Interference: Contracts and Authority with Insecure Communication*

Patrick Legros† Ecares, ULB, Liège, and CEPR

Andrew F. Newman Columbia University

December 7, 1999

Abstract

We develop a theory of mechanism design when agents are able to interfere with each others' communication channels. We develop a kind of revelation principle, the "Noninterference Principle" which permits representation of arbitrary mechanisms by direct ones. The incentives to interfere will depend on the mechanism chosen; interference thus constrains contractual design. For instance, authority emerges as a governance mechanism which may economize on the costs of securing channels, particularly when the organization needs to be flexible and there is diversity in its members' preferences. We also show that there are environments in which the possibility of interference actually facilitates full implementation by providing a means of "protest" in undesired equilibria.

1 Introduction

In a recently publicized case [10], an employee of Morgan Stanley who had been dismissed allegedly for expense account abuses claimed that on the contrary the cause of his dismissal was racism and homophobia. The firm tried to produce testimony from a witness who was claiming that the employee was in touch with a computer hacker to plant racist and homophobic e-mail in the computers of the company. As a result, the employee was arrested on

*Very preliminary and incomplete. Prepared for the ESSET, Gersenzee, Switzerland, June 1999.

†This author benefited from the financial support of the Communauté française de Belgique (projet ARC 98/03-221).

charges of forgery, coercion and "computer trespass". The lawyers for the employee admitted that he payed the hacker \$200 to plant the phony e-mail. However, it later appeared that the witness had received a \$10,000 payment from the firm and the charges against the employee were dropped. What is distinctive about this case is not that agents can process limited sets of signals or that signals arrive at their targets with noise. Rather, the distortions in the signals arise because competing agents *interfere* with each other's attempts to transmit information. It is the effect of this aspect of imperfect communication on contracting that we wish to explore in this paper.

There seems to be little practical reason to believe that the messages transmitted during play of an arbitrary mechanism are immune to interference, or that they can be made secure at zero cost. Typically mechanisms rely on high powered incentives to play a certain way; a player who deviates stands to punished heavily and the others rewarded. But this creates an incentive to make other players "look bad." If a player can effectively make it appear that another player deviated, he will. The assumption that this cannot happen, implicit in mechanism design, is extreme. A more natural starting assumption would be that messages reach their target only imperfectly, depending on actions that other players take and perhaps on costly investments in relevant technologies. Mechanisms must be designed bearing in mind the endogeneity of interference.

One way in which this endogenous form of interference differs from other forms of limited communication is that it does not seem subject to a technological fix. There have been vast improvements in the transmission and processing of information over the last century. But this may do little to diminish the benefits from forging a signature, for instance. On the contrary, a computer which may be used to encode account information may also be used by a sufficiently motivated hacker to change it. Endogenous interference is thus likely a permanent feature of human existence.

Our approach to studying interference is straightforward. We begin with the conventional mechanism design framework in which preferences are assumed to depend on a set of states of the world. The actual state is common knowledge among the players. Messages must be transmitted from the players to a mediator who, implements a decision based on the messages he receives.

We depart from the standard set-up by taking the transmission of these messages to be problematic: they may be interfered with by the players of the mechanism.[1] Each player is assumed to have a "channel" along which he

[1]Interference may also arise from physical sources. We leave discussion of this "exogenous" interference to another paper.

transmits his message; one player interferes with a second player by (probabilistically) substituting the second player's own message with a different which he transmits along the second player's channel (thus assuming his identity).[2]

Now in practice the process by which players conduct interference might be quite complicated, involving for example delicate extensive-form modeling issues (you break into my office and steal my letterhead; later, I break into yours and find a letter you have written in an imitation of my handwriting on my letterhead; I replace this with a letter of my own and while I'm there substitute a fake version of the letter you wrote on your own behalf but forgot to mail, etc.). More generally, there are a many ways in which people might interfere, ranging from rhetorical ploys to electronic jamming to outright forgery. No doubt the ease of accomplishing interference will depend in part on individual skills, on the physical nature of the messages themselves, on encryption technology, on the availability of physical or institutional means of sending secure messages, etc. A complete understanding of these issues certainly falls out of the domain of competence of the economist (at least *these* economists!)

We therefore take an agnostic approach, and simply posit that there is a technology of interference with certain simple properties. Specifically, each player transmits a message along *every* channel; the message that the mediator receives on channel i may have been sent by player i or by some other player j. The probability that i's message in the guise of j arrives on j's channel can be influenced by costly "efforts" that the players exert in securing their own channels and accessing others. We make one important independence restriction: the probabilities and effort costs are independent of the messages sent and of the "physical" nature of the messages.

This "black box" approach to modeling the interference technology, though crude, seems the natural starting point for investigating what allocations are feasible in the presence of interference. A distinct advantage is that it yields a generalization of the revelation principle, which we dub the "Noninterference Principle" that allows representation of the potentially enormous set of feasible mechanisms by the much smaller set of direct mechanisms. These direct mechanisms involve having each player send a "signed" message on behalf of himself and on behalf of each of the other players. The principal implements different allocations depending not only on the full set of messages but also on the signatures. Any equilibrium of a mechanism with

[2]This model corresponds precisely to the internet practice of "spoofing." Of course, "channel" should be interpreted broadly; it can indicate for instance one player's vocal chords and the mediator's ear. Interference may involve shouting or more clever and deceptive rhetoric

possible interference can be represented by a game in which the messages consist of the true state of the world and the true identities of the players who sent them.[3] Thus, with all feasible allocations generated by a relatively simple set of mechanisms, the study of mechanism design with interference is amenable to optimization and/or equilibrium techniques.

We then go on to apply this apparatus in to a simple contracting environment in which agents must take a common production decision; the total surplus maximizing decision depends on the state of the world. Different agents have different state dependent costs and benefits associated with the decision. We show that if their preferences are consonant, then even with the possibility of interference the first best can be achieved with no costs of securing channels. However, when preferences differ enough, and when the organization needs to be "flexible" (specifically, there is a large variance in states and the optimal decision therefore also varies widely across states), this will no longer be feasible. The second best mechanism may involve giving "authority" or effective decision making power to one person. In the direct game this is accomplished by securing his channel perfectly; in equilibrium all other player's messages are ignored. When all channels cost the same to secure, this person is the one whose preferences are most consonant with the organization objective.

Authority begins to assume the richer meaning it has in everyday parlance. Not only does it convey the notion of having the power to decide on things which affect others. But also the idee that one's word carries a lot of weight relative to others'. Our model shows that these two ideas may be closely connected: having the right to decide is equivalent to being the only one with a perfectly secure channel. This idea broaches the question raised in the recent debate on incomplete contracts by Maskin-Tirole: why can't message games substitute for authority? Our answer is that authority can be modelled as a *kind* of message game (with interference) and can arise endogenously as an optimal mechanism in certain environments.

1.1 Literature

Our work is related to a number of papers in organization theory. One strand of the literature considers *exogenous* constraints on communication: agents have limited ability to transmit or process information. Part of this literature (e.g., Bolton-Dewatripont [3], Radner [17], Radner-van Zandt [18],

[3] Although all players are sending messages along all channels in the direct game, they are being honest and "up front" about it (they are signing with their true identities) and in this sense they are not interfering with each other.

Segal [19]) ignores incentive problems and conceives of organizations essentially as communication networks designed to overcome these limitations. The design of the organization itself does not affect the ability of an individual agent to communicate. Other papers (such as Green-Laffont [5] and Melumad-Mookherjee-Reichelstein [14]) do incorporate incentives and examine how standard incentive schemes or how the design of communication structures may be altered when communication is costly. By contrast, the framework presented in this paper is not actually based on the inability for agents to process information: messages in our set-up may be arbitrarily complex and/or processed arbitrarily quickly. But neither is it inconsistent with this approach. Indeed, a deeper examination of the technology underlying the interference probabilities might very well take these considerations into account.

The other strand of literature is that in which communication is *endogenously* limited by the design of the organization itself. Here the literature is perhaps more sparse. Fudenberg-Tirole's theory of signal jamming is close in spirit to our notion of interference. Aside from the context in which the idea is applied, there is a methodological difference between our approach and theirs in that the signal jamming structure itself is not endogenously determined by the agents in the model. Perhaps more closely related in spirit are "influence activities" (Milgrom-Roberts [15]). However, their focus is on the weakening of incentive schemes in response to influence costs; they too take as given the communication and decision making structures.

Finally, the notion of authority has appeared in a number of recent papers, beginning with Grossman-Hart [6] (see for instance Hart [7], Hart-Moore [8], Aghion-Tirole [2]). These papers start from the assumption that allocations cannot be completely specified by contracts; decisions are then necessarily made by one of the parties without having been specified in advance, which is interpreted as power or authority. We certainly think this is a useful notion of authority. There have been a number of recent criticisms of this approach (Maskin-Tirole [13]) based on the fact that incompleteness can be filled by appropriate message games. Both sides of the debate miss the connection to the security of communication channels, and therefore the idea of an authority as one who is influential as well as powerful. Moreover, the scope for authority arises endogenously here: agents are perfectly capable of *conceiving* of writing arbitrarily complex contracts. But interference problems limit what could actually be carried out via contract.

2 A Model of Interference

We consider an environment with 2 agents. The extension to n agents is straightforward, but introduces some minor complications that we wish to avoid at present. The state space is Θ and the set of decisions is D. At time $t = 1$, a realization of the state $\theta \in \Theta$ is observed by all the n agents but not by outside parties. Agents have state contingent preferences on decisions and we assume that these preferences are represented by vNM utility functions u^i. $u^i(d, \theta)$ denotes the utility of agent i if decision $d \in D$ is chosen in state θ; when d is a lottery, $u^i(d, \theta)$ denotes the expected utility of the lottery d. At time $t = 0$, a "contract" is signed: a contract is here a mechanism, i.e., $(\{M_i\}, g)$ where M_i is the set of messages available to agent i and $g : \times_{i=1}^2 M_i \to D$ is a message contingent decision rule. Assume that the agents have at the time of contracting common beliefs F about the distribution of the states of the world. We assume that F has a continuous and positive density and has a bounded support $[\underline{\theta}, \overline{\theta}]$. Traditional mechanism design assumes that each agent has access to a perfectly "secure" channel of communication: when message m_i^i is received on channel i, it is known that agent i sent this message. We depart from this assumption and assume that after contracting but *before the realization of the state*[4], agents can exert efforts $\mathbf{e}$ that enable them to secure their channels and to interfere on other agents' channels. To keep things tractable, we use an indirect representation and assume that efforts translate into a probability distribution over messages received on each channel. Given a vector of efforts $\mathbf{e} = (\mathbf{e}^1, \mathbf{e}^2)$, there is a probability $\mu_i(\mathbf{e})$ that a message that he sends on his channel arrives safely and a probability $1 - \mu_i(\mathbf{e})$ that the interference of the other agent succeeds.[5] The (private) cost of effort is $c^i(\mathbf{e}^i)$. Note that this formulation

[4]The extension of the analysis to the case where the effort levels are exerted after the realization of the state is straightforward but somewhat more complex.

[5]In general, and agent can spend effort not only to secure his channel but also to interfere on other channels, i.e., $i's$ effort is a vector $\mathbf{e}^i = (e_1^i, e_2^i)$ and the probability that player i succeeds on channel j is $\mu_j^i(\mathbf{e})$, where $\mathbf{e} = (\mathbf{e}^1, \mathbf{e}^2)$. In the application we assume that $\mu_i^i(\mathbf{e})$ depends only on e_{ii}.

The general formalism can allow different interesting situations. For instance, suppose that $\mu_i^i(\mathbf{e})$ is differentiable. In general, we should expect that more effort on one's channel increases security, $\partial\mu_i^i/\partial e_i^i > 0$, that more effort by the other agent on one's channel decreases security, $\partial\mu_i^i/\partial e_i^j < 0$. Of interest are the cross effects of one's effort on the security of the other agent's channel . For instance, if $\frac{\partial^2\mu_1^1}{\partial e_1^2 \partial e_2^2} < 0$ by increasing his effort to secure his own channel (e_2^2) player 2 finds it easier (needs a lower e_1^2) to interfere on the other's channel. A simple example of this could be learning-by-doing: learning how to secure an internet network makes you learn how to get into other people's network (and reciprocally).

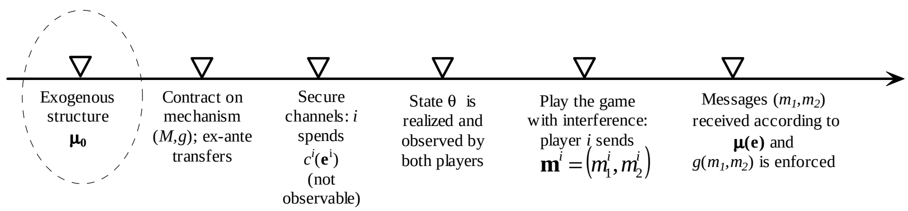

Figure 1:

allows for *public* investments in security before contracting (e.g., creation of secured court proceedings, or of a postal service which ensures against mail fraud); public investments could make channels secure even in the absence of private investment, that is $\boldsymbol{\mu}(\mathbf{0}) = \boldsymbol{\mu}_0 > 0$. Hence, while there might be a social cost associated with the public investment, they might also make contracts more efficient once $\boldsymbol{\mu}_0$ is sunk. Note that even if $\boldsymbol{\mu}_0 = \mathbf{1}$, interference by the agents might generate $\boldsymbol{\mu} < \mathbf{1}$ in some contracting environments.

Interference by agent i on channel j takes the form of agent i sending a message m_j^i on channel j. Let $\mathbf{m}^i = (m_1^i, m_2^i)$ be the vector of messages that agent i sends on the two channels. Communication is then summarized by a vector $\boldsymbol{\mu} = (\mu_1, \mu_2)$. For instance, $\boldsymbol{\mu} = (1, 0)$ represents a situation in which channel 1 is secure but channel 2 is not, $\boldsymbol{\mu} = (1, 1)$ is a case where each channel is secure and $\boldsymbol{\mu} = (0, 0)$ is a case where each channel is not secure. Note that the two last cases are equivalent from a mechanism design point of view: in $\boldsymbol{\mu} = (0, 0)$, each agent has in fact access to a secure channel: the channel of the *other* agent! For a given $\boldsymbol{\mu}$, if agents play $\mathbf{m}^1, \mathbf{m}^2$, the outcome will be a lottery $g^{\boldsymbol{\mu}}(\mathbf{m}^1, \mathbf{m}^2)$:

$$\begin{aligned} & g\left(m_1^1, m_2^1\right) \text{ with probability } \mu_1 (1 - \mu_2) \\ & g\left(m_1^1, m_2^2\right) \text{ with probability } \mu_1 \mu_2 \\ & g\left(m_1^2, m_2^1\right) \text{ with probability } (1 - \mu_1)(1 - \mu_2) \\ & g\left(m_1^2, m_2^2\right) \text{ with probability } (1 - \mu_1) \mu_2. \end{aligned} \tag{1}$$

Hence, agents play an "extended two-stage game": at the first stage they choose to secure their communication channels and then they play the message game that they contracted upon.

To summarize, the sequence of events is as follows.

In this extended game the strategy of each agent consists of an effort level $\mathbf{e}^i \in \mathbb{R}^2$ and of a message sending strategy $\sigma^i : \Theta \times \mathbb{R}^2 \to (M^1 \times M^2)$, i.e., $\sigma^i(\theta, \mathbf{e}^i)$ is the interference strategy of agent i in state θ when his effort choice was $\mathbf{e}^i$.

It will be convenient to denote the probability that agent i succeeds on channel 1 and agent j succeeds on channel 2 by $\pi_{ij}(\mathbf{e})$. From (1), we have $\pi_{11}(\mathbf{e}) = \mu_1(\mathbf{e})(1-\mu_2(\mathbf{e}))$, etc. From g and $\mathbf{m}$, we can also define four functions $h_{ij}(\mathbf{m}) = g\left(m_1^i, m_2^j\right)$, $i = 1,2$, $j = 1,2$. Let $\boldsymbol{\pi}(\mathbf{e}) = (\pi_{11}(\mathbf{e}), \pi_{12}(\mathbf{e}), \pi_{21}(\mathbf{e}), \pi_{22}(\mathbf{e}))$ and $\mathbf{h} = (h_{11}, h_{12}, h_{21}, h_{22})$. Then, if agents use strategies $\mathbf{m}$ the outcome is the lottery $\boldsymbol{\pi}(\mathbf{e}) \cdot \mathbf{h}(\mathbf{m})$ that selects outcome $h_{ij}(\mathbf{m})$ with probability $\pi_{ij}(\mathbf{e})$.

An equilibrium is defined in the usual way: for given $\mathbf{e}$, the strategies σ^i must form an equilibrium and for each i, $\mathbf{e}^i$ must be optimal given $\mathbf{e}^{-i}$ and σ:

$$\mathbf{e}^i \in \arg\max_{\hat{\mathbf{e}}^i} \int u^i\left(\pi\left(\mathbf{e}^{-i}, \hat{\mathbf{e}}^i\right) \cdot \mathbf{h}\left(\sigma^{-i}\left(\theta, \mathbf{e}^{-i}\right), \sigma^i\left(\theta, \hat{\mathbf{e}}^i\right)\right), \theta\right) dF(\theta) - c^i\left(\hat{\mathbf{e}}^i\right), \text{ for all } i \tag{2}$$

$$\sigma^i\left(\theta, \mathbf{e}^i\right) \in \arg\max_{\mathbf{m}^i} u^i\left(\pi\left(\mathbf{e}^{-i}, \mathbf{e}^i\right) \cdot \mathbf{h}\left(\sigma\left(\theta, \mathbf{e}^{-i}\right), \mathbf{m}^i\right), \theta\right), \text{ for all } \theta, \text{ for all } i.$$

We can define our concept of implementation.

Definition 1 *A decision rule $f : \Theta \to D$ is implementable at cost c if there exists a mechanism (M, g) and an equilibrium $(\mathbf{e}, \boldsymbol{\sigma})$ of the extended game such that $g^{\boldsymbol{\mu}(\mathbf{e})}(\sigma(\theta)) = f(\theta)$ and $c^1(\mathbf{e}^1) + c^2(\mathbf{e}^2) = c$.*

We will sometimes make the following assumption.

Assumption There exists a "unanimously worst outcome" d^0 : for any state θ, and any agent i, $u^i(d, \theta)$ is minimized at $d = d^0$. We normalize utilities in such a way that $u^i(d^0, \theta) = 0$ for all θ, all i.

3 Example : Facilitating Coordination

There are two agents, indexed by $i = 1, 2$ and two states, θ and ϕ. The set of possible decisions is $\{a, b, c\}$. Interference probabilities are exogenous (i.e., $\boldsymbol{\mu}$ is given). Payoffs are as follows ($\varepsilon \in [0, 1)$).

	a	b	c
θ	$2,2$	$3,\varepsilon$	$0,0$
ϕ	$\varepsilon,3$	$2,2$	$0,0$

Hence, agent 1 always prefer b to a to c while agent 2 always prefer a to b to c. Note that the surplus maximizing decision rule $f(\theta) = a$, $f(\phi) = b$ is not Nash implementable (failure of monotonicity) and is not subgame perfect implementable (failure of a "test pair")

As we will show shortly, it is nevertheless easy to (weakly) implement f in Nash: the direct game $g(\theta,\theta)=a$, $g(\phi,\phi)=b$, $g(\theta,\phi)=g(\phi,\theta)=c$ has indeed a truthful equilibrium. However, the direct game has also other equilibria. In particular, it is possible that a is always the equilibrium outcome.

With the possibility of interference, the truthful equilibrium is fragile. Note that since agent 2 gains 1 if a is chosen instead of b in state ϕ, as long as the cost of interference is less than 1, agent 1 will like to interfere and have a as the equilibrium outcome. Similarly for agent 2 who would like to interfere and generate outcome b in state θ.

But if *both* agents can interfere, then it is possible to fully implement f by the direct revelation game. Moreover, in case $\varepsilon=0$, *any* symmetric interference technology $\mu_1=\mu_2=\mu\in(0,1)$ will make the direct game fully implement f. The result, while somewhat surprising, is actually intuitive. In state θ, agent 1 has a gain of 1 if he successfully interferes and obtains b rather than a; however, since a is the surplus maximizing decision, the other agent, agent 2, loses *more* than 1 when decision b is chosen instead of a : therefore, agent 2 has even more incentive than 1 to interfere if she believes that agent 1 will not be truthful.

Consider a symmetric interference technology and the direct game (Θ,g), where g has been defined above and $\varepsilon=0$. We claim that in state θ the unique equilibrium outcome is a. By symmetry, this shows that in state ϕ the unique equilibrium outcome is b.

Observation Consider $\varepsilon=0$ and the direct game (Θ,g). Truth-telling is the unique equilibrium of the game for any $\boldsymbol{\mu}=(\mu,\mu)$ with $\mu\in(0,1)$.

Proof. Consider state θ. Assume that agent 1 sends θ_1^1 on his channel and θ_2^1 on the other channel: agent 1 uses the strategy $\left(\theta_1^1,\theta_2^1\right)$ in the "extended game". Suppose that agent 2 sends θ_2^2 on his channel and θ_1^2 on $1'$s channel, i.e., uses $\left(\theta_1^2,\theta_2^2\right)$ in the extended game. Then since an agent succeeds on his channel with probability μ and succeeds on the other agent's channel with probability $1-\mu$, there are four possible outcomes

$$\begin{aligned}
&g\left(\theta_1^1,\theta_2^1\right) \text{ with probability } \mu(1-\mu)\\
&g\left(\theta_1^1,\theta_2^2\right) \text{ with probability } \mu^2\\
&g\left(\theta_1^2,\theta_2^1\right) \text{ with probability } (1-\mu)^2\\
&g\left(\theta_1^2,\theta_2^2\right) \text{ with probability } (1-\mu)\mu.
\end{aligned}$$

Note that if a agent sends (θ,θ), the best response of the other agent is to send (θ,θ). For instance, if 1 sends (θ,θ), agent 2 gets the maximum payoff of 2 in sending (θ,θ) : any other strategy will generate a lower probability of getting decision a. If agent 2 sends (θ,θ), since the outcome is c in cases of conflicts, if agent 1 sends (ϕ,ϕ), he obtains a payoff of $5\mu(1-\mu)$ which is less than 2 since $\mu\in(0,1)$. If agent 1 sends (θ,ϕ) or (ϕ,θ) his payoff is less than 2 since he will either agree with agent 2 on state θ or will disagree with positive probability. Hence truth-telling is indeed an equilibrium play.

Consider the other possible pure strategies of agent 1. Let $u^2\left(\left(\theta_1^2,\theta_2^2\right)\middle|\left(\theta_1^1,\theta_2^1\right)\right)$ be the payoff to agent 2 of using $\left(\theta_1^2,\theta_2^2\right)$ when agent 1 uses $\left(\theta_1^1,\theta_2^1\right)$. Suppose that $\left(\theta_1^1,\theta_2^1\right)=(\theta,\phi)$; then

$$\begin{aligned}
u^2\left((\theta,\theta)|\,(\theta,\phi)\right) &= 2\left(\mu^2+\mu(1-\mu)\right)\\
u^2\left((\theta,\phi)|\,(\theta,\phi)\right) &= 0\\
u^2\left((\phi,\theta)|\,(\theta,\phi)\right) &= 2\mu^2\\
u^2\left((\phi,\phi)|\,(\theta,\phi)\right) &= 0,
\end{aligned}$$

and the unique best response of agent 2 is (θ,θ) since $\mu\in(0,1)$. But then by the previous remark, (θ,ϕ) is not a best response to (θ,θ). The same reasoning applies for all other strategies $\left(\theta_1^1,\theta_2^1\right)$ of agent 1. Hence, the unique pure equilibrium strategy in state θ is truth-telling. Note that truth-telling is in fact a strict Nash equilibrium. It then follows that truth-telling is also the unique equilibrium in mixed strategies. ∎

Remark 2 *We have so far assumed that agents cannot exert effort and change the security of their channels. Consider the case* $\varepsilon=0$. *Letting* $\mu_0>0$ *be the initial (symmetric) security of the channels, say due to existing public investments, etc., assume that each agent can increase the security on his channel to* $\mu>\mu_0$ *at cost* $c(\mu-\mu_0)$. *Assume that agent 2 does not invest in security, hence that* $\mu_2=\mu_0$. *If agent 1 has security* $\mu_1\geq\mu_0$ *on his channel, then, given that agent 2 is truthful, agent 1 obtains by deviating in state* θ *a payoff*

$$u^1\left((\phi,\phi),(\theta,\theta)\right)=3\mu_1(1-\mu_0)+2\mu_0(1-\mu_1)$$

agent 1 finds it beneficial to deviate only if $u^1\left((\phi,\phi),(\theta,\theta)\right)\geq u^1\left((\theta,\theta),(\theta,\theta)\right)=2$, *i.e., if*

$$\mu_1(3-5\mu_0)\geq 2(1-\mu_0).$$

If $\mu_0\geq\frac{3}{5}$, *agent 1 cannot gain from deviating, even if his channel is more secure* $(\mu_1>\mu_0)$ *and since he does not want to interfere, agent 1 has no incentive to invest in more security. Similarly, for* $\mu_0\in\left[\frac{1}{3},\frac{3}{5}\right)$, *the condition is* $\mu_1\geq\frac{2(1-\mu_0)}{3-5\mu_0}>1$ *which is impossible.*

If $\mu_0 < \frac{1}{3}$ then agent 1 gains in state θ when $\mu_1 \geq \frac{2(1-\mu_0)}{3-5\mu_0}$. Assuming that each state has equal probability, agent 1's utility gain is equal to $\frac{1}{2}\delta(\mu_1) - c(\mu_1 - \mu_0)$, where $\delta(\mu_1) = \mu_1(3-5\mu_0) - 2(1-\mu_0)$ is the utility gain in state θ. Note that $\delta(\mu_1)$ is increasing and is equal to zero at $\mu_1 = \frac{2(1-\mu_0)}{3-5\mu_0}$, which is strictly greater than $\frac{2}{3}$ since $\mu_0 \geq 0$. Hence, in order to replicate *the outcome without deviation, agent 1 needs to exert an effort that will increase the security on his channel by at least $\frac{1}{3}$, and the smaller μ_0 is, the larger the effort that agent 1 has to exert to replicate the outcome without deviation, and the more costly it is. Clearly, if $c\left(\frac{2}{3}\right)$ is large enough, even if μ_0 is close to* 0, *agent* 1 *will not want to deviate. For instance, let $c(\mu_1 - \mu_0) = \mu_1 - \mu_0$ for $\mu_1 \geq \mu_0$, the marginal gain of agent* 1 *is*

$$\frac{1}{2}\delta'(\mu_1) - c'(\mu_1) = 2 - 5\mu_0.$$

If $\mu_0 < \frac{2}{5}$, the marginal gain is positive and agent 1 maximizes his payoff from deviating to $\mu_1 = 1$, which yields a net gain of $-2\mu_0 < 0$. If $\mu_0 \in \left(\frac{2}{5}, \frac{1}{3}\right)$, the maximum from deviation is attained at $\mu_1 = \mu_0$. Hence, independently of the initial public investment $\mu_0 \in (0,1)$, the unique equilibrium of the extended game is for agents not to modify the security of their channel, to be truthful and the surplus maximizing decision rule is fully implemented by the direct game.

Obviously, if the marginal cost $c'(\mu_1 - \mu_0)$ is "small" around 0, *agent 1 will in general gain from deviating if μ_0 is small. For instance, if $c(\mu_1 - \mu_0) = \frac{(\mu_1-\mu_0)^2}{2}$, it is possible to show that agent 1 will not deviate only if $\mu_0 > \frac{5}{7} - \frac{4}{35}\sqrt{15}$. If μ_0 is smaller than this bound, then both agents will invest in additional security.*

Remark 3 *The assumption that the set of decisions is $\{a, b, c\}$ is not innocuous. If one extends D to include lotteries over these outcomes,*[6] *then the decision rule $f(\theta) = a$ and $f(\phi) = b$* is monotonic. *For instance, letting π denote the lottery $(\pi_{22}, \pi_{11}, 1 - \pi_{11} - \pi_{22})$, agent 1 prefers a to π in state θ but prefers π to a in state ϕ and monotonicity is not violated for f. As it is well known, monotonicity is far from being sufficient when there are only two agents ([16]). Nevertheless, for the example at hand, if lotteries can be used, full implementation is obtained by the following game (when channels are secure). Let $h : \Theta^2 \to \{a, b, c\}$ be such that $h(\theta_1, \theta_2) = c$ if $\theta_1 \neq \theta_2$ and $h(\theta, \theta) = f(\theta)$. The messages are $M_1 = M_2 = \Theta^2$, and the outcome function is the lottery $g\left(\left(\theta_1^1, \theta_2^1\right), \left(\theta_1^2, \theta_2^2\right)\right)$ that selects $h\left(\theta_1^i, \theta_2^j\right)$ with probability π_{ij}. If lotteries cannot be used as outcomes, then our model suggests that the*

[6] For another use of lotteries in implementation, see Abreu-Matsushima [1].

design of communication structures that are not perfectly secure *can create the desired randomization.*

The result of Observation 3 generalizes to any $\varepsilon \in [0, 2/3)$: one can always find a symmetric interference structure (μ, μ) such that the efficient decision is the unique equilibrium outcome of the extended game corresponding to the direct game (Θ, h) where $h(\theta, \phi) = h(\phi, \theta) = d^0$ and $h(\theta, \theta) = a$, $h(\phi, \phi) = b$. The proof of the following proposition mimics the proof of Observation 3: the bounds on μ are found by imposing that for any strategy (θ_1^i, θ_2^j) of agent i, agent $-i$'s unique best response is (θ, θ). Note that as ε is close to 0, μ can be chosen as small as we want but that as ε approaches $2/3$, μ must be chosen close to $\frac{1}{2}$. For $\varepsilon \geq 2/3$, it is not possible to rule out b in state θ as an outcome of this mechanism.

Proposition 4 *Suppose that* $\varepsilon \in \left[0, \frac{2}{3}\right)$. *Choose* $\mu \in \left(\frac{1}{2} - \sqrt{\frac{2-3\varepsilon}{8+4\varepsilon}}, \frac{1}{2} + \sqrt{\frac{2-3\varepsilon}{8+4\varepsilon}}\right)$. *Then the unique equilibrium of the extended game corresponding to the direct game is the efficient decision rule.*

There are two main lessons to draw from this example. First, interference structures can facilitate the implementation of desired decision rules since by creating "voice" on other agents' channels, interference creates the possibility for agents to signal their disagreement. This intuition is very similar to the type of construction used in the implementation literature. In fact, as we noted in our Remark 3, the logic is in fact equivalent since there exists, in a perfectly secured environment, a mechanism in which each agent sends an ordered pair of states and the decision rule depends on the results of a lottery that selects a pair among all possible pairs. We show in the next section that while we do not insist on full implementation here, the revelation principle by which we can describe the set of allocations is based on the same logic. Second, the example suggests that there is a relationship between the stakes (measured by $2-\varepsilon$) and the minimal security of channels required for implementation. The agent who gains by imposing the inefficient decision gains 1 but the other agent loses $2-\varepsilon$. As ε increases, the other agent loses less from a deviation and therefore has less incentives to signal the deviation or to secure his channel. This explains why as ε increases it is *more difficult* to obtain full implementation with a given security structure. In general, the logic that agents have less incentives to interfere when the stakes are lower seems rather intuitive. We will show in an application how this logic is articulated and how our approach can provide some foundations for the emergence of authority relationships in some contracting environments and the emergence of "consensual" relationships in others.

4 The Noninterference Principle

The previous example proved that a direct game satisfying truth-telling can implement the surplus maximizing decision rule. Readers familiar with the revelation principle would not be surprised by this observation. However, it turns out that some care must be taken before generalizing this observation.

Indeed, a necessary condition for (Θ, h) to represent the outcome of an equilibrium σ of (M, g) is that $h(\theta, \theta) = g^{\mu}(\sigma^1(\theta), \sigma^2(\theta))$ where $\boldsymbol{\mu} = \boldsymbol{\mu}(\mathbf{e})$, $\mathbf{e}$ is the equilibrium choice of effort levels. Suppose that there is a worst outcome and that $h(\theta, \phi) = h(\phi, \theta) = d^0$. If agent 1 uses strategy (ϕ, ϕ) in the direct game and if agent 2 uses the strategy (θ, θ), the outcome is $h(\phi, \phi) = g^{\mu}(\sigma^1(\phi), \sigma^2(\phi))$ with probability π_{11} and $h(\theta, \theta) = g^{\mu}(\sigma^1(\theta), \sigma^2(\theta))$ with probability π_{22}. Incentive compatibility requires that

$$\pi_{11} u^1\left(g^{\mu}\left(\sigma^1(\phi), \sigma^2(\phi)\right)\right) \leq (1 - \pi_{22})\, u^1\left(g^{\mu}\left(\sigma^1(\theta), \sigma^2(\theta)\right)\right). \tag{3}$$

But the equilibrium conditions in (M, g) require that

$$u^1\left(g^{\mu}\left(\sigma^1(\phi), \sigma^2(\theta)\right)\right) \leq u^1\left(g^{\mu}\left(\sigma^1(\theta), \sigma^2(\theta)\right)\right) \tag{4}$$

and there is no immediate relationship between the two conditions. In the direct game, when agent 1 interferes, he can in fact change (with some probability) the *strategy* that agent 2 uses in the initial mechanism. To avoid this problem, one needs to be able to distinguish whether a message received on channel 2 was sent by agent 1 or by agent 2. Obviously, the structure does not allow us to have this information and therefore agents must "self-signal" their identity. To induce agents to signal their identity, it is then necessary to adjust the outcome with respect to the believed origin of the message. Hence, if one receives on channel 1 a message labeled "1" and on channel 2 a message labeled "1", the outcome will be a function h_{11} while if one receives on channel 1 a message labeled "2" and on channel 2 a message labeled "1", the outcome will be a function h_{21}. In the example above, all the h_{ij} functions were taken to be the same, but this is not a general property. In fact, since it is also necessary to induce the agents to invest in the "right" level of security, it is necessary in general to have different functions h_{ij} (the next example will be such an example).

A way to think of our revelation principle is that we first create for each initial channel i, two "virtual" channels, one for each agent . Incentive compatibility requires that each agent "tells the truth": announces the true state *and* his true identity on his virtual channel. On each initial channel i, the probability that the message sent on the virtual channel j succeeds is equal to μ_i if $i = j$ and is equal to $1 - \mu_i$ if $j \neq i$. The figure below is a representation of the principle.

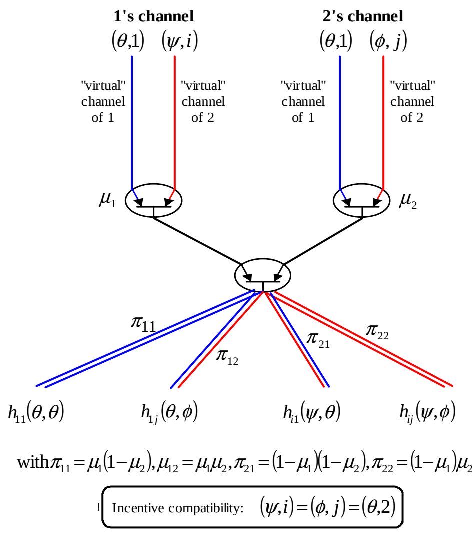

Figure 2:

Hence, we have here a "two-dimensional" incentive compatibility problem and it is this additional dimension that enables us to replicate in the direct game the independence of strategies that exists in the initial game.

A *direct* game is defined by message sets $M_1 = \Theta \times \{1,2\}$, $M_2 = \Theta \times \{1,2\}$ and an outcome function $h((\theta,i),(\phi,j)) = h_{ij}(\theta,\phi)$. Fixing $\boldsymbol{\mu}$, we define a *truth-telling equilibrium* in the extended game by the strategy $\gamma^i(\theta) = ((\theta,i),(\theta,i))$.[7]

Proposition 5 *Consider an equilibrium* $(\mathbf{e},\sigma)$ *of the extended game corresponding to the mechanism* (M,g) *and consider the decision rule* $f(\theta) = g^{\boldsymbol{\mu}(\mathbf{e})} \circ \sigma(\theta)$. *There exists a direct game* $(\Theta \times \{1,2\}, h)$ *such that, if* γ *is the truth-telling strategy,* $(\mathbf{e},\boldsymbol{\gamma})$ *is an equilibrium of the corresponding extended game and such that* $f(\theta) = h^{\boldsymbol{\mu}(\mathbf{e})} \circ \gamma(\theta)$.

Proof. Appendix. ■

When there is a worst outcome, incentive compatibility of agent 1 takes two forms, depending on the nature of his deviation. Agent 1 can either tell the truth about the state but lie about his identity or he can lie about the state but tell the truth about his identity. Additional conditions are that $\mathbf{e}$ is indeed an equilibrium, i.e., that (2) holds for the direct game.

[7] Note that we require truth-telling only when agent i has taken his equilibrium effort. If agent i deviates from his equilibrium effort, his play in the extended game will be different from truth-telling.

5 Authority as a Solution to the Interference Problem

Consider a firm with n agents. The agents have to decide on which production technology to use, and we index the technology by a parameter $q \in \mathbb{R}^+$ that we will also call "decision". Once the technology is settled upon, an output is realized that is equal to $Y(q) = Y - \frac{1}{2}(q - q^*)^2$; [8] Agents might disagree on which technology is best, and this disagreement might be a function of the state of the world. For instance, engineers and marketing people might disagree on the degree of quality that the production line should produce: the larger the quality, the slower the pace of the production. Engineers value quality but marketing people value volume. Among engineers, some might prefer to produce sport cars while others might prefer to produce station wagons. We represent the preferences of the agents by $\alpha^i(\theta)\,q$, where $\theta \in \left[\underline{\theta}, \overline{\theta}\right] \subset \mathbb{R}^+$ indexes the state of the world and each α^i is differentiable. Total surplus is maximized at

$$q^*(\theta) = q^* + \sum \alpha^i(\theta).$$

A contract specifies an income $I^i(\theta)$ and a decision $q(\theta)$ for each announced state θ. We impose budget balancing:

$$\sum I^i(\theta) = Y(q(\theta)).$$

The first best is attained when each agent does not invest in securing his channel and "tells the truth". We will prove that as long as the preferences of the agents are consonant, even if an agent can perfectly interfere with other agents' channels, he will still prefer not to interfere and to tell the truth.

Trivially, if for each θ the sum $\sum \alpha^i(\theta)$ is constant (for instance, $\boldsymbol{\alpha}$ is an element of the $n-1$ simplex), there is a uniform optimal decision and agents can contract ex-ante on this decision. (There is no need for message games to reveal the state unless one is also worried about ex-post utility levels.)

When the sum of the benefits varies across states, implementation is less obvious. Starting from $\boldsymbol{\mu} = (0, 0)$, if agent 1 deviates and secure his channel perfectly, he can decide on the outcome since $\pi_{11} = 1$. A sufficient condition for such a deviation not to be beneficial is when agent 1 will in fact choose the outcome that would have been chosen under $\boldsymbol{\mu} = (0, 0)$. This sufficient condition amounts to show that the initial contract would make agent 1 be

[8] By risk neutrality and the absence of moral hazard in production, this might as well be the mean of a stochastic variable. We will consider this interpretation when we introduce monotone mechanisms.

truthful even if he could choose the outcome. The same reasoning applies to agent 2. Conditions for an agent to be "classically" incentive compatible are that for each $\theta, \hat{\theta}$:

$$\begin{aligned} I^i(\theta) + \alpha^i(\theta) q^*(\theta) &\geq I^i\left(\hat{\theta}\right) + \alpha^i(\theta) q^*\left(\hat{\theta}\right) \\ I^i\left(\hat{\theta}\right) + \alpha^i\left(\hat{\theta}\right) q^*\left(\hat{\theta}\right) &\geq I^i(\theta) + \alpha^i\left(\hat{\theta}\right) q^*(\theta). \end{aligned}$$

Adding the two inequalities yield

$$\left[\alpha^i(\theta) - \alpha^i\left(\hat{\theta}\right)\right]\left[\sum\left(\alpha^j(\theta) - \alpha^j\left(\hat{\theta}\right)\right)\right] \geq 0. \tag{5}$$

Moreover, dividing both sides by $\left(\theta - \hat{\theta}\right)^2$ and taking the limit as $\theta \to \hat{\theta}$ yield

$$\alpha^{i\prime}(\theta) \sum \alpha^{j\prime}(\theta) \geq 0. \tag{6}$$

If (5) does not hold, then an agent who can interfere perfectly with other agents' messages has an incentive to deviate from truth-telling. If the cost of interfering perfectly goes to zero, (6) is in fact a necessary condition for implementation of the first best. It is not sufficient in general (this has to do with global versus local incentive compatibility). We provide a sufficient condition below.

Definition 6 *Let* α^i *be the derivative of* α^i *and* $\boldsymbol{\alpha}' = \left(\alpha^{1\prime}, \ldots, \alpha^{n\prime}\right)$. *The agents have consonant preferences if* $\boldsymbol{\alpha}' \geq 0$ *or if* $\boldsymbol{\alpha}' \leq 0$.

Linear case $\alpha^i(\theta) = \alpha^i\theta$. In this case, consonance occurs when all α^i are of the same sign.

Proposition 7 *Suppose that agents are consonant. Then the first best can be implemented.*

Proof. Incentive compatibility implies that $I^i(\theta)$ is differentiable[9] and that

$$I^{i\prime}(\theta) = -\alpha^i(\underline{\theta}) \sum \alpha^{j\prime}(\theta).$$

[9]Indeed, combing the two incentive compatibility conditions, we obtain

$$\alpha^i\left(\hat{\theta}\right)\left(q^*\left(\hat{\theta}\right) - q^*(\theta)\right) \geq I^i(\theta) - I^i\left(\hat{\theta}\right) \geq \alpha^i(\theta)\left(q^*\left(\hat{\theta}\right) - q^*(\theta)\right)$$

.Since $q^*\left(\hat{\theta}\right) - q^*(\theta) = \sum\left(\alpha^j\left(\hat{\theta}\right) - \alpha^j(\theta)\right)$ and since α^j is differentiable, by dividing all sides by $\theta - \hat{\theta}$ yields the result.

It follows that

$$I^i(\theta) = I^i(\underline{\theta}) - \int_{\underline{\theta}}^{\theta} \alpha^i\left(\tilde{\theta}\right) \sum \alpha^{j\prime}\left(\tilde{\theta}\right) d\tilde{\theta}. \tag{7}$$

Choose $\{I^i(\underline{\theta})\}$ in such a way that,

$$\sum I^i(\underline{\theta}) = Y - \frac{1}{2}\left(\sum \alpha^j(\underline{\theta})\right)^2.$$

It follows that[10]

$$\begin{aligned}\sum I^i(\theta) &= Y - \frac{1}{2}\left(\sum \alpha^j(\underline{\theta})\right)^2 - \int_{\underline{\theta}}^{\theta} \left(\sum \alpha^i\left(\tilde{\theta}\right)\right)\left(\sum \alpha^{j\prime}\left(\tilde{\theta}\right)\right) d\tilde{\theta} \\ &= Y - \frac{1}{2}\left(\sum \alpha^j(\theta)\right)^2.\end{aligned}$$

and budget balancing is satisfied everywhere.

We prove our claim if we show that the local incentive compatibility condition implies the global incentive compatibility condition. Here, the argument is familiar from the regulation literature. Let $u^i\left(\hat{\theta}, \theta\right) = I^i\left(\hat{\theta}\right) + \alpha^i(\theta) q^*\left(\hat{\theta}\right)$ be the utility of i if he interferes in state θ and succeeds in obtaining the outcome of state $\hat{\theta}$. Suppose that there exist $\hat{\theta}$ and θ such that $u^i\left(\hat{\theta}, \theta\right) - u^i(\theta, \theta) > 0$. Standard arguments show that failure of global incentive compatibility implies that

$$\int_{\theta}^{\hat{\theta}} \int_{x}^{\theta} \frac{\partial^2 u^i(x,y)}{\partial x \partial y} dy dx > 0.$$

However by consonance,

$$\frac{\partial^2 u^i(x,y)}{\partial x \partial y} = \alpha^{i\prime}(x) \sum \alpha^{j\prime}(y) \geq 0$$

and we obtain a contradiction (for instance, if $\hat{\theta} > \theta$, then $x > \theta$ and the integrand is negative). ■

Example 8 *In the linear case, if $\sum \alpha^j \neq 0$ and agents are consonant, the first best decision rule $q^*(\theta)$ can be implemented with a mechanism with a linear sharing rule $I^i(\theta) = s_i Y(q^*(\theta))$, where $s_i = \frac{\alpha^i}{\sum \alpha^j}$.*

[10] Indeed, note that the primitive of the function $\sum \alpha^j(\theta) \sum \alpha^{j\prime}(\theta)$ is $\frac{1}{2}\left(\sum \alpha^j(\theta)\right)^2$.

We now come to the case in which the first best decision is state contingent and in which agents are not consonant.

Here we appeal to our noninterference principle. We assume that $\mu_i(\mathbf{e})$ depends only on $i's$ effort level on his own channel, in particular, it is not possible for agent 1 to change the security level on channel 2. This enables us to identify the effort level of i by μ_i and to define the cost of effort by $c(\mu_i)$ (we assume symmetric cost functions). We restrict attention to the linear case: $\alpha^i(\theta) = \alpha^i\theta$ for all i and all θ.

5.1 Two Agents

Assume that $\alpha^1(\theta) = \alpha\theta$, that $\alpha^2(\theta) = \beta\theta$ and that $\alpha + \beta > 0$ while $\alpha > 0 > \beta$. To simplify notation, let $q^* = 0$. The first best is

$$q^*(\theta) = (\alpha + \beta)\theta.$$

Note that agent 1 would like q to increase (since $\alpha < 0$) and agent 2 would like q to decrease (since $\beta > 0$) with respect to the first best situation. However, agent 1 has the same direction of preferences as the first best while agent 2 has an opposite direction of preferences. We are interested in finding contracts that maximize the ex-ante total surplus of the relationship. The underlying assumption is that agents can make ex-ante transfers in order to select such a contract. As we can anticipate from our previous work ([11], [12]), the conclusions (in terms of information structure) will be quite different if agents have limited means of transferring money ex-ante, say because they have limited wealth and/or because the financial market is imperfect.

Suppose that there can be two levels of effort: high effort makes the channel perfectly secure ($\mu_i = 1$) and low effort makes the channel insecure ($\mu_i = 0$). The cost of high effort is $c > 0$ and the cost of low effort is zero. We consider only pure strategy equilibria in effort levels.

There are four possible communication structures since $\boldsymbol{\mu} \in \{0,1\}^2$. In the case $\boldsymbol{\mu} = (1,0)$, agent 1 can select the outcome and agent 2 cannot interfere. We view this situation as a situation of "authority". The case $\boldsymbol{\mu} = (0,1)$ is the mirror case but we will show that this case is ex-ante dominated by the previous case: *if* an agent has authority, it should be the agent who has the same direction of preferences as in the first best. This result generalizes nicely to the n person case.

Whether or not authority emerges as the optimal communication structure depends on the implied cost that the incentive compatibility constraints impose on the ex-ante contract *and* on the stakes that are present in the first best contract. Remember that two sets of conditions must be satisfied: first, truth-telling must hold given the equilibrium structure $\boldsymbol{\mu}$, second, no agent must want to change his effort level. For instance, agent 2 does not want to increase μ_2 and agent 1 does not want to decrease μ_1 when respect to $\boldsymbol{\mu} = (1, 0)$.

5.1.1 A Necessary Condition for Authority

When $\boldsymbol{\mu} = (1, 0)$, agent 1 can effectively select the messages that will be received since $\boldsymbol{\pi} = (1, 0, 0, 0)$. Incentive compatibility of agent 2 is immediate since agent 2 has "no voice". By sending a message $\left(\left(\theta_1^1, i\right), \left(\theta_2^1, j\right)\right)$ agent 1 effectively selects the decision $h_{ij}\left(\theta_1^1, \theta_2^1\right)$. Incentive compatibility means that agent 1 prefers to send the message $((\theta, 1), (\theta, 1))$ in state θ. This implies two conditions:

- Agent 1 does not want to misrepresent the state.
- Agent 1 does not want to lie about his identity.

The second condition is easily satisfied by assuming that $h_{ij} = d^0$ as long as $ij \neq 1$. The first condition is easier to satisfy if $h_{11}(\theta, \phi) = d^0$ whenever $\theta \neq \phi$. Since by sending $((\phi, 1), (\phi, 1))$ agent 1 selects the decision $h_{11}(\phi, \phi)$, the first condition is satisfied when for each state θ, agent 1 prefers $h_{11}(\theta, \theta)$ to $h_{11}(\phi, \phi)$ for any ϕ. Since the worst decision is chosen if the two announced states disagree, we might simply write $h_{11}(\theta)$ to denote $h_{11}(\theta, \theta)$ and we have just stated that h_{11} (considered as a function of one variable) must be incentive compatible in the "classical" sense. From our previous observations, this implies that in $h_{11} = (I_{11}, q_{11})$, q_{11} is increasing in θ and I_{11} is differentiable and satisfies $I_{11}'(\theta) = -\alpha\theta q_{11}'(\theta)$.

Now, we have to verify that agents indeed want to exert the right efforts in terms of securing their channels. If agent 1 does not exert effort, then $\boldsymbol{\pi} = (0, 0, 1, 0)$, i.e., each agent succeeds on the *other* channels. Since 2 is incentive compatible and sends $((\theta, 2), (\theta, 2))$, if 1 sends $\left(\left(\theta_1^1, i\right), \left(\theta_2^1, j\right)\right)$, the final message that is received is $\left((\theta, 2), \left(\theta_2^1, j\right)\right)$. However, since $h_{2j} = d^0$, agent 1 does not gain by exerting a low level of effort as long as his expected equilibrium utility is positive (i.e., there is enough surplus in the relationship to compensate 1 for his cost of effort c), which is what we will assume from now on.

Suppose now that agent 2 deviates and exerts the high level of effort, i.e., generates a security structure $(1, 1)$ and $\boldsymbol{\pi} = (0, 1, 0, 0)$. Since agent 1 is incentive compatible and send $((\theta, 1), (\theta, 1))$ in state θ, if agent 2 sends $\left(\left(\theta_1^2, i\right), \left(\theta_2^2, j\right)\right)$ in state θ, the final message will be $\left((\theta, 1), \left(\theta_2^2, j\right)\right)$. Since

$h_{1j}\left(\theta, \theta_2^2\right) = d^0$ whenever $j \neq 1$ or $\theta_2^2 \neq \theta$, agent 2 can either replicate the decision $h_{11}(\theta, \theta)$ or obtain the worst outcome d^0. Hence, agent 2 cannot be made strictly better off. As long as $c > 0$, agent 2 will not exert the high level of effort.

Clearly, incentives to exert the high effort level are strongest for agent 1, the larger is his payoff under $\boldsymbol{\mu} = (1, 0)$. It follows that we might as well take h_{11} to be the efficient decision rule, subject to the incentive compatibility condition (7).

Lemma 9 *Let W^* be the first best level of total ex-ante surplus and assume that $W^* > c$. A lower bound on the total welfare is that obtained with 1-authority and is equal to $W^* - c$.*

Proof. Using $q_{11}(\theta, \theta) = q^*(\theta)$ and $I_{11}(\theta, \theta) = I_{11}(\underline{\theta}) - \alpha(\alpha + \beta)\left(\int_{\underline{\theta}}^{\theta} \tilde{\theta} dF\left(\tilde{\theta}\right)\right)$, agent 1 is incentive compatible given $\boldsymbol{\mu} = (1, 0)$ and the outcome function h_{11}. The arguments in the text conclude the proof. ■

What about $2's$ having authority? Replicating the reasoning in the text, if the equilibrium is $\boldsymbol{\mu} = (0, 1)$—agent 2 has authority— h_{22} must be incentive compatibility. However, incentive compatibility for 2 implies that q_{22} is *non-increasing* in θ; this is a simple illustration of 2 having a different direction of preferences than the first best. It is straightforward then that the surplus maximizing contract—conditional on $\mu = (0, 1)$ and truth-telling being an equilibrium— is for $q_{22}(\theta, \theta)$ to be a constant. Maximizing welfare subject to a constant decision yields $q_{22}(\theta, \theta) = (\alpha + \beta)\left(\int \tilde{\theta} dF\left(\tilde{\theta}\right)\right)$, i.e., the constant decision should be equal to the first best decision at the average state. It follows that authority to 2 yields a lower total surplus than the combination of insecure channels and a constant decision (since c is saved in the later case). Letting E denote the expectation with respect to F, and var the variance, total welfare is then equal to

$$
\begin{aligned}
Y + \frac{1}{2}(\alpha + \beta)^2 (E[\theta])^2 &= W^* - \frac{1}{2}\left\{E\left[\theta^2\right] - (E[\theta])^2\right\} \\
&= W^* - \frac{var[\theta]}{2}.
\end{aligned}
$$

The following is then immediate.

Proposition 10 *1-authority dominates a constant decision rule if and only if the cost of securing a channel is such that $c \leq \frac{(\alpha+\beta)^2}{2} var[\theta]$.*

From Proposition 7, if $\beta = 0$, then 1 and 2 are consonant and the first best can be attained with $\boldsymbol{\mu} = (0, 0)$. Hence, as long as β is not "too negative", the best contract compatible with $\boldsymbol{\mu} = (0,0)$ dominates the best contract with $\boldsymbol{\mu} = (1,0)$. When $\beta = -\alpha$, the optimal decision is the decision $q = 0$. Hence, as long as β is "not too different" from $-\alpha$, the structure $\boldsymbol{\mu} = (0,0)$ will again be optimal.

Hence, if 1-authority is optimal, it will be optimal for intermediate values of β. We show indeed below that 1-authority is optimal when β takes intermediate values. But we also show that 1-authority is more likely to emerge as the *variance* in the first best decisions is important. This shows how the *stakes* created in the efficient decision rule (stakes that are proportional to the variance of the state) influence the incentives of agents to interfere and how these stakes make the cost c spent by agent 1 for securing his channel the exact bound on contract efficiency. This bound is attained only when agent 1 has authority.

5.1.2 A Sufficient Condition for authority

Assume that $\boldsymbol{\mu} = (0,0)$ is part of an equilibrium, i.e., that both agents exert the low effort. In this case, each agent has access to a secure channel: the channel of the other agent. The easiest way to satisfy incentive compatibility is to set $h_{ij}(\theta,\phi) = d^0$ when $ij \neq 21$ or when $\theta \neq \phi$. This is enough to make *any* decision rule incentive compatible *once* $\boldsymbol{\mu} = (0,0)$ is given.

Now, each agent can unilaterally change the security structure and once he has done so he can select in each state θ the outcome in the range of h_{21} that he prefers. Remember that if agent 1 secures his channel, then with probability one his message will succeed on both channels ($\boldsymbol{\pi} = (1,0,0,0)$) and by using the strategy $((\phi,2),(\phi,1))$ he can effectively "select" the outcome $h_{21}(\phi,\phi)$. Because only the deviations in which the same message is sent on each channel can be beneficial, we simplify notation and write $h(\phi)$ instead of $h_{21}(\phi,\phi)$; no ambiguity should arise. We also write the shares and the decision rule that are chosen under $h_{21}(\theta,\theta)$ as $I^i(\theta)$ and $q(\theta)$. It is convenient to write the decision rule in terms of a shift with respect to the first best decision

$$q(\theta) = (\alpha + \beta) d(\theta) \tag{8}$$

Note that the first best decision rule corresponds to $d(\theta) = \theta$ for each θ.

Denote the expected utility of agent i when the decision corresponding to state ϕ is implemented while the state is θ by $u^i(\phi,\theta)$

$$u^i(\phi,\theta) = I^i(\phi) + \alpha^i \theta q(\phi).$$

If i unilaterally exerts the high effort, then in state θ he will send messages in order to maximize $u^i(\phi,\theta)$. Let $\phi^i(\theta)$ be a solution to the problem

$\max_\phi u^i(\phi,\theta)$. Then agent i will decide not to exert the high effort level only if

$$E\left[u^i\left(\phi^i(\theta),\theta\right)\right]-c\leq E\left[u^i(\theta,\theta)\right]. \tag{9}$$

The surplus maximizing mechanism conditional on agents exerting the low effort will solve the problem

$$\begin{cases} \max_h W(h)=E\left[W(\theta)\right] \\ \quad \text{s.t. } (9),\ i=1,2. \end{cases} \tag{10}$$

where $W(\theta)$ is the total surplus in state θ given the mechanism h. By budget balancing, $W(\theta)=Y(q(\theta))+(\alpha+\beta)\theta q(\theta)$.

We note that a revealed preference argument implies that $q\left(\phi^1(\theta)\right)$ and $u^1\left(\phi^1(\theta),\theta\right)$ are increasing in θ and that $q\left(\phi^2(\theta)\right)$ and $u^2\left(\phi^2(\theta),\theta\right)$ are decreasing in θ (see the proof in the appendix for details). We will not attempt here to fully characterize the solution to (10), but rather we will derive conditions under which the optimal welfare in (10) is less than W^*-c, which proves that 1-authority is indeed optimal.

Consider again the moral hazard constraint for agent 2. Clearly, for each θ, and each $\hat{\theta}$,

$$\begin{aligned} u^2\left(\phi^2(\theta),\theta\right) &\geq u^2\left(\hat{\theta},\theta\right) \\ &= I^2\left(\hat{\theta}\right)+\beta(\alpha+\beta)\theta d\left(\hat{\theta}\right) \end{aligned}$$

Hence, taking expectations with respect to $\hat{\theta}$ we have,

$$u^2\left(\phi^2(\theta),\theta\right)\geq E\left[I^2\left(\hat{\theta}\right)\right]+\beta(\alpha+\beta)\theta E\left[d\left(\hat{\theta}\right)\right].$$

Taking expectations over θ, we finally obtain

$$E\left[u^2\left(\phi^2(\theta),\theta\right)\right]\geq E\left[I^2(\theta)\right]+\beta(\alpha+\beta)E[\theta]E[d(\theta)].$$

The expected equilibrium payoff to agent 2 is

$$E\left[u^2(\theta,\theta)\right]=E\left[I^2(\theta)\right]+\beta(\alpha+\beta)E[\theta d(\theta)]$$

Since by (9) $E\left[u^2\left(\phi^2(\theta),\theta\right)\right]\leq E\left[u^2(\theta,\theta)\right]+c$, it follows that a necessary condition for (9) is (using the fact that $\beta<0$)

$$E[\theta d(\theta)]\leq E[\theta]E[d(\theta)]-\frac{c}{\beta(\alpha+\beta)}. \tag{11}$$

Total welfare is $W(h) = Y - \frac{(\alpha+\beta)^2}{2} E\left[d(\theta)^2\right] + (\alpha+\beta)^2 E[\theta d(\theta)]$ and must be greater than $W^* - c = Y + \frac{(\alpha+\beta)^2}{2} E\left[\theta^2\right] - c$. Therefore, if 1-authority *is not optimal,* we must have

$$E[\theta d(\theta)] \geq \frac{E\left[\theta^2\right] + E\left[d(\theta)^2\right]}{2} - \frac{c}{(\alpha+\beta)^2}. \tag{12}$$

Combining (11) and (12), we obtain the condition

$$\begin{aligned} \frac{E\left[\theta^2\right] + E\left[d(\theta)^2\right]}{2} - E[\theta] E[d(\theta)] &\leq c\frac{-\alpha}{\beta(\alpha+\beta)^2} \\ &\Longleftrightarrow \\ \frac{var[\theta] + var[d(\theta)]}{2} + \frac{1}{2}(E[\theta] - E[d(\theta)])^2 &\leq c\frac{-\alpha}{\beta(\alpha+\beta)^2}. \end{aligned} \tag{13}$$

Note that the right hand side is positive since $\beta < 0$ and is greater than $\frac{c}{(\alpha+\beta)^2}$ since $-\beta < \alpha$. The left hand side is minimized when the variance of $d(\theta)$ is equal to zero and when the expectation of θ and of $d[\theta]$ are equal. This happens for the (best) constant decision that we described in the previous section. Whenever the mechanism in (10) *is not* the constant decision, the left hand side is greater than $\frac{var[\theta]}{2}$. Therefore (13) is violated whenever[11]

$$\frac{var[\theta]}{2} > \frac{c}{(\alpha+\beta)^2}\frac{-\alpha}{\beta}.$$

Proposition 11 *1-authority is the optimal interference structure when* $\frac{var[\theta]}{2} > \frac{c}{(\alpha+\beta)^2}\frac{-\alpha}{\beta}$.

The condition in Proposition 11 is intuitive. The first best decision is increasing with the state and the variance of the first best decision is proportional to the variance of the state (is equal to $(\alpha+\beta)^2 var[\theta]$). For *any mechanism*, if agent 2 was able to select the decision (in the range of the decisions available in the mechanism), he would do so in such a way that the decision is a *decreasing* function of the state. The larger the variance in the decision rule of the mechanism, the greater is the benefit for agent 2 to select his preferred decisions. The variance of the first best decision depends both on the variance in the state *and* on the difference in the preference parameters of the two agents. As long as c, the cost of interfering is not too large, a large initial variance of the state and a large initial first best weight will create incentives for agent 2 to interfere.

[11]The bound on the right hand side achieves its minimum value of $\frac{27c}{4\alpha^2}$ for $\beta = -\frac{1}{3}\alpha$. The condition can therefore be satisfied for some values of β when $var[\theta] > \frac{27c}{2\alpha^2}$.

5.2 n Agents

The basic logic of the case $n = 2$ extends to the case $n \geq 3$. Consider the situation in which the $\{\alpha^i\}$ are decreasing in i, $\sum \alpha^i > 0$ and there exists k such that $\alpha^k > 0 > \alpha^{k+1}$. In this case, agents $\{1, \ldots, k\}$ have consonant preferences with the first best and agents $\{k+1, \ldots, n\}$ have dissonant preferences with the first best. Our principle extends readily to the n agent case. It is now necessary to define for each channel i a vector $\boldsymbol{\mu}_i = \left(\mu_i^j, j = 1, \ldots, n\right)$ in the $n-1$ simplex where μ_i^j denotes the probability with which agent j succeeds on channel i).

To simplify, continue to assume that agents can either secure their channel or not. If channel i is secured, i.e., $e^i = 1$, then $\mu_i^i = 1$ (only player i can succeeds on channel i). If channel i is not secured, i.e., $e^i = 0$, then $\mu_i^j = \frac{1}{n-1}$ (each player has an equal chance of succeeding on channel i). Suppose that $\mathbf{e} = \mathbf{0}$, i.e., that $\mu_i^j = \frac{1}{n-1}$ for $j \neq i$. As before, as long as the security structure does not change, any mechanism is incentive compatible in our sense. Suppose that agent j exerts high effort $e^1 = 1$ (at cost $c > 0$) and deviates from truth-telling. Since other agents are telling the truth, as long as $h = d^{\mathbf{0}}$ when announcements about the state disagree, agent j will be able to select a decision $q(\phi) \neq q(\theta)$ with probability $p = \left(\frac{1}{n-1}\right)^{n-1}$ and with probability $1 - p$ will select decision $d^{\mathbf{0}}$. Clearly for n large, the probability of selecting $d^{\mathbf{0}}$ is close to one and no agent can gain by exerting high effort.[12]

For small values of n, an agent might still benefit from exerting high effort and deviating from truth telling. Replicating the reasoning that we made for the case $n = 2$, we can show that as long as $var[\theta]$ is large enough, the best mechanism with $\mathbf{e} = \mathbf{0}$ is dominated by authority. In the $n = 2$ case, authority should be given to the agent who has the same direction of preferences as the first best, i.e., to agent 1 since $\alpha^1 > 0$. However, since agent 1 must be incentive compatible in the classical sense, it is necessary that the sharing rule for agent 2 is a non trivial function of the state. In the case $n \geq 3$, if there exists j such that $\alpha^j = \sum \alpha^i$, then we can give authority to j, pay all other agents a fixed wage and make agent j the full residual claimant for the revenue. This is suggestive of the fact that authority should be given to agents whose preferences coincide with the social preferences. This also suggests that we should observe authority figures who are somehow "consensual" rather than extremists.

[12] We might question however the relevance of the assumption that a large number of agents have common information while outsiders do not.

6 Discussion

Our examples illustrate that the consideration of interference can lead to new insights about contracting and organization. The noninterference principle makes the program tractable. Already it leads to a new interpretation and account of authority in organizations: in environments where communication cannot be secured at no cost, and when the stakes are sufficiently high, it may be optimal to give to one person authority in its dual sense: the power to decide and the power of influence.

There are a number of obvious extensions to the model that we have not yet considered, including continuous effort choices. More interesting perhaps is a comparison of voting with authority: even if a median voter *would* make the same decision as the representative player with authority, the differences in the costs of securing channels may make authority desirable. On the other hand if the distribution is skewed so that the representative and the median are very different, authority will not perform well and majority rule may dominate.

We outline some issues on the agenda.

One observation is that a person who is already influential may end up getting decision power because it is very cheap to secure his channel. For instance, reputation could be an important source of channel security; thus a bankrupt Donald Trump can get control over new real estate development projects. Similarly, if we relax the assumption that there is no binding limit to ex-ante transfers, considering the surplus maximizing outcome is no longer justified. With wealth effects, for instance, we might expect the wealthiest to get authority, even if this isn't optimal, simply because they are able to afford secure channels (Legros-Newman [11], [12]). Someone with a great reputation may get authority (decision power) for the same reason.

Authority as we have modeled it can help to explain why corporations are frequently personified by one individual. It can also help us understand why the one agent paradigm might be reasonable, especially when the standard mechanism desing approach would predict a strong discontinuity between the one and the two agent models. For instance, as should be well known, if two agents have correlated information in a firm, they will not be subject to borrowing constraint while a one agent firm would be. Once one allows for interference, contracts that would relax the borrowing constraint in the two agent case may no longer be feasible.

7 Appendix

7.1 Proof of Proposition

Step 1: Consider an equilibrium $(\mathbf{e}, \sigma)$ of the extended game corresponding to (M, g). Consider sets $\hat{M}^i$ that are disjoint and that contain *two different "copies"* of Θ. For instance, let $\hat{M}_1 = \hat{M}_2 = \Theta \times \{1, 2\}$. Below we fix $\mathbf{e}$ at his equilibrium value and let $\sigma^i(\theta)$ stands for $\sigma^i(\theta, \mu_i)$, $.\pi_{ij}$ stands for $\pi_{ij}(\mathbf{e})$, etc. We want to show that there exists a function $h : \hat{M}_1 \times \hat{M}_2 \to D$ such that in the extended game corresponding to the direct mechanism $\left(\hat{M}, h\right)$, agents find it optimal to "tell the truth" on *each channel*: hence agent i finds it optimal to send $((\theta, i), (\theta, i))$ in state θ when the other agent is telling the truth. Moreover, we want to argue that with such mechanism, each agent also finds it optimal to play his equilibrium $\mathbf{e}^i$ when the other agent is expected to do the same.

We define h as follows.

$$h((\theta, i), (\phi, j)) = \begin{cases} g\left(\sigma_1^i(\theta), \sigma_2^j(\theta)\right) & \text{if } \theta = \phi \\ d^0 & \text{otherwise.} \end{cases}$$

We claim that the strategy γ where $\gamma^i(\theta) = ((\theta, i), (\theta, i))$ is an equilibrium of the extended game corresponding to $\left(\hat{M}, h\right)$. Note that $h\left(\gamma_1^i(\theta), \gamma_2^j(\theta)\right) = g\left(\sigma_1^i(\theta), \sigma_2^j(\theta)\right)$ and therefore that $h^\mu \circ \gamma = g^\mu \circ \sigma$. To show the equilibrium property of γ assume that in state θ agent 1 decides to use a strategy $\left(\left(\theta_1^1, i\right), \left(\theta_2^1, j\right)\right)$. Since agent 2 uses the strategy $\gamma^2(\theta) = ((\theta, 2), (\theta, 2))$, the outcome is the lottery

$$\begin{aligned} &h\left(\left(\theta_1^1, i\right), \left(\theta_2^1, j\right)\right) \text{ with probability } \pi_{11} \\ &h\left(\left(\theta_1^1, i\right), (\theta, 2)\right) \text{ with probability } \pi_{12} \\ &h\left((\theta, 2), \left(\theta_2^1, j\right)\right) \text{ with probability } \pi_{21} \\ &h((\theta, 2), (\theta, 2)) \text{ with probability } \pi_{22} \end{aligned}$$

Note that the outcome $h((\theta, 2), (\theta, 2))$ arises with probability π_{22} independently of the strategy of agent 1. Hence, only the first three outcomes depend on the strategy of agent 1. If $\theta_1 \neq \theta_2 \neq \theta$ the agent obtains a lottery in which he gets d^0 with probability $1 - \pi_{22}$ and $h((\theta, 2), (\theta, 2))$ with probability π_{22}; hence agent 1 cannot be strictly better off than when he plays $\gamma^1(\theta)$. Consider now the case $\theta_1 = \theta_2 = \phi$. Using the definition of h, the above lottery is

$$\begin{array}{l} g\left(\sigma_1^i(\phi), \sigma_2^j(\phi)\right) \text{ with probability } \pi_{11} \\ g\left(\sigma_1^i(\phi), \sigma_2^2(\theta)\right) \text{ with probability } \pi_{12} \\ g\left(\sigma_1^2(\theta), \sigma_2^j(\phi)\right) \text{ with probability } \pi_{21} \\ g\left(\sigma_1^2(\theta), \sigma_2^2(\theta)\right) \text{ with probability } \pi_{22} \end{array}$$

If agent 1 prefers strictly this lottery to the lottery when there is truth-telling, this implies that he prefers to play the strategy $\left(\sigma_1^i(\phi), \sigma_2^j(\phi)\right)$ in state θ in the initial game. But this contradicts the equilibrium property of σ in the initial extended game.

Suppose now that agent 1 want to choose $\hat{\mathbf{e}}^1 \neq \mathbf{e}^1$ at the first stage. Let $\hat{\boldsymbol{\mu}} = \boldsymbol{\mu}(\hat{\mathbf{e}}^1, \mathbf{e}^2)$ be the new vector describing the security of the channels. When the agents play, agent 2 does not obtain information about the effort level of 1 and continues to use his truthful strategy. Suppose that agent 1 uses a strategy $\hat{\gamma}^1(\theta)$ and has a higher ex-ante utility:

$$\mathcal{E}u^1\left(h^{\hat{\mu}}\left(\hat{\gamma}^1(\theta), \gamma^2(\theta)\right), \theta\right) - c^1\left(\hat{\mathbf{e}}^1\right) > \mathcal{E}u^1\left(h^{\mu}\left(\gamma^1(\theta), \gamma^2(\theta)\right), \theta\right) - c^1\left(\mathbf{e}^1\right).$$

We note that since $h_{ij}(\phi, \theta) = d^{\mathbf{0}}$, for any θ,

$$\max_{\hat{m}^1 \in \hat{M}_1 \times \hat{M}_2} u^1\left(h^{\hat{\mu}}\left(\hat{m}^1, \gamma^2(\theta)\right), \theta\right) \leq \max_{m^1 \in M_1 \times M_2} u^1\left(g^{\hat{\mu}}\left(m^1, \sigma^2(\theta)\right), \theta\right).$$

Indeed, if the maximum on the right hand side is attained at m^1 such that $m^1 = \left(\sigma_1^i(\theta_1), \sigma_2^j(\theta_2)\right)$ for some ij and some θ_1, θ_2, then agent 1 by using $\hat{m}^1 = ((\theta_1, i), (\theta_2, j))$ will either obtain the same payoff or will obtain a strictly lower payoff if θ_1 or θ_2 is not equal to θ. It follows that $\max_{\hat{\sigma}^1} \mathcal{E}u^1\left(g^{\hat{\mu}}\left(\hat{\sigma}^1(\theta), \sigma^2(\theta)\right), \theta\right) - c^1\left(\hat{\mathbf{e}}^1\right) > \mathcal{E}u^1\left(g^{\mu}\left(\sigma^1(\theta), \sigma^2(\theta)\right), \theta\right) - c^1\left(\mathbf{e}^1\right)$, which contradicts the assumption that $(\mathbf{e}, \boldsymbol{\sigma})$ is an equilibrium of the extended game (M, g).

Note that we can dispense of the assumption of a worst outcome quite easily by defining

$$h((\theta, i), (\phi, j)) = g\left(\sigma_1^i(\theta), \sigma_2^j(\phi)\right)$$

and going through the same steps as before. Finally, defining $h_{ij}(\theta, \phi) = h((\theta, i), (\phi, j))$ establishes the proposition.

7.2 Proof of Monotonicity

We first show that $q\left(\phi^i(\theta)\right)$ and $u^i\left(\phi^i(\theta), \theta\right)$ are increasing in θ when $i = 1$ and are decreasing in θ when $i = 2$. Consider two states θ and $\hat{\theta}$. By definition of ϕ^i,

$$\begin{aligned}
I^i\left(q\left(\phi^i\left(\theta\right)\right)\right)+\alpha^i\theta q\left(\phi^i\left(\theta\right)\right) &\geq I^i\left(q\left(\phi^i\left(\hat{\theta}\right)\right)\right)+\alpha^i\theta q\left(\phi^i\left(\hat{\theta}\right)\right)\\
I^i\left(q\left(\phi^i\left(\theta\right)\right)\right)+\alpha^i\hat{\theta} q\left(\phi^i\left(\hat{\theta}\right)\right) &\geq I^i\left(q\left(\phi^i\left(\theta\right)\right)\right)+\alpha^i\hat{\theta} q\left(\phi^i\left(\theta\right)\right)
\end{aligned}$$

which implies that

$$\alpha^i\left(\theta-\hat{\theta}\right)\left(q\left(\phi^i\left(\theta\right)\right)-q\left(\phi^i\left(\hat{\theta}\right)\right)\right)\geq 0.$$

When $\theta > \hat{\theta}$ implies that $q\left(\phi^i\left(\theta\right)\right)\geq q\left(\phi^i\left(\hat{\theta}\right)\right)$ when $\alpha^i\geq 0$ and implies $q\left(\phi^i\left(\theta\right)\right)\geq q\left(\phi^i\left(\hat{\theta}\right)\right)$ when $\alpha^i\leq 0$. Note that the above conditions can be written

$$\alpha^i\left(\theta-\hat{\theta}\right)q\left(\phi^i\left(\theta\right)\right)\geq u^i\left(\phi^i\left(\theta\right),\theta\right)-u^i\left(\phi^i\left(\hat{\theta}\right),\hat{\theta}\right)\geq\alpha^i\left(\theta-\hat{\theta}\right)q\left(\phi^i\left(\hat{\theta}\right)\right).$$

As $q\geq 0$, the result follows.

8 References

References

[1] Abreu, Dilip and Matsushima, Hitoshi (1992), "Virtual Implementation in Iteratively Undominated Strategies: Complete Information", *Econometrica*, 60, 993-1008.

[2] Aghion, Philippe and Tirole, Jean (1997), "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105, 1-29.

[3] Bolton, Patrick and Dewatripont, Mathias (1994), "The Firm as a Communication Network," *Quarterly Journal of Economics*, 109, 809-839.

[4] Fudenberg, Drew and Jean Tirole (1986), "A "Signal-Jamming" Theory of Predation," *Rand Journal of Economics*, 17, 366-376.

[5] Green, Jerry R. and Laffont, Jean-Jacques (1986), "Incentive Theory with Data Compression", in *Essays in honor of Kenneth J. Arrow, volume 3: Uncertainty, information, and communication.*, W.P. Heller, R.M. Starr, D.A. Starrett editors, (Cambridge; New York and Sydney: Cambridge University Press), pages 239-53.

[6] Grossman, Sanford and Hart, Oliver (1986), "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration", *Journal of Political Economy*, 94, 691-714.

[7] Hart, Oliver (1995) *Firms, Contracts, and Financial Structure* (Oxford: Oxford University Press).

[8] Hart, Oliver and Moore, John (1990), "Property Rights and the Nature of the Firm", *Journal of Political Economy*, 98, 1119-1158.

[9] Hart, Oliver and Moore, John (1999), "Foundations of Incomplete Contracts", *Review of Economic Studies,* 66, 115–138.

[10] "Morgan Stanley Sued for $1.35B," *Las Vegas Sun*, May 20, 1999 (http://www.lasvegassun.com/stories).

[11] Legros, Patrick and Newman, Andrew (1996), "Wealth Effects, Distribution, and the Theory of Organization," *Journal of Economic Theory* 70:2 312-341.

[12] Legros, Patrick and Newman, Andrew (1996), "Competing for Ownership," Working Paper, Columbia University.

[13] Maskin, Eric and Tirole, Jean (1999), "Unforseen Contingencies and Incomplete Contracts", *Review of Economic Studies,* 66, 83-114.

[14] Melumad, Nahum, Mookerjee, Dilip and Reichelstein, Stefan (1992) "A Theory of Responsibility Centers," *Journal of Accounting and Economics*, 15, 445-484.

[15] Milgrom, Paul and Roberts, John (1986), "Bargaining Costs, Influence Costs, and the Organization of Economic Activity," in *Perspectives on Positive Political Economy*, J. Alt and K. Shepsle, eds. (Cambridge: Cambridge University Press).

[16] Moore, John and Repullo, Rafael (1990), "Nash Implementation: A Full Characterization", *Econometrica*, 58(5), 1083-1099.

[17] Radner, Roy (1993), "The Organization of Decentralized information Processing," *Econometrica, 62, 1109-1146.*

[18] Radner, Roy and van Zandt, Timothy (1992), "Information Processing and Returns to Scale," *Annales d'Economie et de Statistique*, 25/26, 265-298.

[19] Segal, Ilya (1996), "Communication, Complexity and Coordination by Authority, " mimeo.

[20] Segal, Ilya (1999), "Complexity and Renegotiation: a Foundation for Incomplete Contracts", *Review of Economic Studies, 66, 57-82.*

[21] Williamson, Oliver E. (1975). *Markets and Hierarchies, Analysis and Antitrust Implications: a Study in the Economics of Internal Organizations.* New York: Free Press.