e., higher realizations of θ are less likely than low realizations when $\eta = \eta^*$. (ii) is satisfied for instance when θ has a uniform prior and there are just two possible signal values satisfying the standard monotone likelihood ratio property. Condition (iii) says that costs are high in the sense that the support of the cost distribution is shifted sufficiently to the right. This condition is equivalent to

$\frac{\theta}{f(\theta|\eta^*)} + \frac{F(\theta|\eta^*)-1}{f(\theta|\eta^*)} > 0$, which guarantees the CVC function (given η^*) is positive-valued on Θ . Finally (iv) is a condition on the benefit function, which is satisfied if V is exponential ($V = V_0[1 - \exp(-rQ)]$; $V_0 > 0, r > 0$).²⁴

The proof develops necessary conditions for the second best to be EAC feasible given the distributional properties (i) and (ii). If the outputs are second-best, they must be a monotone decreasing function of the (ironed) virtual cost $\hat{h}(\theta | \eta)$ in the second-best setting. If they also satisfy the coalition incentive constraints, they must be monotone in CVC $z(\theta | \eta)$. These conditions imply the existence of a monotone transformation from \hat{h} to z , and enable M's ex post rent to be expressed as a function of \hat{h} alone. Conditions (iii) and (iv) are used to show that this rent function is strictly convex; combined with (i) and (ii) this implies that the expected rents of M must be strictly higher (hence strictly positive) in state η^* than any other state. Then M must earn positive rents in state η^* , which ensures the second best cannot be achieved.

5. Extensions

5.1. Side contracts designed by a third party, and alternative allocations of bargaining power

We now explain how the preceding results extend when the side contract is designed not by M, but instead by a third-party that manages the coalition and assigns arbitrary welfare weights to the payoffs of M and A respectively. Such a formulation has been used by a number of authors to model collusion, such as Laffont and Martimort (1997, 2000), Dequiedt (2007) and Celik and Peters (2011). An advantage of this approach is that it enables us to examine effects of varying the allocation of bargaining power between colluding partners.

Our results extend to such a setting, under the following formulation of side contracts designed by a third party. We assume the third-party's objective is to maximize a weighted sum of M and A's interim payoffs. In the subgame C3 following choice of a grand contract by P, the third party designs the side contract after learning the realization of η .²⁵ Both M and A have the option to reject the side contract; if either of them does, they play the grand contract noncooperatively. Otherwise the side contract mechanism is executed.

The notion of EACP allocations is extended as follows. Letting $\alpha \in [0, 1]$ denote the welfare weight assigned by the third-party to A's payoff, the side contract design problem reduces to selecting randomized message $\tilde{m}(\theta | \eta)$ and A's payoff $\tilde{u}_A(\theta, \eta)$ to (using the same notation for the formulation $P(\eta)$ of side contracts in Section 3.4):

$$\max E[(1 - \alpha)\{X(\tilde{m}(\theta | \eta)) - \theta Q(\tilde{m}(\theta | \eta)) - \tilde{u}_A(\theta, \eta)\} + \alpha \tilde{u}_A(\theta, \eta) | \eta]$$

subject to $\tilde{m}(\theta | \eta) \in \Delta(\mathcal{M})$,

$$\tilde{u}_A(\theta, \eta) \geq u_A(\theta, \eta)$$

for all $\theta \in \Theta(\eta)$,

$$\tilde{u}_A(\theta, \eta) \geq \tilde{u}_A(\theta', \eta) + (\theta' - \theta)Q(\tilde{m}(\theta' | \eta))$$

for any $\theta, \theta' \in \Theta(\eta)$, and

$$E[X(\tilde{m}(\theta | \eta)) - \theta Q(\tilde{m}(\theta | \eta)) - \tilde{u}_A(\theta, \eta) | \eta] \geq E[u_M(\theta, \eta) | \eta].$$

Besides modifying the objective function, this formulation adds a participation constraint for M. We refer to this as problem $TP(\eta; \alpha)$. The definition of EACP can be extended to EACP(α) by requiring the null side contract to be optimal in $TP(\eta; \alpha)$ for every η . Further details concerning the justification of this solution concept is provided in the online Appendix.²⁶

We now show that the set of EACP(α) allocations is independent of α . This implies that all our preceding results extend to side contracts designed by a third party.²⁷

Proposition 6. *The set of EACP(α) allocations is independent of $\alpha \in [0, 1]$.*

Despite the existence of asymmetric information within the coalition, the Coase Theorem applies. The reasoning is straightforward. The EACP criterion amounts to the absence of incentive compatible deviations that are Pareto improving for the coalition: this property does not vary with the precise welfare weights. Consider any $\alpha \in (0, 1)$. A given allocation is EACP(α) if and only if there is no other allocation attainable by some non-null side contract which satisfies the incentive

²⁴ This benefit function does not satisfy the Inada conditions assumed in the model. However, the only purpose of imposing the Inada conditions was to ensure that optimal allocations would always involve strictly positive quantities procured in all states. This will be the case here if V_0 is large enough.

²⁵ This assumption can be dropped without affecting the results, since it can be shown the third-party can use cross-reporting of η by M and A to learn its true value.

²⁶ In particular, we show a similar justification applies as in the case where side contracts are offered by M, and we do not need any assumptions regarding off-equilibrium-path beliefs or acceptance of side contracts beyond those incorporated in the PBE(c) notion.

²⁷ FLM provide an analogous result for the case of interim collusion.

constraint for A, and which Pareto-dominates it (for A and M) with at least one of them strictly better off. The same characterization applies to any interior $\alpha' \in (0, 1)$, implying that the set of EACP(α) allocations is independent of $\alpha \in (0, 1)$. The transferability of utility can then be used to show that the set of EACP allocations for interior welfare weights are also the same at the boundary.

5.2. Altruistic expert

Now consider a different variant, where M offers a side-contract to A, but M is altruistic towards A rather than just concerned with his own income. Suppose M's payoff is $u_M = X_M + t + \alpha[X_A - t - \theta q]$, where $\alpha \in [0, 1]$ is the weight M places on A's payoff. A's payoff function remains the same as in the previous section: $u_A = X_A - t - \theta q$.

Our analysis extends as follows. The expression for coalition shadow cost is now modified to

$$\pi_\alpha(\theta|\eta) \equiv \theta + (1 - \alpha) \frac{F(\theta | \eta) - \Lambda(\theta | \eta)}{f(\theta | \eta)}$$

instead of $\pi(\theta|\eta)$ in Definition 4. When P delegates to M, the corresponding expression for the cost of procuring one unit from M is modified from $h(\theta | \eta)$ to $h_\alpha(\theta | \eta) = \theta + (1 - \alpha) \frac{F(\theta|\eta)}{f(\theta|\eta)}$. As long as $\alpha < 1$, this is strictly higher than θ , so delegation will still continue to result in a lower profit than no monitor case (NM). The proof that M is valuable under centralized contracting also goes through *in toto*.

It is interesting to examine the effect of changes in the degree of altruism on P's payoffs. When P delegates to M, an increase in α lowers M's shadow cost of output $h_\alpha(\theta | \eta)$, which benefits P. This is intuitive: the DMR problem becomes less acute with a more altruistic expert. Note that with perfect altruism $\alpha = 1$, and the DMR problem disappears: delegation then becomes equivalent to NM.

On the other hand, an increase in altruism cannot benefit P in centralized contracting. The set of EACP allocations can be shown to be non-increasing in α . Take any EACP allocation corresponding to α : the following argument shows that it is a EACP allocation corresponding to any $\alpha' < \alpha$. Let $z(\theta | \eta)$ be the CVC function that is associated with the allocation at α , i.e., it is the ironed version of $\pi_\alpha(\theta|\eta)$ corresponding to some function $\Lambda_\alpha(\cdot|\eta)$ satisfying the stipulated requirements in the definition of CSC functions on $[\underline{\theta}(\eta), \bar{\theta}(\eta)]$. We can then select

$$\Lambda_{\alpha'}(\theta | \eta) = \frac{\alpha - \alpha'}{1 - \alpha'} F(\theta | \eta) + \frac{1 - \alpha}{1 - \alpha'} \Lambda_\alpha(\theta | \eta)$$

when the altruism parameter is α' , which satisfies the stipulated requirements since $\alpha > \alpha'$. This ensures that the same CSC and CVC function is available when the altruism parameter is α' , since by construction $\pi_\alpha(\theta|\eta) = \pi_{\alpha'}(\theta|\eta)$. Hence the allocation satisfies the sufficient condition for EACP when the altruism parameter is α' .

Finally, if $\alpha = 1$, the CSC function π_α coincides with the identity function θ , the cost of the agent in NM. We thus obtain

Proposition 7. *In centralized contracting, P's optimal payoff is non-increasing in α . When P delegates to M, P's optimal payoff is increasing in α . When $\alpha = 1$, P's optimal payoffs in delegation, centralized contracting coincide and equal that in NM, so M is not valuable in that case.*

5.3. Procuring an indivisible good

In the case where the good to be procured is indivisible, it is possible to extend the analysis and obtain more detailed properties of optimal contracts with ex ante collusion, including a comparison with optimal contracts in the interim collusion setting. Details of this case are available in the online Appendix. We report the main assumptions and results here.

The quantity to be procured is $q \in \{0, 1\}$. P obtains a zero gross benefit if $q = 0$, and a benefit of $V > 0$ if $q = 1$. A's cost of delivering the good is θ which is distributed on an interval $[\underline{\theta}, \bar{\theta}]$. M receives either of two possible signals $\eta_i, i = L, H$, where η_L (resp. η_H) is a low (resp. high) cost signal in the following sense. $F_i(\theta)$ denotes the cdf of θ conditional on signal η_i , with a density $f_i(\theta)$ which is throughout positive and continuous on $[\underline{\theta}, \bar{\theta}]$, where the likelihood ratio $\frac{f_L(\theta)}{f_H(\theta)}$ is decreasing in θ .²⁸ $\kappa_i \in (0, 1)$ denotes the prior probability of signal η_i , so P's prior beliefs are given by the cdf $F = \kappa_L F_L + \kappa_H F_H$ with density $f = \kappa_L f_L + \kappa_H f_H$. We impose some regularity conditions to simplify the analysis: the functions $H(\theta) = \theta + \frac{F(\theta)}{f(\theta)}$, $h_i(\theta) = \theta + \frac{F_i(\theta)}{f_i(\theta)}$, $l_i(\theta) = \theta - \frac{1 - F_i(\theta)}{f_i(\theta)}$ ($i = L, H$) are increasing, and $l_L(\theta)$ is convex. In addition $h_L(\bar{\theta}) > V > \underline{\theta}$ to ensure there is positive supply and M is valuable in the second-best situation.²⁹

These assumptions hold in the case of linear conditional densities $F_L = 2d\theta - (2d - 1)\theta^2$, $F_H = 2(1 - d)\theta + (2d - 1)\theta^2$ on $[0, 1]$, with $d \in (1/2, 1)$, $\kappa_L = \kappa_H = \frac{1}{2}$ which imply a uniform prior, and with $0 < V < h_L(\bar{\theta}) = 1 + \frac{1}{2(1-d)}$. For this family we

²⁸ This also implies that $F_L(\theta) > F_H(\theta)$ on $(\underline{\theta}, \bar{\theta})$ and $\frac{f_L(\theta)}{f_L(\bar{\theta})} > \frac{f_H(\theta)}{f_H(\bar{\theta})}$ on $(\underline{\theta}, \bar{\theta})$.

²⁹ In the second best situation without collusion, optimal price (denoted by p_i^{SB}) maximizes $F_i(p_i)(V - p_i)$ in η_i . Then p_i^{SB} satisfies $h_i(p_i^{SB}) = V$ for $V \in [\underline{\theta}, h_i(\bar{\theta})]$ and $p_i^{SB} = \bar{\theta}$ for $V > h_i(\bar{\theta})$.

are able to compute optimal contracts both with ex ante and interim collusion, properties of which are shown below for the case of $d = 0.99$.

A coalitional contract in this setting is represented by aggregate payments $b + X_0$, X_0 paid by P to the coalition of M and A in the event that the good is, and is not delivered. The payoff function for A is represented by ‘prices’ p_H, p_L offered by P to A in the two states H, L respectively. In the optimal contract, P ends up offering the same price p_i to A in η_i . Hence following a joint report of (θ, η_i) which results in $q = 1$ (resp. $q = 0$), the grand contract pays A p_i (resp. 0) while M receives $b + X_0 - p_i$ (resp. X_0). X_0 can be interpreted as a lump-sum payment made by P to the coalition irrespective of the quantity delivered, and b an additional bonus in the event the good is delivered. Under ex ante collusion, the coalitional participation constraint reduces to $X_0 \geq 0$. In interim collusion, X_0 is permitted to be negative, i.e., where the coalition ends up paying P in the event the good is not delivered.

Under ex ante collusion, the following properties of the optimal contract can be derived analytically. We use superscripts SB, E, I to describe optimal contracts in the second-best, ex ante collusion and interim collusion settings respectively.

- (a) $p_L^{SB} \leq p_L^E < p_H^E \leq p_H^{SB}$ and $X_0^E > 0, b^E < p_i^E$ for some $i = L, H$ if M is valuable in ex ante collusion.
- (b) For V sufficiently small (but larger than $\underline{\theta}$), the second-best outcome is achievable under both ex ante and interim collusion.
- (c) For V sufficiently large (but smaller than $h_L(\bar{\theta})$), M is not valuable in ex ante collusion.
- (d) Whenever the optimal ex ante and interim collusion contracts diverge, the latter contract involves $b^I > p_H^I > p_L^I, X_0^I < 0$, and can be implemented via delegation.

The first part of Property (a) states that the dispersion of prices offered to A is smaller in ex ante collusion compared to the second-best setting unless the latter is achievable. This illustrates how the prospect of collusion limits the extent to which prices offered to A vary with M’s signal. In the second-best setting, a low cost signal motivates P to offer A a lower price ($p_L^{SB} < p_H^{SB}$), owing to a lower price elasticity of supply (since $\frac{\theta f_L(\theta)}{F_L(\theta)} < \frac{\theta f_H(\theta)}{F_H(\theta)}$ for all θ). A large gap between these prices potentially motivates A to bribe M to report a high rather than low cost signal. Hence prices vary less with the cost signal in the presence of collusion. The second part of (a) states that M is provided lower-powered incentives to deliver the good compared with A following some signal η_i – i.e., the bonus paid to the coalition as a whole is lower than the price offered to A, so M is worse off *ex post* if A decides to deliver the good. To compensate for this loss, the coalition is paid a positive amount in the event that the good is not delivered $X_0^E > 0$. In other words, ‘countervailing’ incentives are used by P to control collusion: M and A have opposite *ex post* delivery incentives in some states of the world.

Property (b) states that the second-best is attainable with ex ante (and therefore interim) collusion if P’s value V is small. Intuitively the reason is that the associated dispersion in second-best prices is small enough that no incentives for collusion are actually generated. Property (c) states that for high enough values of V consulting M is not valuable. With high enough V the stakes of collusion become ‘too large’ and P ends up with a pooled contract where prices offered to A do not vary with M’s signal. This result differs from what we found when the good is divisible and P’s benefit function satisfies Inada conditions, where P additionally has scope to vary the quantity of the good delivered with respect to submitted cost reports. With an indivisible good, the quantity of the good is either zero or one, so M’s cost report is less valuable.

Finally, (d) states that when the optimal contract with interim collusion differs from that under ex ante collusion, it ends up providing higher powered delivery incentives to M compared to A following each possible signal, whence M has a stronger stake in delivering the good. At the same time the constant payment X_0^I is negative, so M ends up paying P in the event the good is not delivered. This resembles a ‘franchise’ contract where M pays a fixed fee to P and ends up worse off than non-participation in the mechanism if the good is not delivered owing to a high cost state for A. This franchise contract is implementable via delegation to M, and is obviously not feasible with ex ante collusion.

The optimal contracts in the case of a uniform prior and linear conditional densities described above for the case of $d = 0.99$ are illustrated in Figs. 1–3 below. Fig. 1 contrasts the second-best contract with optimal ex ante collusion contract. For a range of low values of V , the two coincide. As V rises, the dispersion in prices offered to A rises. For an intermediate range of values of V from slightly below 1 until $V = 2$, the two contracts diverge. The dispersion between prices offered to A tends to narrow in the optimal ex ante collusion contract, as the price offered in the low signal state rises faster than the one in the high signal state. At the same time the aggregate bonus b^E rises faster, thereby narrowing the extent of ‘countervailing’ or ‘opposing’ incentives between M and A. The dispersion in prices converges to zero as V converges to 2; thereafter M ceases to be valuable.

The corresponding optimal contract with interim collusion is shown in Figs. 2 and 3. Fig. 2 plots the optimal prices offered to A against values of V : these coincide with the optimal prices under ex ante collusion for values of V below some threshold between 1.5 and 2. Beyond this threshold, the optimal contract with interim collusion switches discontinuously to one with higher dispersion in prices offered to A, and the maximum price of 1 offered following report of a high cost signal, which guarantees delivery of the good. Fig. 3 shows that the bonus b^I jumps upwards at the threshold above the prices offered to A. M continues to be valuable for all values of V shown under interim collusion. Effectively, P combats interim collusion with either of two contrasting contracts. One is a ‘low powered’ contract with low b and positive X_0 . The other is a ‘high powered’ contract with high b and negative X_0 , resembling a franchise contract. The switch of the optimal contract from the former to the latter arises at some threshold of V , which is reflected in the jumps that appear in Figs. 2




Fig. 1. Optimal Prices and Bonus with Ex Ante Collusion.




Fig. 2. Optimal Prices with Interim Collusion.




Fig. 3. Optimal Bonus in Ex Ante and Interim Collusion.

and 3. Only the low powered contract is feasible in ex ante collusion; hence the optimal contracts diverge between the ex ante and interim collusion settings when P obtains a high benefit from the good.

6. Concluding comments

Our results have interesting implications for hierarchical contract design. They provide a rationale for the widespread practice of consulting third party experts in the design of incentive contracts, even when they have ‘prior connections’ with the agent that could facilitate ex ante collusion. In such circumstances, unconditional delegation is suboptimal; the mechanism must allow the agent to communicate directly with the Principal. The Principal could appoint the expert as a ‘manager’, i.e., contract only with the expert and delegate subcontracting with the agent on the equilibrium path, but the agent must be provided the option to bypass the manager and register an ‘appeal’ with the Principal, prompting the latter to intervene directly. The existence of such off-equilibrium-path options is essential for the Principal to be able to control the prospect of ex ante collusion sufficiently to make it profitable to consult the expert. Within firms, it explains the role of worker rights to appeal the evaluations reported by their managers to higher level managers or an ombudsman ap-

pointed for this purpose, thereby formalizing Williamson’s (1975) claim that such dispute settlement procedures constitute an advantage of hierarchies over market relationships.³⁰

The result concerning effects of altruism of M towards A implies that the Principal ought to appoint ‘external’ experts rather than ‘insiders’ likely to be altruistic towards the agent. In the context of corporate governance, for instance, this is an argument in favor of appointing ‘outsiders’ rather than ‘insiders’ to a company’s Board of Directors.³¹ In the context of regulation, it confirms the normal intuition in favor of preventing any direct conflict of interest for the regulator (e.g., who should not have a financial stake in the agent’s fortunes, nor have any social or personal connections with the agent).

Extensions of the model to bilateral asymmetric information within the coalition (e.g., if A does not observe M’s signal) and to discrete type spaces are examined respectively in Tsumagari (2016a, 2016b). Mookherjee and Tsumagari (2017) show that the allocation of bargaining power between M and A does matter in the case of ‘strong’ collusion, e.g., where the side contract includes commitments regarding subsequent actions by one partner if the other refuses it.

Our analysis is subject to a number of shortcomings. We excluded the possibility of other coalitions that may co-exist with the M-A coalition, a topic studied by Ortner and Chassang (2018). We also ignored the enforcement of side contracts within the coalition; modeling self-enforcing collusion via a relational contract in a side game between colluding parties seems to be an interesting extension that could be pursued in future research.

Acknowledgments

For helpful comments and discussions we thank three referees and the Advisory Editor of this journal, Jacopo Bizotto, Yeon-Koo Che, Christian Hellwig, Bart Lipman, Juan Ortner, Yutaka Suzuki and seminar participants at Universities of Cambridge, New South Wales, Washington, Kellogg School of Management, Sloan School of Management, the Theory Workshop at Boston University and the June 2014 Ravello Workshop. Tsumagari acknowledges funding from Keio Gijuku Academic Development Funds.

Appendix A. Proofs

Proof of Proposition 1. *Proof of Necessity.*

Suppose (u_A, u_M, q) is EAC feasible. Then it obviously satisfies interim participation constraints of A and M, since each has the option of unilaterally opting out of the mechanism. Hence we need to show that it is also a EACP allocation. Suppose not. Then there exists $\eta \in \mathcal{N}$ such that $(\tilde{m}(\theta | \eta), \tilde{u}_A(\theta, \eta)) = ((\theta, \eta), u_A(\theta, \eta))$ does not solve the side-contracting problem $P(\eta)$. Suppose that $(\tilde{m}^*(\theta | \eta), \tilde{u}_A^*(\theta, \eta))$ is a solution of $P(\eta)$. Defining

$$\tilde{u}_M^*(\theta, \eta) \equiv X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - \tilde{u}_A^*(\theta, \eta),$$

we must have

$$E[\tilde{u}_M^*(\theta, \eta) | \eta] > E[u_M(\theta, \eta) | \eta]$$

and

$$\tilde{u}_A^*(\theta, \eta) \geq u_A(\theta, \eta)$$

for any $\theta \in \Theta(\eta)$. Since (u_A, u_M, q) is an achievable allocation, this means there exists a grand contract GC and an associated PBE(c) which results in this allocation. Let this PBE involve beliefs $b(\eta)$ and non-cooperative equilibrium $c(\eta) = (m_A(\cdot, \eta), m_M(\eta))$ of GC based on M’s beliefs $b(\eta)$ resulting if A rejects the side contract $SC(\eta)$ offered on the equilibrium path. The payoff accruing to A in this noncooperative equilibrium then cannot exceed $u_A(\theta, \eta)$ in any state (θ, η) , otherwise A would find it optimal to deviate from the behavior prescribed in the PBE(c).

For $\tilde{m}^*(\theta | \eta) \in \Delta(K \cup e)$, there exists $\tilde{m}^c(\theta, \eta) \in \Delta(\mathcal{M}_A \times \mathcal{M}_M)$ such that

$$(X(\tilde{m}^*(\theta | \eta)), Q(\tilde{m}^*(\theta | \eta))) = (X_A(\tilde{m}^c(\theta, \eta)) + X_M(\tilde{m}^c(\theta, \eta)), q(\tilde{m}^c(\theta, \eta))).$$

Given GC and η , consider the side-contract $SC^c(\eta)$ in which the report to P is selected according to $\tilde{m}^c(\theta', \eta)$ on the basis of A’s report of $\theta' \in \Theta(\eta)$, associated with the transfer to A:

$$t_A^c(\theta', \eta) = \tilde{u}_A^*(\theta', \eta) - [X_A(\tilde{m}^c(\theta', \eta)) - \theta' q(\tilde{m}^c(\theta', \eta))].$$

Now construct a different Perfect Bayesian Equilibrium (PBE) which differs from the previous one only in state η , where on the equilibrium path M offers instead $SC^c(\eta)$, and this is accepted by all types of A. Rejection of this offer results in the

³⁰ It also relates to Hirschman’s (1970) depiction of the value of ‘voice’ within organizations over and above exit options.

³¹ See Harris and Raviv (2008) for a model based on limited commitment by P where this result may not hold in some settings.

same noncooperative equilibrium $c(\eta)$ of the grand contract. What occurs in the continuation of any other side contract offer remains the same as in the previous PBE. To check this is a PBE, note that it is optimal for A to accept $SC^c(\eta)$, and then report truthfully. Moreover, given that this side contract is accepted by all types of A, it is optimal for M to offer it (since offering $SC(\eta)$ was optimal in state η in the previous PBE).

Hence $(\tilde{u}_A^*(\theta, \eta), \tilde{u}_M^*(\theta, \eta))$ can be realized as a PBE outcome. Since M is better off without making any type of A worse off, it contradicts the hypothesis that (u_A, u_M, q) is realized as the outcome of a PBE(c).

Proof of Sufficiency.

Step 1: Construction of grand contract.

Suppose that (u_A, u_M, q) is a EACP allocation satisfying interim participation constraints. We show that there exists a grand contract which achieves (u_A, u_M, q) as a PBE(c) outcome.

The grand contract is constructed as follows:

$$GC = (X_A(m_A, m_M), X_M(m_A, m_M), q(m_A, m_M) : \mathcal{M}_A, \mathcal{M}_M)$$

where

$$\mathcal{M}_A = K \cup \{e_A\}$$

$$\mathcal{M}_M = \mathcal{N} \cup \{e_M\}$$

$$X_A(e_A, m_M) = X_M(e_A, m_M) = q(e_A, m_M) = X_M(m_A, e_M) = 0$$

for any (m_A, m_M) .

- $(X_A((\theta_A, \eta_A), \eta_M), q((\theta_A, \eta_A), \eta_M)) = (u_A(\theta_A, \eta_M) + \theta_A q(\theta_A, \eta_M), q(\theta_A, \eta_M))$
- $X_M((\theta_A, \eta_A), \eta_M) = u_M(\theta_A, \eta_A)$ for $\eta_M = \eta_A$ and $X_M((\theta_A, \eta_A), \eta_M) = -T$ for $\eta_M \neq \eta_A$, with T sufficiently large
- $(X_A((\theta_A, \eta_A), e_M), q((\theta_A, \eta_A), e_M)) = (X(\tilde{m}^*(\theta_A)), Q(\tilde{m}^*(\theta_A)))$ where $\tilde{m}^*(\theta)$ maximizes $X(\tilde{m}) - \theta Q(\tilde{m})$ subject to $\tilde{m} \in \Delta(K \cup \{e\})$ and the definition of $(X(\tilde{m}), Q(\tilde{m}))$ is provided in Section 3.4.

Step 2: Non-cooperative equilibrium.

First we argue $(m_A(\theta, \eta), m_M(\eta)) = ((\theta, \eta), \eta)$ is a non-cooperative equilibrium of the grand contract based on prior beliefs (denoted by $b_\phi(\eta)$) for η . EACP and A's participation constraint imply that A always has an incentive to participate and report truthfully: $m_A(\theta, \eta) = (\theta, \eta)$. Since M's interim participation constraint ($E[u_M(\theta, \eta) | \eta] \geq 0$) holds, taking A's strategy $m_A(\theta, \eta) = (\theta, \eta)$ as given, M also has an incentive to participate and report truthfully.

This equilibrium results in allocation $(u_A(\theta, \eta), u_M(\theta, \eta), q(\theta, \eta))$. By offering a null side-contract, M can always realize the allocation $(u_A(\theta, \eta), u_M(\theta, \eta), q(\theta, \eta))$ and achieve interim payoff $E[u_M(\theta, \eta) | \eta]$, because M's beliefs regarding A's type will remain unchanged in that case. Therefore M would have an incentive to offer a non-null side-contract only if the deviation results in a higher payoff. We show that there exists a PBE of (C3) following the GC constructed above, in which M's interim payoff from any deviating side contract offer cannot exceed $E[u_M(\theta, \eta) | \eta]$.

Consider any deviating side contract offer in state η , and let $b(\eta)$ denote beliefs of M regarding θ which result following rejection of this side contract by A. M and A then play GC noncooperatively with beliefs $b(\eta)$. By construction, A has an incentive to report truthfully and participate in GC irrespective of what M does, i.e., irrespective of the beliefs $b(\eta)$ held by M (as well as irrespective of the particular deviating side contract offered). If T is sufficiently large, it is then a best response for M to report truthfully, conditional on participating. We focus on PBEs satisfying these two properties following rejection by A of any deviating side contract.

In what follows, there are two cases to consider. (a) $E_{b(\eta)}[u_M(\theta, \eta)] \geq 0$, in which case it is a best response for M to participate (and report truthfully) in GC when it is played noncooperatively with beliefs $b(\eta)$. We refer to this as the T case. (b) $E_{b(\eta)}[u_M(\theta, \eta)] < 0$, whereby M exits from GC following rejection of the side contract and attains zero payoff. We refer to this as the E case.

Step 3: Side-contract choice.

Now we argue that without loss of generality, the choice of deviating side contract can be limited to those where in every state (θ, η) : either A and M both participate and submit consistent reports $\eta_A = \eta_M$, or they both exit. That they should submit consistent reports conditional on joint participation, follows if T is sufficiently large. Suppose there is some state in which the side contract prescribes an exit for M alone. Given the construction of the grand contract, for any $(m_A, m_M) = ((\theta, \eta), e_M)$, there exists $\tilde{m}' \in \Delta(\mathcal{M}_A \times \mathcal{M}_M \setminus \{(\theta, \eta), e_M\})$ such that

$$(X_A(m_A, m_M) + X_M(m_A, m_M), q(m_A, m_M)) = (X_A(\tilde{m}') + X_M(\tilde{m}'), q(\tilde{m}')),$$

given

$$(X_A((\theta, \eta), e_M) + X_M((\theta, \eta), e_M), q((\theta, \eta), e_M)) = (X(\tilde{m}^*(\theta_A)), Q(\tilde{m}^*(\theta_A)))$$

and the definition of (X, Q) . Therefore \tilde{m}' and $(m_A, m_M) = ((\theta, \eta), e_M)$ generate the same total payment and output target for the coalition. A similar argument ensures that outcomes involving exit for A alone can be eliminated without loss of generality, since $(m_A, m_M) = (e_A, \eta)$ generates the same outcome $X_A = X_M = q = 0$ in the GC as $(m_A, m_M) = (e_A, e_M)$. With the abuse of the notation, let $\hat{\mathcal{M}}$ represent $\{((\theta, \eta), \eta) \mid (\theta, \eta) \in K\} \times \{e_A, e_M\}$.

Step 4: Continuation payoffs following non-null side-contract.

Suppose that M offers some non-null side-contract SC for η , which is described as $(\tilde{m}(\theta, \eta), \tilde{u}_A(\theta, \eta))$ which satisfies

$$\tilde{u}_A(\theta, \eta) \geq \tilde{u}_A(\theta', \eta) + (\theta' - \theta)q(\tilde{m}(\theta', \eta))$$

for any $\theta, \theta' \in \Theta(\eta)$ and $\tilde{m}(\theta, \eta) \in \Delta(\hat{\mathcal{M}})$. Let $\kappa^*(\theta) \in [0, 1]$ denote the probability that $\theta \in \Theta(\eta)$ accepts SC. We focus on PBE's with the property that A reports truthfully to M conditional on accepting the SC. The inequality above ensures that this is optimal for A. In any such PBE, the payoff resulting for M when A accepts the SC equals (in state (θ, η)):

$$X_A(\tilde{m}(\theta, \eta)) + X_M(\tilde{m}(\theta, \eta)) - \theta q(\tilde{m}(\theta, \eta)) - \tilde{u}_A(\theta, \eta).$$

If A rejects SC, A and M play the grand contract non-cooperatively with belief $b^*(\eta)$, which is consistent with Bayes rule as required in a PBE. Sequential rationality of A's participation decision $\kappa^*(\theta)$, given beliefs $b^*(\eta)$ and the non-cooperative equilibrium associated with $b^*(\eta)$, implies the following. In the T-case, $\kappa^*(\theta) = 0$ (or 1 or $\in [0, 1]$) if and only if $u_A(\theta, \eta) >$ (or $<$ or $=$) $\tilde{u}_A(\theta, \eta)$. A ends up with payoff

$$\max\{u_A(\theta, \eta), \tilde{u}_A(\theta, \eta)\},$$

and M's interim payoff is

$$E[\kappa^*(\theta)\{X_A(\tilde{m}(\theta, \eta)) + X_M(\tilde{m}(\theta, \eta)) - \theta q(\tilde{m}(\theta, \eta)) - \tilde{u}_A(\theta, \eta)\} + (1 - \kappa^*(\theta))u_M(\theta, \eta) \mid \eta].$$

Conversely, in the E-case, $\kappa^*(\theta) = 0$ (or 1 or $\in [0, 1]$) if and only if

$$X(\tilde{m}^*(\theta_A)) - \theta Q(\tilde{m}^*(\theta_A)) > (\text{or } < \text{ or } =) \tilde{u}_A(\theta, \eta).$$

A's payoff is

$$\max\{X(\tilde{m}^*(\theta_A)) - \theta Q(\tilde{m}^*(\theta_A)), \tilde{u}_A(\theta, \eta)\}$$

while M's interim payoff is

$$E[\kappa^*(\theta)\{X_A(\tilde{m}(\theta, \eta)) + X_M(\tilde{m}(\theta, \eta)) - \theta q(\tilde{m}(\theta, \eta)) - \tilde{u}_A(\theta, \eta)\} \mid \eta].$$

Step 5: Upper bound on M's interim payoff in continuation play following non-null side-contract.

Here we establish an upper bound of M's interim payoff in PBE of the continuation game for non-null side-contract.

(i) T-Case.

Consider the following problem: select $(\hat{m}(\theta, \eta), \hat{u}_A(\theta, \eta))$ to

$$\max E[X_A(\hat{m}(\theta, \eta)) + X_M(\hat{m}(\theta, \eta)) - \theta q(\hat{m}(\theta, \eta)) - \hat{u}_A(\theta, \eta) \mid \eta]$$

subject to $\hat{m}(\theta, \eta) \in \Delta(\hat{\mathcal{M}})$,

$$\hat{u}_A(\theta, \eta) \geq \hat{u}_A(\theta', \eta) + (\theta' - \theta)q(\hat{m}(\theta', \eta))$$

for any $\theta, \theta' \in \Theta(\eta)$ and

$$\hat{u}_A(\theta, \eta) \geq u_A(\theta, \eta)$$

for any $\theta \in \Theta(\eta)$.

This is equivalent to problem $P(\eta)$ used to characterize EACP allocations. The EACP property implies that $(\hat{m}(\theta, \eta), \hat{u}_A(\theta, \eta)) = ((\theta, \eta), \eta), u_A(\theta, \eta))$ solves this problem and the maximum value is $E[u_M(\theta, \eta) | \eta]$.

We now show that this is an upper bound on M's interim payoff from the deviating side contract in the T-case. Suppose that non-null side-contract $(\tilde{m}(\theta, \eta), t(\theta, \eta))$ is associated with acceptance probability $\kappa^*(\cdot)$ and the T-case applies. Select $(\hat{m}(\theta, \eta), \hat{u}_A(\theta, \eta))$ as follows:

$$\hat{m}(\theta, \eta) = \kappa^*(\theta)\tilde{m}(\theta, \eta) + (1 - \kappa^*(\theta))I(\theta, \eta)$$

and

$$\hat{u}_A(\theta, \eta) = \max\{\tilde{u}_A(\theta, \eta), u_A(\theta, \eta)\}$$

where $I(\theta, \eta)$ is the probability measure concentrated on (θ, η) . In this allocation, A earns exactly the same payoffs as in the continuation following offer of side-contract $(\tilde{m}(\theta, \eta), t(\theta, \eta))$. Hence the agent's incentive constraint is satisfied, and so is the participation constraint by construction. Hence the continuation play following offer of side-contract $(\tilde{m}(\theta, \eta), t(\theta, \eta))$ results in an interim payoff for M which cannot exceed $E[u_M(\theta, \eta) | \eta]$.

(ii) E-Case.

Now consider the following problem: select $(\hat{m}(\theta, \eta), \hat{u}_A(\theta, \eta))$ to

$$\max E[X_A(\hat{m}(\theta, \eta)) + X_M(\hat{m}(\theta, \eta)) - \theta q(\hat{m}(\theta, \eta)) - \hat{u}_A(\theta, \eta) | \eta]$$

subject to $\hat{m}(\theta, \eta) \in \Delta(\mathcal{M})$,

$$\hat{u}_A(\theta, \eta) \geq \hat{u}_A(\theta', \eta) + (\theta' - \theta)q(\hat{m}(\theta', \eta))$$

and

$$\hat{u}_A(\theta, \eta) \geq X(\tilde{m}^*(\theta)) - \theta Q(\tilde{m}^*(\theta)).$$

In order to derive the solution of this problem, consider the problem of maximizing

$$X_A(\hat{m}) + X_M(\hat{m}) - \theta q(\hat{m})$$

subject to $\hat{m} \in \Delta(\mathcal{M})$. Denoting its solution by $\hat{m}^*(\theta)$, we have

$$X_A(\hat{m}^*(\theta)) + X_M(\hat{m}^*(\theta)) - \theta q(\hat{m}^*(\theta)) = X(\tilde{m}^*(\theta)) - \theta Q(\tilde{m}^*(\theta)),$$

because of the definition of (X, Q) and $\tilde{m}^*(\theta)$. Therefore in the above problem, an upper bound of objective function is given by

$$E[X_A(\hat{m}^*(\theta)) + X_M(\hat{m}^*(\theta)) - \theta q(\hat{m}^*(\theta)) - \{X(\tilde{m}^*(\theta)) - \theta Q(\tilde{m}^*(\theta))\} | \eta] = 0$$

This upper bound can be achieved by selecting

$$(\hat{m}(\theta, \eta), \tilde{u}_A(\theta, \eta)) = (\hat{m}^*(\theta), X(\tilde{m}^*(\theta)) - \theta Q(\tilde{m}^*(\theta))).$$

Since this also satisfies all the constraints of the problem, this is a solution of the problem. It follows that the maximum value is equal to zero.

Next check that this maximum value provides an upper bound on M's payoff in the continuation play following the offer of the deviating side contract $(\tilde{m}(\theta, \eta), t(\theta, \eta))$ in which the E-case arises. Select $(\hat{m}(\theta, \eta), \hat{u}_A(\theta, \eta))$ as follows:

$$\hat{m}(\theta, \eta) = \kappa^*(\theta)\tilde{m}(\theta, \eta) + (1 - \kappa^*(\theta))\tilde{m}^*(\theta)$$

and

$$\hat{u}_A(\theta, \eta) = \max\{\tilde{u}_A(\theta, \eta), X(\tilde{m}^*(\theta)) - \theta Q(\tilde{m}^*(\theta))\}.$$

This generates the same payoffs for A as in the continuation play following the offer of the deviating side contract $(\tilde{m}(\theta, \eta), t(\theta, \eta))$, and is therefore feasible in the maximization problem above. Hence zero is an upper bound to M's interim expected payoff when the E-case applies.

Step 6: PBE in collusion game.

We can construct a PBE in the overall collusion game as follows. If M offers null side-contract, he receives $E[u_M(\theta, \eta) | \eta]$. If M offers any non-null side-contract, it follows from Step 5 that his subsequent continuation payoff is not larger than $E[u_M(\theta, \eta) | \eta]$. Since $E[u_M(\theta, \eta) | \eta] \geq 0$, there exists a PBE in which M offers a null side-contract on the equilibrium path, resulting in allocation $(u_A(\theta, \eta), u_M(\theta, \eta), q(\theta, \eta))$.

Step 7: Check PBE(c) property.

Finally check that PBE constructed in the above argument is also a PBE(c). Otherwise there would exist a PBE resulting in a Pareto superior allocation for the coalition. This would violate the EACP property of the allocation we started with. \square

Proof of Proposition 2. At the first step, note that the optimal side contract problem for M in DM involves an outside option for A which is identically zero. This reduces to a standard problem of contracting with a single agent with adverse selection and an outside option of zero, where M has a prior distribution $F(\theta|\eta)$ over the agent's cost θ in state η . The expected procurement cost incurred by M is then $\hat{h}(\theta|\eta)$.

Given this, P's contract with M in DM is effectively a contracting problem for P with a single supplier whose unit supply cost is $\hat{h}(\theta|\eta)$. P's prior over this supplier's cost is given by distribution function $G(h) \equiv \Pr(\theta|\eta | \hat{h}(\theta|\eta) \leq h)$ for $h \geq \min_{\eta} \hat{h}(\theta|\eta) = \underline{\theta}$ and $G(h) = 0$ for $h < \underline{\theta}$. Let $G(h|\eta)$ denote the cumulative distribution function of $h = \hat{h}(\theta|\eta)$ conditional on η : $G(h|\eta) \equiv \Pr(\theta | \hat{h}(\theta|\eta) \leq h, \eta)$ for $h \geq \hat{h}(\underline{\theta}|\eta) (= \underline{\theta}(\eta))$ and $G(h|\eta) = 0$ for $h < \underline{\theta}(\eta)$. Then $G(h) = \sum_{\eta \in \mathcal{N}} p(\eta)G(h|\eta)$. Since $\hat{h}(\theta|\eta)$ is continuous and nondecreasing on $\Theta(\eta)$, $G(h|\eta)$ is strictly increasing in h on $[\underline{\theta}(\eta), \hat{h}(\bar{\theta}|\eta)]$. However, $G(h|\eta)$ may fail to be left-continuous.

Hence P's problem in DM reduces to $\max E_h[V(q(h)) - X(h)]$ subject to $X(h) - hq(h) \geq X(h') - hq(h')$ for any $h, h' \in [\underline{\theta}, \bar{h}]$ and $X(h) - hq(h) \geq 0$ for any $h \in [\underline{\theta}, \bar{h}]$ where the distribution function of h is $G(h)$ and $\bar{h} \equiv \max_{\eta \in \mathcal{N}} \hat{h}(\bar{\theta}|\eta)$. The corresponding problem in NM is $\max E_{\theta}[V(q(\theta)) - X(\theta)]$ subject to $X(\theta) - \theta q(\theta) \geq X(\theta') - \theta q(\theta')$ for any $\theta, \theta' \in \Theta$ and $X(\theta) - \theta q(\theta) \geq 0$ for any $\theta \in \Theta$. The two problems differ only in the underlying cost distributions of P: $G(h)$ in the case of DM and $F(\theta)$ in the case of NM. Since $\theta < \hat{h}(\theta|\eta)$ for $\theta > \underline{\theta}(\eta)$,

$$G(h|\eta) \equiv \Pr(\theta | \hat{h}(\theta|\eta) \leq h, \eta) < \Pr(\theta | \theta \leq h, \eta) = F(h|\eta)$$

for $h \in (\underline{\theta}(\eta), \hat{h}(\bar{\theta}|\eta) | \eta)$, implying $G(h) = \sum_{\eta \in \mathcal{N}} p(\eta)G(h|\eta) < \sum_{\eta \in \mathcal{N}} p(\eta)F(h|\eta) = F(h)$ for any $h \in (\underline{\theta}, \bar{h})$. Therefore the distribution of h in DM (strictly) dominates that of θ in NM in the first order stochastic sense. $\bar{h} > \bar{\theta}$, since $\hat{h}(\bar{\theta}|\eta) | \eta > \bar{\theta}(\eta)$ for any η .

It remains to show that this implies that P must earn a lower profit in DM. We prove the following general statement. Consider two contracting problems with a single supplier which differ only in regard to the cost distributions G_1 and G_2 , where $G_1(h) < G_2(h)$ for any $h \in (\underline{h}, \bar{h})$ and $G_2(h) = 1$ on $h \in [\bar{\theta}, \bar{h}]$. Standard arguments imply the problem can be reduced to selecting $q(h)$ to maximize the expected value of $V(q(h)) - hq(h) - \int_h^{\bar{h}_1} q(y)dy$ (where $\bar{h}_1 \equiv \bar{h}$ and $\bar{h}_2 \equiv \bar{\theta}$) with respect to distribution G_i , subject to the constraint that $q(\cdot)$ is nonincreasing. Let the maximized profit of P with distribution G be denoted $W(G)$. We will show $W(G_1) < W(G_2)$.

Let $q_1(h)$ denote the optimal solution of the problem based on $G_1(h)$. If $q_1(h)$ is constant on (\underline{h}, \bar{h}) with $q_1(h) = q > 0$, $W(G_1) = V(q) - \bar{h}q$. It is feasible for P to select this output schedule when the cost distribution is G_2 , generating expected profit $V(q) - \bar{\theta}q$. Then $W(G_2) \geq V(q) - \bar{\theta}q > W(G_1)$ since $\bar{h} > \bar{\theta}$. We henceforth focus on the case where $q_1(h)$ is not constant.

(i) First we show that $V'(q_1(h)) < h$ does not hold for any set of values of h with positive measure. Suppose otherwise that there exists some interval over which $V'(q_1(h)) < h$. Then we can replace the portion of $q_1(h)$ with $V'(q_1(h)) < h$ by $q^*(h)$ with $V'(q^*(h)) = h$, without violating the constraint that $q(h)$ is non-increasing. It raises the value of the objective function, since $V(q_1(h)) - hq_1(h) < V(q^*(h)) - hq^*(h)$ for h where $q_1(h)$ is replaced by $q^*(h)$, and $\int_h^{\bar{h}} q(y)dy$ decreases with this replacement. This is a contradiction.

(ii) Define

$$\Phi(h) \equiv V(q_1(h)) - hq_1(h) - \int_h^{\bar{h}} q_1(y)dy.$$

We claim that $\Phi(h)$ is left-continuous. First we show that our attention can be restricted to the case that $q_1(h)$ is left-continuous. Otherwise, there exists $h' \in (\underline{h}, \bar{h})$ such that $q_1(h'-) > q_1(h')$. Now consider $\tilde{q}_1(h)$ (which is left-continuous at h') such that $\tilde{q}_1(h') = q_1(h'-)$ and $\tilde{q}_1(h) = q_1(h)$ for any $h \neq h'$. Defining $\tilde{\Phi}(h) \equiv V(\tilde{q}_1(h)) - h\tilde{q}_1(h) - \int_h^{\bar{h}} \tilde{q}_1(y)dy$, observe that $\tilde{\Phi}(h) = \Phi(h)$ for $h \neq h'$ and $\tilde{\Phi}(h) > \Phi(h)$ when $h = h'$. Then

$$\int_{[\underline{h}, \bar{h}]} \tilde{\Phi}(h)dG(h) = \int_{[\underline{h}, \bar{h}] \setminus h'} \tilde{\Phi}(h)dG(h) + \tilde{\Phi}(h')[G(h'+) - G(h'-)]$$

$$\geq \int_{[\underline{h}, \bar{h}] \setminus h'} \tilde{\Phi}(h) dG(h) + \Phi(h') [G(h'+) - G(h'-)] = \int_{[\underline{h}, \bar{h}]} \Phi(h) dG(h)$$

with strict inequality if $G(h)$ is discontinuous at $h = h'$. When $q_1(h)$ is left-continuous, $\Phi(h)$ is also so.

(iii) We claim that $\Phi(h)$ is non-increasing in h and is not constant on (\underline{h}, \bar{h}) . To show the former, note that for any h , we have

$$\begin{aligned} & \lim_{\epsilon \rightarrow 0^+} \frac{\Phi(h + \epsilon) - \Phi(h)}{\epsilon} \\ &= \lim_{\epsilon \rightarrow 0^+} (1/\epsilon) [V(q_1(h + \epsilon)) - (h + \epsilon)q_1(h + \epsilon) - \int_{h+\epsilon}^{\bar{h}} q_1(y) dy - [V(q_1(h)) - hq_1(h) - \int_h^{\bar{h}} q_1(y) dy]] \\ &= [V'(\hat{q}(h)) - h] \lim_{\epsilon \rightarrow 0^+} \frac{q_1(h + \epsilon) - q_1(h)}{\epsilon} - q_1(h+) + \lim_{\epsilon \rightarrow 0^+} (1/\epsilon) \int_h^{h+\epsilon} q_1(y) dy \\ &= [V'(\hat{q}(h)) - h] \lim_{\epsilon \rightarrow 0^+} \frac{q_1(h + \epsilon) - q_1(h)}{\epsilon} \end{aligned}$$

for some $\hat{q}(h) \in [q_1(h+), q_1(h)]$. This is non-positive since $V'(\hat{q}(h)) \geq V'(q_1(h)) \geq h$ and $\lim_{\epsilon \rightarrow 0^+} \frac{q_1(h+\epsilon) - q_1(h)}{\epsilon} \leq 0$. Because of left-continuity of $\Phi(h)$, it implies that $\Phi(h)$ is non-increasing in h . Next to show $\Phi(h)$ is not constant, suppose otherwise. Then $\Phi(h) = \Phi(\bar{h}-) = V(q_1(\bar{h}-)) - \bar{h}q_1(\bar{h}-)$ which must be equal to $W(G_1)$. It means that $W(G_1)$ is attainable with constant output schedule ($q_1(h) = q_1(\bar{h}-)$ for any $h \in (\underline{h}, \bar{h})$). We obtain a contradiction, since we supposed that $q_1(h)$ is not constant.

(iv) Now consider the contracting problem with cost distribution $G_2(h)$. Since $q_1(h)$ is non-increasing in h , it is feasible for P to select this output schedule when the cost distribution is G_2 . Hence $W(G_2) \geq \int_{\underline{h}}^{\bar{h}} \Phi(h) dG_2(h)$. Therefore if $\int_{\underline{h}}^{\bar{h}} \Phi(h) dG_2(h) > \int_{\underline{h}}^{\bar{h}} \Phi(h) dG_1(h) = W(G_1)$, it follows that $W(G_2) > W(G_1)$. Since $G_i(h)$ ($i = 1, 2$) is right-continuous and $\Phi(h)$ is left-continuous, we can integrate by parts:

$$\int_{\underline{h}}^{\bar{h}} \Phi(h) dG_i(h) + \int_{\underline{h}}^{\bar{h}} G_i(h) d\Phi(h) = \Phi(\bar{h}).$$

Hence

$$\int_{\underline{h}}^{\bar{h}} \Phi(h) dG_2(h) - \int_{\underline{h}}^{\bar{h}} \Phi(h) dG_1(h) = \int_{\underline{h}}^{\bar{h}} [G_1(h) - G_2(h)] d\Phi(h).$$

By (iii) and $G_2(h) > G_1(h)$ for $h \in (\underline{h}, \bar{h})$, this is positive. \square

Proof of Proposition 3. The proof of the necessity part has already been described in the text. So we establish the sufficiency part. Consider any state η . Suppose there is a CSC function $\pi(\cdot | \eta) \in Y(\eta)$ which is ironed to yield the CVC function $z(\cdot | \eta)$ such that $(u_M(\theta, \eta), u_A(\theta, \eta), q(\theta, \eta))$ satisfies all the conditions in the proposition. Define $(X(m), Q(m))$ on $\hat{\mathcal{M}} \equiv K \cup \{e\}$ such that

$$(X(\theta, \eta), Q(\theta, \eta)) = (u_M(\theta, \eta) + u_A(\theta, \eta) + \theta q(\theta, \eta), q(\theta, \eta))$$

and $(X(e), Q(e)) = (0, 0)$, and extend this to $(X(\tilde{m}), Q(\tilde{m}))$ on $\Delta(\hat{\mathcal{M}})$ in the obvious manner. Consider the problem $P(\eta)$. For $\tilde{u}_A(\theta, \eta)$ which is feasible, we have

$$\int_{[\theta(\eta), \bar{\theta}(\eta)]} [\tilde{u}_A(\theta, \eta) - u_A(\theta, \eta)] d\Lambda(\theta | \eta) \geq 0.$$

Then

$$\begin{aligned} & E[X(\tilde{m}(\theta | \eta)) - \theta Q(\tilde{m}(\theta | \eta)) - \tilde{u}_A(\theta, \eta) | \eta] \\ & \leq E[X(\tilde{m}(\theta | \eta)) - \theta Q(\tilde{m}(\theta | \eta)) - \tilde{u}_A(\theta, \eta) | \eta] \end{aligned}$$

$$+ \int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} [\tilde{u}_A(\theta, \eta) - u_A(\theta, \eta)] d\Lambda(\theta | \eta).$$

Now consider the problem of maximizing the right hand side of this inequality, subject to the constraint that $Q(\tilde{m}(\theta | \eta))$ is non-increasing in θ . Using the same steps in the proof of the necessity part, this can be expressed as a problem of selecting $\tilde{m}(\theta | \eta)$ to maximize the Lagrangean (8) subject to the constraint that $Q(\tilde{m}(\theta | \eta))$ is non-increasing in θ . Conditions (i)-(iii) imply that the right-hand-side is maximized at $\tilde{m}(\theta | \eta) = (\theta, \eta)$ and $\tilde{u}_A(\theta, \eta) = u_A(\theta, \eta)$. Since

$$\int_{[\underline{\theta}(\eta), \bar{\theta}(\eta)]} [\tilde{u}_A(\theta, \eta) - u_A(\theta, \eta)] d\Lambda(\theta | \eta) = 0$$

when $\tilde{u}_A(\theta, \eta) = u_A(\theta, \eta)$, this shows that the left hand side of the above inequality is also maximized at $\tilde{m}(\theta | \eta) = (\theta, \eta)$ and $\tilde{u}_A(\theta, \eta) = u_A(\theta, \eta)$. Hence $(\tilde{m}(\theta | \eta), \tilde{u}_A(\theta, \eta)) = ((\theta, \eta), u_A(\theta, \eta))$ solves $P(\eta)$. \square

Proof of Proposition 4. Step 1:

Define

$$A(\theta | \eta) \equiv \frac{F(\theta | \eta)}{f(\theta | \eta)} / \frac{F(\theta)}{f(\theta)} \equiv \frac{\int_{\underline{\theta}(\eta)}^{\theta} f(y)a(\eta|y)dy}{a(\eta|\theta)F(\theta)}.$$

We show that there exist $\eta^{**} \in \mathcal{N}$ and a closed interval $I = [\theta', \theta'']$ with positive measure (conditional on η^{**}) such that $A(\theta | \eta^{**})$ is increasing in θ over I .

Evidently this holds for η such that $\underline{\theta} < \underline{\theta}(\eta)$, since $A(\underline{\theta}(\eta) | \eta)$ is continuous, $\lim_{\theta \rightarrow \underline{\theta}(\eta)} A(\theta | \eta) = 0$ and $A(\theta | \eta) > 0$ for $\theta > \underline{\theta}(\eta)$. Suppose otherwise; then $\underline{\theta}(\eta) = \underline{\theta}$ for all η . Using l'Hôpital's rule, $\lim_{\theta \rightarrow \underline{\theta}} A(\theta | \eta) = 1$. If $A(\theta | \eta)$ is non-increasing in θ for all η , $A(\theta | \eta) \leq 1$ or equivalently $\int_{\underline{\theta}}^{\theta} f(y)a(\eta|y)dy \leq a(\eta|\theta)F(\theta)$ for all $(\theta, \eta) \in K$. Since

$$\int_{\underline{\theta}}^{\theta} f(y)a(\eta|y)dy = \int_{\underline{\theta}}^{\theta} a(\eta|y)F(y)dy = F(\theta)$$

for all θ , $A(\theta | \eta) = 1$ for all $(\theta, \eta) \in K$. Then $h(\theta | \eta) = H(\theta)$ for any $(\theta, \eta) \in K$. This is a contradiction, since η is informative about θ .

For η^{**} and I , we choose $\lambda > 0$, closed intervals $\Theta_L = [\underline{\theta}^L, \bar{\theta}^L] \subset I$ and $\Theta_H = [\underline{\theta}^H, \bar{\theta}^H] \subset I$ with $\bar{\theta}^L < \underline{\theta}^H$ such that

$$\frac{F(\theta)}{f(\theta)} / \frac{F(\theta | \eta^{**})}{f(\theta | \eta^{**})} < \lambda < \frac{F(\bar{\theta})}{f(\bar{\theta})} / \frac{F(\bar{\theta} | \eta^{**})}{f(\bar{\theta} | \eta^{**})} \text{ for } \bar{\theta} \in \Theta_L, \theta \in \Theta_H.$$

These conditions are equivalent to

$$H(\theta) - (1 - \lambda)\theta - \lambda h(\theta | \eta^{**}) > 0 \text{ for } \theta \in \Theta_L$$

and

$$H(\theta) - (1 - \lambda)\theta - \lambda h(\theta | \eta^{**}) < 0 \text{ for } \theta \in \Theta_H.$$

Step 2: Construction of $z(\cdot | \eta)$.

Now let us construct $z(\cdot | \eta)$ which satisfies the following conditions.

- (A) For $\eta \neq \eta^{**}$, $z(\theta | \eta) = \theta$ for any $\theta \in \Theta(\eta)$.
- (B) For η^{**} , $z(\cdot | \eta^{**}) \in Z(\eta^{**})$ which satisfies
 - (i) $z(\theta | \eta^{**}) = \theta$ for any $\theta \notin \Theta_H \cup \Theta_L$;
 - (ii) For $\theta \in \Theta_L$, $z(\theta | \eta^{**})$ satisfies (a) $z(\theta | \eta^{**}) \leq \theta$ with strict inequality for some subinterval of Θ_L of positive measure, and (b) $H(z) - (1 - \lambda)z - \lambda h(\theta | \eta^{**}) > 0$ for any $z \in [z(\theta | \eta^{**}), \theta]$;
 - (iii) For $\theta \in \Theta_H$, $z(\theta | \eta^{**})$ satisfies (a) $z(\theta | \eta^{**}) \geq \theta$ with strict inequality for some subinterval of Θ_H of positive measure, (b) $z(\theta | \eta^{**}) < h(\theta | \eta^{**})$ and (c) $H(z) - (1 - \lambda)z - \lambda h(\theta | \eta^{**}) < 0$ for any $z \in [\theta, z(\theta | \eta^{**})]$;
 - (iv) $E[(z(\theta | \eta^{**}) - h(\theta | \eta^{**}))Q^*(z(\theta | \eta^{**})) + \int_{z(\theta | \eta^{**})}^{\bar{\theta}(\eta^{**})} Q^*(z)dz | \eta^{**}] = 0$.

It is shown in the online Appendix that there exists $z(\cdot | \eta^{**}) \in Z(\eta^{**})$ which satisfies (B(i)-(iv)).

Step 3:

For $z(\cdot | \eta)$ constructed in Step 2, consider the following allocation (u_A, u_M, q) :

$$\begin{aligned} q(\theta, \eta) &= Q^*(z(\theta | \eta)) \\ u_A(\theta, \eta) &= \int_{\theta}^{\bar{\theta}(\eta)} Q^*(z(y | \eta)) dy + \int_{\bar{\theta}(\eta)}^{\bar{\theta}} Q^*(y) dy \\ u_M(\theta, \eta) &= X^*(z(\theta | \eta)) - \theta Q^*(z(\theta | \eta)) - \int_{\theta}^{\bar{\theta}(\eta)} Q^*(z(y | \eta)) dy - \int_{\bar{\theta}(\eta)}^{\bar{\theta}} Q^*(y) dy \end{aligned}$$

where

$$X^*(z(\theta | \eta)) \equiv z(\theta | \eta) Q^*(z(\theta | \eta)) + \int_{z(\theta | \eta)}^{\bar{\theta}} Q^*(z) dz.$$

In the online Appendix it is shown that (u_A, u_M, q) is a EAC feasible allocation.

Now we show that this allocation generates a higher payoff to P than the optimal allocation in NM. Define $\Phi_P(z)$ and $\Phi_M(z, \theta)$ as

$$\Phi_P(z) \equiv V(Q^*(z)) - zQ^*(z) - \int_z^{\bar{\theta}} Q^*(\tilde{z}) d\tilde{z}$$

and

$$\Phi_M(z, \theta) \equiv (z - h(\theta | \eta^{**})) Q^*(z) + \int_z^{\bar{\theta}(\eta^{**})} Q^*(\tilde{z}) d\tilde{z}.$$

P's resulting expected payoff conditional on η^{**} is $E[\Phi_P(z(\theta | \eta^{**})) | \eta^{**}]$. P's expected payoff conditional on η^{**} in the optimal allocation in NM is $E[\Phi_P(\theta) | \eta^{**}]$. By the definition of $\Phi_M(z, \theta)$ and $E[u_M(\theta, \eta^{**}) | \eta^{**}] = 0$, $E[\Phi_M(z(\theta | \eta^{**}), \theta) | \eta^{**}] = E[\Phi_M(\theta, \theta) | \eta^{**}] = 0$. Then the difference between two payoffs is

$$\begin{aligned} & E[\Phi_P(z(\theta | \eta^{**})) | \eta^{**}] - E[\Phi_P(\theta) | \eta^{**}] \\ &= E[\Phi_P(z(\theta | \eta^{**})) + \lambda \Phi_M(z(\theta | \eta^{**}), \theta) | \eta^{**}] - E[\Phi_P(\theta) + \lambda \Phi_M(\theta, \theta) | \eta^{**}] \\ &= E\left[\int_{\theta}^{z(\theta | \eta^{**})} \{\Phi'_P(z) + \lambda \partial \Phi_M(z, \theta) / \partial z\} dz \mid \eta^{**} \right] \\ &= E\left[\int_{\theta}^{z(\theta | \eta^{**})} [V'(Q^*(z)) - \{(1 - \lambda)z + \lambda h(\theta | \eta^{**})\}] Q^{*'}(z) dz \mid \eta^{**} \right] \\ &= E\left[\int_{\theta}^{z(\theta | \eta^{**})} [H(z) - \{(1 - \lambda)z + \lambda h(\theta | \eta^{**})\}] Q^{*'}(z) dz \mid \eta^{**} \right]. \end{aligned}$$

The last equality follows from $V'(Q^*(z)) = H(z)$. From the construction of $z(\theta | \eta^{**})$ in Step 2 and $Q^{*'}(z) < 0$, this is positive. We have thus found an EAC feasible allocation which generates a higher payoff to P compared to the optimal allocation in NM. \square

Proof of Proposition 5. Conditions (i) and (ii) imply $h(\theta | \eta)$ satisfies the following properties:

- $h(\theta | \eta^*) = \hat{h}(\theta | \eta^*)$ is strictly increasing and continuously differentiable in θ

- $\hat{h}(\theta | \eta^*) > \hat{h}(\theta | \eta)$ for $\theta \in (\underline{\theta}, \bar{\theta}]$ and $\hat{h}(\underline{\theta} | \eta^*) = \hat{h}(\underline{\theta} | \eta) = \underline{\theta}$ for any $\eta \neq \eta^*$
- Define $G(h | \eta) \equiv \int_{\{\theta | \hat{h}(\theta | \eta) \leq h\}} f(\theta | \eta) d\theta$. Then $G(h | \eta^*)$ is a mean-preserving spread of $G(h | \eta)$ for any $\eta \neq \eta^*$

The first one is evident, since $f(\theta | \eta^*)$ is decreasing in θ and $h(\theta | \eta^*)$ is increasing in θ . Then $q^*(\hat{h}(\theta | \eta^*))$ is also continuously differentiable and strictly decreasing in θ . By the second property, the range of \hat{h} conditional on η^* (which is denoted by H) includes the range of \hat{h} conditional on η . The proof of the second and third properties are provided in the online Appendix.

Suppose the result is false, and the second best allocation $(u_A^{SB}, u_M^{SB}, q^{SB})$ is EAC feasible. Then Proposition 3 implies the existence of $\pi(\cdot | \eta) \in Y(\eta)$ such that for any $(\theta, \eta), (\theta', \eta') \in K$,

$$\begin{aligned} q^{SB}(\theta, \eta) &= q^*(\hat{h}(\theta | \eta)) \\ X^{SB}(\theta, \eta) - z(\theta | \eta)q^{SB}(\theta, \eta) &\geq 0 \\ X^{SB}(\theta, \eta) - z(\theta | \eta)q^{SB}(\theta, \eta) &\geq X^{SB}(\theta', \eta') - z(\theta | \eta)q^{SB}(\theta', \eta') \end{aligned}$$

where $z(\cdot | \eta)$ is an ironing transformation of $\pi(\cdot | \eta)$ based on $F(\theta | \eta)$, and

$$X^{SB}(\theta, \eta) \equiv u_A^{SB}(\theta, \eta) + u_M^{SB}(\theta, \eta) + \theta q^{SB}(\theta, \eta).$$

Step 1: Existence of $(\tilde{X}(h), \phi(h))$.

Consider $(\theta, \eta), (\theta', \eta')$ such that $\hat{h}(\theta | \eta) = \hat{h}(\theta' | \eta')$. Then $q^{SB}(\theta, \eta) = q^{SB}(\theta', \eta')$. The coalitional incentive constraint implies $X^{SB}(\theta, \eta) = X^{SB}(\theta', \eta')$, since otherwise the coalition would misrepresent a state with higher payment in the other state where the same output is produced. It guarantees the existence of $\tilde{X}(h)$ such that $X^{SB}(\theta, \eta) = \tilde{X}(\hat{h}(\theta | \eta))$ for any (θ, η) .

Next suppose that $\hat{h}(\theta'' | \eta'') = \hat{h}(\theta' | \eta')$ and $z(\theta'' | \eta'') > z(\theta' | \eta')$ for some $(\theta', \eta'), (\theta'', \eta'')$. The ironing procedure ensures $z(\theta | \eta)$ and $\hat{h}(\theta | \eta)$ are continuous and non-decreasing for θ on Θ . Since $\hat{h}(\underline{\theta} | \eta) = \underline{\theta} < \bar{\theta} < \hat{h}(\bar{\theta} | \eta)$, $\hat{h}(\theta | \eta)$ is not constant on Θ . Then by adjusting θ' and θ'' , we can find $(\tilde{\theta}', \tilde{\theta}'')$ such that $\hat{h}(\tilde{\theta}'' | \eta'') < \hat{h}(\tilde{\theta}' | \eta')$ and $z(\tilde{\theta}'' | \eta'') > z(\tilde{\theta}' | \eta')$. We obtain a contradiction, since the coalitional incentive constraint implies that whenever $z(\tilde{\theta}'' | \eta'') > z(\tilde{\theta}' | \eta')$, $q^{SB}(\tilde{\theta}'', \eta'') \leq q^{SB}(\tilde{\theta}', \eta')$ or equivalently $\hat{h}(\tilde{\theta}'' | \eta'') \geq \hat{h}(\tilde{\theta}' | \eta')$. Hence there exists $\phi(h)$ which satisfies $z(\theta | \eta) = \phi(\hat{h}(\theta | \eta))$ for any (θ, η) . Since $\hat{h}(\theta | \eta^*)$ and $z(\theta | \eta^*)$ are continuous and non-decreasing for θ , $\phi(h)$ is continuous and non-decreasing on H .

The coalitional incentive constraint implies that for any $h \in H$, h maximizes $\tilde{X}(h') - \phi(h)q^*(h')$ subject to $h' \in H$. By the continuity of $\phi(h)$ and the differentiability of $q^*(h)$, we obtain the differentiability of $\tilde{X}(h)$ and the first order condition $\tilde{X}'(h) = \phi(h)q^*(h)$.

Step 2: Properties of $\phi(h)$.

Here we show that (a) $\phi(h) \geq 0$ on H and (b) $h - \phi(h)$ is non-negative and increasing in h .

Since $q^{SB}(\theta, \eta^*) = q^*(\hat{h}(\theta | \eta^*))$ is strictly decreasing in θ , the pooling region $\Theta(\pi(\cdot | \eta^*), \eta^*)$ must be empty. Hence it must be the case that

$$z(\theta | \eta^*) = \phi(\hat{h}(\theta | \eta^*)) = \theta + \frac{F(\theta | \eta^*) - \Lambda(\theta | \eta^*)}{f(\theta | \eta^*)}.$$

Since $\phi(\hat{h}(\theta | \eta^*))$ is non-decreasing in θ and $\Lambda(\theta | \eta^*) \leq 1$,

$$\phi(\hat{h}(\theta | \eta^*)) \geq \phi(\underline{\theta}) \geq \underline{\theta} - 1/f(\underline{\theta} | \eta^*) > 0$$

by property (iii), which implies (a). The above equality can be rewritten as

$$\frac{\Lambda(\theta | \eta^*)}{f(\theta | \eta^*)} = \hat{h}(\theta | \eta^*) - \phi(\hat{h}(\theta | \eta^*)).$$

The LHS is non-negative and increasing in θ , since $f(\theta | \eta^*)$ is decreasing in θ and $\Lambda(\theta | \eta^*)$ is non-negative and non-decreasing in θ . It implies (b).

Step 3: M earns positive rent.

Define $L(h) \equiv \tilde{X}(h) - hq^*(h)$. M's interim payoff is

$$E[X^{SB}(\theta, \eta) - h(\theta | \eta)q^{SB}(\theta, \eta) | \eta] = E[L(\hat{h}(\theta | \eta)) | \eta],$$

utilizing a property of the ironing transformation. If the second best allocation is EAC feasible, $E[L(\hat{h}(\theta | \eta)) | \eta] = 0$ holds for any η . The first derivative of $L(h)$ is

$$L'(h) = (\phi(h) - h)q^{*'}(h) - q^*(h).$$

Since $q^*(h)$ is continuously differentiable and $\phi(h)$ is continuous and almost everywhere differentiable, $L'(h)$ is continuous and also differentiable almost everywhere and

$$L''(h) = (\phi'(h) - 1)q^{*'}(h) + (\phi(h) - h)q^{*''}(h) - q^{*'}(h).$$

By using $V'(q^*(h)) = h$, we can show that $V'''(q) \leq 0$ implies $q^{*''}(h) \leq 0$, and $0 < V'''(q) \leq \frac{(V''(q))^2}{V'(q)}$ implies $q^{*''}(h) > 0$ and $hq^{*''}(h) + q^{*'}(h) < 0$. By $\phi'(h) - 1 < 0$ and $0 \leq \phi(h) \leq h$, it follows that $L''(h) > 0$.

By the strict convexity of L and the mean-preserving spread property of $G(h | \eta^*)$,

$$E[L(\hat{h}(\theta | \eta^*)) | \eta^*] = \int L(h)dG(h | \eta^*) > \int L(h)dG(h | \eta) = E[L(\hat{h}(\theta | \eta)) | \eta] \geq 0$$

for any $\eta \neq \eta^*$. Therefore M must earn a positive rent in state η^* . This is a contradiction. \square

Proof of Proposition 6. We show that the set of EACP(α) is independent of $\alpha \in [0, 1]$. Suppose otherwise that $(u_A(\theta, \eta), u_M(\theta, \eta), q(\theta, \eta))$ is a EACP(α) allocation, but not a EACP(α') ($\alpha \neq \alpha'$) allocation. It implies that for some η , $(\tilde{m}(\theta | \eta), \tilde{u}_A(\theta, \eta)) = ((\theta, \eta), u_A(\theta, \eta))$ is not the solution of $TP(\eta; \alpha')$ defined for $(u_A(\theta, \eta), u_M(\theta, \eta), q(\theta, \eta))$. If $(\tilde{m}^*(\theta | \eta), u_A^*(\theta, \eta)) \neq ((\theta, \eta), u_A(\theta, \eta))$ is a solution of $TP(\eta; \alpha')$, it satisfies all constraints of $TP(\eta; \alpha')$ and realizes a higher payoff to the third party than in the choice of $(\tilde{m}(\theta | \eta), \tilde{u}_A(\theta, \eta)) = ((\theta, \eta), u_A(\theta, \eta))$:

$$\begin{aligned} & E[(1 - \alpha')[X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - u_A^*(\theta, \eta)] + \alpha' u_A^*(\theta, \eta) | \eta] \\ & > E[(1 - \alpha')[X(\theta, \eta) - \theta Q(\theta, \eta) - u_A(\theta, \eta)] + \alpha' u_A(\theta, \eta) | \eta]. \end{aligned}$$

It also satisfies A and M's participation constraints:

$$u_A^*(\theta, \eta) \geq u_A(\theta, \eta)$$

and

$$E[X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - u_A^*(\theta, \eta) | \eta] \geq E[u_M(\theta, \eta) | \eta].$$

On the other hand, since $(\tilde{m}(\theta | \eta), \tilde{u}_A(\theta, \eta)) = ((\theta, \eta), u_A(\theta, \eta))$ solves $TP(\eta; \alpha)$,

$$\begin{aligned} & E[(1 - \alpha)[X(\theta, \eta) - \theta Q(\theta, \eta) - u_A(\theta, \eta)] + \alpha u_A(\theta, \eta) | \eta] \\ & \geq E[(1 - \alpha)[X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - u_A^*(\theta, \eta)] + \alpha u_A^*(\theta, \eta) | \eta] \end{aligned}$$

Let us consider three cases: (i) $\alpha \in (0, 1)$, (ii) $\alpha = 1$ and (iii) $\alpha = 0$.

(i) $\alpha \in (0, 1)$.

The last three inequalities imply

$$u_A^*(\theta, \eta) = u_A(\theta, \eta)$$

and

$$E[X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - u_A^*(\theta, \eta) | \eta] = E[u_M(\theta, \eta) | \eta].$$

But this is not compatible with the first inequality. We obtain a contradiction.

(ii) $\alpha = 1$.

With $\alpha = 1$, the above four inequalities imply

$$E[u_A(\theta, \eta) | \eta] = E[u_A^*(\theta, \eta) | \eta]$$

and

$$E[X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - u_A^*(\theta, \eta) | \eta] > E[u_M(\theta, \eta) | \eta].$$

But for sufficiently small $\epsilon > 0$, the choice of

$$(\tilde{m}(\theta | \eta), \tilde{u}_A(\theta, \eta)) = (\tilde{m}^*(\theta, \eta), u_A^*(\theta, \eta) + \epsilon)$$

(instead of $((\theta, \eta), u_A(\theta, \eta))$) in $TP(\eta; \alpha = 1)$ generates a higher value of the objection function without violating any constraint. We obtain a contradiction.

(iii) $\alpha = 0$.

With $\alpha = 0$, the four inequalities imply

$$E[u_M(\theta, \eta) | \eta] = E[X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - u_A^*(\theta, \eta) | \eta]$$

and

$$E[u_A^*(\theta, \eta) | \eta] > E[u_A(\theta, \eta) | \eta].$$

Since $u_A^*(\theta, \eta) \geq u_A(\theta, \eta)$ for any θ , there is a subset of θ with the positive measure such that $u_A^*(\theta, \eta) > u_A(\theta, \eta)$. Consider a modified problem of $TP(\eta; \alpha = 0)$ such that the constraint $\tilde{u}_A(\theta, \eta) \geq u_A(\theta, \eta)$ is replaced by $\tilde{u}_A(\theta, \eta) \geq u_A^*(\theta, \eta)$ in $TP(\eta; \alpha = 0)$. Since the optimal solution $(\tilde{m}(\theta | \eta), \tilde{u}_A(\theta, \eta)) = ((\theta, \eta), u_A(\theta, \eta))$ in $TP(\eta; \alpha = 0)$ violates the constraint, the maximum value of the objective function in the modified problem would become lower. On the other hand, $(\tilde{m}^*(\theta | \eta), u_A^*(\theta, \eta))$ satisfies all the constraints of the modified problem, and brings

$$E[X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - u_A^*(\theta, \eta) | \eta].$$

The argument implies

$$E[u_M(\theta, \eta) | \eta] > E[X(\tilde{m}^*(\theta | \eta)) - \theta Q(\tilde{m}^*(\theta | \eta)) - u_A^*(\theta, \eta) | \eta].$$

We obtain a contradiction. \square

Appendix B. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.geb.2020.04.011>.

References

- Baliga, S., Sjostrom, T., 1998. Decentralization and collusion. *J. Econ. Theory* 83 (2), 196–232.
- Baron, B., Myerson, R.B., 1982. Regulating a monopolist with unknown costs. *Econometrica* 50 (4), 911–930.
- Celik, G., 2009. Mechanism design with collusive supervision. *J. Econ. Theory* 144 (1), 69–95.
- Celik, G., Peters, M., 2011. Equilibrium rejection of a mechanism. *Games Econ. Behav.* 73 (2), 375–387.
- Che, Y.K., Condorelli, D., Kim, J., 2018. Weak cartels and collusion-proof auctions. *J. Econ. Theory* 178, 398–435.
- Che, Y.K., Kim, J., 2006. Robustly collusion-proof implementation. *Econometrica* 74 (4), 1063–1107.
- Che, Y.K., Kim, J., 2009. Optimal collusion-proof auctions. *J. Econ. Theory* 144, 565–603.
- Dequiedt, V., 2007. Efficient collusion in optimal auctions. *J. Econ. Theory* 136 (1), 302–323.
- Faure-Grimaud, A., Laffont, J.J., Martimort, D., 2003. Collusion, delegation and supervision with soft information. *Rev. Econ. Stud.* 70 (2), 253–279.
- Fudenberg, D., Tirole, J., 1991. Perfect Bayesian equilibrium and sequential equilibrium. *J. Econ. Theory* 53 (2), 236–260.
- Harris, M., Raviv, A., 2008. A theory of board control and size. *Rev. Financ. Stud.* 21 (4), 1797–1832.
- Hirschman, A., 1970. *Exit, Voice and Loyalty: Responses to Decline in Firms, Organizations, and States*. Harvard University Press, Cambridge/Mass.
- Jullien, B., 2000. Participation constraints in adverse selection models. *J. Econ. Theory* 93 (1), 1–47.
- Laffont, J.J., Martimort, D., 1997. Collusion under asymmetric information. *Econometrica* 65 (4), 875–911.
- Laffont, J.J., Martimort, D., 2000. Mechanism design with collusion and correlation. *Econometrica* 68 (2), 309–342.
- Laffont, J.J., Tirole, J., 1993. *A Theory of Incentives in Procurement and Regulation*. MIT Press, Cambridge, MA.
- Mookherjee, D., Motta, A., Tsumagari, M., 2018. Bypassing Intermediaries via Vertical Integration, a Transaction-Cost-Based Theory. Working paper. Institute for Economic Development, Boston University. Downloadable from <http://people.bu.edu/dilipm/wkppap/BypassIntermediaries.pdf>.
- Mookherjee, D., Tsumagari, M., 2004. The organization of supplier networks: effects of delegation and intermediation. *Econometrica* 72 (4), 1179–1219.
- Mookherjee, D., Tsumagari, M., 2017. Hierarchical Control Rights and Strong Collusion. Working paper. Department of Economics, Boston University. Downloadable from <http://people.bu.edu/dilipm/wkppap/index.html>.
- Motta, A., 2009. Collusion and selective supervision. Working paper. Downloadable from <https://sites.google.com/site/albertomottaeconomics/research>.
- Myerson, R., 1981. Optimal auction design. *Math. Oper. Res.* 6, 58–73.
- Ortner, J., Chassang, S., 2018. Making corruption harder: asymmetric information, collusion, and crime. *J. Polit. Econ.* 126 (5), 2108–2133.
- Pavlov, G., 2008. Auction design in the presence of collusion. *Theor. Econ.* 3 (3), 383–429.
- Quesada, L., 2004. Collusion as an Informed Principal Problem. University of Wisconsin. Working paper.
- Strausz, R., 2006. Deterministic versus stochastic mechanisms in principal-agent models. *J. Econ. Theory* 128 (1), 306–314.
- Tirole, J., 1986. Hierarchies and bureaucracies: on the role of collusion in organizations. *J. Law Econ. Organ.* 2, 181–214.
- Tirole, J., 1995. Collusion and the theory of organizations. In: Laffont, J.J. (Ed.), *Advances in Economic Theory: Volume 2: Sixth World Congress of the Econometric Society*. Cambridge University Press.
- Tsumagari, M., 2016a. On the role of collusive supervisor: case of bilateral asymmetric information. Working paper. <https://sites.google.com/site/masatoshitsumagari/home/working-papers>.
- Tsumagari, M., 2016b. Weak Ex Ante collusion and value of supervisor: discrete type model. Working paper. <https://sites.google.com/site/masatoshitsumagari/home/working-papers>.
- Williamson, O.E., 1975. *Markets and Hierarchies: Antitrust Analysis and Implications*. The Free Press, New York.