MECHANISM DESIGN WITH COMMUNICATION CONSTRAINTS[1]

Dilip Mookherjee and Masatoshi Tsumagari[2]

This Version: October 12 2012

Abstract

We consider the problem of designing mechanisms for contexts including auctions, internal organization and regulation in which multiple agents (whose outputs need to be coordinated) are privately informed about their respective production costs, and have limited capacity to communicate their private information owing to time or resource costs of sending detailed messages. This implies centralized revelation mechanisms become prohibitively expensive if agents' information is sufficiently detailed. We characterize feasible as well as optimal mechanisms subject to incentive and communicational constraints, without imposing *ad hoc* restrictions on the number of rounds of communication. Optimal mechanisms are shown to generally involve dynamic, interactive communication between agents and delegation of production decisions.

KEYWORDS: communication, mechanism design, auctions, decentralization, incentives, organizations

[1]We thank Ingela Alger, Richard Arnott, Hideshi Itoh, Ulf Lilienfeld, Bart Lipman, Michael Manove, Preston McAfee, Margaret Meyer, Alberto Motta, Andy Newman, Roy Radner, Stefan Reichelstein, Ilya Segal, Hideo Suehiro and seminar participants at Boston College, Brown, Columbia, Northwestern, Hitotsubashi, Keio, Kobe, Oxford, Nagoya, Padua, Paris, Stanford, Toulouse and Warwick for useful comments.

[2]Boston University and Keio University, respectively

1 Introduction

Most practical contexts of economic mechanisms such as auctions, regulation or internal organization of firms involve dynamic, interactive communication. Participants exchange bids, cost reports, or budgets through dynamic, time-consuming processes which feed into ultimate decisions concerning production, exchange and financial transfers. These are in sharp contrast with revelation mechanisms commonly studied in the theoretical literature, in which agents make reports of everything they know in a simultaneous, static process.

Part of the reason is that the dimensionality of private information is considerably richer than what can be feasibly be communicated to others in real time. This observation was made by Hayek (1945) in his famous critique of Lange-Lerner socialist resource allocation mechanisms, in which he argued that communication constraints provided much of the justification for a decentralized market economy coordinating decisions through price signals. Hayek's observation motivated a large literature on resource allocation mechanisms that economize on communication costs. Examples are the message space literature (Hurwicz (1960, 1972), Mount and Reiter (1974)) and the theory of teams (Marschak and Radner (1972)).[3] This early literature on mechanism design ignored incentive problems.[4] The more recent literature on mechanism design on the other hand focuses only on incentive problems, ignoring communication costs. As a result the space of relevant mechanisms studied are revelation mechanisms, where agents communicate reports of their entire private information in a single, instantaneous step.

A few recent papers have explored the implications of co-existence of communication costs and incentive problems. Van Zandt (2007) and Fadel and Segal (2009) pose the question of the extent to which incentive problems increase communicational complexity of mechanisms that implement a desired allocation rule, in a general setting. Other authors have sought to characterize optimal incentive mechanisms in settings with restricted message spaces and the standard assumptions of single-dimensional outputs and single-crossing preferences (Green and Laffont (1986, 1987), Melumad, Mookherjee and Reichelstein (1992, 1997), Laffont and Martimort (1998), Blumrosen, Nisan and Segal (2007), Blumrosen and Feldman (2006), Kos (20011a, 2011b)).[5] For reasons of tractability, these authors have severely restricted the communication protocols as well as the range of mechanisms considered. With the exception of Kos (2011b), most authors focus on mechanisms with a single round of communication (with restricted message spaces).

[3]Segal (2006) surveys recent studies of informationally efficient allocation mechanisms.

[4]A notable exception is Reichelstein and Reiter (1988), who examined implications of strategic behavior for communicational requirements of mechanisms implementing efficient allocations.

[5]Battigali and Maggi (2002) study a model of symmetric but nonverifiable information where there are costs of writing contingencies into contracts. This is in contrast to the models cited above which involve asymmetric information with constraints on message spaces.

However dynamic, interactive communication is valuable when agents' opportunities to express themselves are intrinsically limited.

The key analytical problem in incorporating dynamic communication protocols into models with strategic agents is to find a suitable characterization of incentive constraints. From the standpoint of informational efficiency, it is valuable to allow multiple rounds of communication that permit agents to learn messages sent by other agents before they respond with additional messages. This can create complications in the presence of incentive problems. Dynamic mechanisms enlarge the range of possible deviations available to participants, over and above those typically characterized by incentive compatibility constraints in a static revelation mechanism. Apart from the possibility of mimicking equilibrium strategies of other types of the same agent, additional deviations are possible in a dynamic communication process. Van Zandt (2007) observes that this is not a problem when the solution concept is *ex post incentive compatibility (EPIC)*, where agents do not regret their strategies even after observing all messages sent by other agents. When we use the less demanding concept of a (perfect) Bayesian equilibrium, dynamic communication protocols do impose additional incentive constraints. Then there is a potential trade-off between informational efficiency and incentive problems.

The problem in studying this trade-off is that a precise characterization of incentive constraints for dynamic protocols is not available in existing literature. In a very general setting Fadel and Segal (2009) provide different sets of sufficient conditions that are substantially stronger than necessary conditions. In this paper we restrict attention to contexts with single dimensional outputs and single-crossing preferences for each agent, and obtain a set of conditions that are both necessary and sufficient for Bayesian implementation in arbitrary dynamic communication protocols (see Proposition 1).

This implies that the mechanism design problem reduces to selecting an output allocation rule which maximizes a payoff function of the Principal (modified to include the cost of incentive rents paid to agents in a standard way with 'virtual' types replacing actual types) subject to communication feasibility restrictions alone (Proposition 2). This extends the standard approach to solving for optimal mechanisms with unlimited communication (following Myerson (1981)), while providing a convenient representation of the respective costs imposed by incentive problems and communicational constraints.

A number of implications of this result are then derived. The first concerns the value of delegating production decisions to agents.[6] In contexts of unconstrained communication where the Revelation Principle applies, it is well known that centralized decision-making

[6]Earlier literature such as Melumad, Mookherjee and Reichelstein (1992, 1997) and Laffont and Martimort (1998) have focused on a related but different question: the value of decentralized contracting (or subcontracting) relative to centralized contracting. Here we assume that contracting is centralized, and examine the value of decentralizing production decisions instead.

can always (trivially) replicate the outcome of any delegation mechanism.[7] This is no longer so when communication constraints prevent agents from reporting everything they know to the Principal (as argued by Hayek). In the presence of incentive problems, however, delegation can generate costs owing to opportunistic behavior, which have to be traded off against the benefits from enhanced informational efficiency. In our context, it turns out the benefits of improved information always outweigh attendant incentive costs, provided it is feasible for agents to select their own outputs independently. Under this condition, in addition to a host of regularity conditions on the production function (smoothness, concavity and Inada conditions ensuring interiority of optimal productions), Proposition 3 shows that any centralized mechanism is strictly dominated by some mechanism with decentralized production choices made by agents. This vindicates Hayek's arguments in favor of decentralized mechanisms. Practical implications include the superiority of taxes over quantitative controls, and of firm organizations which delegate production decisions to workers (e.g., Aoki (1990)).

A second set of implications concern the design of optimal communication protocols. Proposition 2 has strong implications for how communication processes ought to be structured: to maximize the amount of information exchanged by agents. We show that if communication costs either involve material costs which are linear in the length of messages sent and in the size of the communication channel (defined by the maximum length of messages sent), or delay which is linear in the size of the communication channel, then communication should take place over multiple rounds in each of which agents are assigned a small message set (consisting of all messages not exceeding unit length) in each round. Such dynamic protocols enable agents to exchange maximal information subject to the communication constraints. If communication costs consist only of delay, agents must report simultaneously in each round. But if they consist of only material costs, it is optimal for agents to alternate in sending messages across successive rounds.

The paper is organized as follows. Section 2 provides some examples to help explain the informational benefit of dynamic communication protocols, and then the incentive problems they give rise to. Section 3 introduces the general model. Section 4 is devoted to characterizing feasible allocations. Section 5 uses this to represent the design problem as maximizing the Principal's incentive-rent-modified welfare function subject to communicational constraints alone. Section 6 uses this to compare centralized and decentralized mechanisms, while Section 7 describes implications for design of optimal communication protocols. Section 8 concludes.

[7]For a formal statement and proof, see Myerson (1982).

2 Examples and Related Literature

Example 1

We start by providing an example illustrating the informational value of dynamic, interactive communication protocols when there are limitations on the amount of information that can be communicated by any agent in any given round, and incentive issues are ignored.

Suppose two agents jointly produce a common output q for the Principal, which takes three possible levels $0, 1, 2$. The corresponding gross revenues $V(q)$ earned by the Principal are $0, 38$ and 50 respectively. Agent i incurs a unit cost θ_i of production. For Agent 1, this cost takes two possible values $0, 10$ which are equally likely *ex ante*. On the other hand θ_2 can take three possible values $0, 30, 100$ where the prior probability of costs 0 and 100 are $\frac{1}{4}$ each and the probability of cost 30 is a half.

Suppose the Principal is concerned about the efficiency measured by the expected value of $V(q) - (\theta_1 + \theta_2)q$. The first best allocation (without any consideration of communication or incentive constraints) is that $q = 2$ if $\theta_1 + \theta_2 < V(2) - V(1) = 12$, $q = 1$ if $V(2) - V(1) = 12 \leq \theta_1 + \theta_2 \leq 38 = V(1)$, $q = 0$ if $V(1) = 38 < \theta_1 + \theta_2$. The corresponding first-best outputs $q^{FB}(\theta_1, \theta_2)$ are $q^{FB}(0, 100) = q^{FB}(10, 100) = q^{FB}(10, 30) = 0$, $q^{FB}(0, 30) = 1$ and $q^{FB}(0, 0) = q^{FB}(0, 10) = 2$. This first best allocation is shown in case (a) of Figure 1.

Now introduce a constraint on communication: each agent can send only a binary message $m_i \in \{0, 1\}$ only once. Agent 2 who has three types then cannot report his true type. Ignoring incentive issues, suppose agents follow instructions provided by the Principal regarding what message to send in different contingencies. In what follows we focus on *threshold reporting strategies*, in which the type space of each agent is partitioned into a number of subintervals (which equals the number of feasible messages), and they report the subinterval in which their type belongs. It is easy to check that this is without any loss of generality, as the outcomes of other kinds of reporting strategies are dominated by those of threshold reporting strategies.

Consider first a protocol where the two agents send a binary report simultaneously to the Principal. There are two possible threshold reporting strategies, shown in (b) and (c) of Figure 1. The Principal now does not have the information available to implement the first-best efficient allocation, lacking information of Agent 2's true type. Conditional on the information available, the constrained efficient allocations are also shown in (b) and (c) of this figure. In the case of simultaneous reporting, the Principal's information is represented by a 'rectangular' partition of the type space, where she knows whether Agent 1's cost is high (10) or low (0), and whether Agent 2's cost is high or low (the precise definition of which depends on the particular threshold strategy used by Agent 2).

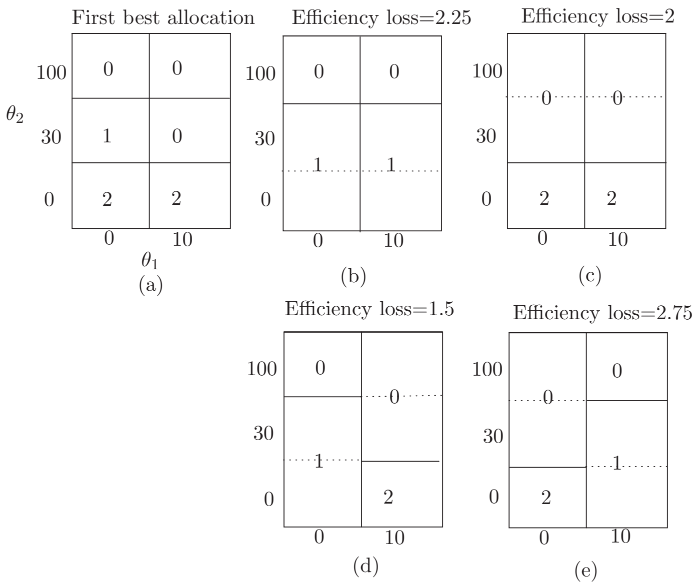

Figure 1: Example 1

The threshold for high cost of Agent 2 is 100 in case (b) and is 30 in case (c). The latter turns out to generate a smaller loss of efficiency compared with the first-best allocation.

Now consider a sequential protocol where Agent 1 sends a binary report first, to which Agent 2 responds with a binary report. Agent 2 can now condition his threshold on the report sent by Agent 1. Cases (d) and (e) show two different reporting strategies used by Agent 2 where the threshold does vary with Agent 1's report. Note also that cases (b) and (c) continue to be available here, since Agent 2 can also use a strategy in which the threshold does not vary with Agent 1's message. It turns out the smallest efficiency loss is incurred in case (d). Compared to the simultaneous reporting protocol, a sequential protocol allows more information to be communicated to the Principal, within the constraints allowed by the communication technology. Of course this requires Agent 2 be granted access to the reports filed by Agent 1. If there are no incentive problems there is no loss associated with letting Agent 2 acquire this information, while production assignments are chosen on the basis of richer information.

If agents type spaces included a larger number of possible types, with agents constrained to only send binary messages in any round, further gains could be achieved by adding more rounds of communication. With enough communication rounds, the agents would be able to report their true types. However, if each communication round involves time delays or other resource costs, the costs of having enough rounds that permit full revelation would become prohibitive if agents' private information were rich enough. With a limit on the number of rounds of binary communication which prevents agents from reporting everything they know, it is intuitively clear that the greatest amount of information could be communicated if the agents were to release information gradually, as this would enable them to tailor their reporting strategies in later rounds on the messages exchanged in previous rounds.

Providing more information to agents regarding messages sent by other agents in previous rounds may of course generate incentive problems. That is the problem that we address in this paper. To illustrate the nature of the incentive problem in dynamic communication protocols, we turn to the next example.

Example 2

We now present a different example which illustrates in the simplest possible way the special problems that arise in analyzing incentives in dynamic protocols.

Suppose that Agent 1's type θ_1 takes three possible values: 0, 1 and 2 with equal probability, while Agent 2's type θ_2 is takes two possible values 0 and 1 with equal probability where θ_i is unit cost of production for $i \in \{1, 2\}$. Consider a communication protocol with three rounds of binary messages. We abstract from Agent 2's reporting incentives by assuming he reports truthfully, and focus only on Agent 1's reporting incentives.

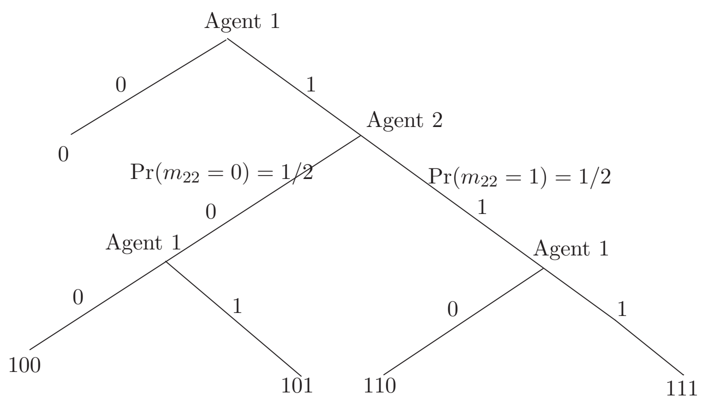

Figure 2: Example 2 (Extensive Form of Communication)

Round 1 Agent 1 sends a binary message $m_{11} \in \{0, 1\}$ in the first round. If $m_{11} = 0$, the communication ends and the mechanism specifies allocation $(t_1(0), q_1(0))$ for agent 1 where t_1 is a transfer to Agent 1 and q_1 is a quantity to be supplied by Agent 1. If Agent 1 chose $m_{11} = 1$, the game proceeds to Round 2.

Round 2 Agent 2 sends a binary message $m_{22} \in \{0, 1\}$, and the game proceeds to Round 3.

Round 3 Agent 1 sends a binary message $m_{13} \in \{0, 1\}$. The mechanism then specifies a transfer and quantity to be supplied by agent 1 as a function of messages exchanged at rounds 2 and 3: $(t_1(1, m_{22}, m_{13}), q_1(1, m_{22}, m_{13}))$.

At the beginning of each round, each agent knows messages sent by the other agent until the end of the previous round. Figure 2 describes the extensive form, which has five terminal nodes $h \in H \equiv \{0, 100, 101, 110, 111\}$.

Consider the question of incentive compatibility of the following communication strategies $(c_1(\theta_1), c_2(\theta_2))$, which specifies messages to be taken by agents for each decision node, recommended by the Principal:

Round 1 Agent 1 sends $m_{11} = 0$ if $\theta_1 = 0$, and $m_{11} = 1$ if $\theta_1 \in \{1, 2\}$.

Round 2 Agent 2 sends $m_{22} = 0$ if $\theta_2 = 0$, and $m_{22} = 1$ if $\theta_2 = 1$

Round 3 If $m_{22} = 0$, agent 1 sends $m_{13} = 0$ if $\theta_1 = 1$ and $m_{13} = 1$ if $\theta_1 = 2$. If $m_{22} = 1$, agent 1 sends $m_{13} = 0$ if $\theta_1 = 1$ and $m_{13} = 1$ if $\theta_1 = 2$.

	$(P(0), P(100), P(101), P(110), P(111))$
type 0	$(1, 0, 0, 0, 0)$
type 1	$(0, 1/2, 0, 1/2, 0)$
type 2	$(0, 0, 1/2, 0, 1/2)$

Figure 3: Example 2 (Parsimonious Protocol)

Given a quantity allocation $\{q_1(h) : h \in H\}$, the question is whether there exist transfers $\{t(h) : h \in H\}$ that ensure Agent 1's incentives to follow these recommendations, assuming Agent 2 follows them. Note that the message choices of type 0 of agent 1 in Round 3 (if he happened to have chosen $m_{11} = 1$ in Round 1) have not been specified. It will be necessary to consider possible choices that this type could make, in order to determine whether the preceding strategies are incentive compatible for agent 1.

The recommended strategies determine the probability with which various terminal nodes are achieved: let $P(h)$ denote the probability that terminal node h is reached. Figure 3 describes these probabilities for each type. Note that every terminal node is reached with positive probability. In other words, all message options provided at various rounds are used with positive probability. The protocol is *parsimonious* relative to communication strategies $(c_1(\theta_1), c_2(\theta_2))$, to use the terminology of Van Zandt (2007).

A necessary condition for incentive compatibility of these strategies for Agent 1 for some set of transfers, is that no type should want to deviate to the strategy prescribed for another type. Had the Principal been able to use a one-shot revelation mechanism in which each agent independently sends a type report, these conditions would also have been sufficient to ensure incentive compatibility. The reason is that the size of the message space is equal to the set of possible types, so the range of possible deviations available to any type is precisely the range of strategies utilized by other types. Following standard arguments, these conditions reduce to a single monotonicity condition on the expected quantity produced by Agent 1 with respect to his type report (when using Bayesian equilibrium as the solution concept):

$$q_1(0) \geq (1/2)q_1(100) + (1/2)q_1(110) \geq (1/2)q_1(101) + (1/2)q_1(111). \tag{1}$$

In the dynamic mechanism above, however, each type has a larger range of deviations available. Agent 1 has eight possible communication strategies to choose from, corresponding to various combinations of messages sent at Round 1, and those sent in Round

3 following the message sent by Agent 2 in Round 2. The question is whether the necessary condition for incentive compatibility described above suffices to ensure that no type of Agent 1 has a profitable deviation.

Van Zandt (2007) poses a related question which can be viewed as the converse of the Revelation Principle: does incentive compatibility in the static revelation mechanism ensure incentive compatibility in the dynamic mechanism (after the latter has been pruned to eliminate unused messages, i.e., the dynamic mechanism is parsimonious)? He points out the answer is yes, if the solution concept used is *ex post incentive compatibility (EPIC)*. This concept imposes the requirement that no type should regret his strategy choice at the end of the game, after learning the messages sent by the other agent. In the static revelation mechanism, in our example this requires type 0 to prefer $h = 0$ to both $h = 100$ and $h = 101$, nodes that could have been reached upon mimicking type 1's strategy. It also requires type 0 to prefer $h = 0$ to $h = 110$ and $h = 111$, nodes that it could have reached upon mimicking type 2's strategy. Hence type 0 weakly prefers the node $h = 0$ to all the four other terminal nodes. Similarly type 1 should not prefer either $h = 0$ or $h = 101$ to the node reached $h = 100$ when agent 2 sends message $m_{22} = 0$ at round 2. Nor should he prefer either $h = 0$ or $h = 111$ to the node $h = 110$ reached when agent 2 sent message $m_{22} = 1$ instead. An analogous set of inequalities for type 2 completes the set of requirements for incentive compatibility in the revelation mechanism.

It is obvious that these suffice to ensure incentive compatibility as well in the dynamic protocol. Essentially, the condition of 'no-regret after learning messages sent by other agents' implies that letting agents learn about messages sent by other agents does not disturb their incentive to follow the recommended strategies. Hence dynamic protocols do not entail additional incentive constraints when the solution concept is EPIC. In our context there is a simple necessary and sufficient condition for EPIC-incentivizability of a given quantity allocation, involving monotonicity of the assigned quantity with respect to the reported type of the agent, for every possible type reported by the other agent:

$$q_1(0) \geq q_1(100) \geq q_1(101) \tag{2}$$

and

$$q_1(0) \geq q_1(110) \geq q_1(111). \tag{3}$$

This argument does not extend to the case of Bayesian incentive compatibility. Here the incentives of agents to follow the recommended strategies may depend on their lack of knowledge of messages sent by other agents. In our example each type of agent 1 has available deviations which do not constitute strategies chosen by any other type. This is despite the fact that the protocol is parsimonious. For instance, type 0 could deviate to selecting $m_{11} = 1$ in Round 1, followed by $m_{13} = 0$ if $m_{22} = 0$ and $m_{13} = 1$ if $m_{22} = 1$.

This is a strategy not selected by any type of Agent 1. The condition that type 0 does not benefit by deviating to a strategy chosen by some other type, does not automatically ensure that it would not benefit by deviating to some other strategy.

To see that the dynamic protocol imposes additional incentive compatibility constraints, consider the question of characterizing quantity assignments $\{q_1(h) : h \in H\}$ that are Bayesian-incentivizable by some set of transfers. In a static revelation mechanism, incentivizability requires only (1), whereas in the dynamic protocol, the following additional constraints are required:

$$q_1(100) \geq q_1(101) \tag{4}$$

and

$$q_1(110) \geq q_1(111) \tag{5}$$

to ensure that in Round 3 types 1 and 2 do not want to imitate one another's message, having heard the message reported by Agent 2 in Round 2.

Hence the added flexibility of production assignments in a dynamic protocol may come at the expense of increasing the number of incentive constraints. This is the key trade-off involved in comparing dynamic with static protocols. In order to make progress with studying design of optimal mechanisms with limited communication, we need to characterize the precise set of incentive constraints associated with any given dynamic communication protocol.

Such a characterization is not available in existing literature. Fadel and Segal (2009) provide a set of sufficient conditions for Bayesian incentive compatibility in the dynamic protocol, which are stronger than the necessary conditions represented by the combination of (1), (4) and (5). Their Proposition 6 observes that EPIC-incentivizability implies BIC-incentivizability. Hence conditions (2) and (3) are sufficient for BIC-incentivizability. But these are stronger than the combination of (1), (4) and (5). These happen to be automatically satisfied in the informationally efficient allocation (incorporating only the communicational constraints) in contexts involving a single round of communication in many contexts (Melumad, Mookherjee and Reichelstein (1992, 1997), Blumrosen and Feldman (2006), Blumrosen, Nisan and Segal (2007)), as well as in some dynamic communication contexts (Kos (2011b)). In such contexts, thus, the informationally efficient allocation ends up being incentivizable, so the incentive constraints do not impose any additional cost. But this property is not true in general. For instance, it is not satisfied in the constrained efficient allocation in case (d) of Example 1.[8]

In this paper we show that conditions (1), (4) and (5) are collectively both necessary *and* sufficient for Bayesian-incentivizability in the dynamic protocol. The idea for the sufficiency argument is illustrated in Section 4. This property holds in general. It is a

[8] Agent 1's production increases from 1 to 2 as his cost rises from 0 to 10, in the case where the other agent's cost equals 0.

key step that allows us to pose the mechanism design problem as selection of a communication protocol and a contract (represented by quantities and transfers) to maximize the Principal's payoff subject to the constraints on the protocol and these incentive compatibility constraints.

3 Model

There are a Principal (P) and two agents 1 and 2. Agent $i = 1, 2$ produces a one-dimensional nonnegative real valued input q_i at cost $\theta_i q_i$, where θ_i is a real-valued parameter distributed over an interval $\Theta_i \equiv [\underline{\theta}_i, \bar{\theta}_i]$ according to a positive-valued, continuously differentiable density function f_i and associated c.d.f. F_i. The distribution satisfies the standard monotone hazard condition that $\frac{F_i(\theta_i)}{f_i(\theta_i)}$ is nondecreasing, implying that the 'virtual cost' $v_i(\theta_i) \equiv \theta_i + \frac{F_i(\theta_i)}{f_i(\theta_i)}$ is strictly increasing.[9] θ_1 and θ_2 are independently distributed, and these distributions F_1, F_2 are common knowledge among the three players.

The inputs of the two agents combine to produce a gross return $V(q_1, q_2)$ for P. In some contexts there may be technological restrictions on the relation between the outputs of the two agents. For instance the agents may work in a team that produce a joint output q, in which case $q_1 = q_2 = q$. Or the Principal may organize a procurement auction between two suppliers that produce perfect substitutes, to procure a given total quantity of the common good. Normalizing this desired quantity to 1, we can set $q_2 = 1 - q_1 \in [0, 1]$. In this case of course the constraint $q_2 = 1 - q_1$ is not a technological restriction, as the production function can be written as $V = \min\{q_1 + q_2, 1\}$. To accommodate the possibility of joint production, let $Q \subset R_+ \times R_+$ denote a set of technologically feasible combination of input supplies (q_1, q_2).

Additional restrictions will be imposed on V in Sections 6 and 7 which focus on the optimality of delegation and of optimal communication protocols. There we will assume that it is feasible for the two agents to select their outputs independently ($Q = R_+ \times R_+$), V is twice continuously differentiable, strictly concave, and satisfies the Inada condition that $V_i(q_1, q_2)$ tends to ∞ if q_i tends to 0. Our result concerning optimal communication protocols will additionally rely on the assumption that the marginal product of each input always depends on the quantity of the other input: $V_{12}(q_1, q_2) \neq 0$ for any $(q_1, q_2) \neq 0$.

The Principal makes transfer payments t_i to i. The payoff of i is $t_i - \theta_i q_i$. Both agents are risk-neutral and have autarkic payoffs of 0. The Principal's objective takes the form

$$V(q_1, q_2) - \lambda_1(t_1 + t_2) - \lambda_2(\theta_1 q_1 + \theta_2 q_2) \tag{6}$$

[9] Our results can be extended in the absence of this assumption, employing the 'ironing' technique developed by Myerson (1981) and Baron and Myerson (1982).

where $\lambda_1 \geq 0, \lambda_2 \geq 0$ and $(\lambda_1, \lambda_2) \neq 0$ respectively represent welfare weights on the cost of transfers incurred by the Principal and cost of production incurred by the agents. Applications include:

Application 1: Internal organization/procurement

The Principal is the owner of a firm composed of two divisions whose respective outputs combine to form revenues $V = V(q_1, q_2)$. If there is a single joint output, $Q = \{(q_1, q_2) | q_1 = q_2\}$. If the two divisions produce different inputs independently, $Q = R_+ \times R_+$. In this case there are no production externalities across the two agents. The principal seeks to maximize profit, hence $\lambda_1 = 1$ and $\lambda_2 = 0$. The same applies when the two agents correspond to external input suppliers.

Application 2: Regulation

The Principal is a regulator seeking to control outputs or abatements q_i of two firms $i = 1, 2$ with $Q = R_+ \times R_+$. $V(q_1 + q_2)$ is the gross social benefit, and θ_i is the firm i's unit cost. Consumer welfare equals $V - (1 + \lambda)F$ where F is the total tax revenue collected from consumers and λ is the deadweight loss involved in raising these taxes. The revenue is used to reimburse transfers t_1, t_2 to the firms. Social welfare equals the sum of consumer welfare and firm payoffs, which reduces to (6) with $\lambda_1 = \lambda, \lambda_2 = 1$.

Application 3: Allocating private goods

The Principal allocates a fixed quantity q of a private good. In this environment, θ_i is negative, with $-\theta_i$ representing agent i's valuation of one unit of product, and $-t_i$ the amount paid by agent i. $Q = \{(q_1, q_2) \in R_+ \times R_+ \mid q_1 + q_2 \leq q\}$. $V(q_1, q_2) = 0$, $\lambda_1 = 1$ and $\lambda_2 = 0$ represents the case where the Principal is solely concerned about revenue, while $V(q_1, q_2) = 0$, $\lambda_1 = 0$ and $\lambda_2 = 1$ represents the case where the Principal is concerned about efficiency. The auction of a single item corresponds to the case of $Q = \{(q_1, q_2) \in \{0, 1\} \times \{0, 1\} \mid q_1 + q_2 \leq 1\}$.

Application 4: Public good decisions

Here $q = q_1 = q_2$ represents the level of a public good, whose valuation by agent i is $-\theta_i > 0$. Here nonrivalry and nonexcludability implies $Q = \{(q_1, q_2) \mid q_1 = q_2\}$. $V(q, q) = -C(q)$ is interpreted as the cost of producing the public good. These costs are covered by taxes raised from consumers, which involve a deadweight loss of λ. Social welfare corresponds to (6) with $\lambda_1 = 0, \lambda_2 = \frac{1}{1+\lambda}$.

4 Communication and Contracting

4.1 Timing

The mechanism is designed by the principal at an ex-ante stage ($t = -1$). It consists of a *communication protocol* (explained further below) and a set of contracts to each agent. There is enough time between $t = -1$ and $t = 0$ for all agents to read and understand the offered contracts.

At $t = 0$, each agent i privately observes the realization of θ_i, and independently decides whether to participate or opt out of the mechanism. If either agent opts out the game ends; otherwise they enter the planning or communication phase which lasts until $t = T$.

Communication takes place in a number of successive rounds $t = 1, \ldots, T$. We will abstract from mechanisms in which the Principal seeks to limit the flow of information across agents, either by appointing mediators, regulators or scrambling devices. Later we argue that the optimal allocation is implemented with this communication structure, i.e., it is not profitable to restrict or garble the flow of information across agents. Hence this restriction will turn out to entail no loss of generality. This simplifies the exposition considerably.

The Principal is assumed to be able to verify all messages exchanged between agents. Equivalently, an exact copy of every message sent by one agent to another is also sent to the Principal. This rules out collusion between the agents, and allows the Principal to condition transfers *ex post* on messages exchanged. Given that agents exchange messages directly with one another and the absence of any private information possessed by the Principal, there is no rationale for the Principal to send any messages to the agents. In what follows we will not make the Principal's role explicit in the description of the communication protocol, and will focus on the exchange of communication between the agents.[10]

At the end of round T, each agent $i = 1, 2$ (or the Principal) selects production level q_i, depending on whose choice variable q_i is (an issue discussed further below).

Finally, after production decisions have been made, payments are made according to the contracts signed at the *ex ante* stage, and verification by the Principal of messages exchanged by agents and outputs produced by them.

[10] As mentioned above, any mechanism in which agents send some messages to the Principal but not to each other, will end up being weakly dominated by a mechanism in which these messages are also sent to other agents. Hence there is no need to consider mechanisms where agents communicate privately with the Principal.

4.2 Communication Protocol

A communication protocol is a rule defining T the number of rounds of communication, and the message set M_i of each agent i in any given round, which may depend on the history of messages exchanged in previous rounds. If some agents are not supposed to communicate anything in any round, their message sets are null in those rounds. This allows us to include protocols where agents take turns in sending messages in different rounds. Other protocols may involve simultaneous reporting by all agents in each round.

The *vocabulary* of any agent $i \in I \equiv \{1, 2\}$ is a message set $\mathcal{M}^i$, which contains all messages m_i that i can feasibly send in a single round. This incorporates restrictions on the language that agents use to communicate with one another. Specific assumptions concerning such restrictions are introduced below.

The message set M_i assigned to agent i in any round is a subset of the vocabulary of that agent. Message histories and message sets are defined recursively as follows. Let m_{it} denote a message sent by i in round t. Given a history h_{t-1} of messages exchanged (sent and received) by i until round $(t-1)$, it is updated at round t to include the messages exchanged at round t: $h_t = (h_{t-1}, \{m_{it}\}_{i \in I})$. And for every i, $h_0 = \emptyset$. The message set for i at round t is then a subset of $\mathcal{M}^i$ which depends on h_t, unless it is null.

Formally, the *communication protocol* specifies the number of rounds T, and for every round $t \in \{1, \ldots, T\}$ and every agent i, a message set $M_i(h_{t-1}) \subseteq \mathcal{M}^i$ or $M_i(h_{t-1}) = \emptyset$ for every possible history h_{t-1} until the end of the previous round.[11]

4.3 Communication Costs

We now describe communication costs. We focus only on costs of writing or composing messages by senders, and assume that receiving (reading or listening) a message does not entail any time delay or any other cost. This applies to verbal exchanges where speaking is time-consuming and the audience listens at the same time that the speaker speaks. In any case, this is the implicit assumption in most existing literature on costly communication.

We allow agents the option of not sending any message at all in any given round: hence the null message $\phi \in \mathcal{M}_i$. Let $l(m_i)$ denote the *length* of message $m_i \in \mathcal{M}_i$, which is an integer. It is natural to assume $l(\phi) = 0$, and positive-valued for any other message. For example if messages are binary-encoded, $l(m_i)$ could denote the total number of 0's and 1 bits included in m_i.

The following assumption imposes the fundamental limitation on communication, stating

[11] We depart from Fadel and Segal (2009) and Van Zandt (2007) insofar as their definition of a protocol combines the extensive form game of communication as well as the communication strategy of each agent.

that vocabularies consisting of messages not exceeding any given length are finite:[12]

Assumption 1 *For any $k < \infty$, there exists an integer $n < \infty$ such that $\#\{m_i \in \mathcal{M}_i \mid l(m_i) < k\} < n$.*

We focus on two possible dimensions of communication costs: material costs and time delays. Examples of the former are costs of telecommunication: faxes, telephones, electronic mail or videoconferences. Some of these vary with the length of actual messages sent. Others are fixed costs of maintaining a communication channel of a certain capacity (defined by the maximum length of messages that could be sent). Formally, the communication capacity of each agent i depends on the message set M_i assigned to that agent in any given round, and is defined as the longest message contained in M_i: defined for any given message set $M_i \subset \mathcal{M}_i$ as follows: $\bar{l}(M_i) \equiv \max_{m_i \in M_i} l(m_i)$. We assume that the material cost of communication incurred by the organization in any given round in connection with agent i who selects m_i from her message set M_i is given by a function $\Phi_m(l(m_i), \bar{l}(M_i))$ which is increasing in the length of the message (representing the variable cost component) and a fixed cost which depends on the channel capacity. In particular, when material costs of communication exist, we assume that the marginal cost of increasing capacity is bounded away from zero:

Assumption 2 *If there are material costs of communication, there exists a lower bound $\phi_f > 0$ to the increase in cost Φ_m corresponding to a unit increase in communication capacity $\bar{l}(M_i)$.*

This ensures that some material costs are incurred even if an agent sends a null message, as long as the agent had the option to send some non-null message. In other words, sending null messages is costly whenever it is informative.[13] In deriving results concerning optimal communication protocols we will impose additional assumptions on the communication cost function pertaining to the nature of variable costs.

The other component of communication cost is time delay. Production decisions are deferred to the end of the communication phase. Hence a longer communication phase

[12] $\#A$ denotes the number of elements in a finite set A.

[13] In the absence of such an assumption, it is possible for agents to communicate a lot of information using the timing of non-null messages at negligible cost. Suppose that communication costs depend only on the length of messages sent while communication capacity is costless. Then communication cost is zero whenever a null message is sent. Consider some arbitrary non-null message m^*, such as a single bit of information. Then it is possible to set $T = \infty$ and agent i can communicate any rational number at a total cost equal to the cost of sending m^* once, by selecting a single round t when he sends message m^* and a null message at every other round. This is not possible under Assumption 2 since each round will entail a positive capacity cost, thereby imposing an upper bound to the number of communication rounds for any finite communication budget.

delays production which can result in foregone sales or loss in quality of goods produced. The time taken to communicate a message increases with its length. Let this be represented by the function $\delta(l)$. The length of any communication round can be set equal to the time taken to communicate the longest message by either agent in that round. Hence the delay in any round is

$$\Phi_d(M_i, M_j) = \delta(\max\{\bar{l}(M_i), \bar{l}(M_j)\}),$$

independent of the length of messages actually sent in that round. This is because message recipients have to wait till the end of the round to read the entire message sent by the sender: the same delay has to be incurred if agents send shorter messages than the longest one available.[14] Analogous to assumption 2 above, we assume here that

Assumption 3 *If there are costs of time delay, there exists a lower bound $\phi_d > 0$ to the increase in delay $\delta(l)$ corresponding to a unit increase in message length l.*

While the organization can choose its communication cost (budget and/or delay) optimally, it suffices for us to consider the optimal mechanism designed corresponding to any given (finite) cost, as the properties we will be focusing on will not vary with the precise cost.[15] Let the fixed budget for material costs and maximum delay be denoted by K and D respectively. These constraints impose restrictions on feasible choice of communication protocols.

We need the following notation to describe these constraints. For a given communication protocol, let $\bar{H}$ be the set of possible sequences of histories $(h_0, h_1 ...., h_T)$ attained in each round along communication path. $(h_0, .., h_T) \in \bar{H}$ also specifies the message that i sends for a given history h_{t-1} in round t, which is denoted by $m_i(h_{t-1}, h_t) \in M_i(h_{t-1})$ as a function of h_{t-1} and h_t. Then the communication constraints are defined as follows:

Definition 1 *(i) A communication protocol is feasible under a budget for material cost K if and only if*

$$\max_{(h_0,...,h_T)\in\bar{H}} \Sigma_{t=1}^T \Sigma_{i\in\{1,2\}} \Phi_m(m_i(h_{t-1}, h_t), M_i(h_{t-1})) \leq K.$$

(ii) A communication protocol is feasible under total delay D if and only if

$$\max_{(h_0,...,h_T)\in\bar{H}} \Sigma_{t=1}^T \Phi_d(M_i(h_{t-1}), M_j(h_{t-1})) \leq D.$$

[14] Note that each agent has the option of sending some message in the early part of a round, remaining silent for some time, and then may or may not send an additional message for the remainder of the round. The other agent would have to wait till the very end of the round, to see which of these options were chosen by the sending agent.

[15] In other words, we can view as the choice of communication cost being made at a second stage, while at the first stage the optimal mechanism conditional on a given cost is decided.

Let a protocol be denoted by p, and the set of feasible protocols given the communication constraints with cost limits (K, D) be denoted by $\mathcal{P}(K, D)$. In what follows, the dependence on (K, D) will be suppressed unless absolutely necessary.

4.4 Communication Plans and Strategies

Given a protocol $p \in \mathcal{P}$, a *communication plan* for agent i specifies for every round t a message $m_{it}(h_{t-1}) \in M_i(h_{t-1})$ for every possible history h_{t-1} that could arise for i in protocol p until round $t-1$. The set of communication plans for i in protocol p is denoted $C_i(p)$. For communication plan $c = (c_1, c_2) \in C(p) \equiv C_1(p) \times C_2(p)$, let $h_t(c)$ denote the history of messages generated thereby until the end of round t. Let $H_t(p) \equiv \{h_t(c) \mid c \in C(p)\}$ denote the set of possible message histories in this protocol until round t. For a given protocol, let $\mathcal{H} \equiv H_T(p)$ denote the set of possible histories at the end of round T.

We now show that our assumptions regarding communication costs imply that agents have available a finite number of communication plans in any protocol subject to finite cost limits.

Lemma 1 *If either $K < \infty$ or $D < \infty$ holds, then $\max_{p \in \mathcal{P}(K,D)} \#C(p)$ has an upper bound.*

The reasoning is straightforward: under Assumption 2 and 3, each communication round entails a cost bounded away from zero, implying that the number of rounds must be finite given any communication budget. Moreover, communication capacity in any given round must be finite. Hence the set of possible communication plans for each agent must also be finite.

Given a protocol $p \in \mathcal{P}$, a *communication strategy* for agent i is a mapping $c_i(\theta_i) \in C_i(p)$ from the set $\Theta_i \equiv [\underline{\theta}_i, \bar{\theta}_i]$ of types of i to the set $C_i(p)$ of possible communication plans for i. In other words, a communication strategy describes the dynamic plan for sending messages, for every possible type of the agent. The finiteness of the set of dynamic communication plans implies that it is not possible for others in the organization to infer the exact type of any agent from the messages exchanged. Non-negligible sets of types will be forced to pool into the same communication plan.

4.5 Production Decisions and Contracts

Many authors in previous literature (Blumrosen and Feldman (2006), Blumrosen, Nisan and Segal (2007) and Kos (2011a, 2011b)) have limited attention to mechanisms where output assignments and transfers are specified as a function of the information communicated by the agents. Decision-making authority is effectively retained by the Principal in this case. This is natural in settings involving auctions or public goods. We shall refer to such mechanisms as *centralized.* A *contract* in this setting specifies a quantity allocation $q(h) \equiv (q_1(h), q_2(h)) : \mathcal{H} \to Q$, with corresponding transfers $t(h) \equiv (t_1(h), t_2(h)) : \mathcal{H} \to \Re \times \Re$. A *centralized mechanism* is then a communication protocol $p \in \mathcal{P}$ and an associated contract $(q(h), t(h)) : \mathcal{H} \to Q \times \Re \times \Re$.

Some authors (Melumad, Mookherjee and Reichelstein (1992, 1997)) have explored mechanisms where the Principal delegates decision-making to one of the two agents, and compared their performance with centralized mechanisms. This is a pertinent question in procurement, internal organization or regulation contexts. They consider mechanisms where both contracting with the second agent as well as production decisions are decentralized (while restricting attention to communication protocols involving a single round of communication). Here we focus attention on mechanisms where the Principal retains control over the design of contracts with both agents, while decentralizing decision-making authority to agents concerning their own productions. This is feasible only if the production decisions of the two agents can be chosen independently, i.e., there are no technical complementarities or jointness restrictions on their outputs. We refer to such mechanisms as *decentralized.* The potential advantage of decentralizing production decisions to agents is that these decisions can be based on information possessed by the agents which is richer than what they can communicate to the Principal. Transfers can then be based on output decisions as well as messages exchanged.

Formally, when the feasible output space is $Q = \Re_+ \times \Re_+$, a *decentralized mechanism* is a communication protocol p and a pair of contracts for the two agents, where the contract for agent i is a transfer rule $t_i(q_i, h) : \Re_+ \times \mathcal{H} \to \Re$. Such a mechanism induces a quantity allocation $q_i(\theta_i, h) : \Theta_i \times \mathcal{H} \to \Re_+$ which maximizes $[t_i(q_i, h) - \theta_i q_i]$ with respect to choice of $q_i \in \Re_+$.[16] To simplify exposition we specify the quantity allocation as part of the decentralized mechanism itself.

A centralized mechanism can be viewed as a special case of a decentralized mechanism in which $q_i(\theta_i, h)$ is measurable with respect to h, i.e., does not depend on θ_i conditional on h. It corresponds to a mechanism in which the Principal sets an output target for each agent (based on the messages communicated) and then effectively forces them to meet these targets with a corresponding incentive scheme. We can therefore treat every mechanism as decentralized, in a formal sense.

In view of this, say that a mechanism is *truly decentralized* if it is not centralized. We

[16] Since i infers the other's output q_j $(j \neq i)$ only through h, we can restrict attention to contracts where the payments to any agent depends only on his own output without loss of generality. Specifically, if t_i were to depend on q_j, the expected value of the transfer to i can be expressed as a function of q_i and h, since agent i's information about q_j has to be conditioned on h.

shall in due course evaluate the relative merits of centralized and truly decentralized mechanisms.

4.6 Feasible Production Allocations

The standard way of analysing the mechanism design problem with unlimited communication is to first characterize production allocations that are feasible in combination with some set of transfers, and then use the Revenue Equivalence Theorem to represent the Principals objective in terms of the production allocation alone, while incorporating the cost of the supporting transfers. To extend this method we seek to characterize feasible production allocations.

A *production allocation* is a mapping $q(\theta) \equiv (q_1(\theta), q_2(\theta)) : \Theta_1 \times \Theta_2 \to Q$. Restrictions are imposed on production allocations owing both to communication and incentive problems.

Consider first communication restrictions. A production allocation $q(\theta)$ is said to be *communication-feasible* if: (a) the mechanism involves a communication protocol p satisfying the specified constraints on communication, and (b) there exist communication strategies $c(\theta) = (c_1(\theta_1), c_2(\theta_2)) \in C(p)$ and output decisions of agents $q_i(\theta_i, h) : \Theta_i \times \mathcal{H} \to \Re_+$, such that $q(\theta) = (q_1(\theta_1, h(c(\theta))), q_2(\theta_2, h(c(\theta))))$ for all $\theta \in \Theta$. Here $h(c)$ denotes the message histories generated by the communication strategies c in this protocol.

The other set of constraints pertain to incentives. A communication-feasible production allocation $q(\theta)$ is said to be *incentive-feasible* in a mechanism if there exists a Perfect Bayesian Equilibrium (PBE) of the game induced by the mechanism which implements the production allocation.[17] In other words, there must exist a set of communication strategies and output decision strategies satisfying condition (b) above in the requirement of communication-feasibility, which constitutes a PBE.

4.7 Characterization of Incentive Feasibility

We now proceed to characterize incentive-feasible production allocations. Using the single-dimensional output of each agent and the single crossing property of agent preferences, we can obtain as a necessary condition a monotonicity property of expected outputs with respect to types at each decision node. To describe this condition, we need the following notation.

It is easily checked (see Lemma 2 in the Appendix) that given any strategy configuration $c(\theta) \equiv (c_1(\theta_1), c_2(\theta_2))$ and any history h_t until the end of round t in a communication protocol, the set of types (θ_1, θ_2) that could have generated the history h_t can be expressed

[17] This requires both incentive and participation constraints be satisfied.

as the Cartesian product of subsets $\Theta_1(h_t), \Theta_2(h_t)$ such that

$$\{(\theta_1, \theta_2) \mid h_t(c(\theta_1, \theta_2)) = h_t\} = \Theta_1(h_t) \times \Theta_2(h_t). \tag{7}$$

A necessary condition for incentive-feasibility of a production allocation $q(\theta)$ which is communication-feasible in a protocol p and supported by communication strategies $c(\theta)$ is that for any $t = 0, \ldots, T$, any $h_t \in H_t$ and any $i = 1, 2$:

$$E[q_i(\theta_i, \theta_j) \mid \theta_j \in \Theta_j(h_t)] \text{ is non-increasing in } \theta_i \text{ on } \Theta_i(h_t), \tag{8}$$

where H_t denotes the set of possible histories until round t generated with positive probability in the protocol when $c(\theta)$ is played, and $\Theta_i(h_t)$ denotes the set of types of i who arrive at h_t with positive probability under the communication strategies $c(\theta)$.

The necessity of this condition follows straightforwardly from the dynamic incentive constraints which must be satisfied for any history h_t on the equilibrium path. Upon observing h_t, i's beliefs about θ_j are updated by conditioning on the event that $\theta_j \in \Theta_j(h_t)$. Any type of agent i in $\Theta_i(h_t)$ will have chosen the same messages up to round t. Hence any type $\theta_i \in \Theta_i(h_t)$ has the opportunity to pretend to be any other type in $\Theta_i(h_t)$ from round $t+1$ onward, without this deviation being discovered by anyone. A PBE requires that such a deviation cannot be profitable. The single-crossing property then implies condition (8).

As noted earlier, the existing literature has provided sufficient conditions for incentive-feasibility that are stronger than (8). Fadel and Segal (2009) in a more general framework (with abstract decision spaces and no restrictions on preferences) provide two sets of sufficient conditions. One set (provided in their Proposition 6) of conditions is based on the observation that the stronger solution concept of ex post incentive compatibility implies Bayesian incentive compatibility. In our current context ex post incentive compatibility is assured by the condition that for each $i = 1, 2$:

$$q_i(\theta_i, \theta_j) \text{ is globally non-increasing in } \theta_i \text{ for every } \theta_j \in \Theta_j \tag{9}$$

Another set of sufficient conditions (Proposition 3 in Fadel-Segal (2009)) imposes a no-regret property with respect to possible deviations to communication strategies chosen by other types following every possible message history arising with positive probability under the recommended communication strategies. This is applied to every pair of types for each agent at nodes where it is this agent's turn to send a message. In the context of centralized mechanisms (which Fadel and Segal restrict attention to), this reduces to the condition that for any any $i = 1, 2$ and any $h_t \in H_t, t = 0, \ldots, T-1$ where it is i's turn to move (i.e., $M_i(h_t) \neq \emptyset$):[18]

$$E[q_i(\theta_i, \theta_j)|\theta_j \in \Theta_j(h_t)] \text{ is globally non-increasing in } \theta_i. \tag{10}$$

Our first main result is that the necessary condition (8) is also sufficient for incentive feasibility, provided the communication protocol prunes unused messages. Suppose that p is a communication protocol in which communication strategies used are $c(\theta)$. Then p is *parsimonious relative to communication strategies* $c(\theta)$ if every possible history $h \in \mathcal{H}$ in this protocol is reached with positive probability under $c(\theta)$.

Proposition 1 *Condition (8) is sufficient for incentive-feasibility of a production allocation $q(\theta)$ which is communication-feasible in a protocol p and supported by communication strategies $c(\theta)$, provided the protocol is parsimonious with respect to $c(\theta)$.*

Parsimonious protocols have the convenient feature that Bayes rule can be used to update beliefs at every node, and off-equilibrium-path deviations do not have to be considered while checking incentive feasibility. Restricting attention to such protocols entail no loss of generality since any protocol can be pruned by deleting unused messages under any given set of communication strategies, to yield a protocol which is parsimonious with respect to these strategies. Hence it follows that condition (8) is both necessary and sufficient for incentive-feasibility.

The proof of Proposition 1 is provided in the Appendix. The main complication arises for the following reason. In a dynamic protocol with more than one round of communication, no argument is available for showing that attention can be confined to communication strategies with a threshold property. Hence the set of types $\Theta_i(h_t)$ pooling into message history h_t need not constitute an interval. The monotonicity property for output decisions in (8) holds only 'within' $\Theta_i(h_t)$, which may span two distinct intervals. The monotonicity property may therefore not hold for type ranges lying between the two intervals. This complicates the conventional argument for construction of transfers that incentivize a given output allocation.

The proof is constructive. Given a production allocation satisfying (8) with respect to set of communication strategies in a protocol, we first prune the protocol to eliminate unused messages. Then incentivizing transfers are constructed as follows. We start by defining a set of functions representing expected outputs of each agent following any given history h_t at any stage t, expressed as a function of the type of that agent. Condition (8) ensures the expected output of any agent i is monotone over the set $\Theta_i(h_t)$. These

[18]As Fadel and Segal point out, it suffices to check the following condition at the last node of the communication game at which it is agent i's turn to move. Note also that this condition is imposed on nodes of the communication game, and not at nodes where agents make output decisions in the case of a decentralized mechanism.

are the types of i that actually arrive at h_t with positive probability on the equilibrium path. The proof shows it is possible to extend this function over all types of this agent (not just those that arrive at h_t on the equilibrium path) which is globally monotone, in a way that agrees with the actual expected outputs on the set $\Theta_i(h_t)$, and which maintains consistency across histories reached at successive dates. This amounts to assigning outputs for types that do not reach h_t on the equilibrium path, which can be thought of as outputs they would be assigned if they were to deviate somewhere in the game and arrive at h_t. Since this extended function is globally monotone, transfers can be constructed in the usual way to incentivize this allocation of expected output. The construction also has the feature that the messages sent by the agent after arriving at h_t do not affect the expected outputs that would thereafter be assigned to the agent, which assures that the agent does not have an incentive to deviate from the recommended communication strategy. Moreover, interim participation constraints are satisfied.

Consider the following example which illustrates the construction of transfers that incentivize an allocation satisfying the necessary condition (8). Agent i's cost is uniformly distributed over $[0, 1]$. There are three rounds of communication. In round 1, agent 1 reports $m_{11} \in \{E, C\}$. In round 2, agent 2 reports $m_{22} \in \{U, D\}$. In round 3, agent 1 reports $m_{13} \in \{0, 1\}$. The mechanism is decentralized, with each agent choosing their respective outputs at the end of round 3. The recommended communication strategies (on the equilibrium path) are the following. In round 1, agent 1 reports $m_{11} = E$ if $\theta_1 \in [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$, and C otherwise. If $m_{11} = E$ then in round 2, agent 2 reports $m_{22} = U$ if $\theta_2 \in [\frac{1}{3}, 1]$ and D otherwise. And if $m_{11} = C$, agent 2 reports $m_{22} = U$ if $\theta_2 \in [\frac{2}{3}, 1]$ and D otherwise. In round 3, agent 1 has a null message set if he reported C in round 1. If he reported E in round 1, in round 3 he reports $m_{13} = 0$ if $\theta_1 \in [0, \frac{1}{3}]$ and $m_{13} = 1$ if $\theta_1 \in [\frac{2}{3}, 1]$.

We focus on constructing transfers for agent 1 so as to induce this agent to follow the recommended strategy, while assuming that agent 2 follows his. Hence we check only the reporting incentives for agent 1 in rounds 1 and 3. These depend on outputs that agent 1 is expected to select at the end of round 3, as a function of messages sent in the first three rounds, besides the true type of agent 1. These outputs are represented by functions $e(\theta_1, m_{22})$ and $c(\theta_1, m_{22})$ corresponding to first round announcements of E and C respectively, and are shown in Figure 4.

Define first-round expected outputs $m(\theta_1)$ as follows:

$$m(\theta_1) \equiv (1/3)e(\theta_1, D) + (2/3)e(\theta_1, U)$$

for $\theta_1 \in [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$, and

$$m(\theta_1) \equiv (2/3)c(\theta_1, D) + (1/3)c(\theta_1, U)$$

for $\theta_1 \in (\frac{1}{3}, \frac{2}{3})$. This is also shown in Figure 4.

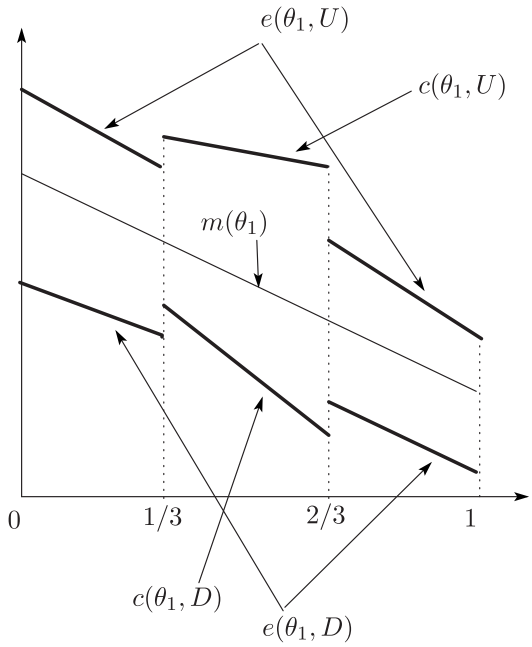

Figure 4: Outputs for Agent 1

Then the necessary condition for incentive feasibility is that the following monotonicity properties of first-round and third-round expected outputs of agent 1 are satisfied:

- $m(\theta_1)$ is non-increasing in $\theta_1 \in [0,1]$
- $e(\theta_1, D)$ and $e(\theta_1, U)$ are non-increasing in θ_1 on $[0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$
- $c(\theta_1, D)$ and $c(\theta_1, U)$ are non-increasing in θ_1 on $(\frac{1}{3}, \frac{2}{3})$

If we now define the output functions $q_U(\theta_1) = e(\theta_1, U)$ if $\theta_1 \notin [\frac{1}{3}, \frac{2}{3}]$ and equal to $c(\theta_1, U)$ otherwise, and analogously $q_D(\theta_1) = e(\theta_1, D)$ if $\theta_1 \notin [\frac{1}{3}, \frac{2}{3}]$ and equal to $c(\theta_1, D)$ otherwise, it is evident from Figure 4 that these functions are not globally monotone in θ_1.

As a first step in our construction, we extend the third round output functions e, c to $\tilde{e}, \tilde{c}$ to the entire type space $[0,1]$ as shown in Figures 5, 6 respectively. This can be done to maintain the following two properties: (a) the extended functions are globally monotone in θ_1, and (b) their average equals $m(\theta_1)$ for all θ_1. These can be interpreted as outputs corresponding to off-equilibrium messages.

Let $\tilde{q}_1(\theta_1, h_3)$ denote the round-3 output of agent 1 corresponding to type θ_1 and message vector h_3 consisting of messages sent in the first three rounds, based on these extended functions. The function $\tilde{q}_1$ is well-defined for all θ_1 and all possible message histories till round 3, both on and off the equilibrium path. Now define transfers to agent 1 as follows:

$$t(\tilde{q}_1(\theta_1, h_3), h_3) = \theta_1 \tilde{q}_1(\theta_1, h_3) + \int_{\theta_1}^{1} \tilde{q}_1(y, h_3) dy \tag{11}$$

This ensures that every type θ_1 will optimally select output of $\tilde{q}_1(\theta_1, h_3)$ at the end of the communication phase following any given history h_3, both on and off the equilibrium path. Moreover, by construction, after conditioning on the output $\tilde{q}_1$ that agent 1 expects to choose at the end of round 3, its transfer does not depend on the message m_{13} sent in that round.[19] Hence no type of Agent 1 has an incentive to deviate from the recommended communication strategy on the equilibrium path in round 3.

Now consider reporting incentives in round 1. At this stage Agent 1 does not yet know the message to be sent by Agent 2 in round 2. By construction of the extended output functions, the expected output for an agent with true type θ_1 is $m(\theta_1)$, irrespective of whether or not it deviates from the recommended reporting strategy in round 1. Its expected transfer in round 1 therefore do not depend on the messages it sends in round 1, implying it has no incentive to deviate in round 1 either.

[19]Specifically, note that $\tilde{q}_1(\theta_1, h_3)$ equals $\tilde{e}(\theta_1, m_{22})$ for h_3 such that $m_{11} = E$, and equals $\tilde{c}(\theta_1, m_{22})$ for h_3 such that $m_{11} = C$.

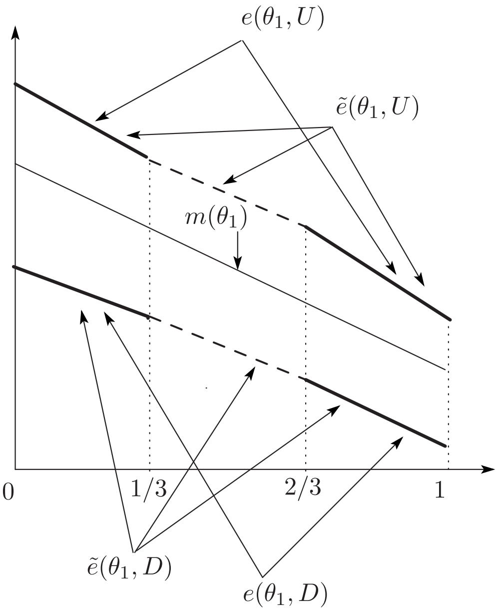

Figure 5: Construction of $\tilde{e}$

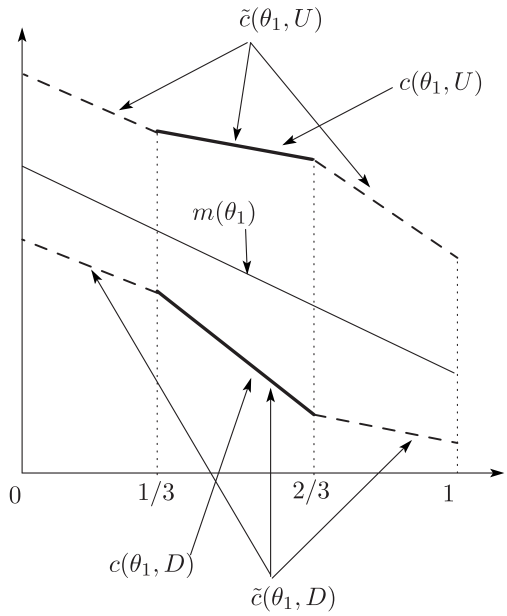

Figure 6: Construction of $\tilde{c}$

Note finally that taking expectations of (11) with respect to types of Agent 2 combined with the equilibrium reporting strategies, the ex ante expected transfers paid by the Principal to Agent 1 depends only on the first round expected output function $m(\theta_1)$, i.e., do not depend on the particular way in which the round 3 output functions were extended.

5 Characterizing Optimal Mechanisms

Having characterized feasible allocations, we can now restate the mechanism design problem as follows.

Note that the interim participation constraints imply that every type of each agent must earn a non-negative expected payoff from participating. Agents that do not participate do not produce anything or receive any transfers. Hence by the usual logic it is without loss of generality that all types participate in the mechanism. The single crossing property ensures that expected payoffs are nonincreasing in θ_i for each agent i. Since $\lambda_1 \geq 0$ it is optimal to set transfers that incentivize any given output allocation rule $q(\theta)$ satisfying (8) such that the expected payoff of the highest cost type $\bar{\theta}_i$ equals zero for each i. The expected transfers to the agents then equal (using the arguments in Myerson (1981) to establish the Revenue Equivalence Theorem):

$$\Sigma_{i=1}^{2} E[v_i(\theta_i) q_i(\theta_i, \theta_j)]$$

where $v_i(\theta_i) \equiv \theta_i + \frac{F_i(\theta_i)}{f_i(\theta_i)}$. Consequently the expected payoff of the Principal is

$$E[V(q_i(\theta_i, \theta_j), q_j(\theta_i, \theta_j)) - w_i(\theta_i) q_i(\theta_i, \theta_j) - w_j(\theta_j) q_j(\theta_i, \theta_j)] \tag{12}$$

where $w_i(\theta_i) \equiv (\lambda_1 + \lambda_2)\theta_i + \lambda_1 \frac{F_i(\theta_i)}{f_i(\theta_i)}$.

This enables us to state the problem in terms of selecting an output allocation in combination with communication protocol and communication strategies. Given the set $\mathcal{P}$ of feasible communication protocols defined by the communication constraints, the problem is to select a protocol $p \in \mathcal{P}$, communication strategies $c(\theta)$ in p and output allocation $q(\theta)$ to maximize (12), subject to the constraint that (i) there exists a set of output decision strategies $q_i(\theta_i, h), i = 1, 2$ such that $q(\theta) = (q_1(\theta_1, h(c(\theta))), q_2(\theta_2, h(c(\theta))))$ for all $\theta \in \Theta$, and (ii) the output allocation satisfies condition (8).

Condition (i) is essentially a communication-feasibility constraint, which applies even in the absence of incentive problems. Condition (ii) is the additional constraint represented by incentive problems. Note that the above statement of the problem applies since attention can be confined without loss of generality to protocols that are parsimonious with respect to the assigned communication strategies. To elaborate, note that conditions

(i) and (ii) are both necessary for implementation. Conversely, given an output allocation, a communication protocol, and communication strategies in the protocol that satisfy conditions (i) and (ii), we can prune that protocol by deleting unused messages to obtain a protocol that is parsimonious with respect to the given communication strategies. Then Proposition 1 ensures that the output allocation can be implemented as a PBE in the pruned protocol with suitably constructed transfers, which generate an expected payoff (12) for the Principal while ensuring all types of both agents have an incentive to participate.

Now observe that constraint (ii) is redundant in this statement of the problem. If we consider the relaxed version of the problem stated above where (ii) is dropped, the solution to that problem must automatically satisfy (ii), since the monotone hazard rate property on the type distributions F_i ensure that $w_i(\theta_i)$ is an increasing function for each i. This generates our main result.

Proposition 2 *The mechanism design problem can be reduced to the following. Given the set $\mathcal{P}$ of feasible communication protocols defined by the communication constraints, select a protocol $p \in \mathcal{P}$, communication strategies $c(\theta)$ in p and output allocation $q(\theta)$ to maximize (12), subject to the constraint of communication feasibility alone, i.e., there exists a set of output decision strategies $q_i(\theta_i, h), i = 1, 2$ such that*

$$q(\theta) = (q_1(\theta_1, h(c(\theta))), q_2(\theta_2, h(c(\theta)))), \forall \theta \in \Theta. \tag{13}$$

In the case of unlimited communication, this reduces to the familiar property that an optimal output allocation can be computed on the basis of unconstrained maximization of expected payoffs (12) of the Principal which incorporate incentive rents earned by the agents. With limited communication additional constraints pertaining to communication feasibility have to be incorporated. In the absence of incentive problems, the same constraint would apply: the only difference would be that the agents would not earn incentive rents and the objective function of the Principal would be different (w_i would be replaced by $\tilde{w}_i = (\lambda_1 + \lambda_2)\theta_i$).

Proposition 2 thus shows how costs imposed by incentive considerations are handled differently from those imposed by communicational constraints. The former is represented by the replacement of production costs of the agents by their incentive-rent-inclusive virtual costs in the objective function of the Principal, in exactly the same way as in a world with costless, unlimited communication. The costs imposed by communicational constraints are represented by the restriction of the feasible set of output allocations, which must now vary more coarsely with the type realizations of the agents. This can be viewed as the natural extension of Marschak-Radner (1972) characterization of optimal team decision problems to a setting with incentive problems. In particular, the same computational techniques can be used to solve these problems both with and without incentive problems: only the form of the objective function needs to be modified to replace actual production costs by virtual costs. The 'desired' communicational strategies can be rendered incentive compatible at zero additional cost.

Van Zandt (2007) and Fadel and Segal (2009) discuss the question of 'communication cost of selfishness', posed as follows. They take an arbitrary social choice function (production allocation in our notation) and ask whether the communicational complexity of implementing it is increased by the presence of incentive constraints. Van Zandt shows this is not the case when using the solution concept of ex post incentive compatibility, while Fadel and Segal provide examples where this is the case when using the solution concept of Bayesian incentive compatibility.

In our context we fix communicational complexity and select an allocation to maximize the Principal's expected payoff. The analogous question in our context may be whether the optimal mechanism in the presence of communication constraints alone is also optimal in the presence of incentive constraints as well. The answer to this question depends on λ_1. If the Principal is solely concerned with efficiency and $\lambda_1 = 0$, the objective function is the same with and without incentive constraints.[20] Then the optimal mechanism in the absence of any incentive constraints is also optimal in the presence of incentive constraints. On the other hand, if $\lambda_1 > 0$ and the Principal seeks to limit transfers to the agents, the objective function with and without incentive constraints differ. Then the optimal allocation in the absence of incentive constraints will typically not be optimal when incentive problems are present.

6 Implications for Decentralization versus Centralization of Production Decisions

We now examine implications of Proposition 2 for the value of truly decentralized mechanisms compared with centralized ones. Consider a setting where there is no technological restriction on the relation between outputs of the two agents, i.e., $Q = \Re_+ \times \Re_+$, so their outputs can be chosen independently. If production decisions are made by the Principal, outputs are measurable with respect to the history of exchanged messages. If production decisions are delegated to the agents, this is no longer true, since they can be decided by the agents on the basis of information about their own true types, which is richer than what they managed to communicate to the Principal. Unlike settings of unlimited communication, centralized mechanisms can no longer replicate the outcomes of decentralized ones. Contracts are endogenously incomplete, thus permitting a nontrivial comparison of centralized and decentralization decision rights.

[20]Van Zandt and Fadel and Segal do not incorporate the costs of incentivizing transfers in posing the implementation problem, so this is the appropriate case to consider when comparing with their result.

The typical tradeoff associated with delegation of decision rights to better informed agents compares the benefit of increased flexibility of decisions with respect to the true state of the world, with the cost of possible use of discretion by the agent to increase his own rents at the expense of the Principal. Proposition 2 however shows that once the incentive rents that agents will inevitably earn have been factored in the Principal's objective, incentive considerations can be ignored. The added flexibility that delegation allows then ensures that delegation is the superior arrangement. We show below that delegation strictly outperforms centralization provided the gross benefit function of the Principal is smooth, strictly concave and satisfies Inada conditions.

Proposition 3 *Suppose that (i) outputs of the two agents can be chosen independently ($Q = \Re_+ \times \Re_+$); and (ii) $V(q_1, q_2)$ is twice continuously differentiable, strictly concave and satisfies the Inada condition $\frac{\partial V}{\partial q_i} \to \infty$ as $q_i \to 0$. Then given any feasible centralized mechanism there exists a corresponding truly decentralized mechanism which generates a strictly higher payoff to the Principal.*

The outline of the argument is as follows. The finiteness of the set of feasible communication plans for every agent implies the existence of non-negligible type intervals over which communication strategies and message histories are pooled. Consequently if decisions are centralized, the production decision for i must be pooled in the same way. Instead if production decisions are left to agent i, the production decision can be based on the agent's knowledge of its own true type. Under the assumptions of Proposition 3 which ensure that the optimal outputs are always interior, this added 'flexibility' will allow an increase in the Principal's objective (12) while preserving communication feasibility. The result then follows upon using Proposition 2.

7 Implications for Choice of Communication Protocol

Proposition 2 also has useful implications for the ranking of different communication protocols. Given any set of communication strategies in a given protocol, in state (θ_i, θ_j) agent i learns that θ_j lies in the set $\Theta_j(h(c_i(\theta_i), c_j(\theta_j)))$, which generates an information partition for agent i over agent j's type.

Say that a protocol $p_1 \in \mathcal{P}$ is *more informative* than another $p_2 \in \mathcal{P}$ if for any set of communication strategies in the former, there exists a set of communication strategies in the latter which yields (at round T) an information partition to each agent over the type of the other agent which is more informative in the Blackwell sense in (almost) all states of the world.

It then follows that a more informative communication protocol permits a wider choice of communication feasible output allocations. Hence Proposition 2 implies that the

Principal prefers more informative protocols, and would not benefit by restricting or scrambling the flow of communication among agents.

This is the reason we assumed that all messages are addressed to everyone else in the organization. If the transmission and processing of messages entail no resource or time costs, this ensures maximal flow of information between agents. In contrast much of the literature on informational efficiency of resource allocation mechanisms (in the tradition of Hurwicz (1960, 1972) or Mount and Reiter (1974)) has focused on centralized communication protocols where agents send messages to the Principal rather than one another. Such protocols restrict the flow of information among agents. Marschak and Reichelstein (1998) have extended this to network mechanisms where agents communicate directly with one another, and examine the consequences of such decentralized 'network' mechanisms for communication costs (in the absence of incentive problems). In our approach the Principal plays no active role in the communication process. If the only costs of communication involve writing or sending messages this is without loss of generality, since the Principal has no private information to report to the agents, and any messages that an agent sends to the Principal which are in turn sent to the other agent could be sent directly to the latter at no additional cost.

Within the class of decentralized communication protocols, more can be said about the nature of optimal protocols, depending on the precise nature of communication costs. We turn to this now.

In the remainder of this section we assume that communication costs are linear in length of messages and in communication capacity. Hence material costs Φ_m take the form

$$\Phi_m = \phi_v l(m_i) + \phi_f \bar{l}(M_i) \tag{14}$$

for some constants $\phi_v \geq 0, \phi_f > 0$, while time delay takes the form

$$\Phi_d = \phi_d \max\{\bar{l}(M_1), \bar{l}(M_2)\} \tag{15}$$

for some $\phi_d > 0$. More generally, the same results obtain as long as there are no increasing returns to scale with respect to length of messages or communication capacity. Note that the same parameters ϕ_v, ϕ_f, ϕ_d apply to both agents; in this sense they have similar communication abilities.

We also limit attention to agent vocabularies consisting of *letters* or messages of unit length, in which longer messages are *words* which are combinations of letters. Hence if there are L_i letters of unit length in agent i's vocabulary, then there are at most L_i^k words or messages of length not exceeding k, for any integer k. For instance, if the agents communicate using binary code, there are two letters or unit bits 0 and 1, and any longer message consists of a string of unit bits, with the length of the message identified by the number of bits. The same is true for most languages which have an alphabet of letters, words are composed of a string of letters and the length of a word is measured by the number of letters contained in that word. In what follows, we use M_i^* to denote the set of letters in i's vocabulary in conjunction with the null message, i.e., $M_i^* \equiv \{m_i \in \mathcal{M}_i | l(m_i) \leq 1\}$.

Our first result shows that under the above assumptions, information ought to be released 'slowly' by agents across multiple rounds of communication. If any agent has a 'large' message set in any given round, the agent can communicate more information at the same cost by breaking this up a sequence of smaller messages in successive rounds. Suppose for instance that communication is in binary code, and an agent has the following message set in some round: $\{\phi, 0, 1, 00, 01, 10, 11\}$. This round can be broken up into two successive rounds in each of which the agent is given the message set $\{\phi, 0, 1\}$. The agent can communicate at least as much information across these two rounds as she could previously (e.g., a null message in both rounds corresponds to a null message previously, a null message in one round combined with a single-bit message 0 (or 1) in the other corresponds to a previous message of 0 (or 1), and so on). Communication costs do not increase since capacity costs are the same: the maximal length of a message was 2 previously with a single round, while it is now 1 in each of the two rounds. The aggregate length of messages remains the same in every state of the world. The agent now has a total of nine possible message combinations across the two rounds, as against seven possible messages previously. Hence the agent can now send strictly more information, e.g., she has the choice of the order in which a null message is sent in one round and a single-bit message in the other. This allows a strict improvement in the Principal's payoff, provided the gross benefit function of the Principal is smooth, strictly concave and satisfies Inada conditions.

Proposition 4 *Suppose that (i) agent vocabularies satisfy the property mentioned above, (ii) assumptions (14, 15) hold, (iii) $Q = \Re_+ \times \Re_+$, (iv) $V(q_1, q_2)$ is twice continuously differentiable, strictly concave and satisfies the Inada condition $\frac{\partial V}{\partial q_i} \rightarrow \infty$ as $q_i \rightarrow 0$, and (v) $V_{12}(q_1, q_2) \neq 0$ for every $(q_1, q_2) >> 0$. Then if communication is constrained either by material costs ($K < \infty$) or total permissible delay ($D < \infty$), any non-null message set assigned to any agent in any round following any history arising with positive probability in any optimal protocol must consist of letters (messages of unit length) alone, i.e., $M_i(h_{t-1}) = M_i^*$ if it is non-null.*

Our final result concerns the contrast between material costs and time delay formulations of communication cost for the nature of optimal protocols.

Proposition 5 *Suppose the same conditions as in Proposition 4. In addition*

(i) suppose that communication is constrained only by total material cost (i.e., $\phi_d = 0$,

$\phi_f > 0$). Then there exists an optimal protocol with the feature that only one agent sends messages in any given communication round.

(ii) Suppose that communication is constrained only by the total time delays (i.e., $\phi_v = \phi_f = 0 < \phi_d$). Then every optimal protocol involves a number of communication rounds equal to the largest integer not exceeding D/ϕ_d, and both agents send messages simultaneously in each round.

The reasoning is the following. If communication entails only material costs, any round with simultaneous communication by both agents (from the set of messages of unit length or less) can be broken down into two successive rounds in which the agents alternate in sending messages from this set. Each agent has the option of sending the same message in this round when it is their turn to report. The agent now moving second has the additional option of conditioning his message on the message just sent by the other agent moving first (while restricted to sending a message of the same or shorter length as he did previously). The rest of the protocol is left unchanged. Material costs of communication are unchanged, as the communication capacity of each remains the same and the length of messages sent do not increase. Hence the Principal's payoff weakly increases. Here the total delay of the mechanism is increased, owing to the sequencing of messages across the two agents, but this is not costly by assumption.

In contrast when communication costs consist only of delay, both agents must send messages in every round. Otherwise there would be a round in which one of the agents (i, say) does not send any messages, while the other agent j does (if neither does then the entire round can be dispensed with). Allowing i to select a message from M_i^* in this round allows him to communicate more information than previously. As there are no material costs of communication this does not cause any problem with the communication constraint, so a strict improvement is now possible.

8 Concluding Comments

We now remark on some of the simplifying assumptions in our model, which suggest avenues for further exploration.

Our approach has been based on the notion that the costs of information only incorporate the costs of writing messages. While the costs of transmitting messages to multiple receivers may be trivial with contemporary information technology, the costs of processing information by receivers may be substantial in many contexts. We conjecture that our results concerning optimality of delegating decision-making will continue to be true when reading messages is costly. However the nature of optimal communication protocols is likely to be considerably different. Analyzing implications of incorporating costs of reading messages could be an interesting question for future research.

We also ignored the possibility of delegating responsibility of contracting with other agents to some key agents. A broader concern is that we ignored the communicational requirements involved in contracting itself, by focusing only on communication in the process of implementation of the contract, which takes place after parties have negotiated and accepted a contract. Under the assumption that pre-contracting communication is costless, and messages exchanged between agents are verifiable by the Principle, it can be shown that delegation of contracting cannot dominate centralized contracting if both are equally constrained in terms of communicational requirements. Subcontracting may thus be potentially valuable in the presence of costs of pre-contract communication, or if agents can directly communicate with one another in a richer way than the way they can communicate with the Principal.

Appendix: Proofs

Proof of Lemma 1: Suppose that $K < \infty$. Without loss of generality, our attention can be restricted to a communication protocol p such that $\bar{l}(M_1(h_{t-1}))+\bar{l}(M_2(h_{t-1})) \geq 1$ for any $h_{t-1} \in H_{t-1}(p)$ $(t = 1, .., T)$ on the communication path, since we can delete any communication round with $M_1(h_{t-1}) = M_2(h_{t-1}) = \{\phi\}$ with no effect on communication strategies and without violating communication constraints. Then for any $(h_0, .., h_T) \in \bar{H}$,

$$T \leq \Sigma_{i=1,2}\Sigma_{t=1}^{T}\bar{l}(M_i(h_{t-1})) \leq (1/\phi_f)\Sigma_{t=1}^{T}\Sigma_{i\in\{1,2\}}\Phi_m(m_i(h_{t-1}, h_t), M_i(h_{t-1})) \leq K/\phi_f.$$

Therefore the number of rounds T has an upper bound K/ϕ_f. For any i and any h_{t-1}, $\bar{l}(M_i(h_{t-1})) \leq K/\phi_f$ must hold, since otherwise the budget constraint for the material cost is violated. By Assumption 1, $\#\{m \in \mathcal{M}_i \mid l(m) \leq K/\phi_f\}$ also has an upper bound, which is denoted by A. It implies that $\#M_i(h_{t-1}) \leq A$ for any i and any h_{t-1}. Then $\#H_t(p) \leq A^t$ for any $t \in \{0, .., T-1\}$. Since $\#C_i(p) = \Pi_{t=0}^{T-1}\Pi_{h_t\in H_t(p)}[\#M_i(h_t)]$ (with $\#M_i(h_t) \equiv 1$ if $M_i(h_t)$ is the null set),

$$\#C_i(p) \leq \Pi_{t=0}^{T-1}A^{A^t} \leq A^{1+A+....+A^{K/\phi_f-1}} = A^{\frac{1-A^{K/\phi_f}}{1-A}}.$$

Therefore the number of communication plans has an upper bound for any p.

Next consider the case that the communication is constrained by the total delays ($D < \infty$). Without loss of generality, our attention can be restricted to a communication protocol such that $\max\{\bar{l}(M_1(h_{t-1})), \bar{l}(M_2(h_{t-1}))\} \geq 1$ for any $h_{t-1} \in H_{t-1}(p)$ $(t = 1, .., T)$ on the communication path, by the same reason as above. Then for any $(h_0, .., h_T) \in \bar{H}$,

$$T \leq \Sigma_{t=1}^{T}\max\{\bar{l}(M_1(h_{t-1})), \bar{l}(M_2(h_{t-1}))\} \leq D/\phi_d.$$

Therefore the number of rounds T has an upper bound D/ϕ_d. For any i and any h_{t-1}, $\bar{l}(M_i(h_{t-1})) \leq D/\phi_d$ must hold, since otherwise the total delay constraint is violated. By Assumption 1, $\#\{m \in \mathcal{M}_i \mid l(m) \leq D/\phi_d\}$ has an upper bound, denoted by B. It implies that $\#M_i(h_{t-1}) \leq B$ for any i and any h_{t-1}. Then by the same procedure as the case of $K < \infty$ above, we have

$$\#C_i(p) \leq B^{\frac{1-B^{D/\phi_d}}{1-B}}.$$

Therefore the number of communication plans has an upper bound for any p. ■

Lemma 2 *Consider any communication protocol $p \in \mathcal{P}$. For any $h_t \in H_t(p)$ and any $t \in \{1, \dots, T\}$:*

$$\{c \in C(p) \mid h_t(c) = h_t\}$$

is a rectangle set in the sense that if $h_t(c_i, c_j) = h_t(c_i^{'}, c_j^{'}) = h_t$ for $(c_i, c_j) \neq (c_i^{'}, c_j^{'})$, then

$$h_t(c_i^{'}, c_j) = h_t(c_i, c_j^{'}) = h_t$$

Proof of Lemma 2: The proof is by induction. Note that $h_0(c) = \phi$ for any c, so it is true at $t = 0$. Suppose the result is true for all dates up to $t - 1$, we shall show it is true at t.

Note that

$$h_t(c_i, c_j) = h_t(c_i^{'}, c_j^{'}) = h_t \tag{16}$$

implies

$$h_\tau(c_i, c_j) = h_\tau(c_i^{'}, c_j^{'}) = h_\tau \tag{17}$$

for any $\tau \in \{0, 1, .., t-1\}$. Since the result is true until $t - 1$, we also have

$$h_\tau(c_i^{'}, c_j) = h_\tau(c_i, c_j^{'}) = h_\tau \tag{18}$$

for all $\tau \leq t - 1$. So under any of the configurations of communication plans (c_i, c_j), $(c_i^{'}, c_j^{'})$, $(c_i^{'}, c_j)$ or $(c_i, c_j^{'})$, agent i experiences the same message history h_{t-1} until $t - 1$. Then i has the same message space at t, and (16) implies that i sends the same messages to j at t, under either c_i or $c_i^{'}$.

(17) and (18) also imply that under either c_j or $c_j^{'}$, j sends the same messages to i at all dates until $t - 1$, following receipt on the (common) messages sent by i until $t - 1$ under these different configurations. The result now follows from the fact that messages sent by j to i depend on the communication plan of i only via the messages j receives from i. So i must also receive the same messages at t under any of these different configurations of communication plans. ■

Proof of Proposition 1:

Let $q_i(\theta_i, \theta_j)$ be a production allocation satisfying (8), which is supported by a communication strategy vector $c(\theta)$ in a protocol p which is parsimonious with respect to these strategies. In this protocol all histories are reached with positive probability on the equilibrium path, hence beliefs of every agent with regard to the types of the other agent are obtained by applying Bayes rule.

Define $\hat{q}_i(\theta_i, h_t)$ by

$$\hat{q}_i(\theta_i, h_t) \equiv E[q_i(\theta_i, \theta_j) \mid \theta_j \in \Theta_j(h_t)].$$

for any $h_t \in H_t$ and any $t \in \{0, 1, .., T\}$. Condition (8) requires $\hat{q}_i(\theta_i, h_t)$ to be non-increasing in θ_i on $\Theta_i(h_t)$. Note that

$$\hat{q}_i(\theta_i, h(c(\theta_i, \theta_j))) = E_{\tilde{\theta}_j}[q_i(\theta_i, \tilde{\theta}_j) \mid \tilde{\theta}_j \in \Theta_j(h(c(\theta_i, \theta_j)))] = q_i(\theta_i, \theta_j),$$

since $q_i(\theta_i, \tilde{\theta}_j) = q_i(\theta_i, \theta_j)$ for any $\tilde{\theta}_j \in \Theta_j(h(c(\theta_i, \theta_j)))$.

Step 1: The relationship between $\hat{q}_i(\theta_i, h_t)$ and $\hat{q}_i(\theta_i, h_{t+1})$

Suppose that i observes h_t at the end of round t. Given selection of $m_{i,t+1} \in M_i(h_t)$ where $M_i(h_t)$ is the message set for h_t in protocol p, agent i's history at round $t+1$ is subsequently determined by messages received by i in round t. Let the set of possible histories h_{t+1} at the end of round $t+1$ be denoted by $H_{t+1}(h_t, m_{i,t+1})$. Evidently for $j \neq i$, $\{\Theta_j(h_{t+1}) | h_{t+1} \in H_{t+1}(h_t, m_{i,t+1})\}$ constitutes a partition of $\Theta_j(h_t)$:

$$\cup_{h_{t+1} \in H_{t+1}(h_t, m_{i,t+1})} \Theta_j(h_{t+1}) = \Theta_j(h_t)$$

and

$$\Theta_j(h_{t+1}) \cap \Theta_j(h'_{t+1}) \neq \phi$$

for $h_{t+1}, h'_{t+1} \in H_{t+1}(h_t, m_{i,t+1})$ such that $h_{t+1} \neq h'_{t+1}$. The probability of $h_{t+1} \in H_{t+1}(h_t, m_{i,t+1})$ conditional on $(h_t, m_{i,t+1})$ is represented by

$$\Pr(h_{t+1} \mid h_t, m_{i,t+1}) = \Pr(\Theta_j(h_{t+1})) / \Pr(\Theta_j(h_t)).$$

From the definition of $\hat{q}_i(\theta_i, h_t)$ and $\hat{q}_i(\theta_i, h_{t+1})$, for any $m_{i,t+1} \in M_i(h_t)$ and any $\theta_i \in \Theta_i$,

$$\Sigma_{h_{t+1} \in H_{t+1}(h_t, m_{i,t+1})} \Pr(h_{t+1} \mid h_t, m_{i,t+1}) \hat{q}_i(\theta_i, h_{t+1}) = \hat{q}_i(\theta_i, h_t).$$

Step 2: For any $h_{t+1}, h'_{t+1} \in H_{t+1}(h_t, m_{it+1})$, $\Theta_i(h_{t+1}) = \Theta_i(h'_{t+1}) \subset \Theta_i(h_t)$

By definition

$$\Theta_i(h_{t+1}) = \{\theta_i \mid m_{i,t+1}(\theta_i, h_t) = m_{i,t+1}\} \cap \Theta_{it}(h_t)$$

where $m_{i,t+1}(\theta_i, h_t)$ denotes i's message choice corresponding to the strategy $c_i(\theta_i)$. The right hand side depends only on $m_{i,t+1}$ and h_t. It implies that the set $\Theta_i(h_{t+1})$ does not vary across different $h_{t+1} \in H_{t+1}(h_t, m_{i,t+1})$. To simplify exposition, we denote this set henceforth by $\Theta_i(h_t, m_{i,t+1})$.

Step 3: Construction of $\tilde{q}_i(\theta_i, h_t)$

We construct $\tilde{q}_i(\theta_i, h_t)$ for any $h_t \in H_t$ based on the following Claim 1.

Claim 1:

For arbitrary $q_i(\theta_i, \theta_j)$ satisfying (8), there exists $\tilde{q}_i(\theta_i, h_t)$ for any $h_t \in H_t$ and any $t \in \{0, .., T\}$ so that

(a) $\tilde{q}_i(\theta_i, h_t) = \hat{q}_i(\theta_i, h_t)$ for $\theta_i \in \Theta_i(h_t)$

(b) $\tilde{q}_i(\theta_i, h_t)$ is non-increasing in θ_i on Θ_i

(c) $\Sigma_{h_{t+1} \in H_{t+1}(h_t, m_{i,t+1})} \Pr(h_{t+1} \mid h_t, m_{i,t+1}) \tilde{q}_i(\theta_i, h_{t+1}) = \tilde{q}_i(\theta_i, h_t)$ for any $\theta_i \in \Theta_i$ and any $m_{i,t+1} \in M_i(h_t)$ where $M_i(h_t)$ is the message set for h_t in protocol p.

Claim 1 states that there exists an 'auxiliary' output rule $\tilde{q}_i$ as a function of type θ_i and message history which is globally non-increasing in type (property (b)) following any history h_t, and $\tilde{q}_i(\theta_i, h_t)$ equals the expected value of $\tilde{q}_i(\theta_i, h_{t+1})$ conditional on (h_t, m_{it+1}) for any $m_{it+1} \in M_i(h_t)$ (property (c)).

In order to establish Claim 1, the following Lemma is needed.

Lemma 3 *For any $B \subset R_+$ which may not be connected, let A be an interval satisfying $B \subset A$. Suppose that $F_i(a)$ for $i = 1, ..., N$ and $G(a)$ are real-valued functions defined on A, each of which has the following properties:*

- *$F_i(a)$ is non-increasing in a on B for any i.*
- *$\Sigma_i p_i F_i(a) = G(a)$ for any $a \in B$ and for some p_i so that $p_i > 0$ and $\Sigma_i p_i = 1$.*
- *$G(a)$ is non-increasing in a on A.*

Then we can construct real-valued function $\bar{F}_i(a)$ defined on A for any i so that

- *$\bar{F}_i(a) = F_i(a)$ on $a \in B$ for any i.*
- *$\Sigma_i p_i \bar{F}_i(a) = G(a)$ for any $a \in A$ and for the same p_i*
- *$\bar{F}_i(a)$ is non-increasing in a on A for any i.*

This lemma says that we can construct functions $\bar{F}_i(a)$ so that the properties of functions $F_i(a)$ on B are also maintained on the interval A which covers B.

Proof of Lemma 3:

If this statement is true for $N = 2$, we can easily show that this also holds for any $N \geq 2$. Suppose that this is true for $N = 2$.

$$\Sigma_{i=1}^{N} p_i F_i(a) = p_1 F_1(a) + (p_2 + ... + p_N) F^{-1}(a)$$

with

$$F^{-1}(a) = \Sigma_{i \neq 1} \frac{p_i}{p_2 + ... + p_N} F_i(a).$$

Applying this statement for $N = 2$, we can construct $\bar{F}_1(a)$ and $\bar{F}^{-1}(a)$ which keeps the same property on A as on B. Next using the constructed $\bar{F}^{-1}(a)$ instead of $G(a)$, we can apply the statement for $N = 2$ again to construct desirable $\bar{F}_2(a)$ and $\bar{F}^{-2}(a)$ on A based on $F_2(a)$ and $F^{-2}(a)$ which satisfy

$$\frac{p_2}{p_2 + ... + p_N} F_2(a) + [1 - \frac{p_2}{p_2 + ... + p_N}] F^{-2}(a) = F^{-1}(a).$$

on B. We can use this method recursively to construct $\bar{F}_i(a)$ for all i.

Next let us show that the statement is true for $N = 2$. For $a \in A \backslash B$, define $\underline{a}(a)$ and $\bar{a}(a)$, if they exist, so that

$$\underline{a}(a) \equiv \sup\{a^{'} \in B \mid a^{'} < a\}$$

and

$$\bar{a}(a) \equiv \inf\{a^{'} \in B \mid a^{'} > a\}.$$

It is obvious that at least one of either $\underline{a}(a)$ or $\bar{a}(a)$ exists for any $a \in A \backslash B$.

Let's specify $\bar{F}_1(a)$ and $\bar{F}_2(a)$ so that $\bar{F}_1(a) = F_1(a)$ and $\bar{F}_2(a) = F_2(a)$ for $a \in B$, and for $a \in A \backslash B$ as follows.

(i) For $a \in A \backslash B$ so that only $\underline{a}(a)$ exists,

$$\bar{F}_1(a) = F_1(\underline{a}(a))$$

$$\bar{F}_2(a) = \frac{G(a) - p_1 F_1(\underline{a}(a))}{p_2}$$

(ii)For $a \in A \backslash B$ so that both $\underline{a}(a)$ and $\bar{a}(a)$ exist,

$$\bar{F}_1(a) = \min\{F_1(\underline{a}(a)), \frac{G(a) - p_2 F_2(\bar{a}(a))}{p_1}\}$$

$$\bar{F}_2(a) = \max\{F_2(\bar{a}(a)), \frac{G(a) - p_1 F_1(\underline{a}(a))}{p_2}\}$$

(iii)For $a \in A \backslash B$ so that only $\bar{a}(a)$ exists,

$$\bar{F}_1(a) = \frac{G(a) - p_2 F_2(\bar{a}(a))}{p_1}$$

$$\bar{F}_2(a) = F_2(\bar{a}(a))$$

It is easy to check that $\bar{F}_i(a)$ is non-increasing in a on A for $i = 1, 2$ and

$$p_1 \bar{F}_1(a) + p_2 \bar{F}_2(a) = G(a)$$

for $a \in A$. This completes the proof of the lemma. ■

Proof of Claim 1:

Choose arbitrary $t \in \{0, ..., T\}$ and $h_t \in H_t$. Suppose that $\tilde{q}_i(\theta_i, h_t)$ satisfies (a) and (b) in Claim 1. Then for any $m_{i,t+1} \in M_i(h_t)$, we can construct a function $\tilde{q}_i(\theta_i, h_{t+1})$ for any $h_{t+1} \in H_t(h_t, m_{it+1})$ so that (a), (b) and (c) are satisfied. This result is obtained upon applying Lemma 3 with

$$B = \Theta_i(h_t, m_{i,t+1})$$

$$A = \Theta_i$$

$$a = \theta_i$$

$$G(\theta_i) = \hat{q}_i(\theta_i, h_t)$$

$$F_{h_{t+1}}(\theta_i) = \hat{q}_i(\theta_i, h_{t+1})$$

$$p_{h_{t+1}} = \frac{\Pr(\Theta_j(h_{t+1}))}{\Pr(\Theta_j(h_t))}$$

for any $h_{t+1} \in H_{t+1}(h_t, m_{i,t+1})$ where each element of the set $H_{t+1}(h_t, m_{i,t+1})$ corresponds to an element of the set $\{1, ..., N\}$ in Lemma 3. This means that for $\tilde{q}_i(\theta_i, h_t)$ which satisfies (a) and (b) for any $h_t \in H_t$, we can construct $\tilde{q}_i(\theta_i, h_{t+1})$ which satisfies (a)-(c) for any $h_{t+1} \in H_{t+1}$.

With $h_0 = \phi$, since $\tilde{q}_i(\theta_i, h_0) = \hat{q}_i(\theta_i, h_0)$ satisfies (a) and (b), $\tilde{q}_i(\theta_i, h_1)$ is constructed so that (a)-(c) are satisfied for any $h_1 \in H_1$. Recursively $\tilde{q}_i(\theta_i, h_t)$ can be constructed for any $h_t \in \cup_{\tau=0}^{T} H_\tau$ so that (a)-(c) are satisfied. ■

Step 4

We are now in a position to complete the proof of sufficiency. We focus initially on the case where $Q = \Re_+ \times \Re_+$ and the mechanism is decentralized so agents select their own outputs independently. Later we show how to extend the proof to other contexts.

Given $\tilde{q}_i(\theta_i, h)$ (with $h = h_T$) constructed in Claim 1, construct transfer functions $t_i(q_i, h)$ as follows:

$$t_i(q_i, h) = \hat{\theta}_i(q_i, h)q_i + \int_{\hat{\theta}_i(q_i,h)}^{\bar{\theta}_i} \tilde{q}_i(x, h)dx.$$

for $q_i \in Q_i(h) \equiv \{\tilde{q}_i(\theta_i, h) \mid \theta_i \in \Theta_i\}$, and $t_i(q_i, h) = -\infty$ for $q_i \notin Q_i(h)$ where $\hat{\theta}_i(q_i, h)$ is defined as follows:

$$\hat{\theta}_i(q_i, h) \equiv \sup\{\theta_i \in \Theta_i \mid \tilde{q}_i(\theta_i, h) \geq q_i\}.$$

We show that the specified communication strategies $c(\theta)$ and output choices $(\tilde{q}_i(\theta_i, h), \tilde{q}_j(\theta_j, h))$ constitute a PBE (combined with beliefs obtained by applying Bayes rule at every history). By construction, $\tilde{q}_i(\theta_i, h)$ maximizes $t_i(q_i, h) - \theta_i q_i$ for any $h \in \mathcal{H} \equiv H_T$ and any $\theta_i \in \Theta_i$, where

$$t_i(\tilde{q}_i(\theta_i, h), h) - \theta_i \tilde{q}_i(\theta_i, h) = \int_{\theta_i}^{\bar{\theta}_i} \tilde{q}_i(x, h)dx.$$

Now turn to the choice of messages. Start with round T. Choose arbitrary $h_{T-1} \in H_{T-1}$ and arbitrary $m_{iT} \in M_i(h_{T-1})$. The expected payoff conditional on $\theta_j \in \Theta_j(h_{T-1})$ (i.e., conditional on beliefs given by $\Pr(h \mid h_{T-1}, m_{iT}) = \frac{\Pr(\Theta_j(h))}{\Pr(\Theta_j(h_{T-1}))}$ for $h \in H_T(h_{T-1}, m_{iT})$) is

$$\begin{aligned} & E_h[t_i(\tilde{q}_i(\theta_i, h), h) - \theta_i \tilde{q}_i(\theta_i, h) \mid h_{T-1}, m_{iT}] \\ = & \int_{\theta_i}^{\bar{\theta}_i} E_h[\tilde{q}_i(x, h) \mid h_{T-1}, m_{iT}]dx \\ = & \int_{\theta_i}^{\bar{\theta}_i} \tilde{q}_i(x, h_{T-1})dx. \end{aligned}$$

This does not depend on the choice of $m_{iT} \in M_i(h_{T-1})$. Therefore agent i does not have an incentive to deviate from $m_{iT} = m_{iT}(\theta_i, h_{T-1})$.

The same argument can recursively be applied for all previous rounds t, implying that $m_{i,t+1} = m_{i,t+1}(\theta_i, h_t)$ is an optimal message choice for any $h_t \in H_t$ and any t. It is also evident that at round 0, it is optimal for agent i to accept the contract. This establishes that participation, followed by the communication strategies $c(\theta)$ combined with output choices $(\tilde{q}_i(\theta_i, h), \tilde{q}_j(\theta_j, h))$ constitute a PBE.

The argument extends to the case of the decentralized mechanism where the feasible outputs are constrained to lie in some set Q which is a subset of $\Re_+ \times \Re_+$. From the construction of $\tilde{q}_i(\theta_i, h)$ (with $h = h_T$) in Claim 1, $(\tilde{q}_i(\theta_i, h), \tilde{q}_j(\theta_j, h)) \in Q$ holds for any $(\theta_i, \theta_j) \in \Theta_i(h) \times \Theta_j(h)$. On the other hand, it may not hold for $(\theta_i, \theta_j) \notin \Theta_i(h) \times \Theta_j(h)$. Define $\tilde{Q}_i(h) \equiv \{\tilde{q}_i(\theta_i, h) \mid \theta_i \in \Theta_i(h)\}$, the set of possible outputs for i on the equilibrium path, following h. Then $\tilde{Q}_i(h) \times \tilde{Q}_j(h) \subset Q$, implying that as long as agent i chooses an output in $\tilde{Q}_i(h)$ following message history h, the allocation constraint is not violated. Now construct a new set of transfers $\hat{t}_i(q_i, h)$ as follows: $\hat{t}_i(q_i, h) = t_i(q_i, h)$ for $q_i \in \tilde{Q}_i(h)$, and $\hat{t}_i(q_i, h) = -\infty$ for $q_i \notin \tilde{Q}_i(h)$ where $t_i(q_i, h)$ is the transfer function constructed in the previous argument. In this new mechanism, the agent's expected payoff is preserved on the equilibrium path, while they are not increased off the equilibrium path. Hence the postulated strategies constitute a PBE.

A similar extension works for the case of a centralized mechanism, since this is a special case of the previous mechanism where the assigned outputs $\hat{q}_i(\theta_i, h) = \hat{q}_i(h)$ are measurable with respect to h, i.e., are independent of θ_i conditional on h. Then $\tilde{Q}_i(h) \equiv \{\tilde{q}_i(\theta_i, h) \mid \theta_i \in \Theta_i(h)\} = \hat{q}_i(h)$. Agent i can effectively be forced to choose output $\hat{q}_i(h)$ following history h at the end of the communication phase with a transfer $\hat{t}_i(q_i, h)$. The same argument as above ensures that the assigned communication strategies constitute a PBE. ■

Proof of Proposition 2:

We show that the solution of the relaxed problem where (ii) is dropped satisfies (ii). Suppose not. Let the solution of the relaxed problem be represented by a (parsimonious) communication protocol p, communication strategies $c(\theta)$ and output allocation $(q_1(\theta_1, \theta_2), q_2(\theta_1, \theta_2))$. H_t, $\Theta_i(h_t)$ and $\Theta_j(h_t)$ are well defined for $(p, c(\theta))$. Then there exists $t \in \{0, \ldots, T\}$, $h_t \in H_t$ and $\theta_i, \theta_i' \in \Theta_i(h_t)$ with $\theta_i > \theta_i'$ so that

$$E_{\theta_j}[q_i(\theta_i, \theta_j) \mid \theta_j \in \Theta_j(h_t)] > E_{\theta_j}[q_i(\theta_i', \theta_j) \mid \theta_j \in \Theta_j(h_t)].$$

This implies that at least either one of

$$\begin{aligned} & E[V(q_i(\theta_i', \theta_j), q_j(\theta_i', \theta_j)) - w_i(\theta_i) q_i(\theta_i', \theta_j) - w_j(\theta_j) q_j(\theta_i', \theta_j) \mid \theta_j \in \Theta_j(h_t)] \\ > \; & E[V(q_i(\theta_i, \theta_j), q_j(\theta_i, \theta_j)) - w_i(\theta_i) q_i(\theta_i, \theta_j) - w_j(\theta_j) q_j(\theta_i, \theta_j) \mid \theta_j \in \Theta_j(h_t)] \end{aligned}$$

or

$$\begin{aligned} & E[V(q_i(\theta_i, \theta_j), q_j(\theta_i, \theta_j)) - w_i(\theta_i') q_i(\theta_i, \theta_j) - w_j(\theta_j) q_j(\theta_i, \theta_j) \mid \theta_j \in \Theta_j(h_t)] \\ > \; & E[V(q_i(\theta_i', \theta_j), q_j(\theta_i', \theta_j)) - w_i(\theta_i') q_i(\theta_i', \theta_j) - w_j(\theta_j) q_j(\theta_i', \theta_j) \mid \theta_j \in \Theta_j(h_t)] \end{aligned}$$

holds. This means that if at least one type of either θ_i or θ_i' takes other type of communication plan and output decision rule, P's payoff is improved. This is a contradiction. ■

Proof of Proposition 3:

Consider any communication-feasible centralized mechanism with protocol p and communication strategies $c(\theta)$ that result in an output allocation $q^*(\theta) = q(h(c(\theta)))$. Consider any history h that arises from these communication strategies with positive probability, and let the corresponding set of types be $\Theta_i(h) \times \Theta_j(h)$. Then $q^*(\theta)$ must be constant over $\Theta_i(h) \times \Theta_j(h)$.

For arbitrary q_i, denote

$$E[V(q_i, q_j^*(\theta_i, \theta_j)) \mid \theta_j \in \Theta_j(h)]$$

by $V(q_i, q_j(h))$. Consider the problem of choosing q_i to maximize $V(q_i, q_j(h)) - w_i(\theta_i)q_i$ for any $\theta_i \in \Theta_i(h)$. It is evident that the function $V(q_i, q_j(h))$ is strictly concave in q_i, and satisfies the Inada condition. Given the monotonicity of $w_i(\theta_i)$, the optimal solution to this problem, denoted by $\hat{q}_i(\theta_i, h)$, is strictly decreasing in θ_i on $\Theta_i(h)$. Hence

$$\begin{aligned}
& E[V(\hat{q}_i(\theta_i, h), q_j^*(\theta)) - w_i(\theta_i)\hat{q}_i(\theta_i, h) - w_j(\theta_j)q_j^*(\theta) \mid (\theta_i, \theta_j) \in \Theta_i(h) \times \Theta_j(h)] \\
> \;& E[V(q^*(\theta)) - w_i(\theta_i)q_i^*(\theta) - w_j(\theta_j)q_j^*(\theta) \mid (\theta_i, \theta_j) \in \Theta_i(h) \times \Theta_j(h)]
\end{aligned}$$

Now replace the output allocation $(q_i^*(\theta), q_j^*(\theta))$ by $(\hat{q}_i(\theta_i, h(c(\theta))), q_j^*(\theta))$ over $\Theta_i(h) \times \Theta_j(h)$, while leaving it unchanged everywhere else. This is a decentralized mechanism which is communication-feasible, which attains a strictly higher expected payoff for the Principal compared with the centralized mechanism. ■

Proof of Proposition 4: Suppose there is a round t and history h_{t-1} with $M_i(h_{t-1}) \neq \phi$ and $M_i(h_{t-1}) \neq M_i^*$ for some agent i. Without loss of generality, let $n_i \equiv \bar{l}(M_i(h_{t-1})) \geq n_j \equiv \bar{l}(M_j(h_{t-1}))$, and $n_i \geq 1$ (otherwise both agents have null message sets and the round can be deleted).

Following history h_{t-1}, we replace round t with rounds $t, t+1, \ldots, t+n_i-1$ with message set M_i^* for i in each of these rounds, and message set M_j^* for j in rounds $t, t+1, \ldots, t+n_j-1$. Agent j is assigned a null message set in rounds $t+n_j, \ldots t+n_i-1$ if $n_i > n_j$. Then notice by construction that

$$\bar{l}(M_k(h_{t-1})) = n_k = n_k \bar{l}(M_k^*) \tag{19}$$

for both agents $k = i, j$, implying that aggregate capacity cost or delay will remain unchanged. Moreover for agent i we have

$$\begin{aligned}
\#M_i(h_{t-1}) &\leq \#\{m_i \in \mathcal{M}_i | l(m_i) \leq n_i\} \\
&\leq 1 + L_i + \ldots + (L_i)^{n_i} \\
&< (1 + L_i)^{n_i} = \{\#M_i^*\}^{n_i}
\end{aligned} \tag{20}$$

if $n_i \geq 2$, while

$$\#M_j(h_{t-1}) \leq 1 + L_j + \ldots + (L_j)^{n_j} \leq (1 + L_j)^{n_j} = \{\#M_j^*\}^{n_j}. \quad (21)$$

If $n_i = 1$ then $M_i(h_{t-1})$ is a proper subset of M_i^* and $\#M_i(h_{t-1}) < \#M_i^*$. Hence the set of messages available to each agent is now larger for both, and is strictly larger for agent i. So for either agent $k = i, j$ we can select $\hat{M}_k$ which is a subset of $(M_k^*)^{n_k}$ such that $\#\hat{M}_k = \#M_k(h_{t-1})$ and for agent i it is a proper subset. In other words, there exists $\tilde{m}_i \in (M_i^*)^{n_i} \backslash \hat{M}_i$. For each $k = i, j$ we can select a one-to-one mapping μ_k from $M_k(h_{t-1})$ to $\hat{M}_k$ such that $l(\mu_k(m_k)) = l(m_k)$ for all $m_k \in M_k(h_{t-1})$. Also $l(\tilde{m}_i) \leq n_i = \bar{l}(M_i(h_{t-1}))$, so there exists $\bar{m}_i \in M_i(h_{t-1})$ such that $l(\bar{m}_i) = n_i \geq l(\tilde{m}_i)$.

Given any choice of a subset $\Theta_i^{'}$ of $\Theta_i(h_{t-1}, \bar{m}_i)$, we can construct communication plans for different types of i in rounds $t, \ldots, t + n_i - 1$ as follows:

(a) If $\theta_i \in \Theta_i^{'}$ then type θ_i of i reports $\tilde{m}_i$ instead of $\bar{m}_i$

(b) If $\theta_i \in \Theta_i(h_{t-1}, \bar{m}_i) \backslash \{\Theta_i^{'}\}$, type θ_i reports $\bar{m}_i$, as before

(c) If θ_i does not belong to $\Theta_i(h_{t-1}, \bar{m}_i)$ and θ_i reported $m_i \in M_i(h_{t-1})$ previously, she now selects the vector of reports $\mu_i(m_i) \in \hat{M}_i$ across the new n_i rounds.

We shall describe later in the proof the method for selecting the subset $\Theta_i^{'}$.

The communication strategy for j is adapted to the following. If type θ_j reported $m_j \in M_j(h_{t-1})$ in round t in the previous protocol, she now selects the vector of reports $\mu_j(m_j) \in \hat{M}_j$ in rounds $t, \ldots, t + n_j - 1$.

From round $t+n_i$ onwards, the continuation of the protocol and communication strategies exactly replicates the previous protocol and communication strategies from round $t + 1$ onwards, with the continuation following $\mu_i(m_i), \mu_j(m_j)$ in the new protocol exactly matching the continuation following messages $m_i \in M_i(h_{t-1}), m_j \in M_j(h_{t-1})$ in the old protocol. Moreover, the continuation following $\tilde{m}_i, \mu_j(m_j)$ in the new protocol matches the continuation following messages $\bar{m}_i, m_j$ in the old protocol.

By construction, then, total cost of communication capacity and delay is maintained the same. The variable material cost has not increased (since $l(\tilde{m}_i) \leq l(\bar{m}_i)$ while the length of all other messages has remained the same). On the other hand, the set of available messages has expanded for each agent, and strictly for agent i.

It remains to describe how the set $\Theta_i^{'}$ is chosen. Consider any history h_T till the end of the communication phase which is a continuation of $(h_{t-1}, \bar{m}_i)$ which arises with positive probability in the previous protocol. Following history h_T, agent j's information about θ_i is that it is contained in $\Theta_i(h_T)$ which is a non-degenerate interval of Θ_i, and is a subset of $\Theta_i(h_{t-1}, \bar{m}_i)$. Now for any $\hat{\theta}_j$ in the interior of $\Theta_j(h_T)$, we can find a subset $\Theta_i^{'}$ of $\Theta_i(h_T)$ such that both $\Theta_i^{'}$ and $\Theta_i(h_T) \backslash \{\Theta_i^{'}\}$ are non-degenerate, and

$$E[V_{q_j}(q_i(\theta_i, \hat{\theta}_j), q_j(\theta_i, \hat{\theta}_j)|\theta_i \in \Theta_i^{'}] > E[V_{q_j}(q_i(\theta_i, \hat{\theta}_j), q_j(\theta_i, \hat{\theta}_j)|\theta_i \in \Theta_i(h_T) \backslash \{\Theta_i^{'}\}] \quad (22)$$

since $V_{12} \neq 0$ and q_i is strictly decreasing in θ_i over $\Theta_i(h_T)$. Since this inequality is strict, and since the production decision functions are continuous under the postulated regularity properties on V, it must also hold in a non-degenerate neighborhood of $\hat{\theta}_j$. This implies that optimal production decisions must change with positive probability.

Agent i's information about j's type remains unchanged in the new protocol. And agent j has strictly better information in the new protocol concerning i's type following history h_T. This information is strictly valuable as the agents must change their production decisions with positive probability. Hence the Principal can secure a strict improvement in her expected payoff. ■

Proof of Proposition 5: Given Proposition 4, we can restrict attention to the protocol where any non-null message set assigned to agent i is M_i^* in every round. To show (i), suppose there exists round t and $h_{t-1} \in H_{t-1}$ such that $M_k(h_{t-1}) = M_k^*$ for both agents $k = i, j$. Then consider a new communication protocol $\tilde{p}$ where round t (following history h_{t-1}) is split into two successive rounds with sequential communication: in the first, i has a message set M_i^* while j is assigned a null message set, and in the second j has a message set M_j^* while i is assigned a null message set. Each agent can send the same message as they did in the previous protocol when it is their turn to report. From the next round onwards the rest of the protocol continues as before. This modification does not raise total material cost (although it evidently raises total delay). In this protocol , j can send messages which can depend on m_{it}, something that is not possible in p. Hence it allows a weak improvement in P's payoff.

For (ii), suppose that there exists round t and $h_{t-1} \in H_{t-1}$ such that $M_i(h_{t-1}) = M_i^*$ and $M_j(h_{t-1}) = \{\phi\}$. We can now construct a new communication protocol $\tilde{p}$ with $M_j(h_{t-1}) = M_i^*$ instead of $\{\phi\}$ in round t (with history h_{t-1}). All other components of the communication protocol are preserved. This modification does not raise the total time delay (although it raises the total material cost). Here agent j who was silent in round t following history h_{t-1} in p can now send some messages in this round, thus increasing the amount of information exchanged between the agents. Using the same argument as in the proof of Proposition 4, P's payoff can be strictly improved. ■

References

Aoki, M. (1990), "Toward an Economic Model of the Japanese Firm," *Journal of Economic Literature*, 28(1), 1-27.

Baron, D. P. and R. Myerson (1982), "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50(4), 911-930.

Battigalli, P. and G. Maggi (2002), "Rigidity, Discretion, and the Costs of Writing Contracts," *The American Economic Review*, 92(4),798-817.

Blumrosen, L. and M. Feldman (2006), "Implementation with a Bounded Action Space," Department of Economics, The Hebrew University.

Blumrosen, L., N. Nisan and I. Segal (2007), "Auctions with Severely Bounded Communication," *Journal of Artificial Intelligence Research*, 28, 233-266.

Green, J. and J. Laffont (1986), "Incentive Theory with Data Compression," in W. Heller, R. Starr and D. Starrett (Ed.), *Uncertainty, Information and Communication: Essays in Honor of Kenneth Arrow*, Cambridge: Cambridge University Press.

——— (1987), "Limited Communication and Incentive Compatibility", *Information, Incentives, and Economic Mechanisms: Essays in honor of Leonid Hurwicz*, T. Groves, Radner, R. and Reiter, S. (ed.), Minneapolis: University of Minnesota Press

Fadel R. and I. Segal (2009), "The Communication Cost of Selfishness," *Journal of Economic Theory* 144, 1895-1920.

Hayek, F.A. (1945), "The Use of Knowledge in Society," *American Economic Review*, 35(4), 519-530.

Hurwicz L. (1960), "Optimality and Informational Efficiency in Resource Allocation Processes," in K. Arrow, S. Karlin and P. Suppes (ed.) *Mathematical Methods in the Social Sciences*, Stanford University Press, Stanford, California.

——— (1972), "On the Dimensional Requirements of Informationally Decentralized Pareto-Satisfactory Processes," presented at Decentralization Conference, Northwestern University. Later published in *Studies in Resource Allocation Processes* ed. K. Arrow and L. Hurwicz. Cambridge: Cambridge University Press 1977.

Kos, N. (2011a), "Communication and Efficiency in Auctions," *Games and Economic Behavior*, Forthcoming.

——— (2011b), "Asking Questions," Department of Economics, Bocconi University.

Laffont, J. and D. Martimort (1998), "Collusion and Delegation," *Rand Journal of Economics* 29(2), 280-305.

Marschak, J. and R. Radner (1972), *Economic Theory of Teams*, New Haven: Yale University Press.

Marschak T. and S. Reichelstein (1998), "Network Mechanisms, Information Efficiency and Hierarchies," *Journal of Economic Theory*, 79(1), 106-141.

Melumad, N., D. Mookherjee and S. Reichelstein (1992), "A Theory of Responsibility Centers," *Journal of Accounting and Economics*, 15, 445-484.

————— (1997), "Contract Complexity, Incentives and the Value of Delegation," *Journal of Economics and Management Strategy*, 6(2), 257-289.

Mount, K. and S. Reiter (1974), "The Informational Size of Message Spaces," *Journal of Economic Theory*, 8, 161–191.

Myerson, R. (1981), "Optimal Auction Design", *Mathematics of Operations Research* 6, 58–73.

———- (1982), "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 10(1), 67-81.

Reichelstein, S. and S. Reiter (1988), "Game Forms with Minimal Message Spaces," *Econometrica*, 56(3), 661-692.

Segal I. (2006), "Communication in Economic Mechanisms,"in *Advances in Economics and Econometrics: Theory and Application, Ninth World Congress (Econometric Society Monographs)*, ed. by Richard Blundell, Whitney K. Newey, and Torsten Persson, Cambridge University Press.

Van Zandt, T. (2007), "Communication Complexity and Mechanism Design," *Journal of the European Economic Association*, 5,543-553.