PUBLISHED VERSION

Mechanism design and communication networks

Ludovic Renou
Department of Economics, University of Essex

Tristan Tomala
Department of Economics and Decision Sciences, HEC Paris

This paper studies a mechanism design model where the players and the designer are nodes in a communication network. We characterize the communication networks (directed graphs) for which, in any environment (utilities and beliefs), every incentive compatible social choice function is partially implementable. We show that any incentive compatible social choice function is implementable on a given communication network, in all environments with *either* common independent beliefs and private values *or* a worst outcome, if and only if the network is strongly connected and *weakly 2-connected*. A network is strongly connected if for each player, there exists a directed path to the designer. It is weakly 2-connected if each player is either directly connected to the designer or indirectly connected to the designer through two disjoint paths, *not necessarily directed*. We couple encryption techniques together with appropriate incentives to secure the transmission of each player's private information to the designer.

Keywords. Mechanism design, incentives, Bayesian equilibrium, communication networks, encryption, secure transmission.

JEL classification. C72, D82.

1. Introduction

The *revelation principle* is the cornerstone of mechanism design and its applications. It asserts that the outcome of any communication system can be replicated by a direct revelation mechanism, in which agents directly and privately communicate with a designer, and truthfully report all their information (Gibbard 1973, Dasgupta et al. 1979, Myerson 1979, 1982, Harris and Townsend 1981). As a technical result, the revelation principle is a blessing. It allows one to abstract away from the very details of communication systems and to focus on the social choice functions to be implemented. At

FIGURE 1. Two communication networks.

the same time, it is slightly disturbing, as it implies that no decentralized communication system, however sophisticated, can dominate the centralized (direct) communication system. Yet, real-world organizations (firms, administrations, armies, terrorist networks, organized crime) seldom take the form of centralized communication systems. The aim of this paper is to characterize the communication systems that replicate the incentive properties of centralized communication and, thus, to show that incentive considerations *alone* can already explain the existence of a large variety of real-world organizations.[1]

Communication systems are naturally modeled as networks (graphs), in which the nodes represent the players and the designer. A player can directly communicate with another player if an edge exists from that player to the other. We then associate communication networks with social environments that represent the preferences and beliefs of the players, and we characterize the topology of communication networks for which, in any environment, *every* incentive compatible social choice function is partially implementable. We first focus on acyclic directed networks and then show how our results extend to any network.

The connectivity of communication networks is at the center of our analysis. A directed network is *strongly 1-connected* if for each player, there exists a directed path from this player to the designer. This is a minimal requirement that ensures that the designer can receive information from each player. A directed network is *weakly 2-connected* if each player is either directly connected to the designer or has two disjoint paths to the designer in the associated *undirected* graph. Figure 1 gives two examples of weakly 2-connected networks. Our analysis shows that in a large class of environments, both networks have the very same incentive properties.[2]

Our main results state that any incentive compatible social choice function is partially implementable on a given communication network, in all environments with *either* common independent beliefs and private values *or* a worst outcome, if and only if

---

[1]There is recent literature labeled *algorithmic mechanism design* that focuses on communication complexity and mechanism design (see Nisan et al. 2007 for an excellent exposition and Nisan and Segal 2006 and Van Zandt 2007 for economic applications). Unlike this literature, we abstract from complexity considerations and entirely focus on incentives.

[2]Other features are therefore needed to discriminate among these networks, e.g., their *span of control* (Williamson 1967 and Calvo and Wellisz 1978) or their associated cost of communication (Bolton and Dewatripont 1994 or Radner 1993).

the network is weakly 2-connected and strongly 1-connected. (In the sequel, we omit the condition of strong 1-connectedness.) The intuition for this result is as follows.[3] A social choice function is incentive compatible if no player has an incentive to lie about his own private information when he expects the others to tell the truth. Importantly, players use their prior beliefs to form their expectations. However, in a general communication network, players receive messages from their neighbors and thus their incentives to tell the truth may be altered (since their posterior beliefs may differ from their prior beliefs). To circumvent this problem, we couple encryption techniques and incentives to transfer "securely" each player's private information to the designer through the network. Our encoding technique guarantees that no player learns anything about the types of the other players and, therefore, posterior beliefs are equal to prior beliefs. To illustrate, assume that the network is *strongly* 2-connected, that is, each player is either directly connected to the designer or has two disjoint *directed* paths of communication to the designer. A player can thus send a private "encoding" key to the designer through one path and his type encoded with the key, a "cypher type," through the other (disjoint) path. However, this is not sufficient: players must also have an incentive to truthfully forward the messages they receive. Our technique precisely guarantees this. Last, incentive compatibility ensures that players also have an incentive to truthfully report their own private information. Our connectivity conditions are necessary. If the network is not strongly 1-connected, then there exists at least one player who has no outgoing edges, i.e., this player cannot send information. It is thus impossible to implement a social choice function that depends on this player's type. Alternatively, if the network is not weakly 2-connected, then a pair of players (i, i^*) exists such that all paths from player i to the designer go through player i^*, who has thus the ability to manipulate all the information transmitted by i.

We now offer some motivations for our study. First, as in Bolton and Dewatripont (1994), we implicitly assume that the communication network (the internal organization of the firm) is established in a prior stage and that it is relatively costly to modify. Consequently, if the designer is uncertain about which incentive compatible social choice functions he will actually have to implement, it is optimal to choose a network in the class of weakly 2-connected networks. Alternatively, we can think of our study as a worst-case analysis: If the communication network is not weakly 2-connected, incentive compatible social choice functions exist that cannot be implemented on that network. Second, the previous discussion suggests that the cost of forming a link between any two agents is an important determinant in choosing among different networks (organizations). How costly is it to form such a link? To answer this question, we need to carefully interpret what a link is in our model. A link between two agents is a perfectly secure channel of communication, i.e., no other agent can eavesdrop on, alter, or intercept messages sent over the link, and any message sent is received with certainty. Private face-to-face communication is probably the closest instance of such perfectly secure communication in real life.[4] Such links are relatively costly to establish as argued

---

[3]See the example in Section 2 for an illustration.

[4]E-mails, phone calls, or text messages are not examples of perfectly secure and reliable channels of communication as the recent *News of the World* scandal demonstrates (*Guardian*, 14 July 2009). In fact, if they were, there would be no need for encryption devices.

by computer scientists; see e.g., Beimel and Franklin (1999). Furthermore, Friebel and Raith (2004) argue that even if it were possible to create, at no cost, such perfectly secure communication links between each agent and the designer in an organization, it may not be optimal to do so. In their words, "requiring intra-firm communication to pass through a 'chain of command' can be an effective way of securing the incentives for superiors to recruit and develop the best possible subordinates."

Related literature

The computer science literature on secure transmission of messages is closely related to this paper. Section 4.3 provides an in-depth discussion of this literature and its relationships to our study. The use of coded messages in games of information transmission is common in the cheap talk literature (see Forges 1990, Bárány 1992, Ben-Porath 2003, and Gerardi 2004) and our techniques are akin to those found there. The paper most closely related to our work is Monderer and Tennenholtz (1999), who study a similar problem to ours. Our paper substantially generalizes their results in several dimensions. First, these authors consider *undirected* networks and environments with a worst outcome, common independent beliefs, and private values. They show that 2-connectedness of the network is a sufficient condition for the implementation of all incentive compatible social choice functions. Crucially, in their model, edges are not directed and thus can be used to communicate in both directions. It follows that the 2-connectedness of the undirected network guarantees the existence of directed subnetworks that are strongly 2-connected. Their protocol (mechanisms and strategies) heavily exploits this fact and indeed breaks down if the undirected network does not have an underlying strongly 2-connected network. We show that in environments with common independent beliefs and private values, weak 2-connectedness—a substantially weaker requirement than strong 2-connectedness—is a necessary and sufficient condition (the assumption of a worst outcome is superfluous). Second, we show that in environments with a worst outcome, weak 2-connectedness is again a necessary and sufficient condition; no further assumption on the environment is needed. In particular, there is no need for independent beliefs or private values. We need to resort to different encryption techniques than those used in Monderer and Tennenholtz (1999), which would fail without common independent beliefs even on strongly 2-connected networks. Furthermore, with the very same techniques, we show that strong 2-connectedness and weak 3-connectedness is a sufficient condition for the implementation of all incentive compatible social choice functions in all environments. Again, the techniques of Monderer and Tennenholtz (1999) would fail here.

2. A simple example

We now illustrate our main results within the context of a simple example. There are three players, labeled 1, 2, and 3, two types for player 2, labeled θ and θ', and two alternatives a and b. Player 2's preferences over these alternatives depend on his type (in all examples, preferences are strict). Player 2 prefers a to b if his type is θ and prefers b to a if

FIGURE 2. Communication network $\mathcal{N}_2$.

his type is θ'. Player 1 always prefers a to b, while player 3 always prefers b to a. Note that this is a private values environment: the preferences of players 1 and 3 do not depend on player 2's type. The designer aims to implement the social choice function f^* that selects the preferred alternative of player 2 for each of his type: player 2 is dictatorial.

If player 2 can securely and directly communicate with the designer, f^* is clearly implementable: the designer can simply ask player 2 to directly report his preferred alternative. Suppose now that player 2 cannot directly communicate with the designer and consider the communication network $\mathcal{N}_2$ in Figure 2 (player 0 is the designer).

With communication network $\mathcal{N}_2$, player 2 can indirectly communicate with the designer through player 1. Moreover, player 3 has two disjoint paths of communication to the designer with player 2 on one of them. Consequently, player 2 has two disjoint paths to the designer, but one of them is not directed. The network $\mathcal{N}_2$ is thus *weakly 2-connected*. The idea is then to use the two disjoint paths from 3 to 0 to secure the communication of player 2's type to the designer, without revealing information to the other players. So, suppose that players 1 and 3 believe that player 2's type is θ with probability $\frac{1}{3}$, independently of their own types. The goal is to design a mechanism and an equilibrium such that the designer implements a in state θ and b in state θ'.

The mechanism allows player 3 to send a real number in $[0, 1)$ to player 2 and another real number in $[0, 1)$ to player 0. Similarly, player 2 (resp., player 1) can send a real number in $[0, 1)$ to player 1 (resp., player 0). An informal description of the strategies is as follows. Independently of his type, player 3 draws an "encoding key" y uniformly on $[0, 1)$ and sends it to both players 0 and 2. Player 2 of type θ (resp., θ') draws a "pseudo-type" $\tilde{x}$ uniformly on $[0, \frac{1}{3})$ (resp., $[\frac{1}{3}, 1)$). The pseudo-type thus "reveals" θ, but its unconditional distribution is uniform on $[0, 1)$.[5] Then player 2 encodes his pseudo-type $\tilde{x}$ with the encoding key y received from player 3 to obtain the "cypher type" $x = (\tilde{x} + y) \bmod_{0,1}$.[6] Player 2 sends x to player 1. Player 1 has to correctly forward the message of player 2 to the designer. Let $(\hat{x}, \hat{y})$ be a pair of messages received by the designer. The allocation rule is the following: If $(\hat{x} - \hat{y}) \bmod_{0,1} \in [0, \frac{1}{3})$, the designer implements a and implements b, otherwise.

---

[5]More precisely, $\mathcal{U}_{[0,1/3)}$ (resp., $\mathcal{U}_{[1/3,1)}$) denotes the uniform distribution on $[0, \frac{1}{3})$ (resp., $[\frac{1}{3}, 1)$). The unconditional distribution of $\tilde{x}$ is $\frac{1}{3}\mathcal{U}_{[0,1/3)} + \frac{2}{3}\mathcal{U}_{[1/3,1)} = \mathcal{U}_{[0,1)}$, the uniform distribution on $[0, 1)$.

[6]For a real number r, $r \bmod_{0,1} = r - \lfloor r \rfloor$, with $\lfloor r \rfloor$ the highest integer less or equal to r.

If the players follow the prescribed strategies, $\hat{y} = y$, $\hat{x} = x$, and $(\hat{x} - \hat{y}) \bmod_{0,1} = \tilde{x}$. Thus, the designer correctly learns player 2's type and implements the desired social choice function f^*. In particular, players 1 and 3 expect the designer to implement a with probability $\frac{1}{3}$ and b with probability $\frac{2}{3}$. We now show that the players do not have an incentive to deviate from the prescribed strategies. Suppose that player 1 deviates and sends a message $\hat{x}$ to the designer instead of x. The designer implements the alternative a if $(\hat{x} - y) \bmod_{0,1} \in [0, \frac{1}{3})$ and b, otherwise. Since y is uniformly distributed, so is $(\hat{x} - y) \bmod_{0,1}$ (see Lemma 2 in the Appendix). Accordingly, player 1 expects the designer to implement a with probability $\frac{1}{3}$ and b with probability $\frac{2}{3}$: Player 1's expected payoff does not depend on the message $\hat{x}$ he sends. Player 1 has, therefore, no incentive to deviate. A similar argument applies to player 3. As for player 2, he has no incentive to deviate since f^* is incentive compatible.

It is worth stressing that the essential feature of the network is its weak 2-connectedness. For instance, if in addition to the links shown in Figure 2, player 3 has a link to player 1, the result remains valid (the network remains weakly 2-connected). Indeed, we can construct a "babbling equilibrium" in which player 3 sends an uninformative message to player 1, and player 1 plays independently of player 3's message. Alternatively, and more simply, we may let the message space from player 3 to player 1 be a singleton. In effect, we show that the weak 2-connectedness of the network is a necessary and sufficient condition for the implementation of any incentive compatible social choice functions in environments with independent common beliefs and private values.

A further and important feature of the proposed mechanism and strategies is that players 1 and 3 learn nothing about player 2's type. This is clearly true for player 3, as he does not receive a message from player 2. As for player 1, we prove that the message x (the cypher type) he receives is uniformly distributed on $[0, 1)$ and independent of player 2's type. This feature is crucial for the implementation of incentive compatible social choice functions that depend on the private information of all players. It guarantees that posterior beliefs are equal to prior beliefs and, consequently, that players' incentives to truthfully reveal their own private information are not altered.

Another important aspect is that the mechanism and strategies are tailored to environments with common independent beliefs and private values. First, let us consider the assumption of common independent beliefs. For concreteness, suppose that player 3's belief remains as above, but that player 1 believes that player 2's type is θ with probability $\frac{2}{3}$. Players 1 and 3 have thus *different* beliefs. In the construction above, the partition of $[0, 1)$ into $\{[0, \frac{1}{3}), [\frac{1}{3}, 1)\}$ is such that the Lebesgue measure of each subset exactly matches the prior beliefs of player 3, but differs now from player 1's prior beliefs. Consider a deviation for player 1, whereby he sends the same message, regardless of the message received from player 2. With this deviation, player 1 expects the designer to decode player 2's type as being θ with probability $\frac{1}{3}$, which is different from his prior belief $\frac{2}{3}$. Consequently, player 1's incentive to truthfully report his private information might be altered and this player may profitably deviate.[7] Note that different (interim)

---

[7]For instance, take $\Theta_1 = \Theta_2 = \{\theta, \theta'\}$, three alternatives a, b, c, and $u_1(a, \theta) = \frac{3}{2}$, $u_1(b, \theta) = 1$, and $u_1(c, \theta) = 0$. Consider the social choice function f, which depends only on players 1 and 2's types with

beliefs of players 1 and 3 may derive from a common *correlated* prior on type profiles. Thus, the importance of the common independent belief assumption is that it allows the mechanism to be tuned simultaneously to the beliefs of all players.

Second, to understand the importance of the private value assumption, suppose that player 1 prefers b to a when player 2's type is θ and prefers a to b when player 2's type is θ' (interdependent values). If player 1 truthfully forwards the message x he received from player 2, the alternative a is implemented if and only if player 2's type is θ, and the alternative b is implemented if and only if player 2's type is θ'. However, if he sends a message $\hat{x}$ independently of the message received from player 2, both alternatives a and b are implemented with positive probability, regardless of player 2's type—a profitable deviation for player 1. In sum, the problem with more general environments is not only to guarantee that no information is revealed, but to provide players with incentives to truthfully communicate their private information and the messages they receive.

With more elaborate encryption techniques, our result remains valid in environments with a worst alternative (Theorem 2). The intuition is as follows. Consider again the network $\mathcal{N}_2$. Player 3 draws a large number of independent encoding keys $y_1, \ldots, y_\eta$ and sends them to players 0 and 2. Player 2 privately chooses one of these keys (with equiprobability) and uses it to encrypt his type. He then sends to player 1 the encrypted type and the unused keys, *without telling him which key was used for coding*. Player 1 has to correctly forward player 2's message to the designer. The designer compares the two vectors he receives. If these vectors differ in exactly one component η^*, he infers that the key y_{η^*} transmitted by player 3 was used for coding and he decodes player 2's type accordingly. Otherwise, the designer implements the worst alternative. This encoding technique guarantees that players 1 and 3 learn nothing about player 2's type and allows the designer to detect unilateral deviations with arbitrarily high probability, since the index η^* is the private information of player 2. In turn, the threat to implement the worst alternative upon detection of a deviation deters players from deviating.

3. Definitions

The primitives of the model consist of two essential ingredients: social environments (players, outcomes, and preferences) and communication networks.

A *social environment* $\mathcal{E}$ is a tuple $\langle N, A, (\Theta_i, P_i, u_i)_{i \in N} \rangle$, where $N := \{1, \ldots, n\}$ is the set of players, A is the finite set of alternatives, and Θ_i is the finite set of types of player $i \in N$.[8] Let $\Theta := \times_{i \in N} \Theta_i$ and $\Theta_{-i} := \times_{j \in N \setminus \{i\}} \Theta_j$, with generic elements θ and θ_{-i}, respectively. Each player knows his own type and player i of type θ_i holds a probabilistic belief $P_i(\cdot | \theta_i)$ over Θ_{-i}. Throughout the paper, we assume $P_i(\theta_{-i} | \theta_i) > 0$ for all $(\theta_i, \theta_{-i}) \in \Theta$ and for all $i \in N$. Each player has a preference relation over alternatives that is representable by the type-dependent utility function $u_i : A \times \Theta \to \mathbb{R}$. Players are expected utility maximizers. Three properties of an environment are of particular importance to our analysis.

$f(\theta, \theta) = a$, $f(\theta', \theta) = f(\theta, \theta') = c$, and $f(\theta', \theta') = b$. This is incentive compatible for player 1 at state θ when he believes that player 2's type is θ with probability $\frac{2}{3}$, but not when he believes that player 2's type is θ with probability $\frac{1}{3}$.

[8] In Section 5, we extend our analysis to environments with infinite type spaces.

- The environment has a *common prior* if there exists a probability distribution P on Θ such that $P_i(\theta_{-i}|\theta_i)$ is the conditional distribution of θ_{-i} given θ_i derived from P. The common prior is *independent* if P is the product of its marginal distributions.

- The environment has *private values* if for each player i, his utility function does not depend on the types θ_{-i} of his opponents.

- The environment has a *worst outcome* if there exists an alternative $\underline{a} \in A$ such that for each player i, each type profile θ, and each alternative $a \in A \setminus \{\underline{a}\}$, $u_i(\underline{a}, \theta) < u_i(a, \theta)$.

A social choice function $f : \Theta \to A$ associates with each type profile θ an alternative $f(\theta) \in A$. A social choice function is *incentive compatible* if for each player $i \in N$, for each pair of types (θ_i, θ_i') of player i, we have

$$\sum_{\theta_{-i}} u_i(f(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}) P_i(\theta_{-i}|\theta_i) \geq \sum_{\theta_{-i}} u_i(f(\theta_i', \theta_{-i}), \theta_i, \theta_{-i}) P_i(\theta_{-i}|\theta_i).$$

Note that our definition of a worst outcome is stronger than actually required; it would be enough to consider an alternative worse than any alternative in the range of the social choice function we aim to implement. Exchange economies with free disposal are examples of environments with worst outcome: the zero allocation is a worst outcome if preferences are strictly monotonic and the social choice function selects positive vectors of goods. Similarly, in quasilinear environments, the assumption of a worst outcome is natural.

A *communication network* captures the possibilities of communication between the players and the designer. A communication network is a *directed* graph with $n + 1$ vertices representing the n players and the designer (henceforth, player 0). There is a directed edge from player i to player j, denoted ij, if i can send a message to j. Formally, the network, denoted by $\mathcal{N}$, is defined as a set of edges $\mathcal{N} \subseteq (N \cup \{0\}) \times (N \cup \{0\})$. We let $C(i) = \{j \in N \cup \{0\} : ij \in \mathcal{N}\}$ denote the set of players to whom player i can directly send a message. Similarly, we let $D(i) = \{j \in N \cup \{0\} : ji \in \mathcal{N}\}$ denote the set of players who can directly send a message to player i. A *directed path* in $\mathcal{N}$ is a finite sequence of vertices $(i_1, \ldots, i_m)$ such that $i_k i_{k+1} \in \mathcal{N}$ for each $k = 1, \ldots, m - 1$. A communication network $\mathcal{N}$ is *strongly m-connected* if for each player $i \in N \setminus D(0)$, there exist m disjoint directed paths (i.e., having no common vertex except i and 0) from player i to the designer. By convention, the communication network is strongly n-connected if $N \setminus D(0) = \varnothing$. A network of particular importance is the star network $\mathcal{N}^\star$, where the designer is the center and $D(i) = \varnothing$, $C(i) = \{0\}$ for all player $i \in N$. With the star network, each player communicates directly and privately with the designer; the star network is n-connected.

We make the following assumptions on the network. First, we assume that networks are strongly 1-connected: for each player $i \in N$, there exists a directed path from i to 0. This assumption ensures that the designer can receive information from each player.

Second, we assume for the time being that the graph is *acyclic*, that is, for each $i \in N \cup \{0\}$, there is no path from i to himself. In particular, these two assumptions imply that $C(0) = \varnothing$, i.e., the designer cannot send messages to the players. In other words, as

FIGURE 3. Network $\mathcal{N}_3$ and a consistent extensive form $G_{\mathcal{N}_3}$.

in the classical model of mechanism design, the designer does not communicate with the players: he merely collects information and implements outcomes accordingly.

Now, we describe the interaction between a social environment and a communication network. The important feature of our model is that players can only send messages to players to whom they are directly connected. The interaction (the extensive form) unfolds as follows.

- Each player i "reads" the messages he receives from players in $D(i)$. Then he sends messages to players in $C(i)$ (he may send different messages to different players).

- The designer "reads" the messages he receives from players in $D(0)$ and selects an alternative.

Note that if $\mathcal{N} = \mathcal{N}^\star$, this corresponds to the classical model where each player communicates directly and privately with the designer.

Acyclicity and strong 1-connectedness of the graph imply that the interaction as described above gives rise to a simple extensive form. With acyclicity, the communication rule stating that *a player sends his messages after having received all his messages* generates a well defined timing structure, where each player i is assigned a stage $t(i)$ at which he sends his messages. This statement is proved in the Appendix, Lemma 1. For instance, in Figure 3, player 3 can directly communicate with player 1, but not with player 2 and the designer. In the associated extensive form, player 3 communicates first with player 1 and, after observing player 3's message, player 1 communicates with the designer.

The assumption of directed and acyclic networks makes our problem of implementation the hardest (the designer is silent, and players speak only once and receive no feedback on the messages they send). Yet, the methods we develop for acyclic directed networks *extend to any network*. More specifically, Section 5.1 drops the assumption of acyclicity and shows how to adapt our results to networks with cycles or to undirected networks, i.e., two-way networks where linked players can converse.

A *mechanism* is a pair $\langle (M_{ij})_{ij \in \mathcal{N}}, g \rangle$, where for each edge ij, M_{ij} is the set of messages that player i can send to player j, and $g : \times_{i \in D(0)} M_{i0} \to A$ is the allocation rule. Note that the allocation rule depends only on the messages the designer can receive. The next step is to define the Bayesian game induced by a mechanism, a communication network, and an environment.

Fix an environment $\langle N, A, (\Theta_i, P_i, u_i)_{i \in N} \rangle$, a communication network $\mathcal{N}$, and a mechanism $\langle (M_{ij})_{ij \in \mathcal{N}}, g \rangle$. Define $M_{D(i)} := \times_{j \in D(i)} M_{ji}$ as the set of messages that player i can receive and $M_{C(i)} := \times_{j \in C(i)} M_{ij}$ as the set of messages that player i can send. A *pure strategy* s_i for player i is a mapping from $M_{D(i)} \times \Theta_i$ to $M_{C(i)}$. We denote by S_i the set of player i's pure strategies and by $s_{ij}(m_{D(i)}, \theta_i)$ the message player i sends to player $j \in C(i)$ conditional on receiving the messages $m_{D(i)}$ and being of type θ_i. A *behavioral strategy* σ_i for player i maps $M_{D(i)} \times \Theta_i$ to $\Delta(M_{C(i)})$, the set of probability distributions over $M_{C(i)}$.[9] We denote by $\mathbb{P}_{\sigma, \theta}$ the probability distribution over profiles of messages (i.e., over $\times_{ij \in \mathcal{N}} M_{ij}$) induced by the strategy profile $\sigma = (\sigma_i)_{i \in N}$ at state θ. The Bayesian game $G_{\mathcal{N}}$ induced by an environment, a mechanism, and a network is defined as follows.

- The set of players is N; the set of player i's types is Θ_i and his beliefs are given by P_i.

- The set of strategies of player i is S_i.

- The payoff of player i is his expected utility conditional on his type and given that the outcomes are selected by the allocation rule g.

DEFINITION 1. The social choice function f is partially implementable on the communication network $\mathcal{N}$ if there exist a mechanism $\langle (M_{ij})_{ij \in \mathcal{N}}, g \rangle$ and a Bayesian–Nash equilibrium σ^* of $G_{\mathcal{N}}$ such that for all $\theta \in \Theta$, $g((m_{i0}^*)_{i \in D(0)}) = f(\theta)$ for all profiles of messages $(m_{i0}^*)_{i \in D(0)}$ received by the designer in the support of $\mathbb{P}_{\sigma^*, \theta}$.

Let $F_{\mathcal{N}}(\mathcal{E})$ denote the set of social choice functions partially implementable on the communication network $\mathcal{N}$ when the environment is $\mathcal{E}$. From the revelation principle, $F_{\mathcal{N}}(\mathcal{E}) \subseteq F_{\mathcal{N}^\star}(\mathcal{E})$ for every environment $\mathcal{E}$, and $F_{\mathcal{N}^\star}(\mathcal{E})$ is precisely the set of incentive compatible social choice functions. The aim of this paper is to characterize the communication networks $\mathcal{N}$ for which $F_{\mathcal{N}}(\mathcal{E}) = F_{\mathcal{N}^\star}(\mathcal{E})$ for every environment $\mathcal{E}$.

Before presenting our main results, a final remark is in order. We present our results for the solution concept of Bayesian equilibrium. Yet all our results remain valid with the solution concept of perfect Bayesian equilibrium. Indeed, as will be apparent below (see also the introductory example), the Bayesian equilibria we construct are such that *every* profile of messages a player can receive is in the support of the equilibrium strategies. Moreover, equilibrium strategies are such that we can apply a continuous version of Bayes' rule at every profile of messages a player can receive. We have chosen to present our results for the concept of Bayesian equilibrium, so as to avoid specifying the belief systems, namely the beliefs a player has about the types of his opponents and the messages they have received, at each of his information sets.

---

[9]We also find it convenient to view a behavioral strategy as a measurable mapping from $M_{D(i)} \times \Theta_i \times Y_i$ to $M_{C(i)}$, where $(Y_i, \mathcal{Y}_i, \mu_i)$ is a probability space independent of types and messages, i.e., a private randomization device.

FIGURE 4. Network $\mathcal{N}_4$ is weakly 2-connected; network $\mathcal{N}_4'$ is not.

4. The main results

This section presents our main results regarding the partial implementation of social choice functions on communication networks. We introduce our main connectivity condition. Recall that we consider strongly 1-connected and acyclic networks. An *undirected path* in $\mathcal{N}$ is a finite sequence of vertices $(i_1, \ldots, i_m)$ such that for each $k = 1, \ldots, m-1$, either $i_k i_{k+1} \in \mathcal{N}$ or $i_{k+1} i_k \in \mathcal{N}$.

DEFINITION 2. The communication network $\mathcal{N}$ is *weakly 2-connected* if for each player $i \in N \setminus D(0)$, there exist two disjoint undirected paths from player i to the designer.

In words, a network is weakly 2-connected if for each player not directly connected to the designer, there exist two disjoint paths—directed or undirected—from this player to the designer. For instance, in Figure 4, the network $\mathcal{N}_4$ is weakly 2-connected, while the network $\mathcal{N}_4'$ is not. Note that in both networks, player 2 has a unique directed path to the designer and, therefore, neither network is strongly 2-connected.

Importantly, if a network is not weakly 2-connected, there exist two players, i and i^*, such that all paths (directed or undirected) from player i to the designer go through player i^*. As a consequence, for each player $j \neq i$ who has a path (directed or undirected) to i, all paths (directed or undirected) from j to the designer go through player i^*. Player i^* thus "controls" all the possible messages that player i can use to communicate his private information. Player i^* even controls the messages of all players who are connected, directly or indirectly, to player i. For instance, on the network $\mathcal{N}_4'$, player 1 controls all messages that players 2 and 3 can send. These simple observations suggest that there is no hope to implement all incentive compatible social choice functions on a network that is not weakly 2-connected. We show that this is indeed the case.

4.1 *Common independent beliefs and private values*

We first consider environments with common independent beliefs and private values. This assumption is common in several applications of the theory of mechanism design, e.g., auction theory (Krishna 2002) or contract theory (Salanié 2005). Our first result states that any incentive compatible social choice function is implementable on a network $\mathcal{N}$ for all such environments if and only if $\mathcal{N}$ is weakly 2-connected.

FIGURE 5. Communication network $\mathcal{N}_5$.

THEOREM 1. *Consider an acyclic network $\mathcal{N}$. For all environments $\mathcal{E}$ with common independent beliefs and private values, $F_{\mathcal{N}}(\mathcal{E}) = F_{\mathcal{N}^\star}(\mathcal{E})$ if and only if $\mathcal{N}$ is weakly 2-connected.*

Theorem 1 extends the work of Monderer and Tennenholtz (1999) in several dimensions. Monderer and Tennenholtz consider environments and communication networks with the following properties: (1) types are independently and identically distributed, (2) a player's payoff does not depend on the private information of others (private values), (3) there exists a worst outcome (to abort the protocol), and (4) networks are undirected and repeated communication is allowed, so that each edge is directed in both ways and players may get feedback on the messages they sent. With these assumptions, they show that the 2-connectedness of the communication network is a sufficient condition for the implementation of any incentive compatible social choice function. First, we show that their result extends to weakly 2-connected directed networks and that this condition is necessary. This result requires the construction of a substantially more elaborate protocol (mechanisms and strategies) than that in Monderer and Tennenholtz (1999). Indeed, their construction relies on the existence of an underlying directed subgraph that is strongly 2-connected, so that a player can send his encrypted type on one directed path and send the encryption key on the other disjoint directed path. Unlike Monderer and Tennenholtz, our assumption of weakly 2-connected networks does not guarantee the existence of two disjoint *directed* paths from each player to the designer. Second, we show that the crucial assumptions to extend their result are common independent beliefs and private values. Neither the existence of a worst outcome nor the possibility of multiple rounds of messages is essential. By contrast, Theorem 2 below shows that in environments with a worst outcome, there is no need to assume common and independent beliefs and private values. Moreover, it is important to note that the mechanism and the strategies for Theorem 2 are quite different from those for Theorem 1. Indeed, the mechanism and the strategies for Theorem 1 do not work in more general environments.

The intuition for Theorem 1 is as follows. We consider the network $\mathcal{N}_5$ in Figure 5 and show how to implement the dictatorial social choice function of player 2. Note that player 2 has a directed path of communication to the designer (through player 1) and two disjoint undirected paths of communication to the designer. However, unlike the

network $\mathcal{N}_2$ in Figure 2, there is no player who has a directed path to player 2 and two disjoint directed paths to the designer. This feature is essential and makes the proof of Theorem 1 quite involved for general weakly 2-connected networks (see the Appendix for the general case).

As in Section 2, there are two alternatives a and b, and two types θ and θ' for player 2. Player 2 prefers a to b if his type is θ and prefers b to a if his type is θ'. Suppose that players 1, 3, 4, and 5 share a common prior and believe that player 2's type is θ with probability $\frac{1}{3}$. The designer aims to implement the dictatorial social choice function f^* of player 2.

An informal description of the strategies to implement f^* is as follows. Player 3 draws an encoding key y uniformly on $[0, 1)$ and sends it to players 2 and 4. Simultaneously, player 5 draws another encoding key z uniformly on $[0, 1)$ and sends it to the designer (player 0) and player 4. Then player 4 encrypts the key y received from player 3 with the key z received from player 5 to obtain $w = (z + y) \bmod_{0,1}$, which he sends to player 1. Player 2 of type θ (resp., θ') draws a pseudo-type $\tilde{x}$ uniformly in $[0, \frac{1}{3})$ (resp., $[\frac{1}{3}, 1)$) and sends the encrypted type $x = (\tilde{x} + y) \bmod_{0,1}$ to player 1. Thus, player 1 receives the encrypted type x from player 2 and the modified key w from player 4. Last, player 1 transfers $u = (w - x) \bmod_{0,1}$ to the designer. Let $(\hat{u}, \hat{z})$ be a pair of messages received by the designer. The allocation rule is the following: If $(\hat{z} - \hat{u}) \bmod_{0,1} \in [0, \frac{1}{3})$, the designer implements a and otherwise implements b.

If the players follow the prescribed strategies, then $w = (z + y) \bmod_{0,1}$ and $u = (w - x) \bmod_{0,1} = ((z + y) \bmod_{0,1} - (\tilde{x} + y) \bmod_{0,1}) \bmod_{0,1} = (z - \tilde{x}) \bmod_{0,1}$. The designer thus receives $\hat{u} = u = (z - \tilde{x}) \bmod_{0,1}$ from player 1 and $\hat{z} = z$ from player 5. It follows that $(\hat{z} - \hat{u}) \bmod_{0,1} = \tilde{x}$ and the designer correctly learns player 2's type and implements the desired social choice function f^*. In particular, all players but player 2 expect the designer to implement a with probability $\frac{1}{3}$ and b with probability $\frac{2}{3}$.

We now show that players do not have an incentive to deviate from the prescribed strategies and focus on player 1. From the point of view of player 1, $\tilde{x}$, y, and z are mutually independent and uniformly distributed. It follows that the two messages $(z + y) \bmod_{0,1}$ and $(\tilde{x} + y) \bmod_{0,1}$ received by player 1 are independent and uniformly distributed (see Lemma 2 in the Appendix), and convey no information about z and $\tilde{x}$. Suppose that player 1 deviates and sends the message $\hat{u}$ to the designer instead of $u = (z - \tilde{x}) \bmod_{0,1}$. The designer implements the alternative a if $(z - \hat{u}) \bmod_{0,1} \in [0, \frac{1}{3})$ and b otherwise. Since, conditional on player 1's information, z is uniformly distributed, so is $(z - \hat{u}) \bmod_{0,1}$ (see again Lemma 2 in the Appendix). Accordingly, player 1 expects the designer to implement a with probability $\frac{1}{3}$ and b with probability $\frac{2}{3}$. It follows that player 1's expected payoff does not depend on the message $\hat{u}$ he sends and that player 1 has no incentive to deviate. Similar arguments apply to players 3, 4, and 5. As for player 2, he has no incentive to deviate since f^* is incentive compatible.

The essential difference with the simpler example of Section 2 is that player 3 does not have two disjoint directed paths of communication to the designer. Thus, player 3 cannot give an encryption key to player 2 and send this key to the designer without player 1 learning both the encryption key and player 2's encrypted type. It is precisely at this point that the protocol of Monderer and Tennenholtz fails. The novel idea is then

FIGURE 6. Communication network $\mathcal{N}_6$ is strongly 1-connected.

to let player 4 encrypt the encryption key that player 3 sends to player 2 with the key received from player 5. Accordingly, player 1 receives an encrypted encryption key from player 4 and, therefore, learns nothing about the type of player 2.

The proof of Theorem 1 extends these arguments to any weakly 2-connected network[10] (all proofs are relegated to the Appendix). In particular, we show that if the network is strongly 1-connected and weakly 2-connected, then there exists a protocol such that if all players abide by the protocol, the designer correctly learns the players' types and no player gets additional information about the types of his opponents. In the language of computer science, we construct a protocol for the *secret transmission* of messages. We then show that the existence of such a protocol guarantees the existence of a mechanism and strategies such that players are indifferent between correctly forwarding the messages they receive or lying. Thus, they indeed have an incentive to abide by the protocol. In the language of computer science, our protocol is *reliable*.

Theorem 1 also states that the weak 2-connectedness is a necessary condition to implement *all* incentive compatible social choice functions. To get some intuition for this result, let us consider a simple example. There are two players, 1 and 2, two alternatives, a and b, and two types, θ and θ' for each player. Regardless of his type, player 1 prefers a over b, player 2 of type θ prefers a over b, while player 2 of type θ' prefers b over a. Consider the social choice function f for which player 2 is dictatorial and the communication network $\mathcal{N}_6$ in Figure 6. The issue with this network, and more generally with any communication network that is not weakly 2-connected, is that player 1 controls all the information sent by player 2 and there is no way for the designer to detect a false report by player 1.

Clearly, f is implementable on the star network $\mathcal{N}^\star$, but not on $\mathcal{N}_6$. By contradiction, suppose that f is implementable on $\mathcal{N}_6$ by the mechanism $\langle M_1, M_2, g \rangle$. There must exist an equilibrium message $m_1 \in M_1$ such that $g(m_1) = b$. However, regardless of his type and message received, player 1 has no incentives to send any message m_1 with $g(m_1) = b$, so that f cannot be implemented. The proof of Theorem 1 generalizes this argument to any network that is not weakly 2-connected.

Two further remarks are worth making. First, our encoding technique extends to environments with continuous type spaces (see Section 5.4). Second, the strategies we consider are behavioral strategies. In Section 5.5, we prove that our result does not hold if we restrict ourselves to pure equilibria, a frequently used solution concept in the mechanism design literature.

Before going further, it is worth stressing again that the encoding technique used in the proof of Theorem 1 is tailored to environments with common independent beliefs and does not apply to more general environments (even with private values). See

---

[10]Note that the protocol (mechanism and strategies) of Monderer and Tennenholtz (1999) for undirected networks does not work in general; there is a need to encrypt encryption keys. Their protocol works only if the directed network is strongly 2-connected.

the example in Section 2 for some intuition. With general beliefs, different encoding techniques have to be used: this is the object of the next section.

4.2 *Worst outcome*

In many concrete applications of the theory of mechanism design, players hold different and correlated beliefs about states of the world either because they have received different signals (information) or on purely subjective grounds. Moreover, the payoff of a player often depends on the private information of others. For instance, in auction models, bidders often have different information about the value of the goods for sale (e.g., mineral or oil rights) and the private information of all players influences the valuation for the *good of each player*. To handle these more general beliefs and payoff functions, we resort to a different encoding technique. Our new technique consists of coding the type of each player such that no information is revealed to the other players, and if a player does not truthfully forward the messages he receives, the designer detects it with arbitrarily high probability.

THEOREM 2. *Consider an acyclic network $\mathcal{N}$. For all environments $\mathcal{E}$ with a worst outcome, $F_\mathcal{N}(\mathcal{E}) = F_{\mathcal{N}^\star}(\mathcal{E})$ if and only if $\mathcal{N}$ is weakly 2-connected.*

The main insight provided by Theorem 2 is that assuming a worst outcome allows us to dispense with the assumptions of common independent beliefs and private values.

The intuition for Theorem 2 is as follows. We construct a mechanism such that the true type of player i is transmitted to the designer, no player $j \neq i$ gets information about the type of player i, and a false report by player j is detected with arbitrarily high probability. Consider again the network $\mathcal{N}_5$ and the dictatorial social choice function of player 2.

An informal description of the strategies is the following. Player 3 sends a large number of encoding keys, all uniformly and independently drawn from $[0, 1)$ to players 2 and 4. Simultaneously, player 5 sends another large number of encoding keys, all uniformly and independently drawn from $[0, 1)$ to player 4 and the designer. Player 4 thus receives a large number of keys from both player 3 and player 5. He adds them one-by-one (addition is modulo $[0, 1)$) and sends the resulting vector of keys to player 1. Simultaneously, player 2 selects at random one of the keys received from player 3 and encrypts his type with this key. He then substitutes the selected key with the cypher type and sends it to player 1 along with all the other keys (without telling player 1 which key was used to encrypt his type). Last, player 1 received a large vector of encrypted encryption keys from player 4 and a large vector of encryption keys and the encrypted type from player 2. Player 1 then subtracts these two vectors (subtraction is componentwise modulo $[0, 1)$) and forwards the resulting vector to the designer. The designer can then detect a false report by comparing the two vectors of messages received from players 1 and 3. Namely, if player 1 truthfully forwards the message he receives, the two vectors should differ by exactly one component. In such a case, the designer decodes the type of player 2 according to this component and implements the appropriate outcome. Otherwise, the designer implements the worst outcome. By construction, only player 2 knows

the key selected to encrypt his type. Thus, any deviation by players 1, 3, 4, and 5 induces the worst outcome with arbitrarily high probability: this deters them from lying.

An essential feature of Theorem 2 is the ability to punish a detected deviation with a worst outcome. It is worth stressing, however, that our definition of a worst outcome is stronger than necessary, since it does not depend on the social choice function we aim to implement. It would be enough to find an outcome worse than any outcome in the range of the social choice function.[11]

If such a worst outcome does not exist, the main difficulty for the designer is the choice of an appropriate alternative to implement whenever a false report is detected. A characterization of networks that allows implementation of all incentive compatible social choice functions in all environments is left as an open problem. Yet we provide sufficient conditions in Section 5.3. Naturally, weak 2-connectedness remains a necessary condition.

4.3 *Connections with computer science*

An essential feature of our results is the use of encryption techniques to secure the transmission of messages from players to the designer. As already alluded in the Introduction, our work is closely related to the computer science literature on secure transmission of messages, which we now review. We first discuss two important notions of security that are commonly found in the computer science literature.

Message security   Informally, the transmission of a message from a sender A to a receiver B is *reliable* if A can communicate with B and no adversary, i.e., a potentially malicious third party (a hacker), can tamper with the content of the message. The transmission of a message is *secret* if no adversary can find out the content of the message sent. Information transmission is said to be *secure* if it is both reliable and secret. To discuss the notion of secrecy more precisely, let us assume that A and B have a reliable channel of communication. There are two main approaches to message security in computer science: cryptographic security and information-theoretic security.

A message transmission is cryptographically secure if it is *computationally very hard* (typically NP-hard) for an adversary to find out the content of the message. This approach assumes that the adversary is computationally limited, that is, has no more computational power than a Turing machine. The reader is referred to the seminal papers of Diffie and Hellman (1976) and Rivest et al. (1978; RSA). In particular, classical encryption techniques with *public and private keys* adopt this notion of security. For instance, the RSA encryption scheme with public keys rests on the idea that computing two large prime numbers p and q when their product $n = pq$ is known is computationally very hard.

By contrast, information-theoretic security considers adversaries with unbounded computational power and requires pieces of communication between A and B, which

---

[11]It is also worth noting that Theorem 2 remains true if we consider environments with a bad outcome, i.e., an outcome $\underline{a}$ such that $u_i(f(\theta), \theta) \geq u_i(\underline{a}, \theta)$ for all $i \in N$, for all $\theta \in \Theta$. For completeness, the proof is given in the Appendix, Corollary 3.

may be eavesdropped, to be probabilistically independent of the content of the message. This concept was originally introduced by Shannon (1949) (see also, among others, Goldwasser and Micali 1984, Dolev et al. 1993). A simple method to achieve information-theoretic security is to map the message m to be sent to a number in, say, $\{1, \ldots, n\}$, and to add (modulo n) a uniformly distributed random key X. The encrypted message $(X + m) \bmod n$ is then uniformly distributed and independent of m: it can be publicly disclosed without harming security. The probability of guessing m correctly is $1/n$ and thus can be made arbitrarily small. Our encryption method (Lemma 2) is a continuous version of this method such that the probability of guessing correctly is zero.

As a game-theoretic model, our work follows the latter approach: the agents we consider are unboundedly rational players. They are very similar to the Byzantine adversaries considered in computer science, i.e., malicious players with unbounded computational power. The key difference, however, is that rational players respond to incentives: they do not behave maliciously if it is not optimal for them to do so.

Security in networks   Assume now that the sender A and the receiver B are some distant nodes in a network, so that there is no secure channel of communication between them. The natural question then is how to characterize the networks that guarantee the secure transmission of messages from A to B in the presence of Byzantine adversaries. This is the object of the computer science literature on secure transmission of messages. A seminal contribution is Dolev et al. (1993), who show that if the adversary controls at most t nodes, then $(2t + 1)$-connectedness of the network is a necessary and sufficient condition for the secure transmission of messages from A to B. Dolev et al. assume unicast communication, i.e., a node can send different messages to its neighbors. Alternatively, Franklin and Wright (2000) study broadcast communication: any message sent by a node is automatically sent to all its neighbors. They show that $(2t + 1)$-connectedness is again necessary and sufficient for perfect security.[12]

Unlike our approach, all these results assume undirected graphs and crucially use the possibility of messages going back and forth from the sender to the receiver (repeated communication). Dolev et al. (1993) show that in one-way problems, i.e., if the information flows only from the sender to the receiver, a necessary and sufficient condition for the secure transmission of messages is the $(3t + 1)$-connectedness of the network. Considering directed networks, Desmedt and Wang (2002) show how this bound can be lowered if there are channels of communication from the receiver to the sender. Namely, they show that if for $u \leq t$, there are $2t + 1 - u$ disjoint directed paths from the sender to the receiver and u disjoint directed paths from the receiver to the sender (these u paths are also disjoint from the $2t + 1$ paths from the sender to the receiver), then secure transmission of messages is possible.

---

[12]Franklin and Wright (2000) also consider a weaker notion of security: security is almost perfect when the adversary has an arbitrarily small probability of modifying the message content and learning the content of the message. They show that $(t + 1)$-connectedness is necessary and sufficient for almost-perfect security (see also Renault and Tomala 2008).

Our contribution to information security   The above discussion suggests a reinterpretation of our results in the language of computer science. Starting from a communication network, a social environment, and an incentive compatible social choice function f, we construct a mechanism that implements f as a Bayesian–Nash equilibrium of the induced game. A necessary condition for this result is the possibility to construct a communication protocol with the following properties: (i) the designer correctly learns the profile of types, (ii) no player gets information beyond his own type, and (iii) no player has an incentive to misexecute the communication protocol. Part (ii) corresponds to the computer science requirement of secrecy, while parts (i) and (iii) are the counterparts of reliability.

Before proceeding, it is worth emphasizing that the concept of Bayesian–Nash equilibrium implies that the adversary is a single potential deviant player. Such an adversary has unbounded computational power, responds to incentives, and controls at most one node ($t = 1$). Our main results are then reinterpreted as information transmission against this class of adversaries.

In Theorem 1, we assume common independent belief and private values, and construct a mechanism such that each player forwards the messages he receives and gets the same expected payoff regardless of the messages he forwards (see Section 2 and the proof of Theorem 1). With this in mind, our implementation problem is rephrased as the following problem of information transmission.

P1. *Characterize the networks for which there exists a communication protocol such that if all players abide by the protocol, the designer correctly learns the entire profile of types and no player gets additional information.*

In the presence of a worst outcome, the designer has the ability to punish all players if he detects a deviation, and we construct a protocol such that any tampering with a message is detected with arbitrarily high probability by the designer (see Section 2 and the proof of Theorem 2). The implementation problem thus gives rise to the following problem of information transmission.

P2. *Characterize the networks for which there exists a communication protocol such that no player gets additional information and if all but at most one player abide by the protocol, then the designer either correctly learns the entire profile of types or detects a deviation with arbitrarily high probability.*

Our main contribution to the literature on secure transmission of messages in networks is thus to solve problems P1 and P2 for directed graphs and one-way problems: *the solutions are the weakly-2-connected graphs.* Compared with the computer science literature cited above, our approach through incentives allows us to get a much weaker connectivity requirement. This statement is a by-product of the proofs of our main results, which are structured as follows. We first show that on any weakly 2-connected graph, there exists a communication protocol such that if all players abide by the protocol, the designer correctly learns the entire profile of types and no player gets additional information. Theorem 1 then easily follows: we use the common prior to make players indifferent between all the messages they may forward. The proof of Theorem 2 uses

FIGURE 7. Communication network $\mathcal{N}_7$.

a multiple key technique, akin to authentication schemes (see, e.g., Rabin and Ben-Or 1989), but requires no prior knowledge of any public or private key. To the best of our knowledge, this technique is new.

Finally, let us remark that the use of continuous message spaces, while consistent with mechanism design theory, is unappealing from a computer science perspective. Theorem 1 remains valid with finite message spaces, provided that prior beliefs are rational numbers: encoding keys are then chosen in the integers modulo **n**, with **n** large. Theorem 2 extends to finite messages spaces without restrictions on priors.

5. EXTENSIONS AND ROBUSTNESS

This section discusses various aspects of our problem and offers some generalizations.

5.1 *Active designer and two-way networks*

A salient feature of our model is that the designer is not active in the communication. However, in some situations, it is natural to assume that the designer can communicate with the players. For instance, a CEO has the ability to communicate with his employees either publicly or privately.

So, let us assume that the designer can communicate with some players, so that $C(0) \neq \varnothing$. An important consequence of assuming an active designer is that the network may then contain cycles. We therefore need to relax the assumption of acyclicity. Clearly, the conditions of strong 1-connectedness and weak 2-connectedness remain necessary for the implementation of all incentive compatible social choice functions. The main insight is that these conditions are also sufficient. In other words, our results extend naturally to networks with cycles.

THEOREM 3. *For all environments $\mathcal{E}$ with common independent beliefs and private values or with a worst outcome, $F_{\mathcal{N}}(\mathcal{E}) = F_{\mathcal{N}^\star}(\mathcal{E})$ if and only if $\mathcal{N}$ is weakly 2-connected.*

To get an intuition for this result, consider the network $\mathcal{N}_7$ in Figure 7.

The idea is simply to let the designer play the role of a provider of keys, as in the proof of Theorem 1 or Theorem 2. To be more specific, let us consider the transmission of player 3's private information in the network $\mathcal{N}_7$ when there is a worst outcome. The designer draws a large number of encoding keys and sends them to player 2. Player 2 forwards the encoding keys to player 3, who selects one key at random and uses it to encode his type. He then sends the unused keys and the encoded type to player 1, who should forward this message to the designer. Last, the designer compares the vector of keys he sent to player 2 and the vector of keys he receives from 1, and decodes the type of player 3 accordingly. As in the proof of Theorem 2, any deviation by player 1 or player 2 is detected with arbitrarily large probability, no information about player 3's type is revealed, and the designer correctly learns the type of player 3.

Theorem 3 admits as a special case two-way communication networks where players can exchange messages back and forth along each edge. Such networks are naturally represented by undirected graphs where there is an edge between i and j whenever i and j can converse privately. For this class of networks, strong 2-connectedness and weak 2-connectedness coincide, since one can choose any orientation of the edges. We thus obtain the following corollary.

COROLLARY 1. *For all environments $\mathcal{E}$ with common independent beliefs and private values or with a worst outcome, $F_{\mathcal{N}}(\mathcal{E}) = F_{\mathcal{N}^{\star}}(\mathcal{E})$ if and only if the two-way network $\mathcal{N}$ is 2-connected.*

Finally, let us mention that the assumption of an active designer is important in generalized principal–agents models (Myerson 1982), where players also have to take an action, thus creating a moral hazard problem in addition to the adverse selection problem. In such models, the designer has to "securely recommend" an action to each player. We believe that our results extend to this more general framework. Indeed, if the designer has two disjoint paths of communication to each player (directed or undirected), then he can follow our protocols to privately and reliably make a recommendation to each player. A careful analysis of this issue awaits future research.

5.2 *Direct mechanisms*

Another central feature of our results is the use of encryption techniques to secure the transmission of messages from the players to the designer. This is largely inescapable if we want to implement *all* incentive compatible social choice functions (in Section 5.5, we show that implementing all such functions in pure strategy equilibria is not possible except on $\mathcal{N}^*$).

However, "direct" mechanisms—where players simply announce their types to their neighbors and forward messages—might suffice if we restrict attention to specific environments or to some specific incentive compatible social choice functions. For instance, consider the set of *ex post* incentive compatible social choice functions. A social choice function f is ex post incentive compatible if for all $i \in N$ and $\theta \in \Theta$, $u_i(f(\theta), \theta) \geq u_i(f(\theta_i', \theta_{-i}), \theta)$ for all $\theta_i' \in \Theta_i$.[13]

---

[13]Bergemann and Morris (2005) show that a social choice function is implementable on all type spaces if and only if it is ex post incentive compatible.

FIGURE 8. Communication network $\mathcal{N}_8$.

PROPOSITION 1. *If the communication network $\mathcal{N}$ is strongly 3-connected, then any ex post incentive compatible social choice function is implementable on $\mathcal{N}$ by a direct mechanism.*

The intuition for Proposition 1 is simple. If a social choice function f is ex post incentive compatible, then every player has an incentive to truthfully reveal his private information, even if he were to know the private information of some other players (e.g., his neighbors). There is, therefore, no particular need for encryption techniques: players can simply truthfully report their types on all paths to the designer. In the computer science terminology, secrecy is not an issue. Yet it remains the issue of reliability: players must have the incentive to truthfully forward the messages they receive. However, with three disjoint directed paths of communication from each player $i \in N \setminus D(0)$ to the designer, a simple majority argument guarantees that no player has an incentive to misreport the messages he receives.

Furthermore, it is clear that not all ex post incentive compatible social choice functions are implementable by direct mechanisms on weakly 2-connected networks, even in environments with common independent beliefs and private values or a worst outcome. For a counterexample, we refer the reader to the example in Section 2. So, weak 2-connectedness is not a sufficient condition.

In some environments, however, some ex post incentive compatible social choice functions can be implemented by direct mechanisms, even on strongly 2-connected networks. We illustrate this possibility with the help of two important economic examples: a second-price auction and the provision of a public good.

Consider an auction with three bidders, labeled 1, 2, and 3. There is a single object to be allocated, bidder i values the object at θ_i, and bidder i's payoff is $\theta_i - x_i$ if he is allocated the object at price x_i and is zero otherwise. Consider the strongly 2-connected network $\mathcal{N}_8$ in Figure 8.

The designer aims to allocate the object to the bidder with the highest valuation (if there are several such bidders, choose one randomly). A simple and direct mechanism to implement the social choice function is as follows. Bidder 3 is required to truthfully

report his valuation θ_3 to both bidders 1 and 2. Bidder 1 (resp., bidder 2) has to truthfully report his valuation θ_1 (resp., θ_2) along with bidder 3's valuation θ_3 to the designer. Let $((\hat{\theta}_1, \hat{\theta}_3^1), (\hat{\theta}_2, \hat{\theta}_3^2))$ be a profile of messages received by the designer. The designer computes the bid profile $(\hat{\theta}_1, \hat{\theta}_2, \max(\hat{\theta}_3^1, \hat{\theta}_3^2))$, allocates the object to the highest bidder, and charges a price equal to the second-highest bid: a second-price auction.

Since a second-price auction implements the efficient allocation in weakly dominant strategies (on the star network), no bidder has an incentive to misreport his own valuation, regardless of the reports of the other bidders. We now argue that bidder 1 has no incentive to misreport bidder 3's valuation. (A symmetric reasoning holds for bidder 2.) Clearly, if bidder 1 reports $\hat{\theta}_3^1 < \theta_3$, he does not affect the outcome since $\max(\hat{\theta}_3^1, \theta_3) = \theta_3$. Alternatively, if bidder 1 reports $\hat{\theta}_3^1 > \theta_3$, he does affect the outcome of the auction. However, this is not a profitable deviation: it not only decreases his likelihood of winning the object, but also increases the price paid if he wins.

The second example is about the provision of a public good and is adapted from Bergemann and Morris (2009). Assume that there are three players and that $\Theta_i \subseteq [0, 1)$ for each player $i \in \{1, 2, 3\}$. The utility to player i is $(\theta_i + \gamma \sum_{j \neq i} \theta_j)x_0 + x_i$, where x_0 is the level of public good provided and x_i is the monetary transfer to player i ($\gamma \geq 0$). The cost of providing the level of public good x_0 is $(\frac{1}{2})(x_0)^2$. The designer aims to implement the efficient level of public good, i.e., $(1 + 2\gamma)(\theta_1 + \theta_2 + \theta_3)$, at the type profile $(\theta_1, \theta_2, \theta_3)$. Again, consider the network $\mathcal{N}_8$ in Figure 8. As in the previous example, the players are required to truthfully report their types along with any message they might have received. Let $((\hat{\theta}_1, \hat{\theta}_3^1), (\hat{\theta}_2, \hat{\theta}_3^2))$ be a profile of messages received by the designer. The designer then computes the type profile $(\hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3)$ with $\hat{\theta}_3 := \min(\hat{\theta}_3^1, \hat{\theta}_3^2)$, produces the level $x_0 = (1 + 2\gamma)(\hat{\theta}_1 + \hat{\theta}_2 + \hat{\theta}_3)$ of public good, and establishes the transfer $x_i = -(1 + 2\gamma)[\gamma \hat{\theta}_i \sum_{j \neq i} \hat{\theta}_j + (\frac{1}{2})\hat{\theta}_i^2 - 2\gamma \sum_{j \neq i} \hat{\theta}_j]$ to each player i. Note that up to the term $(1 + 2\gamma)2\gamma \sum_{j \neq i} \hat{\theta}_j$, independent of player i's type, the transfers are identical to the generalized Vickrey–Clarke–Groves transfers of Bergemann and Morris (2009). In particular, they guarantee that the social choice function is ex post incentive compatible (on the star network). However, and unlike the first example, the mechanism does not implement the social choice function in dominant strategies, even on the star network (unless $\gamma = 0$). Player 1 (resp., player 2) might, therefore, have an incentive to misreport his own type whenever his report of player 3's type leads to $\hat{\theta}_3$ being different from player 3's true type.[14] We argue nonetheless that no player has an incentive to misreport in that example. To do so, we compute the difference $\delta_1((\hat{\theta}_1, \hat{\theta}_3^1)|\theta)$ in player 1's ex post payoff between a truthful report (θ_1, θ_3) and the report $(\hat{\theta}_1, \hat{\theta}_3^1)$ at the type profile θ:

$$\delta_1((\hat{\theta}_1, \hat{\theta}_3^1)|\theta) = \tfrac{1}{2}(\theta_1 - \hat{\theta}_1)^2 + [\theta_1 + \gamma(\theta_2 + \theta_3 - \hat{\theta}_1) + 2\gamma](\theta_3 - \hat{\theta}_3),$$

with $\hat{\theta}_3 := \min(\hat{\theta}_3^1, \theta_3)$, the minimum between player 1's report about player 3's type and player 2's (true) report about player 3's type. Since $\hat{\theta}_3 \leq \theta_3$ and $\theta \in [0, 1)^3$, $\delta_1((\hat{\theta}_1, \hat{\theta}_3^1)|\theta) \geq 0$ for all θ, and thus player 1 has no profitable deviation. A similar reasoning applies to player 2. As for player 3, he clearly has no profitable deviation since the social choice function is ex post incentive compatible.

---

[14]Remember that ex post incentive compatibility guarantees that no player has an incentive to misreport his own type for all *truthful* reports of his opponents (but not necessarily for all reports of his opponents).

Both examples generalize to any number of players provided that the communication network is strongly 2-connected. Last, note that a common feature of both examples is the existence of a "sufficient statistic" to aggregate conflicting reports about player 3's type, with the additional property that this aggregate statistic deters players 1 and 2 from lying about player 3's type. We suspect that this property can be generalized and leave it as an open issue.

5.3 *All environments*

We give sufficient conditions on the network for implementing all incentive compatible social choice functions, regardless of the environments.

Recall that a network is strongly m-connected if for each player $i \in N \setminus D(0)$, there exist m disjoint directed paths from player i to the designer. Likewise, a network is weakly m-connected if for each player $i \in N \setminus D(0)$, there exist m disjoint undirected paths from player i to the designer.

THEOREM 4. *If the communication network $\mathcal{N}$ is strongly 2-connected and weakly 3-connected, then $F_{\mathcal{N}}(\mathcal{E}) = F_{\mathcal{N}^\star}(\mathcal{E})$ for all environments $\mathcal{E}$.*

The intuition is the following. We first prove that for each strongly 1-connected and weakly 2-connected network, there exists a mechanism such that any false report of messages is detected with probability 1 and no additional information about the types is revealed (the construction is in the Appendix, Lemma 8).

Next, consider a strongly 2-connected and weakly 3-connected network, and fix a player $i \in N \setminus D(0)$ who wants to transfer his type to the designer. Notice that for each player $j \neq i$, $j \neq 0$, the subnetwork $\mathcal{N} \setminus \{j\}$ (obtained from $\mathcal{N}$ by deleting j) is strongly 1-connected and weakly 2-connected. From the above, there exists a "submechanism" on this subnetwork that detects deviations with probability 1. A simple "majority" argument then ensures that no player has an incentive to lie. More precisely, any unilateral deviation of player $j \neq i$ is almost surely detected, while the submechanism on $\mathcal{N} \setminus \{j\}$ is truthfully executed and allows the designer to correctly decode the type of player i.[15]

5.4 *A continuum of types and alternatives*

Many applications of mechanism design theory, e.g., contract theory and auction theory, assume a continuum of types and alternatives. While we have cast our results in finite settings, they naturally extend to environments with continuous type and alternative sets.[16]

We now explain how to extend Theorem 1. A key feature of the proof of Theorem 1 is that player i transforms his type θ_i into a pseudo-type $\tilde{x}_i$, which reveals his type and is unconditionally uniformly distributed in $[0, 1)$. The pseudo-type is then transmitted through the network by a communication protocol. It is thus enough to show how

[15]We thank Thomas Voice for suggesting this argument to us.

[16]Appropriate measurability and integrability assumptions have to be made.

to construct the pseudo-type in the continuous setup. Let each player's type space Θ_i be a subset of $[0, 1)$ and let types be independently distributed. Let P be the common prior and let G_i be the cumulative distribution function of the marginal P^i over Θ_i. Assume that G_i is continuous. The key observation to make is that $G_i(\theta_i)$ is uniformly distributed on $[0, 1)$ and, therefore, can be used as a pseudo-type. If G_i has atoms, let θ_i^* be an atom of G_i, i.e., $\lim_{\theta_i \uparrow \theta_i^*} G_i(\theta_i) := G_i^-(\theta_i^*) < G_i^+(\theta_i^*) =: \lim_{\theta_i \downarrow \theta_i^*} G_i(\theta_i)$. Let $\hat{G}_i(\theta_i^*)$ be the realization of a uniform draw on $[G_i^-(\theta_i^*), G_i^+(\theta_i^*))$. Let $\hat{G}_i(\theta_i) = G_i(\theta_i)$ if θ_i is not an atom. Then $\hat{G}_i(\theta_i)$ is uniformly distributed (unconditionally on θ_i) and reveals the value of θ_i, thus is a valid pseudo-type. The mechanism construction of Theorem 1 then extends verbatim.

As for Theorem 2, it extends straightforwardly to a continuum of types and alternatives. In sum, all our constructions naturally extend to the continuous case.

5.5 *Pure equilibria*

With the notable exception of Serrano and Vohra (2010), the literature on implementation in Bayesian environments has entirely focused on the implementation of social choice functions in *pure* equilibria (see Jackson 2001 for a survey). By contrast, the recourse to equilibria in mixed strategies is essential for our results. In effect, to transmit their types to the designer securely, it is essential for the players to encrypt their types with randomly generated keys (mixing). Although the use of randomly generated keys seems natural in our context, and indeed is used in daily life (internet banking, online shopping, etc.), we might legitimately wonder whether similar results hold in environments where only pure equilibria are considered. The next theorem states that the set of social choice functions partially implementable on $\mathcal{N}$ in pure equilibria coincides with the set of incentive compatible social choice functions, irrespective of the utility functions, if and only if every player is directly connected to the designer. There is a sharp divide between implementation in pure equilibria and mixed equilibria. Let $F_{\mathcal{N}}^{\text{pure}}(\mathcal{E})$ denote the set of social choice functions (partially) implementable on $\mathcal{N}$ in pure equilibria when the environment is $\mathcal{E}$.

THEOREM 5. *The set $F_{\mathcal{N}}^{\text{pure}}(\mathcal{E}) = F_{\mathcal{N}^\star}^{\text{pure}}(\mathcal{E})$ for all environments $\mathcal{E}$ with common independent beliefs and private values or a worst outcome if and only if each player is directly connected to the designer, i.e., $D(0) = N$.*

The intuition is simple.[17] If player i is not directly connected to the designer and if the social choice function depends on his type, then he must send an informative message to at least one other player, say player j. Given his updated beliefs, player j might then have no incentive to truthfully report his own private information. This reasoning is valid regardless of how many disjoint paths there are from player i to the designer.

While intuitive, Theorem 5 has remarkable implications for the topology of communication networks and implementation in pure equilibria. All but one player, say

---

[17] See Renou and Tomala (2010) for a formal proof.

player 1, might be directly connected to the designer, player 1 might have $n - 1$ disjoint paths of communication to the designer, and yet there exist incentive compatible social choice functions that are not implementable on that network in pure equilibria. While some theorists might feel uncomfortable with equilibria in mixed strategies, the mixing through encoding techniques, as considered in this paper, seems quite natural.

6. Conclusion

This paper completely characterizes the communication networks for which, in any environments (utilities and beliefs) with either common independent priors and private values or with a worst outcome, every incentive compatible social choice function is (partially) implementable. We show that any weakly 2-connected communication network can replicate the incentive properties of the direct revelation mechanism. Importantly, our constructions couple encryption techniques together with incentives to secure the transmission of each player's private information to the designer.

Appendix

A.1 *Timing Structure*

In this section, we prove that the communication rule stating that "a player sends his messages after having received all his messages" generates a well defined timing structure.

LEMMA 1. *Let $\mathcal{N}$ be a strongly 1-connected and acyclic network. There exists an integer T and a timing function $t : N \to \{1, \dots, T\}$ such that $t(i)$ is the stage at which player i sends his messages. Moreover, $ij \in \mathcal{N} \Rightarrow (i) < t(j)$.*

PROOF. Let $V_1 = \{i \in N : D(i) = \varnothing\}$ be the set of players who cannot receive messages. This set is clearly nonempty, for otherwise, there exists a cycle in $\mathcal{N}$. If $V_1 = N$, then $\mathcal{N} = \mathcal{N}^*$ and the proof is complete. If $V_1 \neq N$, let $V_2 = \{i : i \notin V_1 \text{ and } D(i) \subseteq V_1\}$.

CLAIM 1. *If $V_1 \neq N$, then V_2 is nonempty.*

PROOF. Define $W_1 = \bigcup_{i \in V_1} C(i)$ as the set of players with whom the players in V_1 can communicate. By construction, if j is in W_1, then $D(j)$ is nonempty and, therefore, $j \notin V_1$. Consider then a directed path π of maximal length among the directed paths from a player in W_1 to the designer (such a path exists by strong 1-connectedness). Let j be the starting point of this directed path. We claim that j is in V_2. By contradiction, suppose that there exists $k \in D(j)$ with $k \notin V_1$. There exists then a directed path from some point m in V_1 to k, denoted $\tau = m \to l \to \cdots \to k \to j$. It follows that l is in W_1 and $\tau\pi$ contradicts the maximality of π.   ◁

If $V_1 \cup V_2 = N$, the construction ends. If $V_1 \cup V_2 \neq N$, let

$$V_3 = \{i : i \notin V_1 \cup V_2 \text{ and } D(i) \subseteq V_1 \cup V_2\}.$$

We continue this construction by induction. Assume that for some $k \geq 2$, the set V_s has been defined, $s \leq k$. If $\bigcup_{s \leq k} V_s = N$, the construction ends. If $\bigcup_{s \leq k} V_s \neq N$, let

$$V_{k+1} = \left\{ i : i \notin \bigcup_{s \leq k} V_s \text{ and } D(i) \subseteq \bigcup_{s \leq k} V_s \right\}.$$

CLAIM 2. *If $\bigcup_{s \leq k} V_s \neq N$, then V_{k+1} is nonempty.*

PROOF. Let $W_{k+1} = \{ j \notin \bigcup_{s \leq k} V_s : \exists i \in \bigcup_{s \leq k} V_s, j \in C(i) \}$. Since $\bigcup_{s \leq k} V_s \neq N$, W_{k+1} is nonempty. Consider then a directed path π of maximal length among the directed paths from a player in W_{k+1} to the designer (such a path exists by strong 1-connectedness). The starting point j of this path is in V_{k+1}. By contradiction, suppose that there exists $k \in D(j)$, $k \notin \bigcup_{s \leq k} V_s$. There exists then a directed path from some point m in $\bigcup_{s \leq k} V_s$ to k. The follower of m on this path is in W_{k+1} and this contradicts the maximality of π. ◁

The sequence $(\bigcup_{s \leq k} V_s)_k$ is a weakly increasing sequence of sets and is strictly increasing as long as $\bigcup_{s \leq k} V_s \neq N$. Since N is finite, there exists k such that $\bigcup_{s \leq k} V_s = N$. The timing function is then defined as $t(i) = s$ if $i \in V_s$. □

A.2 *Probabilistic encryption*

We present three important properties about the modular manipulations of real numbers in $[0, 1)$. For a real number x, we let $\lfloor x \rfloor$ denote the greatest integer less than or equal to x, and let $x \bmod_{0,1} = x - \lfloor x \rfloor$ denote the fractional part of x. For $(x, y) \in [0, 1) \times [0, 1)$, we denote $x \oplus y = (x + y) \bmod_{0,1}$ and $x \ominus y = (x - y) \bmod_{0,1}$.

LEMMA 2. *(i) For each $(x, y) \in [0, 1) \times [0, 1)$, $(x \oplus y) \ominus y = x$. More generally, $[0, 1)$ is a commutative group for $\oplus$.*

   (ii) Let Y be a random variable in $[0, 1)$ and let $x \in [0, 1)$. If Y is uniformly distributed, then so are $x \oplus Y$ and $x \ominus Y$.

   (iii) Let X, Y be independent random variables in $[0, 1)$. If Y is uniformly distributed, then so are $Z = X \oplus Y$ and $W = X \ominus Y$. Furthermore, (X, Y, Z) (resp., (X, Y, W)) are pairwise-independent.

PROOF. (i) Consider any pair $(x, y) \in [0, 1) \times [0, 1)$. If $x + y \leq 1$, the statement is clear. If $x + y > 1$, then $(x + y) \bmod_{0,1} = x + y - 1$. Thus $(x + y) \bmod_{0,1} - y = x - 1$ and $(x - 1) \bmod_{0,1} = x$.

   (ii) For each $z \in [0, 1)$, we have

$$\mathbb{P}(x \oplus Y \leq z) = \mathbb{P}((x + Y) \leq z, Y \in [0, 1 - x]) + \mathbb{P}(x + Y - 1 \leq z, Y \in (1 - x, 1])$$

$$= \begin{cases} z - x + x & \text{if } z \geq x \\ z + 1 - x - (1 - x) & \text{if } z < x \end{cases}$$

$$= z.$$

FIGURE 9. Blocks attached at 0.

Thus, $X \oplus Y$ is uniformly distributed. Similarly, for each $z \in [0, 1)$,

$$
\begin{aligned}
\mathbb{P}(x \ominus Y \leq z) &= \mathbb{P}(x - Y \leq z, Y \in [0, x]) + \mathbb{P}(x - Y + 1 \leq z, Y \in (x, 1]) \\
&= \begin{cases} x + 1 - (x + 1 - z) & \text{if } z \geq x \\ z + 0 & \text{if } z < x \end{cases} \\
&= z.
\end{aligned}
$$

Thus, $x \ominus Y$ is uniformly distributed.

(iii) We show only that X and Z are independent, the rest being similar. For each $z \in [0, 1)$, $\mathbb{P}(Z \leq z | X = x) = \mathbb{P}(x \oplus Y \leq z) = z$ from (ii). $\square$

A.3 *Information transmission in weakly 2-connected network*

In this section, we describe the structure of directed paths in weakly 2-connected networks and deduce that messages can be *secretly* transmitted from each player to the designer. These results are building blocks for the proofs of our main theorems.

Throughout, all networks (directed graphs) are assumed to be acyclic, strongly 1-connected, and weakly 2-connected. Given a (directed) network $\mathcal{N}$, we let $\mathcal{N}^u$ denote the associated undirected network: $ij \in \mathcal{N}^u$ if and only if $ij \in \mathcal{N}$ or $ji \in \mathcal{N}$.

Our definition of weakly 2-connected networks is closely related to the definition of 2-connectedness for undirected graphs. An undirected graph is 2-connected if for each pair of distinct vertices i and j, there are two disjoint paths from i to j. There are several equivalent statements for 2-connectedness of undirected graphs and the reader is referred to Bollobás (1998, Chap. III.2). For instance, define a *cut vertex* as a vertex i such that deleting i and all its adjacent edges yields a disconnected graph. The graph is 2-connected if and only if there is no cut vertex. Equivalently, for each distinct vertex i, j, and k, there is a path from i to j that does not contain k.

In our model, the designer (player 0) plays a special role, so that the network $\mathcal{N}$ is weakly 2-connected if and only if no player $i \in N$ is a cut vertex of $\mathcal{N}^u$. The designer, however, can be a cut vertex. In such case, let a *block* be a maximal 2-connected subgraph of $\mathcal{N}^u$. The undirected network $\mathcal{N}^u$ is a collection of blocks attached at 0. See Figure 9 for an example. In the sequel, we assume for simplicity that $\mathcal{N}^u$ is the only block, so that $\mathcal{N}^u$ is 2-connected. (If there are several blocks, all our arguments remain valid block-by-block.)

FIGURE 10. Loop $L(a_2, b_2)$ is a successor of $L(a_1, b_1)$.

In the sequel, we use the letters a, b, etc. to denote nodes (players) in the network. This must not be confused with alternatives.

We define a *loop*, denoted $L(a, b)$, in $\mathcal{N}$ as a pair of directed paths with the same origin a and endpoint b, and no vertex in common except for the origin a and the endpoint b. The loop $L(a_2, b_2)$ is a *successor* of the loop $L(a_1, b_1)$ if $a_2 \notin L(a_1, b_1)$, $b_2 \notin L(a_1, b_1)$, and the intersection $L(a_1, b_1) \cap L(a_2, b_2)$ is a path that contains at least one edge and the vertex b_1. See Figure 10 for an example.

We use the following notation: we write $i \to k$ for a directed path $(i_0 = i, i_2, \ldots, i_R = k)$ from player i to player k and write $i \to k \to l$ for a directed path from i to l through k, etc. We say that two directed paths $(i_0 = i, i_2, \ldots, i_R)$ and $(j_0 = i, j_2, \ldots, j_Q)$ *cross each other* if there exist r^* and q^* such that $j_{q^*} = i_{r^*}$.

To prove our main results, we use the following decomposition of directed graphs into successive loops. We assume that there are at least three player (if $n = 2$, the only strongly 1-connected and weakly 2-connected network is such that $D(0) = N$).

PROPOSITION 2. *Let $n \geq 3$. For each $i \in N \setminus D(0)$ and each $j \in C(i)$, there exists a finite sequence of loops $L(a_1, b_1), \ldots, L(a_M, b_M)$ such that the following statements hold.*

(i) *The edge ij belongs to $L(a_1, b_1)$.*

(ii) *For each $m = 1, \ldots, M - 1$, $L(a_{m+1}, b_{m+1})$ is a successor of $L(a_m, b_m)$ and $a_{m+1} \notin \bigcup_{q \leq m} L(a_q, b_q)$.*

(iii) *Endpoint $b_M = 0$.*

See Figure 11 for an illustration.

PROOF OF PROPOSITION 2. This is trivially true if $n = 3$. Assume that $n \geq 4$. The proof rests on several lemmas.

LEMMA 3. *Let $\mathcal{N}^u$ be a 2-connected undirected graph. Let A be a nonempty set of vertices, and let b and c be two distinct vertices that do not belong to A. There exists $a^* \in A$ and a path from a^* to c that has no vertex in $(A \setminus \{a^*\}) \cup \{b\}$.*

FIGURE 11. A sequence of loops.

PROOF. Since $\mathcal{N}^u$ is 2-connected, for each $a \in A$, there exists a path from a to c that does not contain b (otherwise, b would be a cut vertex). This path must leave the set A to reach c, thus the last point a^* in A on this path has the desired properties.　　$\triangleleft$

LEMMA 4. *Letting $i \in N \setminus D(0)$ and $j \in C(i)$, there exists a loop that contains the edge ij.*

PROOF. Remember that for each player $k \in N$, there exists a directed path from k to 0 by strong 1-connectedness and, thus, $C(k) \neq \varnothing$. Consider a player $i \in N \setminus D(0)$ and $j \in C(i)$.

Case 1. If $C(i)$ contains another player $k \neq j$, then there exists a directed path from i to 0 through the edge ij and a directed path from i to 0 through the edge ik. These paths must cross each other (possibly at 0); thus we have found the desired loop.

Case 2. If $C(i) = \{j\}$, let $D_\infty(i)$ denote the set of players who have a directed path to i. From Lemma 3, there exists $k \in D_\infty(i)$ and an undirected path $(k_0 = k, k_1, \ldots, k_R = 0)$ from k to 0 such that no player k_r is in $D_\infty(i) \cup \{i\}$ for $r > 0$. It follows that edge kk_1 is directed from k to k_1. We choose then a directed path from k_1 to 0 to obtain the directed path $k \to k_1 \to 0$, on the one hand, and the directed path $k \to i \to j \to 0$, on the other hand. These paths must cross each other and, therefore, define a loop with origin k. (The first crossing point defines the endpoint of the loop.) The endpoint of the loop cannot be in $D_\infty(i) \cup \{i\}$ since $k_1 \notin D_\infty(i)$. It follows that the edge ij is contained in this loop.　　$\triangleleft$

We now construct the desired sequence of loops. We start with $i \in N \setminus D(0)$ and $j \in C(i)$.

First Step. Let $L(a_1, b_1)$ be a loop containing ij and such that $t(b_1)$ is maximal among all loops that contain ij ($t(\cdot)$ is the timing function constructed in Lemma 1). (Such a loop

exists by the above lemma.) If $b_1 = 0$, the construction ends. If $b_1 \neq 0$, let $c_1 \in C(b_1)$, and let d_1 and e_1 denote the two predecessors of b_1 on each path of $L(a_1, b_1)$.

The construction then proceeds inductively. Assume that $L(a_1, b_1), \ldots, L(a_M, b_M)$ have been constructed for some $M \geq 1$. If $b_M = 0$, the construction ends. If $b_M \neq 0$, let $c_M \in C(b_M)$, and let d_M and e_M denote the two predecessors of b_M on each of the two disjoint directed paths of $L(a_M, b_M)$.

For each subset of players N', let $D_\infty(N')$ denote the set of players j for whom there exists a directed path from j to some player in N'. Clearly, $D_\infty(N' \cup N'') = D_\infty(N') \cup D_\infty(N'')$ and $D_\infty(D_\infty(N')) = D_\infty(N')$.

LEMMA 5. *There exists a loop $L(a_{M+1}, b_{M+1})$ such that $a_{M+1} \notin \bigcup_{q \leq M} L(a_q, b_q) \cup D_\infty(i)$ and that contains either the path $d_M \to b_M \to c_M$ or the path $e_M \to b_M \to c_M$. Furthermore, this loop is disjoint from $\bigcup_{q \leq M-1} D_\infty(L(a_q, b_q)) \cup D_\infty(i)$.*

PROOF. From Lemma 3, there exists $u_M \in \bigcup_{q \leq M} D_\infty(L(a_q, b_q)) \cup D_\infty(i)$ and an undirected path $(\lambda_0 = u_M, \lambda_1, \ldots, \lambda_S = 0)$ from u_M to 0 disjoint from $(\bigcup_{q \leq M} D_\infty(L(a_q, b_q)) \cup D_\infty(i) \cup \{b_M\}) \setminus \{u_M\}$. Assume that $u_M \in D_\infty(L(a_M, b_M))$. There exists a directed path from u_M to b_M that goes either through d_M or through e_M. Without loss of generality, assume that this path goes through d_M. As before, the edge $u_M \lambda_1$ is directed from u_M to λ_1, and we choose a directed path from λ_1 to 0 to obtain the directed path $u_M \to \lambda_1 \to 0$, on one hand, and the directed path $u_M \to d_M \to b_M \to c_M \to 0$, on the other hand. These paths must cross each other and, therefore, define a loop with origin u_M. Since $\lambda_1 \notin \bigcup_{q \leq M} D_\infty(L(a_q, b_q)) \cup D_\infty(i)$, the path $\lambda_1 \to 0$ cannot go through $\bigcup_{q \leq M} D_\infty(L(a_q, b_q)) \cup D_\infty(i)$, and thus the endpoint of the loop is not in $\bigcup_{q \leq M} D_\infty(L(a_q, b_q)) \cup D_\infty(i)$ either. The path $d_M \to b_M \to c_M$ is thus contained in the new loop.

Finally, u_M cannot be in $\bigcup_{q \leq M-1} D_\infty(L(a_q, b_q)) \cup D_\infty(i)$. Otherwise, the construction above provides a loop that contradicts the maximality property of b_m for some $m < M$; that is, since $t(b_{M+1}) > t(b_m)$, the newly constructed loop would have been used at an earlier stage of the induction. Similarly, the origin a_{M+1} of the new loop cannot be in $\bigcup_{q \leq M} D_\infty(L(a_q, b_q)) \cup D_\infty(i)$.   ◁

Inductive Step. Let $L(a_{M+1}, b_{M+1})$ be a loop containing $d_M \to b_M \to c_M$ or $e_M \to b_M \to c_M$ and such that $t(b_{M+1})$ is maximal among all loops that contain $d_M \to b_M \to c_M$ or $e_M \to b_M \to c_M$. If $b_{M+1} = 0$, the construction ends; otherwise, it continues inductively.

By construction, there is a directed path from b_m to b_{m+1}, thus $t(b_m) < t(b_{m+1})$ from the definition of the timing structure. It follows that the construction stops after a finite number of iterations. This completes the proof of Proposition 2.   □

Proposition 2 is a building block for the construction of a protocol (mechanism and strategies) that allows player i to secretly send a message to the designer. Let us summarize our findings. Proposition 2 has the following implications: For each player $i \in N \setminus D(0)$ and $j \in C(i)$, there exists a finite sequence of loops $(L(a_m, b_m))_{m=1}^M$ such that (i) $ij \in L(a_1, b_1)$, (ii) $b_M = 0$, and (iii) the loop $L(a_{m+1}, b_{m+1})$ is a successor

of the loop $L(a_m, b_m)$, $m = 1, \ldots, M - 1$, with the additional property that there exists $u_m \in L(a_m, b_m) \cap L(a_{m+1}, b_{m+1})$ such that the directed path from u_m to b_m in $L(a_m, b_m)$ is part of the directed path from u_m to b_{m+1} in $L(a_{m+1}, b_{m+1})$. Moreover, the sequence of loops defines a directed path from player i to the designer through all players b_1 to b_{M-1}. To see this, note that player i belongs to the loop $L(a_1, b_1)$ from player a_1 to player b_1 and thus, belongs to one directed path to b_1. Similarly, b_1 belongs to the loop $L(a_2, b_2)$ and, thus, has a directed path to b_2. Iterating this argument, we construct a directed path from i to the designer through the players b_1 to b_{M-1}. We will use this directed path to secretly transfer the private information of player i to the designer.

PROPOSITION 3. *Let v be a random variable in $[0, 1)$ that is privately known to player i. There exists a protocol $\mathcal{M}_i$ (i.e., a mechanism and a profile of strategies) on $\mathcal{N}$ such that whenever all players follow the prescribed strategies, the designer correctly learns the value of v. Moreover, the messages received by any player $j \neq i$ are probabilistically independent from v.*

PROOF. If $i \in D(0)$, this is straightforward. Fix $i \in N \setminus D(0)$ and consider the sequence of loops constructed in Proposition 2. We divide players into several categories.

- A player who belongs to one loop is *active*. All other players are inactive. Inactive players do not send or receive messages (their message sets are singletons).

Let us focus now on active players.

- A player a_m who is the origin of a loop is a *provider*.

- A player b_m who is the endpoint of a loop is a *lock-opener*.

- The player u_m who is the first point on the intersection of the two successive loops $L(a_m, b_m)$ and $L(a_{m+1}, b_{m+1})$ is a *lock-closer*.

- Other active players are *transmitters*.

By construction, note that a provider has no active predecessor and exactly two active successors. A lock-opener or a lock-closer has two active predecessors and one active successor. Transmitters have exactly one active predecessor and one active successor. Finally, player i is either a transmitter or a provider. For each loop, we label the path that contains the lock-closer as *left (L)* and label the other path as *right (R)*. The strategies for active players other than player i are as follows.

- Each transmitter truthfully forwards the message received from his active predecessor to his active successor.

- Each provider a_m draws an encryption key X_m uniformly in $[0, 1)$ and sends it to its two active successors.

- Each lock-closer u_m receives two numbers x_m and x_{m+1} from his two predecessors. He computes $z_m = x_m \oplus x_{m+1}$ and sends z_m to his active successor. Note that there is no lock-closer u_{M+1} in the last loop $L(a_M, b_M)$.

FIGURE 12. Providers, lock-closers $\oplus$, and lock-openers $\ominus$.

- Each lock-opener b_m (with $m < M$) receives two numbers x_m^L and x_m^R from his left and right predecessors. He computes $w_m = x_m^L \ominus x_m^R$ and sends w_m to his active successor.

Player i's strategy is as follows.

- If he is a transmitter, player i receives x_1 from his active predecessor and sends $x_1 \oplus v$ to his active successor.

- If he is a provider, player i sends $X_1 \oplus v$ to his active successor on the left path and X_1 to his active successor on the right path.

See Figure 12 for a heuristic illustration of the strategies.

First, we show that this protocol allows the designer to correctly learn the value of v. To this end, let us assume that these strategies are effectively played and compute the messages w_m sent by the lock-openers.

The sequence of loops defines a directed path from player i to the designer. This path contains all lock-openers (b_m) and some lock-closers (u_m), and is uniquely defined if player i is a transmitter. If player i is a provider, we choose the only such path that begins with the left path of the first loop. Along this path, let us attach labels to players. All lock-openers and player i are labeled $\ominus$ and the lock-closers are labeled $\oplus$. For instance, in Figure 12, we have

$$i^{\ominus} \to u_1^{\oplus} \to b_1^{\ominus} \to b_2^{\ominus} \to u_4^{\oplus} \to b_3^{\ominus} \to b_4^{\ominus} = 0.$$

$$X_2 \oplus X_1 \oplus v \ominus X_1$$

$$X_2 \oplus X_1 \oplus v$$

$$X_2$$

$$i : X_1 \oplus v$$

$$a_1 : X_1$$

FIGURE 13. Message w_1 with player i on the left path.

This induces a sequence in the alphabet $\{\ominus, \oplus\}$. Let $\nu(b_m)$ be the number of occurrences of two consecutive $\ominus$ appearing in the sequence before b_m (including b_m). For instance, in the example above, $\nu(b_1) = 0$, $\nu(b_2) = \nu(b_3) = 1$, and $\nu(b_4) = 2$.

LEMMA 6. *If the players follow the above strategies, for each $m = 1, \ldots, M - 1$, we have*

$$w_m = (-1)^{\nu(b_m)} v \oplus X_{m+1}.$$

The two messages received by the designer are X_M and w_{M-1}.

Consequently, the designer can compute the value v of the private information of player i, which is $X_M \ominus w_{M-1}$ if $\nu(b_{M-1})$ is odd and $w_{M-1} \ominus X_M$ if $\nu(b_{M-1})$ is even.

PROOF OF LEMMA 6. We first compute w_1 and then proceed by induction. Consider the loop $L(a_1, b_1)$. Player i is either on the left path of the loop $L(a_1, b_1)$ or on the right path of $L(a_1, b_1)$. In the former case, the left path from i to b_1 is $i^\ominus \to u_1^\oplus \to b_1^\ominus$ and the right path is $i \to b_1$. Player b_1 thus receives $X_2 \oplus X_1 \oplus v$ from the left and X_1 from the right. It follows that $w_1 = (X_2 \oplus X_1 \oplus v) \ominus X_1 = X_2 \oplus v$. Note that in this case $\nu(b_1) = 0$. See Figure 13 for an illustration.

In the latter case, the left path is $a_1 \to u_1 \to b_1$ and the right path is $i^\ominus \to b_1^\ominus$. Player b_1 thus receives $X_2 \oplus X_1$ from the left and $X_1 \oplus v$ from the right. Thus $w_1 = (X_2 \oplus X_1) \ominus (X_1 \oplus v) = X_2 \ominus v$. Note that in this case $\nu(b_1) = 1$. See Figure 14 for an illustration. We have thus proved the lemma for $m = 1$.

We proceed now by induction. Let us assume that for some $m \leq M - 1$, $w_{m-1} = (-1)^{\nu(b_{m-1})} v \oplus X_m$ and compute w_m. Consider the loop $L(a_m, b_m)$. By construction, this loop contains b_{m-1} and u_m, and the left path is the one that contains u_m. Thus, b_{m-1} is either on the left path or on the right path. In the former case, the left path of this loop is $a_m \to b_{m-1}^\ominus \to u_m^\oplus \to b_m^\ominus$ and the right path is $a_m \to b_m$. Since there is also the path $a_{m+1} \to u_m \to b_m$, the message received by b_m from the left is $X_{m+1} \oplus (-1)^{\nu(b_{m-1})} v \oplus X_m$ and the message received from the right is X_m. Thus,

$$w_m = (X_{m+1} \oplus (-1)^{\nu(b_{m-1})} v \oplus X_m) \ominus X_m = X_{m+1} \oplus (-1)^{\nu(b_{m-1})} v.$$

FIGURE 14. Message w_1 with player i on the right path.

FIGURE 15. Message w_m with player b_{m-1} on the left path.

Note that in this case $\nu(b_m) = \nu(b_{m-1})$. See Figure 15 for an illustration.

In the former case, the left path is $a_m \to u_m \to b_m$ and the right path is $a_m \to b_{m-1}^\ominus \to b_m^\ominus$. Since there is also the path $a_{m+1} \to u_m \to b_m$, the message received from the left is $X_{m+1} \oplus X_m$ and the message received from the right is $(-1)^{\nu(b_{m-1})}v \oplus X_m$. Thus $w_m = (X_{m+1} \oplus X_m) \ominus ((-1)^{\nu(b_{m-1})}v) = X_{m+1} \ominus (-1)^{\nu(b_{m-1})}v$. Note that in this case $\nu(b_m) = \nu(b_{m-1}) + 1$. See Figure 16 for an illustration.

Finally, consider the last loop $L(a_M, b_M)$, where $b_M = 0$ is the designer. By construction, this loop does not contain a lock-closer u_{M+1}. One path of this loop goes through b_{M-1}, i.e., we have $a_M \to b_{M-1} \to b_M$, and the other is $a_M \to b_M$. Other players on this loop are transmitters. The designer thus receives w_{M-1} from the first path and X_M from the other. The proof of Lemma 6 is thus complete.   ◁

To complete the proof of Proposition 3, we argue that the message received by each player $j \neq i$ is probabilistically independent from v. This is clearly true for inactive players and for providers. More generally, the only messages that depend on v are those on the directed path from player i to the designer as constructed above, so the statement clearly holds for players outside of this path. Transmitters on this path receive messages

FIGURE 16. Message w_m with player b_{m-1} on the right path.

of the type $X \oplus v$, where X is some random variable independent from v and uniformly distributed. From Lemma 2(iii), this is independent from v. The very same reasoning holds for lock-closers. For lock-openers, this is a consequence of the above computation: since X_m and X_{m+1} are independent and uniformly distributed, so are the two messages received by b_m.   □

COROLLARY 2. *Let $(v_i)_{i \in N}$ be independent random variables such that v_i is known to player i only. A protocol $\mathcal{M}$ exists on $\mathcal{N}$ such that, whenever all players abide by the protocol, the designer correctly learns the value of each v_i. Moreover, the messages received by any player j are probabilistically independent from $(v_i)_{i \neq j}$.*

PROOF. From Proposition 3, for each player i, there exists a protocol (mechanism and strategies) $\mathcal{M}_i$ such that player i can secretly transfer his private information v_i to the designer without revealing information to the other players. The idea is then to concatenate all these protocols "in parallel"; that is, each player j plays a role in each $\mathcal{M}_i$ (inactive, provider, lock-closer, lock-opener, or transmitter) and should play all the corresponding roles simultaneously. For instance, if he is transmitter in several $\mathcal{M}_i$'s, he should forward the corresponding messages on the corresponding links. Moreover, if a player is a provider in one or several $\mathcal{M}_i$'s, the random draws must be mutually independent and independent of messages received.   □

A.4 *Proof of Theorem 1: Sufficiency*

From Corollary 2, there exists a mechanism and a profile of strategies such that if all players follow the prescribed strategies, the designer correctly learns the private information of each player. We now show that in an environment with common independent beliefs and private values, we can indeed provide the players with appropriate incentives to follow the prescribed strategies. Roughly speaking, we make sure that each player is indifferent between all the messages he may send. This is done as follows.

Fix an environment $\mathcal{E}$ with common independent beliefs and private values, and an incentive compatible social choice function f. Let P^i denote the marginal distribution of the common belief P on Θ_i, i.e., this is the common belief of any player $j \neq i$ on Θ_i. Without loss of generality, assume that $\Theta_i := \{1, \ldots, t_i, \ldots, T_i\}$ for each player $i \in N$ and let $\overline{P}^i(t_i) = \sum_{\theta_i \leq t} P^i(\theta_i)$ denote the cumulative distribution function of P^i. Define a partition $\Pi_i = \{\Pi_i(1), \ldots, \Pi_i(T_i)\}$ of $[0, 1)$ into T_i subsets with $\Pi_i(t_i) = [\overline{P}^i(t_i - 1), \overline{P}^i(t_i))$ (with $\overline{P}^i(0) = 0$). Note that if X is uniformly distributed on $[0, 1)$, the event $\{X \in \Pi_i(t_i)\}$ has probability $P^i(t_i)$.

Part I   We first consider the problem of implementing the social choice function f_i^* for which player i is dictatorial, i.e., for any θ_i, define $f_i^*(\theta_i) \in \arg\max_{a \in A} u_i(a, \theta_i)$ and let $f_i^*(\theta_i, \theta_{-i}) = f_i^*(\theta_i)$ for all θ_{-i}. If $i \in D(0)$, f_i^* is clearly implementable. Assume that $i \notin D(0)$. We claim that the protocol $\mathcal{M}_i$ implies the existence of a mechanism and strategies such that player i has an incentive to truthfully reveal his type and no other active player has an incentive to manipulate the transmission of information from player i to the designer.

The mechanism and strategies are as follows.

- Player i of type t_i draws a random number v_i uniformly in $\Pi_i(t_i)$ and transmits it to the designer by the protocol $\mathcal{M}_i$.

- All other active players follow the strategies constructed in $\mathcal{M}_i$.

- Let $\hat{v}_i$ be the message decoded by the designer and denote $\hat{\theta}_i = t_i$ if $\hat{v}_i \in \Pi_i(t_i)$. (See Lemma 6.) The designer implements the alternative $f_i^*(\hat{\theta}_i)$.

First, observe that the protocol $\mathcal{M}_i$ implies that each active player sends a real number in $[0, 1)$. Second, observe that the unconditional distribution of v_i is the uniform distribution on $[0, 1)$. To see this, let $X_i^{t_i}$ denote a random variable uniformly distributed on $\Pi_i(t_i)$ and observe that $v_i = \sum_{t_i=1}^{T_i} \mathbf{1}_{\{\theta_i = t_i\}} X_i^{t_i}$. From Proposition 3, it follows that the designer correctly learns the type of player i if all players abide by the protocol $\mathcal{M}_i$, while no player gets additional information about the type of player i (posterior beliefs are equal to prior beliefs). So the expected payoff of any active player $j \neq i$ of type θ_j is $\sum_{\theta_i} u_j(f_i^*(\theta_i), \theta_j) P^i(\theta_i)$.

Third, we show that no active player has an incentive to deviate. This is clearly true for player i, as f_i^* is incentive compatible. Consider player $j \neq i$ and suppose that j is a transmitter in the loop $L(a_m, b_m)$ for $m = 2, \ldots, M - 1$. There are several cases to consider.

Case 1. Player j is on the right path of the loop $L(a_m, b_m)$ from player a_m to player b_m and moves before the lock-closer u_{m-1}. Under $\mathcal{M}_i$, he receives the message x_m. Suppose that he deviates and sends the message x_m'. It follows that the designer will receive the messages $(-1)^{\nu(b_{M-1})}(v \oplus x_m' \ominus x_m) \oplus X_M$ and X_M under the deviation, so that the decoded message is $v \oplus x_m' \ominus x_m$. Since v is uniformly distributed on $[0, 1)$, it follows that the probability that $v \oplus x_m' \ominus x_m$ is in $\Pi_i(t_i)$ is $P^i(t_i)$, regardless of x_m' (see Lemma 2(ii)). Player j is thus indifferent between sending x_m and x_m'.

Case 2. Player j is on the right path of the loop $L(a_m, b_m)$ from player a_m to player b_m and moves after the lock-closer u_{m-1}, but before the lock-opener b_{m-1}. Under $\mathcal{M}_i$, player j receives the message $x_m \oplus x_{m-1}$ from the lock-closer u_{m-1}. Suppose that he deviates and sends the message x'_m. It follows that the designer will receive the messages $(-1)^{\nu(b_{M-1})}(v \oplus x'_m \ominus x_m \ominus x_{m-1}) \oplus X_M$ and X_M under the deviation. Since all random variables are uniformly distributed on $[0, 1)$, so is their addition $\oplus$ or subtraction $\ominus$ (this follows from Lemma 2), and, consequently, player j is indifferent between sending $x_m \oplus x_{m-1}$ and x'_m.

Case 3. Player j is on the right path of the loop $L(a_m, b_m)$ from player a_m to player b_m, and moves after the lock-closer u_{m-1} and the lock-opener b_{m-1}. Under $\mathcal{M}_i$, player j receives the message $(-1)^{\nu(b_{m-1})}v \oplus x_m$. Note that j does not learn the value of x_m and believes that it is a realization of X_m. Suppose that he deviates and sends the message x'_m. It follows that the designer will receive the messages $(-1)^{\nu(b_{M-1})}(x'_m \ominus x_m) \oplus X_M$ and X_M under the deviation. Since X_m and X_M are uniformly distributed on $[0, 1)$, it follows yet again that player j evaluates the probability of $\hat{v}_i = x'_m \ominus x_m \in \Pi_i(t_i)$ to be $P^i(t_i)$ and, thus, is again indifferent between reporting the truth and deviating.

Case 4. Player j is on the left path of the loop $L(a_m, b_m)$ from player a_m to player b_m and moves before the lock-closer u_m. This case is similar to Case 1.

Case 5. Player j is on the left path of the loop $L(a_m, b_m)$ from player a_m to player b_m and moves after the lock-closer u_m. In that case, player j also belongs to the right path of the loop $L(a_{m+1}, b_{m+1})$ and the same arguments as in Case 1 apply.

Last, a similar reasoning applies if player j is a transmitter in the first or last loop. For instance, if player j is on the right path of the last loop $L(a_M, b_M)$ and moves before the lock-closer u_M, the same reasoning as in Case 1 applies since the designer receives the message $(-1)^{\nu(b_{M-1})}v \oplus x'_M$ and X_M.

Now, suppose that player j is the provider a_m in the loop $L(a_m, b_m)$ $(m < M)$, and suppose that he sends the message x_m^L on the left path of the loop and the message x_m^R on the right path. If all other players abide by the strategies, it follows that the designer receives the messages $(-1)^{\nu(b_{M-1})}(v \oplus x_m^R \ominus x_m^L) \oplus X_M$ and X_M. Since v and X_M are uniformly and independently distributed on $[0, 1)$, it follows that the probability that the decoded type $\hat{v}_i$ is in $\Pi_i(t_i)$ is $P^i(t_i)$ and, thus, player j is indifferent between following the prescribed strategy or deviating.

Similar arguments apply to the lock-closers or lock-openers, so that the prescribed strategies indeed form a Bayesian equilibrium. To summarize, incentive compatibility of the social choice function implies that player i has indeed an incentive to abide by the protocol $\mathcal{M}_i$, while all other active players have no incentive to deviate, since the protocol guarantees the same expected payoff to each active player other than player i, regardless of the message he sends.

Part II   Let f be a social choice function implementable on $\mathcal{N}^\star$, i.e., f is incentive compatible. To implement f, consider the mechanism and strategies implied by the protocol $\mathcal{M}$: each player $i \notin D(0)$ of type t_i draws a random number v_i uniformly in $\Pi_i(t_i)$ and transmits it to the designer according to the protocol $\mathcal{M}_i$, while in his role of an active player in a protocol $\mathcal{M}_j$ $(j \neq i)$, he follows the prescribed strategy.

From Corollary 2, it follows that the designer learns the true profile of types if all players abide by this protocol, while no player gets additional information about the type of his opponents. To complete the proof, note that as in Part I, no player has an incentive to deviate. The expected payoff of a player i is independent of the messages he sends about his opponents (since the assumption of independent beliefs implies that we can consider each deviation as above). Incentive compatibility guarantees that player i has indeed an incentive to abide by the subprotocol $\mathcal{M}_i$. The proof of the sufficiency part of Theorem 1 is thus complete.

A.5 *Proof of Theorem 1: Necessity*

Now, we prove the "only if" part of Theorem 1. The proof proceeds by contradiction. We assume that $\mathcal{N}$ is not weakly 2-connected, and we construct an environment with common independent belief and private values and an incentive compatible social choice function, which is not implementable on $\mathcal{N}$.

If $\mathcal{N}$ is not weakly 2-connected, there exist two distinct players i and i^* such that all paths, directed or undirected, from i to the designer go through i^*. As a consequence, for each player k who has a path to i, directed or undirected, all paths from k to 0 also go through i^*. This implies that player i^* is a cut vertex in the network. In particular, all information regarding the players k who have a path to i is controlled by i^*.

Let us now construct the environment and the social choice function. Assume that all players but player i have a single type and that player i has two types θ_i and θ_i'. Let a and b be two alternatives. The utilities are $u_i(a, \theta_i) = u_{i^*}(a, \cdot) = 1$, $u_i(b, \theta_i) = u_{i^*}(b, \cdot) = 0$ and $u_i(a, \theta_i') = 0$, $u_i(b, \theta_i') = 1$. All other players are indifferent (get a utility of 0) between a and b. Any other alternative gives a utility of -1 to players i and i^* regardless of their types. The common prior is the uniform distribution on the set of types. The social choice function is the dictatorial social choice function of player i.

We claim that for every mechanism on $\mathcal{N}$, there is no equilibrium that implements this social choice function. By contradiction, assume that there exists such an equilibrium σ. Fix a profile of messages $\bar{m}_{i^*} \in M_{D(i^*)}$ for player i^* in the support of $\mathbb{P}_{\theta_i, \sigma}$, i.e., this is a message compatible with θ_i and the equilibrium strategies. Consider the deviation σ_{i^*}' for player i^* that consists of playing $\sigma_{i^*}(\bar{m}_{i^*})$ regardless of his type and messages received.

By construction of the deviation, $\sigma_{i^*}(\bar{m}_{i^*})$ is compatible with the messages sent by players who have no path to player i, i.e.,

$$\operatorname{supp} \mathbb{P}_{\theta, \sigma_{i^*}', \sigma_{-i^*}} \subseteq \operatorname{supp} \mathbb{P}_{\theta_i, \sigma} \quad \forall \theta \in \{\theta_i, \theta_i'\}.$$

Since the strategies are assumed to implement f, it follows that the outcome is almost surely a under the deviation, regardless of the type of player i. Since player i^* prefers a to any other alternative, this deviation is profitable for player i^*.

It is worthwhile noting that weak 2-connectedness is also a necessary condition for Proposition 3 to hold. Indeed, if i^* is a cut vertex and if the designer learns the type of player i, then i^* must learn it as well.

A.6 *Proof of Theorem 2*

The proof of the "only if part" is identical to the necessity part of the proof of Theorem 1 and is omitted. We turn to the "if" part and fix an environment with a worst outcome and an incentive compatible social choice function f. Without loss of generality, let us assume that f does constantly select the worst outcome (if so, the designer just has to choose the worst outcome irrespective of the messages received). Also, without loss of generality, assume that Θ_i is a finite subset of the open interval $(0, 1)$ for each player $i \in N$. In the proof of Theorem 1, we took advantage of the environment to make players indifferent between any message they can send. This is no longer possible in environments with correlated beliefs and/or common values. We thus modify the protocol in such a way that deviations are detected with arbitrarily high probability by the designer. The threat of the worst outcome then deters the players from deviating.

Let η be a large integer. We employ the terminology and notations from Proposition 3 and modify the protocol $\mathcal{M}_i$ as follows.

- Each transmitter forwards the message received from his active predecessor to his active successor.

- Each provider a_m draws an η-vector of keys $\mathbf{X}_m = (X_m^1, \ldots, X_m^\eta)$, the components of which are independently and uniformly distributed in $[0, 1)$, and sends it to its two active successors.

- Each lock-closer u_m receives two vectors $\mathbf{x}_m$ and $\mathbf{x}_{m+1}$ from his predecessors. He computes $\mathbf{z}_m = \mathbf{x}_m \oplus \mathbf{x}_{m+1}$ and sends it to his active successor, where $\oplus$ denotes componentwise addition.

- Each lock-opener b_m receives two vectors $\mathbf{x}_m^L$ and $\mathbf{x}_m^R$ from his predecessors. He computes $\mathbf{w}_m = \mathbf{x}_m^L \ominus \mathbf{x}_m^R$ and sends it to his active successor.

Player i behaves as follows (recall that by construction, player i is either a transmitter or a provider).

- If he is a transmitter, player i who receives $\mathbf{x}_1$ from his active predecessor uniformly draws a random integer η^* in $\{1, \ldots, \eta\}$ and encodes his type θ_i with the encoding key $x_1^{\eta^*}$ to obtain the cypher type $y_1^{\eta^*}(i) = \theta_i \oplus x_1^{\eta^*}$. Player i then sends the vector $(x_1^1, \ldots, x_1^{\eta^*-1}, y_1^{\eta^*}(i), x_1^{\eta^*+1}, \ldots, x_1^\eta)$ to his active successor.

- If he is a provider, player i draws (uniformly) a random vector $\mathbf{X}_1$ and a random integer η^* in $\{1, \ldots, \eta\}$ and computes $Y_1^{\eta^*}(i) = \theta_i \oplus X_1^{\eta^*}$. Player i then sends the vector $(X_1^1, \ldots, X_1^{\eta^*-1}, Y_1^{\eta^*}(i), X_1^{\eta^*+1}, \ldots, X_1^\eta)$ to his left active successor and sends $\mathbf{X}_1$ to his right active successor.

The decision rule of the designer is the following. The designer receives a message $\mathbf{x}_M^R$ from the path $a_M \to b_{M-1} \to b_M = 0$ and a message $\mathbf{x}_M^L$ from the other path of the last loop $a_M \to b_M = 0$.

- If the vectors $\mathbf{x}_M^L$ and $\mathbf{x}_M^R$ differ by exactly one component η^*, the designer decodes $\hat{\theta}_i = x_M^{\eta^*,R} \ominus x_M^{\eta^*,L}$ if $\nu(b_{M-1})$ is even and $\hat{\theta}_i = x_M^{\eta^*,L} \ominus x_M^{\eta^*,R}$ if $\nu(b_{M-1})$ is odd.

- Otherwise, the designer concludes that there was a deviation.

Note that no player $j \neq i$ gains information about θ_i by this modified mechanism. Indeed, player j observes only vectors of uniformly distributed numbers. If all players abide by the mechanism, then the two vectors received by the designer differ only in the component η^*, and the designer correctly decodes the type of player i from Lemma 6. The key argument is that η^* is the private information of player i. Thus, any deviation by an active player is bound to change another component with probability at least $1 - 1/\eta$.

Finally the mechanism for implementing f is the following.

- Each player i transmits his type to the designer using the modified protocol.

- If the designer concludes that there was no deviation, he implements $f(\hat{\theta}_1, \ldots, \hat{\theta}_n)$, where $\hat{\theta}_i$ is the decoded type of player i.

- Otherwise, the designer implements the worst outcome.

Let us check the equilibrium condition. The expected payoff of j under the mechanism is

$$\sum_{\theta_{-j}} u_j(f(\theta_j, \theta_{-j}), \theta_j, \theta_{-j}) P_j(\theta_{-j}|\theta_j) := C.$$

Assume that player j deviates in at least one submechanism. His expected payoff is at most

$$\frac{1}{\eta} W + \left(1 - \frac{1}{\eta}\right) \sum_{\theta_{-j}} u_j(\underline{a}, \theta_j, \theta_{-j}) P_j(\theta_{-j}|\theta_j) := D,$$

where W is an upper bound on player j's payoff. We have

$$C - D = \frac{1}{\eta}(C - W) + \left(1 - \frac{1}{\eta}\right) \sum_{\theta_{-j}} \left(u_j(f(\theta_j, \theta_{-j}), \theta_j, \theta_{-j}) - u_j(\underline{a}, \theta_j, \theta_{-j})\right) P_j(\theta_{-j}|\theta_j).$$

Since $\underline{a}$ is a worst outcome, $u_j(f(\theta_j, \theta_{-j}), \theta_j, \theta_{-j}) - u_j(\underline{a}, \theta_j, \theta_{-j})$ is nonnegative for all type profiles and is strictly positive for at least one type profile, as f is not constantly equal to $\underline{a}$. Recall that we assumed throughout that beliefs have full support, i.e., $P_j(\theta_{-j}|\theta_j) > 0$ for all type profiles. As a consequence, $C - D$ is positive for η large enough and player j has no incentive to deviate. Last, each player i has an incentive to transmit his true type since f is incentive compatible.

A.7 *Proof of Theorem 3*

The proof is very similar to the proofs of Theorems 1 and 2. The proof that the condition is necessary is the same. For sufficiency, the main task is to extend Proposition 3 to weakly 2-connected networks with cycles. Once this is established, Theorem 3 follows similarly as for Theorems 1 and 2, so this part of the proof is omitted.

We now explain how to extend Proposition 3. A important remark is that since the network has cycles, the existence of the timing structure is no longer guaranteed; in

fact, it simply fails. To define a mechanism, one has to specify a timing structure, i.e., who speaks first, who speaks second, and so on. To avoid this difficulty, we associate to the network $\mathcal{N}$, an augmented network $\mathcal{N}^A$, which is strongly 1-connected, weakly 2-connected, and acyclic. Thus, Proposition 3 holds true on $\mathcal{N}^A$. Then we show how the protocol on $\mathcal{N}^A$ induces the desired protocol on $\mathcal{N}$.

Let us fix a strongly 1-connected and weakly 2-connected network $\mathcal{N}$ (but not necessarily acyclic). Recall that a network is a set of edges. A subnetwork is thus a subset of edges.

LEMMA 7. *There exists an acyclic and strongly 1-connected subnetwork $\mathcal{N}^a$ of $\mathcal{N}$.*

PROOF. For each $i \in N$, consider a shortest directed path from i to 0 in $\mathcal{N}$. Such a shortest directed path exists since $\mathcal{N}$ is strongly 1-connected. Let $\mathcal{N}^a$ be the collection of all these paths. We claim that $\mathcal{N}^a$ has the required properties. By construction, it is strongly 1-connected. Let us show that it is acyclic. By contradiction, assume that $\mathcal{N}^a$ contains the cycle $i_1 \to i_2 \to \cdots \to i_K \to i_1$. By construction, $\mathcal{N}^a$ is such that $C(0) = \varnothing$, i.e., there is no edge $0i$ for some $i \in N$ in $\mathcal{N}^a$. It follows that the cycle does not contain the designer (player 0). It then follows that there exists $k \in \{2, \ldots, K\}$ such that the shortest path from i_k to 0 does not follow the cycle (otherwise, 0 cannot be reached, a contradiction with strong 1-connectedness). Thus, the edge $i_k i_{k+1}$ is not on a shortest path from any player j to 0, contradicting the construction of $\mathcal{N}^a$.   ◁

With a slight abuse of notation, let $\mathcal{N}^a$ be a maximal acyclic and strongly 1-connected subnetwork of $\mathcal{N}$ (it exists by the preceding lemma), and let $\mathcal{C} = \mathcal{N} \setminus \mathcal{N}^a$ be the set of edges of $\mathcal{N}$ that do not belong to $\mathcal{N}^a$. Note that every edge of $\mathcal{C}$ belongs to a cycle of $\mathcal{N}$ and that every cycle of $\mathcal{N}$ contains an edge in $\mathcal{C}$. Let $\mathcal{N}^A$ be the network obtained from $\mathcal{N}$ by replacing each edge ij in $\mathcal{C}$ by two edges $i(j)i$ and $i(j)j$, where $i(j)$ is a fictitious player who is a duplicate of player i; that is, for ij in $\mathcal{C}$,

$$i \to j \text{ is replaced by } i \leftarrow i(j) \to j.$$

The edges of $\mathcal{N}^a$ are unchanged. See Figure 17 for an example.

CLAIM 3. *The network $\mathcal{N}^A$ is strongly 1-connected, weakly 2-connected, and acyclic.*

PROOF. Each "regular" player i has a directed path to 0 in $\mathcal{N}^a$ by construction. Since the fictitious player $i(j)$ is directly connected to i, he also has a path to the designer by strong 1-connectedness of $\mathcal{N}$. Weak 2-connectedness is clearly preserved by the transformation. Let us show that $\mathcal{N}^A$ is acyclic. Assume that $\mathcal{N}^A$ contains a cycle. By our construction, each fictitious player has only outgoing edges, thus cannot belong to a cycle. This implies that the cycle was already a cycle in $\mathcal{N}$ and, therefore, it should contain an edge that belongs to $\mathcal{C}$. This is a contradiction because edges in $\mathcal{C}$ no longer appear in $\mathcal{N}^A$.   ◁

Now, we claim that Proposition 3 extends to strongly 1-connected, weakly 2-connected networks with cycles. First, on the network $\mathcal{N}^A$, for each player i, there exists

FIGURE 17. A cyclic network $\mathcal{N}$ and the associated acyclic network $\mathcal{N}^A$.

a protocol with the desired property by Proposition 3. We assume that each fictitious player has no type and a constant payoff function. Second, on the network $\mathcal{N}$, the players can replicate this protocol. The timing of the protocol is that given by the timing structure of $\mathcal{N}^A$, which is well defined since $\mathcal{N}^A$ is acyclic and strongly 1-connected. In particular, each duplicated player i plays only twice: he plays as the fictitious player $i(j)$ the first time and as player i the second time.

Thus, Proposition 3 extends and Theorem 3 follows, similarly as for Theorems 1 and 2.

A.8 *Detection with probability 1*

LEMMA 8. *Let v be a random variable privately known by player i. If the network is weakly 2-connected, there exists a mechanism $\mathcal{M}_i$ on $\mathcal{N}$ such that if all players abide by the mechanism, then the designer learns the value of v, whereas each player $j \neq i$ receives messages that are probabilistically independent from v. Furthermore, the designer detects deviations with probability 1.*

The intuition is as follows. For each integer η, we can devise a test such that any deviation is detected with probability at least $1 - 1/\eta$. We may thus ask the players to pass *all such tests*.[18] There are several ways to construct such a test and we provide a relatively simple one. We modify our protocol M_i as follows. For simplicity, we assume that player i is not a provider.

- *Providers.* Each provider a_m draws two independent infinite sequences $(X_\eta^{m,H}, X_\eta^{m,T})_{\eta \geq 1}$ of independently and identically (i.i.d.) distributed random variables, with uniform distribution on $[0, 1)$ and sends these sequences.

- *Player i.* Independently of his type and of the message he receives, player i draws an infinite sequence of i.i.d. fair coins $c_\eta \in \{H, T\}$. Define $(Y_\eta^H, Y_\eta^T)_{\eta \geq 1}$ as $(Y_\eta^H, Y_\eta^T) = (X_\eta^{1,H} \oplus \theta_i, X_\eta^{1,T})$ if $c_\eta = H$ and as $(Y_\eta^H, Y_\eta^T) = (X_\eta^{1,H}, X_\eta^{1,T} \oplus \theta_i)$ if

---

[18]We thank Sylvain Sorin for suggesting this argument.

$c_\eta = T$. In words, for each η, player i chooses according to the toss of a fair coin whether to encode his type θ_i with $X_\eta^{1,H}$ or with $X_\eta^{1,T}$. Player i then sends the pair of sequences $(Y_\eta^H, Y_\eta^T)_{\eta \geq 1}$ to his active successor.

- *Other players.* The other active players (transmitters, lock-closers, and lock-openers) behave as in the proof of Theorem 2, except that now vectors are sequences.

- *The designer.* The designer receives two pairs of sequences $(x_\eta^{L,H}, x_\eta^{L,T})_{\eta \geq 1}$ and $(x_\eta^{L,H}, x_\eta^{R,T})_{\eta \geq 1}$. If for each η, it holds true that $(x_\eta^{L,H} = x_\eta^{R,H}$ and $x_\eta^{L,T} \neq x_\eta^{R,T})$ or $(x_\eta^{L,H} \neq x_\eta^{R,H}$ and $x_\eta^{L,T} = x_\eta^{R,T})$, the designer concludes that phase 1 of the test succeeds. Then if $x_\eta^{L,T} \neq x_\eta^{R,T}$, he computes $\hat{\theta}_i = x_M^{\eta^*,R,T} \ominus x_M^{\eta^*,L,T}$ if $\nu(b_{M-1})$ is even and $\hat{\theta}_i = x_M^{\eta^*,L,T} \ominus x_M^{\eta^*,R,T}$ if $\nu(b_{M-1})$ is odd. If $x_\eta^{L,H} \neq x_\eta^{R,H}$, he computes $\hat{\theta}_i = x_M^{\eta^*,R,H} \ominus x_M^{\eta^*,L,H}$ if $\nu(b_{M-1})$ is even and $\hat{\theta}_i = x_M^{\eta^*,L,H} \ominus x_M^{\eta^*,R,H}$ if $\nu(b_{M-1})$ is odd. If all $\hat{\theta}_i^\eta$ have the same value $\hat{\theta}_i$, the designer concludes that phase 2 of the test succeeds, and regards $\hat{\theta}_i$ as the correct type of player i. If the test does not succeed in either phase 1 or phase 2, the designer concludes that there was a deviation.

Under these strategies, the decoded type clearly coincides with the true type. It is also clear that no player gets information about the message of player i. The sequence of coins being privately known to player i, each other active player observes only sequences of i.i.d. uniformly distributed variables. Now we claim that any deviation is detected almost surely. Indeed, if some active player $j \neq i$ modifies the sequence, to pass the test in phase 2, he must modify an entry of the double sequence for each η. But then, to succeed in phase 1, he should modify only the component selected by player i. Consequently, the probability of passing the test while changing the message is at most the probability of guessing correctly an infinite sequence of fair coins, which is 0. Any deviation is thus detected with probability 1.

COROLLARY 3. *If the network is weakly 2-connected and if the environment has a bad outcome, i.e., an outcome $\underline{a}$ such that $u_i(a, \theta) \geq u_i(\underline{a}, \theta)$ for all $i \in N$, for all $a \in A$, for all $\theta \in \Theta$, then $F_\mathcal{N}(\mathcal{E}) = F_{\mathcal{N}^\star}(\mathcal{E})$.*

The proof consists in adapting the construction of Theorem 2. Using the above lemma, any deviation brings the bad outcome almost surely and is, therefore, not profitable.

REFERENCES

Bárány, Imre (1992), "Fair distribution protocols or how the players replace fortune." *Mathematics of Operations Research*, 17, 327–340. [492]

Beimel, Amos and Matthew Franklin (1999), "Reliable communication over partially authenticated networks." *Theoretical Computer Science*, 220, 185–210. [492]

Ben-Porath, Elchanan (2003), "Cheap talk in games with incomplete information." *Journal of Economic Theory*, 108, 45–71. [492]

Bergemann, Dirk and Stephen Morris (2005), "Robust mechanism design." *Econometrica*, 73, 1771–1813. [508]

Bergemann, Dirk and Stephen Morris (2009), "Robust implementation in direct mechanisms." *Review of Economic Studies*, 76, 1175–1204. [510]

Bollobás, Béla (1998), *Modern Graph Theory*. Springer, New York. [515]

Bolton, Patrick and Mathias Dewatripont (1994), "The firm as a communication network." *Quarterly Journal of Economics*, 109, 809–839. [490, 491]

Calvo, Guillermo A. and Stanislaw Wellisz (1978), "Supervision, loss of control and the optimal size of the firm." *Journal of Political Economy*, 86, 943–952. [490]

Dasgupta, Partha, Peter Hammond, and Eric Maskin (1979), "The implementation of social choice rules: Some general results on incentive compatibility." *Review of Economic Studies*, 46, 185–216. [489]

Desmedt, Yvo and Yongge Wang (2002), "Perfectly secure message transmission revisited." In *Advances in Cryptology* (Lars Knudsen, ed.), 502–517, Springer, Berlin. [505]

Diffie, Whitfield and Martin E. Hellman (1976), "New directions in cryptography." *IEEE Transactions on Information Theory*, 22, 644–654. [504]

Dolev, Danny, Cynthia Dwork, Orli Waarts, and Moti Yung (1993), "Perfectly secure message transmission." *Journal of the Association for Computing Machinery*, 40, 17–47. [505]

Forges, Françoise (1990), "Universal mechanisms." *Econometrica*, 58, 1341–1364. [492]

Franklin, Matthew and Rebecca N. Wright (2000), "Secure communication in minimal connectivity models." *Journal of Cryptology*, 13, 9–30. [505]

Friebel, Guido and Michael Raith (2004), "Abuse of authority and hierarchical communication." *RAND Journal of Economics*, 35, 224–244. [492]

Gerardi, Dino (2004), "Unmediated communication in games with complete and incomplete information." *Journal of Economic Theory*, 114, 104–131. [492]

Gibbard, Allan (1973), "Manipulation of voting schemes: A general result." *Econometrica*, 41, 587–601. [489]

Goldwasser, Shafi and Silvio Micali (1984), "Probabilistic encryption." *Journal of Computer and Systems Sciences*, 28, 270–299. [505]

Harris, Milton and Robert M. Townsend (1981), "Resource allocation under asymmetric information." *Econometrica*, 49, 33–64. [489]

Jackson, Matthew O. (2001), "A crash course in implementation theory." *Social Choice and Welfare*, 18, 655–708. [512]

Krishna, Vijay (2002), *Auction Theory*. Academic Press, San Diego. [499]

Monderer, Dov and Moshe Tennenholtz (1999), "Distributed games: From mechanisms to protocols." In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, 32–37, American Association for Artificial Intelligence, Menlo Park, California. [492, 500, 502]

Myerson, Roger B. (1979), "Incentive compatibility and the bargaining problem." *Econometrica*, 47, 61–73. [489]

Myerson, Roger B. (1982), "Optimal coordination mechanisms in generalized principal–agent problems." *Journal of Mathematical Economics*, 10, 67–81. [489, 508]

Nisan, Noam and Ilya Segal (2006), "The communication requirements of efficient allocation and supporting prices." *Journal of Economic Theory*, 129, 192–224. [490]

Nisan, Noam, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, eds. (2007), *Algorithmic Game Theory*. Cambridge University Press, Cambridge. [490]

Rabin, Tal and Michael Ben-Or (1989), "Verifiable secret sharing and multiparty protocols with honest majority." In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, 73–85, Association for Computing Machinery, New York. [507]

Radner, Roy (1993), "The organization of decentralized information processing." *Econometrica*, 61, 1109–1146. [490]

Renault, Jérôme and Tristan Tomala (2008), "Probabilistic reliability and privacy of communication using multicast in general neighbor networks." *Journal of Cryptology*, 21, 250–279. [505]

Renou, Ludovic and Tristan Tomala (2010), "Mechanism design and communication networks." Unpublished paper. [512]

Rivest, Ronald, Adi Shamir, and Leonard Adleman (1978), "A method for obtaining digital signatures and public-key cryptosystems." *Communications of the ACM*, 21, 120–126. [504]

Salanié, Bernard (2005), *The Economics of Contracts*, second edition. MIT Press, Cambridge, Massachusetts. [499]

Serrano, Roberto and Rajiv Vohra (2010), "Multiplicity of mixed equilibria in mechanisms: A unified approach to exact and approximate implementation." *Journal of Mathematical Economics*, 46, 775–785. [512]

Shannon, Claude E. (1949), "Communication theory of secrecy systems." *Bell System Technical Journal*, 28, 656–715. [505]

Van Zandt, Timothy (2007), "Communication complexity and mechanism design." *Journal of European Economic Association*, 5, 543–553. [490]

Williamson, Oliver E. (1967), "Hierarchical control and optimum firm size." *Journal of Political Economy*, 75, 123–138. [490]