Communication on networks and strong reliability

Marie Laclau, Ludovic Renou, Xavier Venel

COMMUNICATION ON NETWORKS AND STRONG RELIABILITY

MARIE LACLAU, LUDOVIC RENOU, AND XAVIER VENEL

ABSTRACT. We consider sender–receiver games, where the sender and the receiver are two distant nodes in a communication network. We show that if the network has two disjoint paths of communication between the sender and the receiver, then we can replicate all equilibrium outcomes not only of the direct communication game (i.e., when the sender and the receiver communicate directly with each other) but also of the mediated game (i.e., when the sender and the receiver communicate with the help of a mediator).

KEYWORDS: Cheap talk, direct, mediated, communication, protocol, network.

JEL CLASSIFICATION: C72; D82.

Date: February 22, 2024.

This paper previously circulated under the title: "Robust communication on networks." The authors gratefully thank Françoise Forges, as well as the associate editor and two anonymous referees at JET. The authors also acknowledge the support of the Agence Nationale pour la Recherche under grant ANR CIGNE (ANR-15-CE38-0007-01) and through the ORA Project "Ambiguity in Dynamic Environments" (ANR-18-ORAR-0005). Laclau gratefully acknowledges the support of the ANR through the program Investissements d'Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047) and under grant ANR StratCom (ANR-19-CE26-0010-01). Venel is a member of GNAMPA-INdAM. This research project received partial support from the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets" and the MIUR PRIN 2022 project "Learning in Markets and Society" (project 2022EKNE5K)..

1. INTRODUCTION

We study sender–receiver games in communication networks, where the information from the sender to the receiver may have to flow through self-interested intermediaries. The following question is addressed: When is it possible to emulate direct or even mediated (i.e., with a mediator) communication between the sender and the receiver as indirect but unmediated (i.e., without a mediator) communication?

More precisely, we consider a sender–receiver game and a communication equilibrium of that game, that is, an equilibrium where the sender and the receiver can communicate privately and securely with a trusted mediator (see Forges, 1986; Myerson, 1986). Now, we assume that no trusted mediator exists but that there exists a communication network, with the sender and the receiver as two distant nodes – the other nodes are intermediaries. The communication network models the communication possibilities, i.e., who can send a message to whom, and thus naturally induces a multistage communication game. The following question is addressed: When are we guaranteed to have a perfect Bayesian equilibrium (henceforth, PBE) of the multistage game that replicates the communication equilibrium? By "guaranteed," we mean that the answer should be independent of the preferences of the intermediaries, and of the selected communication equilibrium.

We prove that we are guaranteed to replicate all communication equilibria as PBE of multistage games on communication networks if, and only if, there exist (at least) two disjoint paths of communication between the sender and the receiver. The central insight of our analysis is the tight connection with the concept of strong reliability-a close relative of the concept of reliability in computer science. To define the concept of strong reliability, we consider the following problem: the sender wishes to transmit a message to the receiver through the communication network. The problem is to construct a communication protocol, with the property that the receiver correctly learns the message sent at all histories consistent with at most k intermediaries deviating at every round of the protocol. (Different intermediaries can deviate in different rounds.) We call this property: k-strong reliability. It guarantees that the transmission of the message can tolerate some mistakes, errors or even deliberate disruptions to communication.

We show that 1-strong reliability (for short, strong reliability) is possible on a network if, and only if, there are two disjoint paths of communication between the sender and the receiver. Thus, the PBE implementation of all communication equilibria on networks

and the strong reliability of communication on networks are equivalent, that is, if the former is guaranteed on the network $\mathcal{N}$, so is the latter, and the reverse is also true.

We now provide some intuition for this tight connection between PBE implementation of all communication equilibria and strong reliability. There is a sender and a receiver, two states and two actions, and the sender and the receiver want to match the state. Clearly, there exists an equilibrium of the direct communication game, where the sender truthfully reveals the state and the receiver matches the state. We assume now that the communication between the sender and the receiver is intermediated, with all intermediaries preferring the receiver to un-match the state. The equilibrium distribution (of the direct communication game) is therefore the least preferred outcome for any intermediary, while any other distribution is strictly preferred. Thus, if an intermediary deviates at any stage of the indirect multistage communication game, the deviation must *not* change the distribution; otherwise, this would be a profitable deviation for the intermediary. In any information set that follows the deviation, *sequential rationality* dictates that no other intermediary has a profitable deviation either, i.e., no other intermediary must be able to change the distribution. Therefore, the correct distribution must be implemented not only for all on-path histories but also for histories that are reachable via sequences of unilateral deviations. (We stress that this is a consequence of imposing sequential rationality, which is a novelty of this paper, see the review of the literature.) This is what strong reliability achieves.

Two additional observations are worth noting. First, if we restrict attention to the (unmediated) equilibria of sender–receiver games (i.e., when the sender communicates with the receiver), the same condition on the network is necessary and sufficient. Thus, the connectivity requirement does not increase as we move from the (unmediated) equilibria to the communication equilibria. This is particularly important because some communication equilibria can Pareto-improve over all (unmediated) equilibria. In these instances, indirect communication dominates direct communication. In other words, communicating through layers of intermediaries may actually benefit the sender and the receiver (compared with direct communication). Second, the protocols we construct require several rounds of communication and rich communication possibilities. In particular, players must be able to *broadcast* messages to any subset of their neighbors. Broadcasting a message to a group ensures that group members have common knowledge of the message. It is a very natural assumption: face-to-face meetings and online

meetings via platforms such as Zoom, Microsoft Teams, and WhatsApp groups all make it possible to broadcast to a group.

Our interpretation of the communication network is as follows. In modern organizations, most employees, from top executives to low-level managers, devote a significant fraction of their time to internal communication: they draft and circulate memos, attend and call meetings, write e-mails, etc. The network $\mathcal{N}$ captures these communication possibilities, particularly who can call a meeting with whom. If there is a link between players i and j and between players i and k, player i can communicate with players j and k both privately (face-to-face meetings) and publicly (group meetings). In organizations, meetings serve several functions, from communicating information to making decisions and generating ideas. The former, i.e., meetings used as an information forum, is the closest to the role that meetings play in our analysis. When player i broadcasts a message to players j and k, player i informs players j and k.

We conclude with an insight of our analysis for organization theory. In large organizations, such as public administrations, governments, armed forces and multinational corporations, information typically flows through the different layers of the organization, from engineers, sale representatives, and accountants to top managers and executives. Communication is indirect. While indirect communication is necessary in large organizations, it may harm the effective transmission of valuable information. Indeed, as the objectives of the members of an organization are rarely perfectly aligned, distorting, delaying or even suppressing the transmission of information are natural ways, among others, for members of the organization to achieve their own objectives. Do organizational arrangements exist that mitigate these issues? An insight of this paper is that *matrix organization* is one such arrangement.[1]

Matrix organization (or management) consists of organizing activities along more than one dimension, e.g., function (marketing, accounting, engineering, R&D, etc.), geography (US, Europe, Asia, etc.) or products. From the early 1960s to the present, large corporations such as NASA, IBM, Pearson, Siemens and Starbucks have adopted this mode of management. A central feature of matrix organization is *multiple* reporting lines; that is, low-level employees report to multiple independent managers. For example, an engineer working on a project to be implemented in Europe will report to the manager of

[1]We do not claim that matrix organization was designed with this goal in mind; it is rather a consequence, even if unintended, of its design.

the engineering division, to the manager of the project and to the manager of the European division. In other words, the information can flow via independent channels. One insight of our analysis is that this is the main reason why the matrix organization facilitates the effective communication of valuable information. (See Baron and Besanko, 1996, and Harriv and Raviv, 2002, for economic models of matrix organization and Galbraith, 2009, and Schrötter, 2014, for qualitative analysis.) However, this comes at a cost: frequent meetings and emails. (In our model, this corresponds to the many rounds of communication needed.) The tendency of matrix organization to generate countless meetings was already noted in the late 1970s.

> *Top managers were spending more time than ever before in meetings or in airplanes taking them to and from meetings.* (McKinsey Quarterly Review, "Beyond matrix organization." September 1979.[2])

In summary, this paper makes two novel contributions. We first show how to emulate a mediator with indirect communication; i.e., we replicate the communication equilibria of cheap-talk games through indirect communication. Second, we show how to do so while ensuring sequential rationality. To our knowledge, we are the first to address this challenging problem.

Related literature. This paper is related to several strands of literature. The commonality between these strands is the construction of protocols to *securely* transmit a message from a sender to a receiver on a communication network. The secure transmission of a message requires that (i) the receiver correctly learns the sender's message, and (ii) intermediaries do not obtain additional information about the message while executing the protocol. Reliability (or resiliency) refers to the first requirement, while secrecy refers to the second.

First, there is a large body of literature in computer science that studies the problem of secure transmission of messages on communication networks (see, among others, Beimel and Franklin, 1999; Dolev and al.,1993; Linial, 1994; Franklin and Wright, 2004; Renault and Tomala, 2008; and Renault and al., 2014). This literature provides conditions on the topology of communication networks to enable the secure transmission of a message from a sender to a receiver (see Renault and al., 2014, for a summary of these results). An important assumption of all of these studies is that the adversary

[2]https://www.mckinsey.com/business-functions/organization/our-insights/beyond-the-matrix-organization#

controls a *fixed* set of nodes throughout the execution of the communication protocols. The adversary we consider is stronger in that it can control different sets of nodes in each round of communication (see Section 3.1 for more details). To our knowledge, such an adversary has not been studied in the computer science literature. However, we restrict our attention to singletons, while the computer science literature considers larger sets.

This paper is also linked to the literature on repeated games on networks, where the network models the monitoring structure and/or the communication possibilities (see, among others, Ben-Porath and Kahneman, 1996; Laclau, 2012, 2014; Renault and Tomala, 1998; Tomala, 2011; and Wolitzky, 2015). This literature characterizes the networks for which folk theorems exist. An essential step in obtaining a folk theorem is the construction of protocols that guarantee that, upon observing a deviation, players start a punishment phase. To do so, when a player observes a deviation, he must be able to securely transmit the message "my neighbor has deviated" to all other players.

Except for Wolitzky (2015), none of these studies have imposed sequential rationality while restricting communication on a network. Either the communication is restricted on a network, in which case the solution concept is Nash equilibrium (Laclau, 2012; Renault and Tomala, 1998; Tomala, 2011), or the communication is unrestricted, in which case sequential rationality is imposed (Ben-Porath and Kahneman, 1996; Laclau, 2014). (In the latter case, the network models the structure of observation and/or interaction, but not the communication possibilities.) Wolitzky (2015) imposes sequential rationality and restricts communication on a network. However, he assumes that either a mutual minmax Nash equilibrium exists or that players have access to undifferentiated tokens, which can be freely transferred among neighbors, in addition to cheap talk messages. We make none of these assumptions. (In fact, the existence of a mutual minmax Nash equilibrium imposes restrictions on the preferences of the intermediaries, which is at odds with our analysis.)

This paper is also related to the literature on mediated and unmediated communication in games (regarding mediated communication, see, among others, Aumann, 1974; Ben-Porath, 1998; Forges, 1986, 1990; Forges and Vida, 2013; Myerson, 1986; Renou and Tomala, 2012; and Rivera, 2018). Like in the literature on unmediated communication in games, e.g., Barany (1992), Forges (1990), Forges and Vida (2013), and Gerardi

(2004), we show that we can emulate mediated communication with unmediated communication. The novelty is that communication is restricted to a network, albeit for a particular class of games, i.e., sender–receiver games.

Renou and Tomala (2012) and Rivera (2018) consider mediated communication games, where not all players can communicate directly with the mediator. The mediator is a fixed node in a communication network. These authors characterize the conditions on communication networks that make it possible to replicate all the equilibrium outcomes of the direct communication game, i.e., when the players can communicate directly with the mediator.[3] There are three major differences in our work. First, we do not have a mediator. In fact, we show how we can emulate the presence of the mediator on the network. Second, these authors restrict the communication so that it is unicast, i.e., group meetings are not allowed, while we consider rich communication possibilities. Third, they do not impose *strong* reliability (in the sense that the transmission of messages is reliable *after all histories consistent with unilateral deviations*). While none of their constructions guarantee strong reliability, these authors are nonetheless able to impose sequential rationality. Renou and Tomala (2012) achieve this by restricting attention to games with either independent private values or with strict punishment. Rivera (2018) needs three disjoint paths of communication.

Finally, this paper connects to the large body of literature on cheap talk games, pioneered by Crawford and Sobel (1984) and Aumann and Hart (2003) (see Forges, 2020, for a recent survey). The closest papers to ours are Ivanov (2010) and Ambrus et al. (2014). These papers consider simple communication networks (perfectly hierarchical networks) and restrict attention to a particular class of games. Their emphasis is complementary to ours. We ask when it is possible to PBE implement all communication equilibrium outcomes of sender–receiver games on a communication network, while they ask what the equilibrium outcomes of their fixed game are. Another related paper is Migrow (2021). Building on the work of Ambrus et al. (2014), Migrow shows that indirect communication can improve upon direct communication – suitably designed hierarchies of intermediaries can simulate some of the outcomes that mediated communication would achieve. Again, as in Ambrus et al., a particular class of games and

[3]Renou and Tomala (2012) study pure adverse selection problems, while Rivera (2018) extends the analysis to both adverse selection and moral hazard. A common thread of these papers is the need to construct secret and reliable protocols so that players (the mediator) can transmit their private information (private recommendations) to the mediator (the players).

communication networks is considered. In addition, detailed knowledge of the preferences of the intermediaries is needed.

2. THE SETUP

We start with mathematical preliminaries. Unless indicated otherwise, all sets X are complete separable metric spaces, endowed with their Borel σ-algebra $\mathbb{B}_X$. We write $\Delta(X)$ for the set of probability measures on X. Let X and Y be two complete separable metric spaces. A probability kernel is a function $f : Y \times \mathbb{B}_X \to [0,1]$ such that (i) for all $y \in Y$, $f(y,\cdot) : \mathbb{B}_X \to [0,1]$ is a probability measure, and (ii) for all $B \in \mathbb{B}_X$, $f(\cdot, B) : Y \to [0,1]$ is measurable. Throughout, we abuse notation and write $f : Y \to \Delta(X)$ for the probability kernel $f : Y \times \mathbb{B}_X \to [0,1]$.

2.1. **The problem.** There is a sender and a receiver, labeled S and R, respectively. The sender knows a payoff-relevant state $\omega \in \Omega$, with prior probability $\nu \in \Delta(\Omega)$. The receiver takes action $a \in A$. For all $i \in \{S, R\}$, player i's payoff function is $u_i : A \times \Omega \to \mathbb{R}$, which we assume to be measurable.

Direct communication. In the direct communication game, the sender directly communicates with the receiver; that is, the sender sends a message $m \in M$ to the receiver prior to the receiver choosing an action $a \in A$. A strategy for the sender is a map $\sigma : \Omega \to \Delta(M)$, while a strategy for the receiver is a map $\tau : M \to \Delta(A)$. We denote $\mathcal{E}^d \subseteq \Delta(\Omega \times A)$ the set of (Bayes-Nash) equilibrium distributions over states and actions of the direct communication game. Notably, we may have $\mathcal{E}^d = \emptyset$.[4]

Mediated communication. In the mediated communication game, the sender first sends a message $m \in M$ to a mediator. The mediator then sends a message $r \in R$, possibly randomly, to the receiver, who then takes an action $a \in A$. A strategy for the sender is a map $\sigma : \Omega \to \Delta(M)$, while a strategy for the receiver is a map $\tau : R \to \Delta(A)$. The mediator follows a recommendation rule: $\varphi : M \to \Delta(R)$. A communication equilibrium is a communication device $\langle M, R, \varphi \rangle$, and an equilibrium (σ^*, τ^*) of the mediated game. Thanks to the revelation principle (Forges, 1986 and Myerson, 1986), we can restrict our attention to canonical communication equilibria, where $M = \Omega$, $R = A$, the sender has an incentive to be truthful (to report the true state), and the receiver has an incentive to be obedient (to follow the recommendation). We denote $\mathcal{CE}^d \subseteq \Delta(A \times \Omega)$ as the

[4]Indeed, the existence of a best reply is not guaranteed, as we do not assume that the payoff function is continuous nor that the action space is compact.

set of communication equilibrium distributions over actions and states of the mediated communication game. It is well-known that $\mathcal{E}^d \subseteq \mathcal{CE}^d$.

It is commonly acknowledged that the set of communication equilibrium payoffs might be strictly larger than the set of Nash equilibrium payoffs. In particular, both the sender and the receiver might benefit strictly from mediated communication (see Forges, 1985, and Myerson, 1991). However, this requires the existence of a trusted mediator, which is a rather strong assumption. A message of this paper is that there is a way to organize the communication between the sender and the receiver to emulate the trusted mediator. The communication must be indirect and intermediated. (As we shall see later, the possibility emulating the mediator through indirect communication will be independent of the preferences of the intermediaries.) We now turn to a formal description of indirect communication.

Communication game on a network (indirect communication). To model indirect communication, we assume that the sender and the receiver are two distinct nodes on an (undirected) network $\mathcal{N}$. The set of nodes, other than S and R, is denoted I, which we interpret as a set of n intermediaries. Communication between the sender and the receiver transits through these intermediaries. We let $\mathcal{N}_i$ be the set of neighbors of $i \in I^* := I \cup \{S, R\}$ in the network. Throughout, we assume that the sender and the receiver are not directly connected in the network $\mathcal{N}$.

A communication game on the network $\mathcal{N}$ is a multistage game with $T \leq \infty$ stages, where at each stage players send costless messages to their neighbors and the receiver decides either to take an action (and stop the game) or to continue communicating.

A *communication mechanism*, denoted $\boldsymbol{\mathcal{M}}$, specifies the sets of messages players can send to their neighbors along with their dependence on past messages sent and received. We allow for a rich set of communication possibilities. Communication can be private, e.g., private emails or one-to-one meetings; public, e.g., emails sent to distribution lists or group meetings; or a mix of both. We say that player i *broadcasts* a message to a (nonempty) subset of his neighbors $N \subseteq \mathcal{N}_i$ if (i) all players in N receive the same message and (ii) it is common belief among all players in N that they have received the same message (in other words, the list of recipients of the message is certifiable among them). An example of a communication mechanism is as follows: only public messages to neighbors are allowed (broadcasting to all neighbors only), and each player has two different messages that he can send.

Communication unfolds as follows: at each stage t, players *broadcast* messages to all possible (nonempty) subsets of neighbors. The set of messages player i can broadcast to the subset of neighbors $N \in 2^{\mathcal{N}_i} \setminus \{\emptyset\}$ is $\mathcal{M}_{iN}$. Let $\mathcal{M}_i = \prod_{N \in 2^{\mathcal{N}_i} \setminus \{\emptyset\}} \mathcal{M}_{iN}$ be the set of messages available to player i and $\mathcal{M} = \prod_{i \in I^*} \mathcal{M}_i$ the set of messages available to all players. A few remarks are worth noting here. First, private messages correspond to broadcasts to singletons. Players can thus send private messages to their neighbors. Second, the sets of messages available to a player are independent of both the time and past history of the messages sent and received. The latter has a loss of generality. However, without such an assumption, the model has no bite. Indeed, if the only message player i can transmit upon receiving the message m is the message m itself, an extreme form of history dependence, then we trivially reproduce direct communication with indirect communication.

Finally, two additional elements are needed to obtain the communication game from the communication mechanism. First, we assume that at each stage, the receiver can either take an action $a \in A$ or continue communicating, in which case he sends a message $m_R \in \mathcal{M}_R$. If the receiver takes the action a, the game stops. Second, we need to associate payoffs with terminal histories (*i.e.* histories where the receiver takes an action and histories where the receiver never does). The payoff to player $i \in I^*$ is $u_i(a, \omega)$ when the state is ω and the receiver takes action a, with $u_i : A \times \Omega \to \mathbb{R}$ being a measurable function. (If the receiver never takes an action, the payoff to all players is $-\infty$.) Thus, communication is purely cheap talk. We denote $\Gamma(\boldsymbol{\mathcal{M}}, \mathcal{N})$ the communication game induced by the mechanism $\boldsymbol{\mathcal{M}}$ on the network $\mathcal{N}$.

Strategies and equilibrium. A history of messages received and sent by player i up to (but not including) period t is denoted h_i^t, with H_i^t the set of all such histories. A (pure) strategy for player $i \in I$ is a collection of maps $\sigma_i = (\sigma_{i,t})_{t \geq 1}$, where at each stage t, $\sigma_{i,t}$ maps H_i^t to $\mathcal{M}_i$. A (pure) strategy for the sender is a collection of maps $\sigma_S = (\sigma_{S,t})_{t \geq 1}$, where at each stage t, $\sigma_{S,t}$ maps $\Omega \times H_S^t$ to $\mathcal{M}_S$. A (pure) strategy for the receiver is a collection of maps $\sigma_R = (\sigma_{R,t})_{t \geq 1}$, where at each stage t, $\sigma_{R,t}$ maps H_R^t to $\mathcal{M}_R \cup A$. With a slight abuse of notation, we use the same notation for behavioral strategies. Let $H^t = \times_{i \in I^*} H_i^t$. We write $\mathbb{P}_\sigma(\cdot | h^t)$ for the distribution over terminal histories and states induced by the strategy profile $\sigma = (\sigma_S, \sigma_R, (\sigma_i)_{i \in I})$, conditional on the history $h^t \in H^t$. We write $\mathbb{P}_\sigma$ for the distribution, conditional on the initial (empty) history. It is easy to show that any equilibrium distribution over actions and states is an element of $\mathcal{CE}^d$.

2.2. **PBE implementation of direct (resp., mediated) communication on a network.** We are now ready to define the *PBE implementation of direct communication (resp., mediated communication) on a network*. We first start with an informal description. We say that the PBE implementation of direct communication (resp., mediated communication) is possible on a network if for every direct (resp., mediated) communication game, regardless of the preferences of the intermediaries, we can construct a communication game on the network with the property that for any Bayes-Nash equilibrium (resp., communication equilibrium) distribution over actions and states of the direct communication game (resp., mediated communication game), there exists a PBE of the communication game that replicates that distribution.

Definition 1. PBE implementation of direct communication (resp., mediated communication) *is possible on the network $\mathcal{N}$ if there exists a communication mechanism $\boldsymbol{\mathcal{M}}$ on $\mathcal{N}$ such that for all utility profiles of the sender, receiver and intermediaries, for all distributions $\mu \in \mathcal{E}^d$ (resp., $\mu \in \mathcal{CE}^d$) of all sender–receiver games, there exists a perfect Bayesian equilibrium σ of $\Gamma(\boldsymbol{\mathcal{M}}, \mathcal{N})$ satisfying:*

$$marg_{A\times\Omega}\mathbb{P}_\sigma = \mu.$$

In other words, if PBE implementation of direct communication (resp. mediated communication) is possible on a network, it means that it is possible to replicate any equilibrium outcome of the direct (resp., mediated) game via intermediated (not to be confused with mediated) communication between the sender and the receiver, without the need for a trusted mediator, and regardless of the preferences of the intermediaries.

The solution concept is PBE, i.e., whenever possible, beliefs are consistent with Bayes rule.[5] As we shall see later, none of our arguments relies on "crazy" off-equilibrium path beliefs. Moreover, stronger solution concepts, such as sequential equilibria, are generally not defined for arbitrary games as ours.

3. THE THEOREM

This section characterizes the networks, for which the PBE of direct and mediated communication is possible. With the help of a simple example, we now illustrate some of the

[5]Recall that we only require the sets of states, actions, and messages to be Polish, so that conditioning events often have a null measure. To deal with that issue, we consider a generalized version of PBE, where beliefs are updated using Bayes' rule whenever possible, with the use of regular conditional probabilities (like the distributional strategies in Crawford and Sobel, 1982).

difficulties that the requirement of sequential rationality introduces. (Recall that the literature on unmediated communication on networks has not imposed the requirement of sequential rationality thus far.[6])

We suppose that there are two states and two actions. The sender and receiver want to match the state, while the intermediaries want to unmatch the state. Clearly, there exists an equilibrium in the direct communication game where the sender truthfully reveals the state and the receiver matches the state. The equilibrium distribution is the least preferred outcome of any intermediary-any other distribution is strictly preferred.

Thus, if an intermediary deviates at a stage in the indirect communication game, the deviation must *not* change the distribution to be nonprofitable. At any information set that follows the deviation, sequential rationality dictates that no other intermediary must have a profitable deviation either, i.e., no other intermediary must be able to change the distribution. Therefore, the correct distribution must be implemented not only for all on-path histories but also for histories that are reachable via sequences of unilateral deviations.

This observation motivates the notion of *strong reliability*, which we introduce next. We will prove that strong reliability on a network is in fact equivalent to the PBE implementation of direct, and mediated, communication in that network.

3.1. **Strong reliability.** We consider the following alternative problem: the sender wishes to transmit the message $m \in M$, a realization of the random variable m with distribution $\boldsymbol{\nu}$, to the receiver, through the network $\mathcal{N}$. We want to construct a *protocol*, i.e., a communication mechanism and a profile of strategies, such that the receiver correctly "learns" the message sent at all terminal histories consistent with *unilateral deviations*.

Before formally introducing the concept of strong reliability, we define what we mean by "consistent with unilateral deviations." We fix a strategy profile σ. We define $\Sigma(\sigma)$ as the set of strategy profiles such that $\sigma' \in \Sigma(\sigma)$ if, and only if, there exists a sequence of intermediaries $(i_1, \dots, i_t, \dots)$ such that $\sigma'_t = (\sigma'_{i_t,t}, \sigma_{-i_t,t})$ for all t. Thus, $\Sigma(\sigma)$ consists of all strategy profiles consistent with at most one intermediary deviating at each stage. Note that the same intermediary may deviate at several, or even all, stages. We let

[6]Wolitzky (2015) is an exception, but he needs additional assumptions.

$\mathcal{H}(\sigma)$ be the set of terminal histories consistent with $\Sigma(\sigma)$; that is, $h \in \mathcal{H}(\sigma)$ if there exists $\sigma' \in \Sigma(\sigma)$ such that h is in the support of $\mathbb{P}_{\sigma'}$.

We are now ready to define the concept of *strong reliability on a network*.

Definition 2. *The transmission of messages is strongly reliable on network $\mathcal{N}$ if there exist a protocol σ and a decoding rule $\mathbf{m}_d : H_R^{T+1} \to M$ such that*

$$\mathbb{P}_{\sigma'}\Big(\Big\{h_R^{T+1} : \mathbf{m}_d(h_R^{T+1}) = m\Big\}\Big|\mathbf{m} = m\Big) = 1,$$

for all $\sigma' \in \Sigma(\sigma)$, for all m.

The study of the reliable transmission of messages on networks is not new, see Dolev et al (1993), for an early attempt in computer science. (See Renault et al., 2014, for a summary of the literature.) Computer scientists assume that an adversary controls at most k nodes and they provide conditions on the network for the reliable transmission of messages. An important feature, however, is that the adversary controls the same k nodes throughout the execution of the protocol. This is a natural assumption in computer science, where communication is nearly instantaneous. An adversary would not have the time or capacity to take control of different nodes during the execution of the communication protocol.[7] A distinctive feature of our analysis is that we consider a *dynamic* adversary, i.e., an adversary that controls a different set of nodes in each round of the execution of the protocol. However, we limit our attention to singletons, i.e., $k = 1$. To our knowledge, this approach is new.

The notion of strong reliability has a clear and strong motivation. We want the transmission of messages to be reliable not only in the case of errors and unintentional mistakes but also in the case of intentional manipulations. For instance, without our notion of strong reliability, if a single e-mail were to not reach its recipients, this would entirely disrupt the communication. In addition, if the content of a single e-mail were to be modified, an entirely different message would be transmitted. Strong reliability guarantees that communication is resilient to these events.

The network in Figure 1 illustrates some of the difficulties associated with the requirement of strong reliability. There are three disjoint paths from the sender to the receiver, so it is tempting to use a majority argument. That is, to have the sender transmit his

[7]Formally, the reliable transmission of messages requires that $\mathbb{P}_{(\sigma_i',\sigma_{-i})}\Big(\{h_R^{T+1} : \mathbf{m}_d(h_R^{T+1}) = m\}\Big|\mathbf{m} = m\Big) = 1$ for all σ_i', for all $i \in I$. This is a weaker requirement that strong reliability.

message to intermediaries 1, 2 and 3 and to have all intermediaries forward their messages. If the intermediaries are obedient, then the receiver obtains three identical copies of the message sent and thus learns it. We suppose now that the sender wishes to transmit the message m. If intermediary 1 reports $m' \neq m$ in the first stage and intermediary 5 reports $m'' \neq m$ in the second stage, the receiver then receives the profile of reports (m', m'', m). Thus, we need the receiver to decode it as m. However, the receiver receives the same profile of reports (m', m'', m) when the sender wishes to transmit the message m', intermediary 3 reports m and intermediary 5 reports m''. Since the receiver would still decode it as m, he would learn the wrong message. Such a simple strategy does not work in general.[8] We thus need a more sophisticated construction.

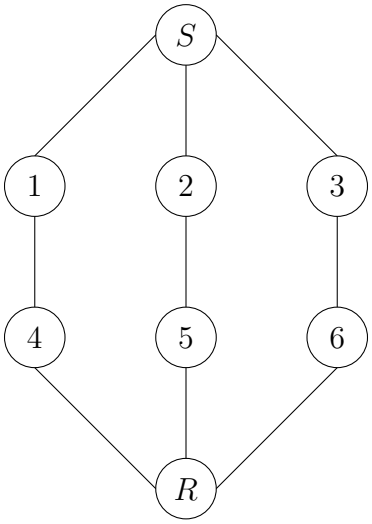

FIGURE 1. Strong reliability: Illustration of the difficulties

We note that there might be some PBE profiles that do not rely on strongly reliable transmission of messages. Indeed, in the previous example, the majority rule is a PBE: since no intermediary observes what others do on the other paths, no intermediary has a profitable deviation, and sequential rationality is guaranteed. Such an argument does not work if intermediaries, say 1 and 5 for instance, are linked, or if there are two paths; most importantly, it does not make it possible to emulate the mediator (see next section with our main result). Even if there are more PBE than the one we construct, we will show that strong reliability implies sequential rationality (see the discussion just before Section 3.4).

3.2. **The main result.** We can now state our main result.

[8]This simple strategy works if there are enough disjoint paths. In that example, we need two additional disjoint paths.

Theorem 1. *The following statements are equivalent.*

(1) PBE implementation of mediated communication is possible on the network $\mathcal{N}$.

(2) PBE implementation of direct communication is possible on the network $\mathcal{N}$.

(3) The transmission of messages is strongly reliable on the network $\mathcal{N}$.

(4) The network $\mathcal{N}$ admits two disjoint paths between the sender and the receiver.

Theorem 1 states that if there are two disjoint paths between the sender and the receiver in the network, then PBE implementation of mediated communication is possible. In other words, we can replicate the mediator through unmediated communication. This implies that the sender and the receiver may be better off communicating through intermediaries, rather than directly (since both the sender and the receiver may be better off with mediated communication than with direct communication). The theorem also states that the connectivity requirement to implement the communication equilibria is no stronger than the one required to implement the Bayes-Nash equilibrium of the underlying sender–receiver game.

3.3. **Proof: the main idea.** The proof of Theorem 1 is constructive and relegated to Appendix B. In what follows, we sketch the main idea.

(1) $\Rightarrow$ (4)**,** (2) $\Rightarrow$ (4) ***and*** (3) $\Rightarrow$ (4)**.** It is easy to see that Statement (4) of Theorem 1 is necessary for all other statements to be satisfied. Indeed, if no two disjoint paths exist between the sender and the receiver, there exists an intermediary that controls all the information transmitted between the sender and the receiver. In graph-theoretic terms, the intermediary is a *cut* of the graph. In games where the sender and the receiver have perfectly aligned preferences, but the intermediaries have opposite preferences, the "cut" can then simulate the histories of messages he would have received in a particular state and behave accordingly. Thus, even in games with perfectly aligned preferences, neither strongly reliable transmission nor PBE implementation of direct communication holds if there is a "cut" (and, a fortiori, neither does PBE implementation of mediated communication hold).

(4) $\Rightarrow$ (3). We prove that if there are two disjoint paths between the sender and the receiver in $\mathcal{N}$, then strong reliability is possible. This is the most important part of the proof, which we now explain in detail.

We suppose that the sender wishes to send the message m to the receiver. We want to find a protocol (a communication mechanism and a strategy profile) such that the receiver correctly learns the message m not only at all on-path histories, but also at all histories consistent with at most one intermediary deviating at each stage of the protocol. We show that such a protocol exists when there are two disjoint paths between the sender and the receiver. Moreover, the receiver learns the message after at most $1+(n^C-3)(2n^C-3)$ stages, where n^C is the number of nodes on the two shortest disjoint paths from the sender to the receiver (including the sender and the receiver). We now illustrate our protocol with the help of the network in Figure 2, where $n^C = 4$. (The protocol we construct is slightly more complicated, but they share the same properties.)

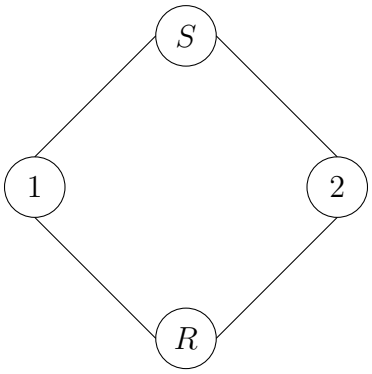

FIGURE 2. Illustration of Theorem 1

A similar example of Nash implementation (not requiring any sequential rationality or strong reliability) was independently studied by Franklin and Wright (2000) and Renault and Tomala (2004, 2008). The basic idea of their protocol is for the sender to broadcast m to intermediaries 1 and 2, who are then supposed to broadcast it, along with authentication keys. Thus, if the receiver observes two identical messages from intermediaries 1 and 2, he decodes it as the correct message. However, if intermediary 1 (resp., 2) broadcasts message $m' \neq m$ (a deviation), the sender then broadcasts 1's authentication key to intermediaries 1 and 2, with intermediary 2 (resp., 1) broadcasting it at the next stage. Thus, upon matching 1's authentication key received from intermediary 1 (at the second stage) and intermediary 2 (at the fourth stage), the receiver correctly learns that intermediary 1 (resp., 2) is the deviator and decodes the message as m. Although simpler than ours, this protocol is not strongly reliable. If intermediary 2 does not correctly broadcast the authentication key of intermediary 1, the receiver does not learn the correct message. In addition, since intermediary 2 (resp., 1) learns whether intermediary 1 (resp., 2) deviated, we cannot guarantee sequential rationality.

In adversarial problems, intermediary 2 would also deviate after observing the deviation of intermediary 1. To our knowledge, none of the previous contributions are able to address that issue or, ultimately, emulate the mediator.

Our communication protocol circumvents this issue by adding repetition: new authentication keys are drawn at each stage, and the receiver will only analyze messages that he has received several times from an intermediary. This latter requirement ensures that evidence supporting the existence of some deviation (via authentication keys) will actually be transmitted by the other intermediary; indeed, at the stages where the first intermediary keeps sending the same false message, then other players are not deviating at these stages, and his authentication key has time to be transmitted through the other disjoint path of the network. Hence, the receiver cannot learn a false message (see Lemma 1 in the Appendix). Moreover, this requirement of analyzing several messages cannot be too large so that the receiver can learn the message; in particular, a deviator cannot distract the protocol with false alerts of deviation (see Lemma 2 in the Appendix).

Formally, our protocol has six stages, which we now describe. First, the sender broadcasts the message m to intermediaries 1 and 2 at stage $t = 1$. At all other stages $t = 2, \dots, 6$, the protocol requires the intermediaries to broadcast the message m and an *authentication key* x_i^t, where x_i^t is the authentication key of intermediary i at stage t, a uniform draw from $[0, 1]$, independent of all messages the intermediary has sent and received. Finally, if the sender observes intermediary i broadcasting message $m' \neq m$ at stage t, the sender broadcasts the triplet (i, t, x_i^t) at stage $t + 1$ (in addition to his other messages). We interpret triplet (i, t, x_i^t) as stating that intermediary i has deviated at stage t and that his authentication key is x_i^t. If an intermediary receives triplet (i, t, x_i^t) at stage $t + 1$, the protocol requires the intermediary to broadcast that triplet at all subsequent stages.

The receiver does not send messages. At the end of the six stages, the receiver decodes the message as follows. If at any stage, the receiver has received the same message from both intermediaries, then he decodes it as the correct message. In all other instances, if the receiver has received the *same* message m_i from intermediary i at some stages

t_1, t_2 and t_3 ($t_1 < t_2 < t_3$), and he has *not* received the triplet $(i, t_1, x_i^{t_1})$ from the other intermediary by stage t_3, then he assumes that the correct message is m_i.[9]

We now argue that the protocol guarantees the strong reliability of the transmission. First, since at most one intermediary deviates at any stage, the protocol guarantees that the receiver obtains at least one sequence of three identical messages from either intermediary 1 or 2 (the stages in the sequence need not be consecutive). This statement represents a simplest version of Lemma 2 (see Appendix A). Moreover, if at any stage, the receiver obtains the same message from both intermediaries, it must be the correct message (since at least one intermediary must be broadcasting the correct message). Therefore, we assume that the receiver obtains message m_i from intermediary i at some stages t_1, t_2 and t_3 ($t_1 < t_2 < t_3$). If $m_i \neq m$ (hence intermediary i deviates at stages t_1, t_2 and t_3), the protocol requires the sender to broadcast the triplet $(i, t_1, x_i^{t_1})$ at stage $t_1 + 1 \leq t_2$. The protocol also requires intermediary $j \neq i$ to broadcast triplet $(i, t_1, x_i^{t_1})$ at all stages after receiving it. Hence, the receiver obtains the triplet at stage t_3 at the latest. Indeed, since intermediary i deviates at the stages (t_1, t_2, t_3), intermediary $j \neq i$ cannot deviate at t_2 and t_3. Since the authentication key received from intermediary j at either t_2 or t_3 matches the key received from intermediary i at t_1, the receiver learns that the message m_i is not correct. The correct message must therefore be the one broadcasted by intermediary j at stage t_1, that is, m. Alternatively, if $m_i = m$, the sender does not broadcast the triplet $(i, t_1, x_i^{t_1})$. Intermediary j may pretend that the sender had sent the triplet $(i, t_1, y_i^{t_1})$ at stage $t_1 + 1$. However, the probability that the reported authentication key y_i^t matches the actual authentication key x_i^t is zero; therefore, the receiver correctly infers that the message is m.

We now preview some secondary aspects of the above construction. First, the protocol is robust to deviations by the sender at all stages except the initial stage (when the sender broadcasts m). Indeed, if the sender deviates at stage $t \geq 2$, the two intermediaries do not, and the receiver then correctly learns the message. Similarly, the protocol is trivially robust to deviations by the receiver at all stages except for the last one. The protocol we constructed shares these two properties, which are key in proving that $(4) \Rightarrow (1)$, as we shall see. Second, the protocol starts with the two immediate successors of the sender on the two disjoint paths to the receiver learning the message m. At the end of the communication protocol, the receiver also learns the message. In general, the

[9]Equivalently, the receiver assumes that the correct message is m_i when he has received m_i from intermediary i at stages t_1, t_2 and t_3, and all triplets $(i, t_1, y_i^{t_1})$ received from the intermediary $3 - i$ by stage t_3 are such that $y_i^{t_1}$ is different from $x_i^{t_1}$, the authentication key received from i at stage t_1.

receiver is not the immediate successor of these intermediaries; these intermediaries have as successors other intermediaries. The key step in our construction is to show that at least one of the immediate successors of these intermediaries correctly learns the message at the end of the first block of communication. Therefore, as the protocol goes through blocks, the receiver eventually learns the message. Moreover, each block has $2n^C - 3$ stages. (Recall that n^C is the total number of nodes on the two disjoint paths from the sender to the receiver, including them.) Thus, the receiver learns the message in at most $1 + (2n^C - 3)(n^C - 3)$ stages, where the two immediate successors of the sender learn m immediately and then each of the remaining $n^C - 3$ players, who do not yet know m learns the message progressively over time.

(4) $\Rightarrow$ (1). The gist of the proof consists of repeatedly using the protocol constructed above to simulate the mediator, with the help of jointly controlled lotteries. To get a flavor of our construction, let us again consider the network in Figure 2. For simplicity, we assume that Ω and A are finite sets. We fix a canonical communication equilibrium $\varphi : \Omega \to \Delta(A)$. For all $\omega \in \Omega$, let A_ω be a partition of $[0, 1]$ into $|A|$ subsets, with the subset $A_\omega(a)$ corresponding to a having Lebesgue measure $\varphi(a|\omega)$.[10]

The protocol has three distinct phases. In the first phase, the sender broadcasts the state ω to the intermediaries 1 and 2. The second phase replicates the communication device. To do so, the sender and intermediary 1 simultaneously choose a randomly generated number in $[0, 1]$. Let x and y be the numbers generated by the sender and intermediary 1, respectively. Players then follow the strongly reliable communication protocol constructed above, which makes it possible for intermediary 1 to strongly reliably transmit y to intermediary 2 since there are two disjoint paths between them on the same circle.[11] Thus, at the end of the second phase, the sender and both intermediaries know ω, x and y, while the receiver only knows y. The third phase starts once the sender and both intermediaries have learned x and y.[12] In the first stage of the third phase, the sender and the intermediaries simultaneously compute $x + y \pmod{[0, 1]}$, output the recommendation a if $x + y \pmod{[0, 1]} \in A_\omega(a)$, and each starts a copy of the communication protocol constructed above to strongly reliably transmit the recommendation to the receiver. Thus in the third phase, the three communication protocols are synchronized

[10]For example, if there are three actions, a_1, a_2 and a_3, a possible partition is $[0, \varphi(a_1|\omega)), [\varphi(a_1|\omega), \varphi(a_1|\omega) + \varphi(a_2|\omega)), [\varphi(a_1|\omega) + \varphi(a_2|\omega), 1]$.

[11]In this subprotocol, intermediary 1 plays the role of the sender, while intermediary 2 plays the role of the receiver.

[12]Notice that even if they do not learn x and y at the same stage, x and y were generated simultaneously.

and start at the very same stage. At the end of the third phase, the receiver learns the recommendations sent by the sender and both intermediaries. Since at most one of them can deviate in the stage where they broadcast their recommendation, the receiver decodes at least two identical recommendations and plays it.

It is then straightforward, albeit tedious and cumbersome, to define belief systems to guarantee the sequential rationality of the equilibria we construct. To see this, let σ be a Nash equilibrium of the communication game we construct. There are four strongly reliable protocols that are run independently of one another; therefore, we can consider them separately. For each such protocol, consider any history h_i^t consistent with unilateral deviations, i.e., there exists $\sigma' \in \Sigma(\sigma)$ such that h_i^t is in the support of $\mathbb{P}_{\sigma'}$. If intermediary i's belief at h_i^t is the "conditioning" of $\mathbb{P}_{\sigma'}$ on h_i^t, then strong reliability implies sequential rationality. Indeed, for all $\sigma' \in \Sigma(\sigma)$, for all $\tilde{\sigma}_i$, the concatenated strategy profile $\langle \sigma', (\tilde{\sigma}_i, \sigma_{-i}) \rangle = ((\sigma'_{t'})_{t'<t}, (\tilde{\sigma}_{i,t'}, \sigma_{-i,t'})_{t' \geq t})$ is consistent with unilateral deviations, i.e., $\langle \sigma', (\tilde{\sigma}_i, \sigma_{-i}) \rangle \in \Sigma(\sigma)$. Strong reliability thus implies that $\mathbb{P}_{\langle \sigma', (\tilde{\sigma}_i, \sigma_{-i}) \rangle} = \mu$, that is, $\mathbb{P}_{(\tilde{\sigma}_i, \sigma_{-i})}(\cdot|h^t)\mathbb{P}_{\sigma'}(h^t) = \mu$ for all $\tilde{\sigma}_i$, for all $\sigma' \in \Sigma(\sigma)$. Intermediary i is therefore indifferent between all his strategies at h_i^t (since his belief about h^t is $\mathbb{P}_{\sigma'}(h^t|h_i^t)$). Since the argument does not rely on the specific $\sigma' \in \Sigma(\sigma)$ we select, we have sequential rationality with respect to any belief system, which is fully supported on the histories consistent with unilateral deviations at h_i^t. In other words, as long as the intermediary believes that at most one player deviated in each of all past stages, we have sequential rationality at h_i^t. Similarly, the equilibria we construct are also robust to deviations by the sender and receiver at all stages except for the first one, where the sender sends the message, and the last one, where the receiver chooses an action. Therefore, we also have sequential rationality at all histories h_S^t and h_R^t since σ has, by construction, the same distribution over actions and states as the communication equilibrium φ. Finally, for all other histories, it is easy to construct beliefs and actions to guarantee sequential rationality. (See Section A.2 for detail.)

(4) ⇒ (2). This follows immediately from the previous step since $\mathcal{E}^d \subseteq \mathcal{CE}^d$ (hence, $(1) \Rightarrow (2)$). Notice, however that a simpler construction is possible since there is no need to emulate the mediator. We can directly use the communication protocol we construct to prove $(4) \Rightarrow (3)$ to (strongly) reliably transmit the message m to the receiver, where m is drawn with probability $\sigma_S^*(m|\omega)$, the equilibrium strategy of the direct communication game. Finally, we bserve that if there are three disjoint paths from the sender to the receiver as in Figure 1, a simpler majority argument works. It suffices for

the sender to send the message m on each path. Since no intermediary observes what others do on the other paths, no intermediary has a profitable deviation, and sequential rationality is guaranteed. Such an argument does not work if there are only two paths and, most importantly, does not make it possible to emulate the mediator, which is our core contribution.

3.4. **Remarks.** We conclude this section with some observations about our analysis. First, we allow for rich communication possibilities, most notably, that players are able to broadcast messages to any subset of neighbors. This is necessary for our results to hold. For instance, if players can only send private messages (unicast communication), then reliable transmission of messages, let alone strong reliability, is impossible, on the network in Figure 2 (see Dolev and al., 1993; or Beimel and Franklin, 1999). Similarly, if players can only send public messages (broadcast communication), reliable transmission of messages, let alone strong reliability, is impossible on the network shown in Figure 3. See Franklin and Wright (2000) and Renault and Tomala (2008).

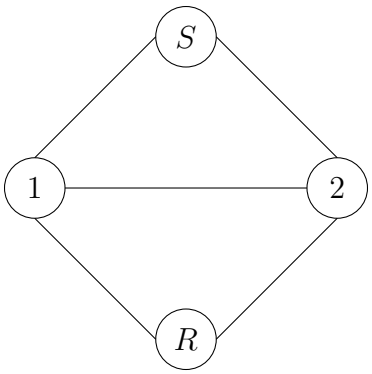

FIGURE 3. Broadcasting to all: Reliable communication is impossible

While formal proofs of these two impossibilities are complicated, the intuition is that the receiver is unable to distinguish between two types of histories: histories where intermediary 1 pretends that the message is m' and intermediary 2 is deviating, and histories where intermediary 2 pretends that the message is m and intermediary 1 is deviating. The key is that an intermediary can simulate fictitious histories, i.e., histories of messages sent and received when the message is any m, and behave accordingly. As is clear from the example, the protocol we construct makes it possible for the receiver to distinguish these two types of histories. If intermediary 1 deviates and pretends that the message is m', the receiver correctly infers that intermediary 2 is not deviating. This requires the sender to broadcast messages to selected subsets of neighbors.

Second, our protocol does not restrict the messages that players can broadcast to any subset of their neighbors. For example, in addition to the messages that our protocol requires the intermediaries to broadcast, the intermediaries can also send private messages to the sender and the receiver. The equilibrium we construct simply treats these additional messages as uninformative (babbling).

Third, detailed knowledge of the communication network is not needed. To execute our protocol, a player on one of the paths from the sender to the receiver needs to know only his two immediate neighbors on the path and the total number of players on the two disjoint paths. Hence, only local information is needed to perform the protocol.

Fourth, we assumed that the sender and the receiver are not directly connected. We now discuss how our results would change if the sender and the receiver can also communicate directly. The equivalence between statements (2) and (3) and a variation of (4) of Theorem 1 extends immediately. More precisely, PBE implementation of direct communication is possible (resp. the transmission of messages is strongly reliable) if, and only if, either the sender and receiver are directly connected or if there are two disjoint paths between the sender and the receiver. The extension of the equivalence to Statement (1) in Theorem 1 is more delicate. First, notice that if the direct link is the only link between the sender and the receiver, then there is no hope to replicate the mediator, as all communication would be direct. We need to be able to use intermediaries in communication. We claim that if there are at least two intermediaries, labeled 1 and 2, such that the sender, the receiver and the two intermediaries are on a "circle," then PBE implementation of mediated communication is possible.[13] An informal proof is given in Appendix B. We do not know whether the condition of having at least two intermediaries such that the sender, the receiver and the two intermediaries are on a "circle," is necessary. We conjecture this is so.

4. CONCLUDING REMARKS

Our analysis extends to communication games with multiple senders. More precisely, we consider a direct communication game, in which senders receive private signals about a payoff-relevant state, send messages to the receiver, and the receiver takes an action. If there exist two disjoint paths of communication from each sender to the receiver, we can then replicate our analysis to implement any PBE distribution of the direct communication game. The key is to have all senders broadcast their messages in the first

[13] A circle is a collection of nodes such that all pairs of nodes have two disjoint paths between them.

stage and then to run copies of our protocol in parallel. Learning from one protocol does not have any impact on other protocols; hence, this guarantees that the receiver learns the correct messages. We stress, however, that it is essential that all senders move simultaneously in the first stage; otherwise, their incentives might change from the direct communication game if they have learned the messages of other senders before sending theirs.

In contrast, extending our analysis to communication games with multiple receivers is very challenging. A first approach would be to increase the connectivity of the network: if there are two disjoint paths from the sender to each receiver, all these paths are disjoints one from another; then, our analysis easily extends to this case. However, this condition requires much connectivity, as there is a need for $2R$ disjoint paths with R receivers. A second direction would be to have a strongly reliable communication protocol with the additional property of secrecy in the sense that intermediaries do not learn the messages sent by the sender while performing the protocol. Our protocol clearly does not have that property. We conjecture that more disjoint paths are necessary to construct such a protocol, but it may be that fewer are needed than in the first approach above. This is a challenging question that is left for future research.

Finally, while we do not consider the possibility of k faults/deviations at each stage, we conjecture that our main ideas apply to this situation. (Naturally, the sender and the receiver cannot be faulty.) To see this, we consider the network in Figure 4, where there are $k+1$ reporting lines between the sender and the receiver.

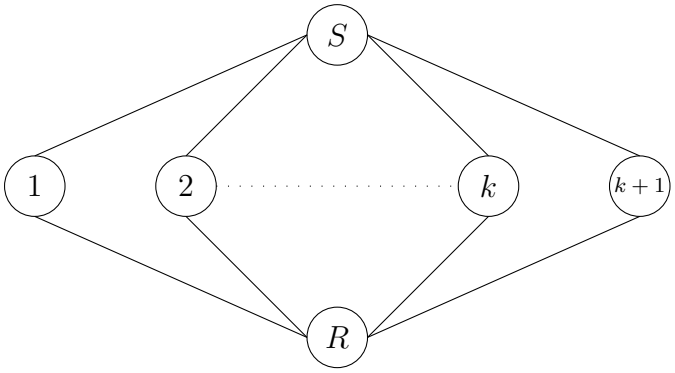

FIGURE 4. $k+1$ reporting lines

We now explain how to adapt our construction. As in Section 3, we continue to assume that the sender and the intermediaries broadcast the message to be transmitted and authentication keys at each stage. In addition, if an intermediary receives triplet (i, t, x_i^t) at a stage, he broadcasts it at all subsequent stages. We also require the receiver to

validate a sequence of three identical messages from intermediary i, received at stages t_1, t_2 and t_3, only if he has *not* received the authentication key $x_i^{t_1}$ at stage t_3 or earlier from intermediary $j \neq i$. Unlike in our original construction, however, we require communication to last $1 + [2(k+1)+1]$ stages. Intuitively, we need to extend the rounds of communication to guarantee that at least one intermediary sends the same message three times. Since at most k intermediaries deviate at a stage, at least one intermediary is transmitting the correct message at each stage; therefore, after $2(k+1)+1$ stages, at least one intermediary has transmitted the correct message three times to the receiver. Thus, as in our original construction, the receiver never validates an incorrect message (since he will receive the correct authentication key from at least one intermediary) and is guaranteed to validate the correct message (since he is guaranteed to receive the correct message three times and not the authentication key). Future research should explore whether we can extend our arguments to all sufficiently connected networks.

APPENDIX

APPENDIX A. PROOF OF THEOREM 1

Statement (1) clearly implies Statement (2) in Theorem 1, as the set of Bayes–Nash equilibria is included in the set of communication equilibria in the direct communication game. Additionally, we do not prove that Statement (2) implies Statement (4), as the proofs follow well-trodden paths, see e.g., Renault and Tomala (2008) or Renault et al. (2014).

A.1. **Proof that Statement (4) implies Statement (3).** We prove here that two disjoint paths between the sender and the receiver is a sufficient condition for the transmission of messages to be strongly reliable. Assume that network $\mathcal{N}$ admits two disjoint paths between the sender and the receiver, and denote the two shortest paths by $S, i_1, \ldots, i_K, R$ and $S, j_1, \ldots, j_{K'}, R$, respectively, for some $K, K' \geq 1$. We let $\mathcal{P}$ be the set of nodes on these two paths, including the sender and the receiver, and let n^C be its cardinality. Throughout, we refer to these two paths as the "circle," with the nodes $\{S, i_1, \cdots, i_K, R\}$ (resp.,$\{S, j_1, \ldots, j_{K'}, R\}$) as the "left side of the circle" (resp., "right side of the circle"). For each player $p \in \mathcal{P} \setminus \{S, R\}$, we call the *successor* of p, denoted p^+, his immediate successor on the path to the receiver. Similarly, we call the *predecessor* of p, denoted p^- his immediate predecessor on the path to the sender. For example, if $p = i_k$ for some $1 \leq k \leq K$, $p^+ = i_{k+1}$ and $p^- = i_{k-1}$, with the convention that $i_0 = S$ and

$i_{K+1} = R$. The sender and the receiver have a special position since they participate in communication on both paths; therefore, their neighbors need to play the roles of both predecessors and successors. In particular, if $p = R$ and $p^- = i_K$ (resp, $j_{K'}$), then $p^+ = j_{K'}$ (resp., $p^- = i_K$). The same reasoning applies for the sender. However, as will be clear later, whenever i_1 (resp., i_K) plays the role of a predecessor, then j_1 (resp., $j_{K'}$) plays the role of the successor and, vice versa. Thus, we mostly focus on messages flowing from the sender to the receiver. However, messages will also need to flow in the other direction.

A.1.1. *The communication protocol* $(\boldsymbol{\mathcal{M}}, \sigma)$. Throughout, we write $[1:T]$ for $\{1, \ldots, T\}$.

The message space. Recall that M is the set of messages in the direct communication game and let $m_0 \notin M$ be an arbitrary message, interpreted as null. The set of messages that player $i \in I^*$ can broadcast to subset $N_i \in 2^{\mathcal{N}_i} \setminus \{\emptyset\}$ is as follows:

$$\begin{aligned}\mathcal{M}_{i,N_i} \;=\; & \Big(M \cup \{m_0\}\Big) \times [0,1] \\ & \bigtimes \Big(\bigtimes_{j \in \mathcal{N}_i} \Big\{ \{\emptyset\} \bigcup \big(\{j\} \times [1:T] \times [0,1]\big) \Big\} \Big) \\ & \bigtimes \Big(\bigtimes_{j \in I^* \setminus \mathcal{N}_i} \Big\{ \{\emptyset\} \bigcup \big(\{j\} \times [1:T] \times [0,1]\big) \Big\}^L \Big),\end{aligned}$$

with $L = 1 + (n^C - 3)(2n^C - 3)$.

The set of messages that player i can send is $\bigtimes_{N_i \in 2^{\mathcal{N}_i} \setminus \{\emptyset\}} \mathcal{M}_{i,N_i}$. In other words, each player can broadcast to any subset of his neighbors a grand message composed of: (i) a message $m \in M$ or the null message m_0, (ii) a number in $[0,1]$ and (iii) a tuple of triplets, each of which is composed of the name of a player, a stage, and a number in $[0,1]$. Crucially, player i and his predecessors following the protocol, validate at most a *single* triplet about each of their neighbors in each stage. However, player i can send several triplets (at most L) about the other players. In other words, at each stage, player i can send a list of triplets to his neighbors N_i, but no list includes more than one triplet about $j \in \mathcal{N}_i$.

The strategies of the players. For any player $p \notin \mathcal{P}$, the protocol requires them to broadcast (uniformly) randomly drawn messages in $\mathcal{M}_{i,N_i}$ at each stage t to each subset of neighbors $N_i \in 2^{\mathcal{N}_i} \setminus \{\emptyset\}$, independently of all the messages received and sent up to stage t. In words, they are babbling.

We now define the strategy for player $p \in \mathcal{P}$. We focus on the messages they broadcast to their neighbors on the circle, i.e., to $\mathcal{N}_p \cap \mathcal{P}$. To all other subsets of neighbors, they

send randomly generated messages, independent of the histories of the messages sent and received, i.e., they babble. In what follows, when we say that player p broadcasts a message, we mean that he broadcasts the message to the subset $\{p^-, p^+\}$. Remember that for S (resp., R), $\{p^-, p^+\} = \{i_1, j_1\}$ (resp., $\{i_K, j_{K'}\}$).

- **Authentication keys:** At each stage $t \geq 1$ of the communication protocol, p broadcasts a uniformly drawn message x_p^t in $[0, 1]$ to his neighbors on the "circle," that is, to the two players in $\mathcal{N}_p \cap \mathcal{P}$: this message x_p^t is called the *authentication key* of player p at stage t.
- **First stage:** At stage $t = 1$, the sender broadcasts the message m to his neighbors on the circle, i.e., to i_1 and j_1, along with his authentication key. All the other players broadcast m_0 to their neighbors on the circle (along with their authentication keys).
- **Subsequent stages:** Starting from stage $t = 2$ onwards, the protocol proceeds in blocks of $2n^C - 3$ stages. We denote these blocks by B_b, where $b = 1, 2 \ldots, \overline{B}$. (We have at most $n^C - 3$ blocks.) For instance, B_1 represents the block that starts at stage $t = 2$, and B_2 represents the block that starts at stage $t = 2 + 2n^C - 3 = 2n^C - 1$, etc. For each $b = 1, 2 \ldots$, let $B_b := \{t_b, \ldots, t_b + 2n^C - 4\}$ where t_b is the first stage of block B_b. In each block B_b, the strategy of p is the following:
  - **Transmission of the sender's message:**
    - If p knows the message m at the beginning of the block, that is, if p is in $\{S, i_1, j_1\}$ or p has *learned* message m at the end of the previous block (see below, where the decoding rule at the end of each block is defined), then p broadcasts the message m to his neighbors p^- and p^+ at all stages $t_b, \ldots, t_b + 2n^C - 4$ of the current block (the neighbors of S are i_1 and j_1).
    - If p does not know the message m at the beginning of the block, then p broadcasts m_0 to his neighbors p^- and p^+ at all stages $t_b, \ldots, t_b + 2n^C - 4$ of the current block.

    (Remember that p also sends an authentication key.)
  - **Detection of deviations:**
    - If p detects his successor p^+ making a false announcement about the message $m \in M$ at some stage $t \in \{t_b, \ldots, t_b + 2n^C - 4\}$, that is,
      - either p knows $m \in M$ and p^+ broadcasts at stage t the message $\tilde{m} \in M \setminus \{m\}$, interpreted as "*player p^+ is broadcasting the false message $\tilde{m}$*,"

· or p does not know the message m and p^+ broadcasts at stage t the message $\tilde{m} \in M$, interpreted as *"player p^+ is broadcasting the message $\tilde{m}$ although he cannot know it,"*

then p broadcasts the triplet $(p^+, t, x^t_{p^+})$ to player p^- and to p^+, where $x^t_{p^+}$ is the true authentication key broadcasted by p^+ at t. Note that if $p = S$ (resp., R), then p^+ is either i_1 or j_1 (resp., i_K or $j_{K'}$).

∗ If p does not detect his successor p^+ making a false announcement about the message $m \in M$ at some stage $t \in \{t_b, \ldots, t_b + 2n^C - 4\}$, then p broadcasts the triplet (p^+, t, y) to p^- and p^+, where y is randomly drawn from $[0, 1]$.

The key observation is that only players p and p^{++} know the true authentication key of p^+ at stage t. Therefore, p^{++} can authenticate whether p^+ deviates at some stage t by cross-checking the authentication key $x^t_{p^+}$ received from p^+ at t with the key broadcasted by p (and having transited on the circle in the opposite direction).

– Transmission of past deviations:

∗ If $p \neq i_1$ is on the left side of the circle and receives at some stage $t \in \{t_b, \ldots, t_b + 2n^C - 4\}$

· from p^+ a message containing the triplet $(p', d, x^d_{p'})$, $t_b \leq d < t$ with p' on the left side of the circle, then p broadcasts the message to p^- and p^+ at stage $t + 1$.

· from p^- a message containing the triplet $(p', d, x^d_{p'})$, $t_b \leq d < t$ with p' on the right side of the circle, then p broadcasts the message to p^- and p^+ at stage $t + 1$.

∗ Similarly, if $p \neq j_1$ is on the right side of the circle and receives at some stage $t \in \{t_b, \ldots, t_b + 2n^C - 4\}$

· from p^+ a message containing the triplet $(p', d, x^d_{p'})$, $t_b \leq d < t$ with p' on the right side of the circle, then p broadcasts the message to p^- and p^+ at stage $t + 1$.

· from p^- a message containing the triplet $(p', d, x^d_{p'})$, $t_b \leq d < t$ with p' on the left side of the circle, then p broadcasts the message to p^- and p^+ at stage $t + 1$.

∗ If $p = i_1$ (respectively $p = j_1$) receives from p^+ a message containing the triplet $(p', d, x^d_{p'})$, $t_b \leq d < t$ with p' on the left (resp. right) side of the circle, then p broadcasts the message to p^- and p^+ at stage $t + 1$.

- If $p = i_1$ (respectively $p = j_1$) receives from $p^- = S$ at some stage $t \in \{t_b, \dots, t_b + 2n^C - 4\}$ a message containing the triplet $(p', d, x^d_{p'})$, with $t_b \leq d < t$ and p' on the right (resp. left) side of the circle, then three cases are possible:
  - (i): If $p' \neq j_1$ (resp., $p' \neq i_1$), then p broadcasts it to both $p^- = S$ and $p^+ = i_2$ (resp., $p^+ = j_2$) at stage $t + 1$.
  - (ii): If $p' = j_1$ (resp., $p' = i_1$) and $p = i_1$ (resp., $p = j_1$) has not received the triplet (p, d, x^d_p), then $p = i_1$ (resp., $p = j_1$) broadcasts the triplet $(j_1, d, x^d_{j_1})$ (resp., $(i_1, d, x^d_{i_1})$) at stage $t + 1$.
  - (iii): If $p' = j_1$ (resp., $p' = i_1$) and $p = i_1$ (resp., $p = j_1$) has received the triplet (p, d, x^d_p), then $p = i_1$ (resp., $p = j_1$) broadcasts the triplet (p', d, y) at state $t + 1$, with y a random draw from $[0, 1]$.

  The intuition is that if p receives from S a message, which reads *"S claims that both i_1 and j_1 deviated at the same stage d,"*, then S must be deviating (since under unilateral deviations, at most one player deviates at each stage); hence, the protocol changes in the fact that i_1 and j_1 are not transmitted the message of the sender since they know the sender is deviating at this stage.
- **Autocorrecting past own deviations:** if p has received the triplet $(p', d, x^d_{p'})$ at stage $t \in \{t_b, \dots, t_b + 2n^C - 5\}$ but did not forward it at stages $t + 1, \dots, t + \Delta$, for some $\Delta \geq 1$, then he forwards it at stage $t + \Delta + 1$. In other words, the protocol requires a player to broadcast the triplet $(p', d, x^d_{p'})$ at stage $t + 1$ upon receiving it at stage t, to broadcast it at $t + 2$ if he fails to broadcast it at stage $t + 1$, to broadcast it at $t + 3$ if he fails to broadcast it at stage $t + 1$ and $t + 2$, etc., so that unless the player deviates at all stages $t' \geq t + 1$, the triplet is broadcast at some stage during the block.

The decoding rule. The decoding rule describes how messages are analyzed at the end of each block. Players not in $\mathcal{P}$ do not analyze their messages. We consider the players in $\mathcal{P}$. First, we say that a player knows a message m in two cases: either he is a neighbor of the sender and receives the sender's message at stage 1 , or he decodes a message as correct according to the protocol in some block.

At the beginning of block B_1, the sender and his two neighbors i_1 and j_1 *know* the message m broadcasted by the sender at stage $t = 1$. At the end of each block, only players

who do not yet know m analyze the message received during the block. Thus, only the players in $\mathcal{P} \setminus \{S, i_1, j_1\}$ analyze messages at the end of the block B_1. (The purpose of our arguments is to show that the set of players who know m at the end of a block is strictly expanding over time and, ultimately, includes the receiver.) Thus, we consider player p, who does not yet know m at the beginning of the block B_b. At the end of the block B_b, he analyzes the messages as follows:

- If during the block, p has received $(n^C - 1)$ times a grand message containing the same message $m \in M$ from his predecessor p^-, let say at stages $d^1, \ldots, d^{n^C-1}$, where $t_b \leq d^1 < d^2 < \cdots < d^{n^C-1} \leq t_b + 2n^C - 4$,
- and if p has not received by stage d^{n^C-1} at the latest from his successor p^+ the message $(p^-, d^1, x_{p^-}^{d^1})$ where $x_{p^-}^{d^1}$ matches the value of the authentication key received by p from p^- at stage $d1$,

then, player p *learns* the message m and starts the next block B_{b+1} as a player who *knows* m. Otherwise, player p does not learn the message. Moreover, once a player knows the message m, he knows it for all the subsequent blocks.

For all other histories, the strategies are left unspecified.

A.1.2. *Two key properties of the protocol.* The protocol we constructed has two key properties. The first property states that no player $p \in \mathcal{P}$ learns incorrectly; that is, if player $p \in \mathcal{P}$ learns a message, the message is indeed the one the sender has sent. Lemma 1 is a formal statement of that property.

Lemma 1. *Let $m \in M$ be the message broadcast by the sender to i_1 and j_1 at stage $t = 1$. If at most one player deviates from the protocol at each stage, then it is not possible for player $p \in \mathcal{P}$ to learn $m' \in \mathcal{M} \setminus \{m\}$.*

Proof of Lemma 1. By contradiction, we assume that player $p \in \mathcal{P}$ learns $m' \in M \setminus \{m\}$ at the end of some block B_b, $b \geq 1$. Without loss of generality, we assume that p is the first player to learn m' on the path from S to R, where p lies.

For player p to learn m', during the block B_b, it must be that player p has received a sequence of $n^C - 1$ grand messages from his predecessor p^-, say at stages $d^1, \ldots, d^{n^C-1}$ with $t_b \leq d^1 < d^2 < \cdots < d^{n^C-1} \leq t_b + 2n^C - 4$, such that (i) the first element of each of the $n^C - 1$ grand messages is m' and (ii) player p did not receive from p^+ the triplet $(p^-, d^1, x_{p^-}^{d^1})$ on or before stage d^{n^C-1}, where $x_{p^-}^{d^1}$ matches the value of the authentication key received from p^- at stage d^1.

Since we assume that $m' \neq m$, it must be that that p^- deviates at all stages $d^1, \ldots, d^{n^C-1}$, as we assume that p is the first player to learn m'. Therefore, since we consider at most one deviation in each stage, all players in $\mathcal{P} \setminus \{p^-\}$ are playing according to σ at all stages $d^1, \ldots, d^{n^C-1}$. It follows that player p^{--}, the predecessor of p^-, broadcasts triplet $(p^-, d^1, x^{d^1}_{p^-})$ to p^- and p^{---} at stage d^2 at the latest and that player p^{---} broadcasts it at stage d^3 at the latest, etc.[14]

Since there are $n^C - 2$ nodes other than p^- and p on the circle, player p^+ broadcasts the triplet $(p^-, d^1, x^{d^1}_{p^-})$ to players p and p^{++} at stage d^{n^C-1} at the latest. Thus, player p does not validate m' with that sequence of grand messages.

Since this is true for any such sequence, player p does not learn m' at block B_b. □

Lemma 1 states that no player on the circle learns an incorrect message. The next Lemma states that at least one new player learns the correct message at the end of each block, which guarantees that the receiver learns the correct message at the latest after $1 + (n^C - 3)(2n^C - 3)$ stages.

Lemma 2. *Let $m \in M$ be the message broadcasted by the sender to i_1 and j_1 at stage $t = 1$. Suppose that all intermediaries i_1 to i_k and j_1 to $j_{k'}$ know the message m at the beginning of the block B_b. If at most one player deviates from the protocol at each stage, then either intermediary i_{k+1} or intermediary $j_{k'+1}$ learns m at the end of block B_b with probability one.*

To be more precise, Lemma 2 states that for all $\sigma' \in \Sigma(\sigma)$, the subset of histories for which either intermediary i_{k+1} or intermediary $j_{k'+1}$ is learning m at the end of block B_b has a probability of one according to $\mathbb{P}_{\sigma'}$.

Proof of Lemma 2. Given a finite set M, we denote its cardinality by $|M|$. For simplicity, let $i := i_k$ and $j := i_{k'}$. We want to prove that either i^+ or j^+ learns the message at the end of the block $B_b = \{t_b, \ldots, t_b + 2n^C - 4\}$. The proof is by contradiction. Therefore, we assume that neither i^+ nor j^+ learns the message at the end of the block.

We fix a strategy profile $\sigma' \in \Sigma(\sigma)$. D^i denotes the stages where player i deviates from σ, D^{i^-} denotes the stages where player i^- deviates, D^j denotes the stages where player

[14]Notice that it is possible for p^{--} to broadcast the message $(p^-, d^1, x^{d^1}_{p^-})$ before stage d^2. For instance it is possible to have a stage $d^{1'}$, with $d^1 < d^{1'} < d^2$, such that (i) p^- is not deviating at stage $d^{1'}$ and sends either m or m_0 depending on whether he knows m and (ii) p^{--} is deviating at stage $d^{1'}$ by not transmitting $(p^-, d^1, x^{d^1}_{p^-})$ to p^- and p^{--}.

j deviates and D^{j^-} denotes the stages where player j^- deviates. According to the definition of $\Sigma(\sigma)$, the sets D^i, D^{i^-}, D^j and D^{j^-} are pairwise disjoints (by definition, we consider histories with unilateral deviations in each stage only). In particular,

$$|D^i| + |D^{i^-}| + |D^j| + |D^{j^-}| \leq |D^i \cup D^{i^-} \cup D^j \cup D^{j^-}| \leq 2n^C - 3. \tag{1}$$

Throughout, for any subset D of B_b, we write $\overline{D}$ for its complement in B_b. By definition, in all stages in $\overline{D^i}$, player i follows σ_i and deviates at all others. Let $\overline{D^i} := \{\overline{d}_1, \ldots, \overline{d}_\ell, \ldots, \overline{d}_{\ell_i^*}\}$, with $\overline{d}_\ell < \overline{d}_{\ell+1}$ for all ℓ. Note that $\ell_i^* = (2n^C - 3) - |D^i|$.

We assume that $\ell_i^* \geq n^C - 1$. Since player i follows the protocol at all stages in $\overline{D^i}$, player i^+ observes at least one sequence of messages such that m is broadcast $n^C - 1$ times by player i. Consider all sequences $(\overline{d}_{\ell_1}, \ldots, \overline{d}_{\ell_{n^C-1}})$ of distinct elements of $\overline{D^i}$ such that all sequences have $n^C - 1$ consecutive elements, that is, if $\overline{d}_\ell$ and $\overline{d}_{\ell'}$ are elements of the sequence, then all $\overline{d}_{\ell''}$ satisfy $\overline{d}_\ell < \overline{d}_{\ell''} < \overline{d}_{\ell'}$. By construction, there are $(\ell_i^* + 1) - (n^C - 1) = n^C - |D^i| - 1$ such sequences. All these sequences have different starting stages and, therefore, different ending stages. Recall that player i broadcasts m at all stages of these sequences.

We fix the sequence $(\overline{d}_{\ell_1}, \ldots, \overline{d}_{\ell_{n^C-1}})$. The protocol specifies that player i^+ learns m if and only if he has not received the correct authentication key $x_i^{d_{\ell_1}}$ from player i^{++} by the stage $\overline{d}_{\ell_{n^C-1}}$. Therefore, player i^+ does *not* learn m only if player i^- broadcasts the correct authentication key in some stage $d > \overline{d}_{\ell_1}$; the other players do not know the authentication key and the probability of guessing it correctly is zero. Moreover, since the protocol requires player i^- to broadcast $x_i^{d_{\ell_1}}$ only if player i broadcasts $m' \neq m$ in stage $\overline{d}_{\ell_1}$, which he does not, player i^- must be deviating. Therefore, $d \in D^{i^-}$.

Remember that player i^- can broadcast at most one authentication key about player i at each stage. Therefore, since there are $n^C - |D^i| - 1$ such sequences, player i^- must deviate at least $n^C - |D^i| - 1$ times for player i^+ to not learn m; that is,

$$|D^{i^-}| \geq n^C - |D^i| - 1. \tag{2}$$

It follows that

$$|D^i| + |D^{i^-}| \geq n^C - 1. \tag{3}$$

Now, we assume that $\ell_i^* < n^C - 1$. We have that $|D^i| = (2n^C - 3) - \ell_i^* > 2n^C - 3 - n^C + 1 = n^C - 2$, hence $|D^i| \geq n^C - 1$. Inequality (3) is also satisfied.

A symmetric argument applies to the pair of players j and j^-, hence,

$$|D^j| + |D^{j^-}| \geq n^C - 1, \tag{4}$$

since player j^+ does not learn m either. Summing Equations (3) and (4), we obtain

$$|D^i| + |D^{i^-}| + |D^j| + |D^{j^-}| \geq 2n^C - 2, \tag{5}$$

contrary to Equation (1). This completes the proof of Lemma 2. □

To conclude this proof, it is sufficient to invoke Lemmas 1 and 2, which guarantee that the receiver learns almost surely the message broadcast by the sender provided there are more than $n^C - 1$ blocks.

A.2. **Proof that Statement (4) implies Statement (1).** The proof is constructive and relies extensively on the use of the strongly reliable communication protocol $(\boldsymbol{\mathcal{M}}, \sigma)$ constructed above. The informal idea is to generate jointly controlled lotteries between the sender and one of the two intermediaries i_1 or j_1 to generate a recommendation, which is then (strongly) reliably transmitted to the receiver.

As explained in the main text, if A is finite, any distribution μ over A can be generated by a jointly controlled lottery. The idea is to partition the interval $[0, 1]$ into $|A|$ subintervals, where the length of the subinterval associated with a is $\mu(a)$. Let $f : [0, 1] \to A$, where $f(r) = a$ if r is in the subinterval associated with a. Note that $f^{-1}(a)$ is a Borel set and has measure $\mu(a)$. We consider two uniform random variables X and Y. Notably, the sum $X + Y \mod [0, 1]$, $x + Y \mod [0, 1]$, and $X + y \mod [0, 1]$ are also uniformly distributed on $[0, 1]$, regardless of the values of x and y. Therefore, if we let $\phi(x, y) = a$ if $x + y \mod [0, 1] \in f^{-1}(a)$, then we indeed generate μ. If Ω is finite, one can repeat the previous construction for each ω to obtain a way to generate any strategy as one that is jointly controlled. We have two technical complications to address. First, the set A is an arbitrary complete and separable metric space and not finite. Second, the set Ω is an arbitrary complete and metric space, hence we need to ensure that this construction can be completed in a measurable way. We first introduce a formal definition extending the notion of a jointly controlled lottery to strategies and then prove its existence in our framework.

Definition 3. *Let $f : \Omega \times \mathbb{B}_A \to [0,1]$. A jointly controlled kernel generating f is a triple (X, Y, ϕ) such that*

- *X is a measurable function from a probability space $(U, \mathcal{U}, \mathbb{P})$ to $([0,1], \mathbb{B}_{[0,1]})$,*
- *Y is a measurable function from a probability space $(U, \mathcal{U}, \mathbb{P})$ to $([0,1], \mathbb{B}_{[0,1]})$,*
- *X and Y are independent,*
- *and ϕ is a probability kernel function on $(\Omega \times [0,1] \times [0,1]) \times \mathbb{B}_A$*

such that for every $O \in \mathbb{B}_A$

(i) $\mathrm{E}[\phi(\omega, X, Y, O)] = f(\omega, O)$,

(ii) for every $x \in [0,1]$, $\mathrm{E}[\phi(\omega, x, Y, O)] = f(\omega, O)$,

(ii) and for every $y \in [0,1]$, $\mathrm{E}[\phi(\omega, Y, y, O)] = f(\omega, O)$,

When Ω is a singleton, μ is the formalization of what we informally described as a jointly controlled lottery.

Proposition 1. *For any probability kernel, $f : \Omega \times \mathbb{B}_A \to [0,1]$, there exists a jointly controlled kernel (X, Y, ϕ) generating f. Moreover, the kernel ϕ is degenerated in the sense that for every $(\omega, x, y) \in \Omega \times [0,1] \times [0,1]$, $\phi(\omega, x, y, .)$ has only one atom.*

Proof of Proposition 1. Let λ be the Lebesgue measure on $[0,1]$ and $f : \Omega \times \mathbb{B}_A \to [0,1]$ be a probability kernel. We assume that A is uncountable which is without loss of generality up to an embedding.

Let us assume first that $A = [0,1]$. We can define a quantile kernel: for every $\omega \in \Omega$, for every $x \in [0,1]$,

$$\theta(\omega, x) = \sup\{t \in \mathbb{Q}, f(\omega,]-\infty, t[) < x\} = \sup_{t \in \mathbb{Q}} t \mathbb{1}_{f(\omega,]-\infty, t[) < x}.$$

By construction for every $x \in [0,1]$, $\theta(., x)$ is measurable on Ω. Moreover, let Z be a uniform r.v over $[0,1]$; then, for every $\omega \in \Omega$, the probability distribution of $\theta(\omega, Z)$ is $f(\omega, .)$.

We can now define

- X is a measurable function from a probability space $(U, \mathcal{U}, \mathbb{P})$ to $([0,1], \mathbb{B}_{[0,1]})$,
- Y is a measurable function from a probability space $(U, \mathcal{U}, \mathbb{P})$ to $([0,1], \mathbb{B}_{[0,1]})$, independent of X

- for every $(\omega, x, y) \in \Omega \times [0,1] \times [0,1]$ and $O \in \mathbb{B}_A$,

$$\phi(\omega, x, y, O) = \begin{cases} 1 \text{ if } \theta(\omega, x + y \ mod\ 1) \in O, \\ 0 \text{ otherwise.} \end{cases}$$

Informally, the strategy computes z as $x + y$ modulo 1 and plays with probability one $\theta(\omega, z)$. The key observations are as follows. First, the sums $X + Y \mod [0,1]$, $x + Y \mod [0,1]$, and $X + y \mod [0,1]$ are all uniformly distributed on $[0,1]$. By construction, the distribution of $\phi(\omega, X, Y)$ over $\mathbb{B}_A$ is the same as $\theta(\omega, Z)$, with Z being a uniform r.v. Hence, $f(\omega, .)$. This concludes the proof for the case where $A = [0,1]$.

Finally, we can show the result for any separable metric space. By assumption, $(A, \mathbb{B}_A)$ is Borel standard; therefore, there is a Borel isomorphism ψ from $(A, \mathbb{B}_A)$ to $([0,1], \mathbb{B}_{[0,1]})$ such that the inverse is also measurable (Theorem 3.3.13, p. 99 in Srivastava, 2008). Hence, we can define the kernel

$$\widetilde{f} : \Omega \times \mathbb{B}_{[0,1]} \to [0,1]$$

by

$$\widetilde{f}(\omega, O) = f(\omega, \psi(O))$$

According to the previous section, we know that there exists a jointly controlled strategy $(X, Y, \widetilde{\phi})$ replicating $\widetilde{f}$. It follows that $(X, Y, \psi^{-1} \circ \widetilde{\phi})$ is a jointly controlled strategy replicating f.

□

We now explain how to PBE implement mediated communication on the network $\mathcal{N}$. Let $\tau^* : \Omega \to \Delta(A)$ be a canonical communication equilibrium of the direct communication game. From Proposition 1, for each ω, there exists a jointly controlled lottery $(X_\omega, Y_\omega, \phi_\omega)$, which generates $\tau^*(\omega)$.

We let $\mathcal{P}$ be the players on the two disjoint paths from the sender to the receiver. We now describe the strategies used in the communication game:

$t = 1$**:** The sender truthfully broadcasts the state ω to the intermediaries i_1 and j_1.

$t = 2$**:** The sender and intermediary i_1 each draw a random number in [0,1]. Let x (resp., y) be the number drawn by the sender (resp., intermediary i_1). The sender and the intermediary i_1 broadcast (simultaneously) their random numbers.

$t = 3, \ldots, 2 + (2n^C - 3)(n^C - 3)$**:** Players $p \in P$ execute the protocol $(\boldsymbol{\mathcal{M}}, \sigma)$ starting from its second stage, with i_1 in the role of the sender and j_1 in the role of the receiver;

the message to be transmitted is y. The first stage of the protocol $(\boldsymbol{\mathcal{M}}, \sigma)$ is not executed because the keys x and y were already broadcasted at $t = 2$. Accordingly, unilateral deviations of i_1 cannot affect the message learned by j_1 at the end of the protocol $(\boldsymbol{\mathcal{M}}, \sigma)$–$j_1$ thus learns the random number y.

$t = 3 + (2n^C - 3)(n^C - 3)$**:** The sender and the intermediaries i_1 and j_1 output a recommendation $a \in A$ according to the jointly controlled kernel (X, Y, ϕ); that is, the recommendation is the unique action in the support of $\phi(\omega, x, y)$ with ω the state broadcasted at $t = 1$. The three of them truthfully broadcast the recommendation.

$t = 4 + (2n^C - 3)(n^C - 3), \ldots, 5 + 2(2n^C - 3)(n^C - 3)$**:** Players execute in parallel and independently three copies of the protocol $(\boldsymbol{\mathcal{M}}, \sigma)$ with S, i_1 and j_1 in the role of the sender, respectively, and the message to be transmitted is the recommendation a. In the last stage, the receiver follows the recommendation made most often, if any. (If there is no majority, then he chooses an arbitrary action.)

Since at stage $t = 3 + (2n^C - 2)(n^C - 3)$, at most one of the three "senders" can deviate, the correct recommendation is sent at least twice. It follows that if the receiver is obedient, the receiver chooses the correct action in all histories consistent with unilateral deviations. Moreover, since the receiver observes neither ω nor x, he has no additional information about the state than in the direct communication game, hence he has an incentive to be obedient.

Clearly, in each iteration of the protocol $(\boldsymbol{\mathcal{M}}, \sigma)$, no intermediary has an incentive to deviate since it would result in the same expected payoff: indeed, since the protocol is strongly reliable, the "receiver" (of the protocol considered) learns the message sent by the "sender" (of the protocol considered) for all histories consistent with unilateral deviations. Additionally, the sender and intermediary i_1 have no incentive to deviate in stage $t = 2$ since this would not change the outcome of the jointly controlled lottery. Additionally, the sender and both intermediaries i_1 and j_1 have no incentive to deviate by broadcasting another recommendation than the one obtained by a jointly controlled lottery, as it would not be followed by the receiver who follows the recommendation made the other two (majority rule). Moreover, the sender has no incentive to deviate in the first stage when sending the state ω since the jointly controlled lottery generates the distribution of a canonical communication equilibrium of the direct game, and conditional on broadcasting ω in the first stage, the receiver receives $\tau^*(\omega)$ in all histories consistent with unilateral deviations, including deviations by the sender. Finally, the receiver has

no incentive to deviate either. If he stops the game earlier, then his expected payoff is weakly lower as a consequence of Blackwell's theorem. Indeed, the only informative message about ω is $\tau^*(\omega)$ and stopping earlier is a garbling of $\tau^*(\omega)$.

Sequential rationality. We now prove that the profile of strategies constructed above satisfies sequential rationality. Notice first that no deviation can be observed at stages $t = 1$ and $t = 2$ since all the messages are in the support of the players' strategies. Sequential rationality must then be proven starting from stage 3, where the remainder of the communication game is composed of a first protocol where i_1 is the sender and j_1 the receiver, then three simultaneous protocols, in which respectively the sender, i_1 and j_1 are the senders of these three protocols, and in all of them, the receiver remains the receiver.

As already argued above, in each of these subprotocols, sequential rationality is guaranteed at all histories consistent with at most one intermediary deviating in each stage of the communication game. We therefore focus our attention on all other histories, i.e., histories not in $\mathcal{H}(\sigma)$.

Each of these four subprotocols is run independently of one another; therefore, we consider them separately. For each subprotocol, we first consider all the intermediaries of this subprotocol $(i_1, \dots, i_K)$ and $(j_1, \dots, j_{K'})$ (with an abuse of notation, as for instance in the first protocol, when i_1 is the sender, the intermediaries are in fact $(i_2, \dots, j_1)$ and (S)). We treat the sender and receiver of this subprotocol separately.

Rebooting strategies. We say that player i reboots his strategy at period t if, from any history $h_i^t \notin \mathcal{H}_i(\sigma)$ onwards, he follows the protocol as if he knows that the message is m_0. That is, at history h_i^t, he broadcasts the message m_0, an authentication key x_i^t, and random triplets $(j, t_j, x_j^{t_j})$, $j \in \mathcal{N}_i$. At all subsequent histories consistent with at most one intermediary deviating from the protocol at each stage, player i continues to follow the protocol. That is, player i continues to broadcast m_0, authentication keys and triplet $(j, t_j, x_j^{t_j})$, as specified by the protocol when a player knows a message (here, it is m_0). At all other histories, player i reboots yet again his strategy, that is, player i continues to broadcast m_0, authentication keys and triplets, as if the multilateral deviation had not taken place.[15]

[15]Since our strongly reliable communication protocol is not defined after multilateral deviation, we cannot simply say that we reboot the strategy only once after a multilateral deviation. Indeed, if a new multilateral deviation occurs after some player has rebooted his strategy, his strategy is then not defined.

Beliefs. In history $h_i^t \notin \mathcal{H}_i(\sigma)$, player i believes that all other players on the same side of the circle reboot their strategies, while players on the other side of the circle as well as the sender and the receiver, continue to follow the protocol. In other words, player i believes that all other players on the same side of the circle have also observed a multilateral deviation, while players on the other side, as well as the sender and the receiver, have observed no deviations.

We now consider the sender. In all histories h_S^t, the sender continues to follow the protocol as if the observed multilateral deviations had not happened. However, he believes that all intermediaries reboot their strategies at period t, while the receiver continues to follow the protocol.

We now consider the receiver. The receiver continues to validate messages as he does in the protocol, i.e., he tests sequences of messages of length $n^C - 1$ received by his two predecessors and validates a message if he has received a sequence of $n^C - 1$ identical copies of the message and has not received the correct authentication on time (see the construction of the protocol for details). To complete the construction of the strategies, we assume that if the receiver validates $m \in M$ and $m_0 \notin M$, then he plays $\tau^*(m)$. Similarly, if he validates two different messages $(m, m') \in M \times M$ or (m_0, m_0) or no messages at all, he plays a best reply to his prior. In all histories, the receiver continues to follow the protocol as if the observed deviations Had not happened. He believes that all intermediaries reboot their strategies, while the sender continues to follow the protocol.

Finally, we assume that beliefs are independent between the different subprotocols, in that if an intermediary observes a multilateral deviation in one subprotocol only, he believes that other players on the same side of the circle reboot their strategies in that subprotocol only and not on the others.

Sequential rationality. In history $h_i^t \notin \mathcal{H}_i(\sigma)$, an intermediary expects the receiver to validate a message $m \in M$ from the other side and to validate the message m_0 from his side. Since the receiver takes the decision $\tau^*(m)$ when validating the messages $m \in M$ and $m_0 \notin M$, the intermediary cannot deviate profitably (as, regardless of his play, the receiver validates m from the other side). Therefore, rebooting the strategy is optimal. Similarly, since the sender expects the intermediaries i_K and $j_{K'}$ to reboot their strategies, he expects the receiver to play a^* and, therefore, cannot profitably deviate. The same applies to the receiver.

APPENDIX B. THE CASE IN WHICH THE SENDER AND THE RECEIVER ARE DIRECTLY CONNECTED.

We assume that the sender and the receiver can also communicate directly. We claim that if there are at least two intermediaries, labeled 1 and 2, such that the sender, the receiver and the two intermediaries are on a "circle," then PBE implementation of mediated communication is possible.[16]

We now present an informal proof. As a preliminary observation, we note that on the circle, it must be that either the two intermediaries are on two disjoint paths from the sender to the receiver or are on the same path. See the two networks in Figure 5 for an illustration.

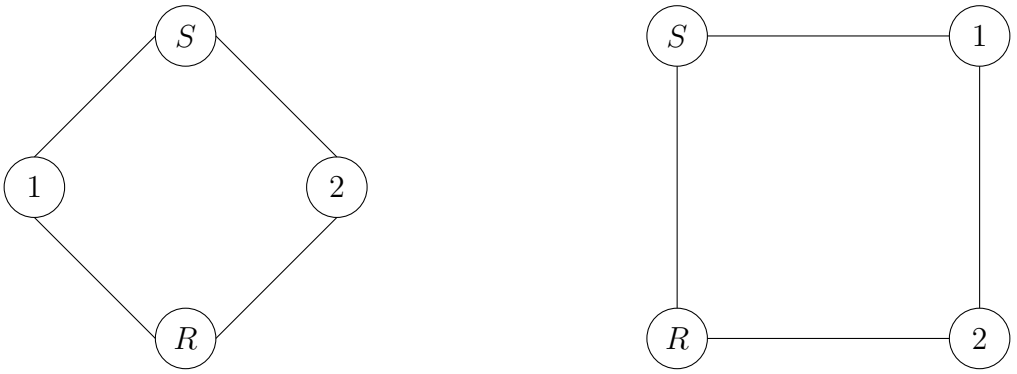

FIGURE 5. $\mathcal{N}$ (left) and $\mathcal{N}^*$ (right)

In the former case (network $\mathcal{N}$), Theorem 1 applies verbatim. In the latter case (network $\mathcal{N}^*$), we need to modify the protocol of Theorem 1 to guarantee that intermediary 2 learns the state ω without the receiver learning it. Once the sender and the two intermediaries know ω, we can then use the second and third phases of the protocol constructed in Section 3 to implement the communication equilibrium.

We modify the first phase as follows. We first let intermediary 1 broadcast an encryption key k to the sender and intermediary 2. The sender then encrypts the state ω with the encryption key k and transmits the encrypted message to intermediary 2. The transmission of the encrypted message is achieved via the strongly reliable protocol constructed in Section 3.

It remains to be argued that at the end of the first phase, the sender and the two intermediaries know ω. It is clear for the sender and intermediary 1. For intermediary 2, at the end of the first phase, he knows the encryption key and the encrypted message and, thus learns the state ω. This argument completes the informal "proof."

[16]A circle is a collection of nodes such that all pairs of nodes have two disjoint paths between them.

We note that either the sender or an intermediary may attempt to reveal the state to the receiver. We cannot exclude this possibility. However, as already explained, the equilibria we construct are such that the receiver does not expect the sender and intermediaries to do so and thus consider their attempts to be gibberish.

REFERENCES

[1] Ambrus, A., and E. Azevedo, and Y. Kamada, Hierarchical Cheap Talk, 2014, Theoretical Economics, 8, pp. 233-261.

[2] Aumann R.J., Subjectivity and Correlation in Randomized Strategies, 1974, Journal of Mathematical Economics, 1, pp. 67-96.

[3] Aumann R.J., and S. Hart, Long Cheap Talk, 2003, Econometrica, 71, pp. 1619-60.

[4] Bárány, I., Fair Distribution Protocols or How The Players Replace Fortune, 1992, Mathematics of Operations Research, 17 (2), pp. 327-340.

[5] Baron D., and D. Besanko, Matrix Organization, 1996, Working Paper, Stanford University.

[6] Beimel, A., and M. Franklin, Reliable Communication Over Partially Authenticated Networks. 1999, Theoretical Computing Science, 220, pp. 185-210.

[7] Ben-Porath E., Correlation Without Mediation: Expanding the Set of Equilibria Outcomes by Cheap Pre-Play, 1998, Journal of Economic Theory, 80 pp. 108-122.

[8] Ben-Porath E., and M. Kahneman, Communication in Repeated Games with Private Monitoring, 1996, Journal of Economic Theory, 70 pp. 281-297.

[9] Crawford, V. P., and J. Sobel, Strategic Information Transmission, 1982, Econometrica. 50, pp. 1431-1451.

[10] Dolev, D., and C. Dwork, and O. Waarts, and M. Yung, Perfectly Secure Message Transmission, 1993, Journal of Association of Computing Machinery, 40, pp. 17-47.

[11] Forges, F., Correlated Equilibria in a Class of Repeated Games with Incomplete Information, 1985, International Journal of Game Theory, 14, pp. 129–149.

[12] Forges, F., An Approach to Communication Equilibria, 1986, Econometrica, 54, pp. 1375-1385.

[13] Forges, F., Universal Mechanisms, 1990, Econometrica, 58, pp. 1341-1364.

[14] Forges, F., Games with Incomplete Information: From Repetition to Cheap Talk and Persuasion, 2020, Annals of Economics and Statistics 137, pp. 3-30.

[15] Forges, F. and P. Vida, Implementation of Communication Equilibria by Correlated Cheap Talk: The Two-Player Case, 2013, Theoretical Economics, 8, pp. 95-123.

[16] Franklin, M., and R.N. Wright, Secure Communication in Minimal Connectivity Models, 2000, Journal of Cryptology, 13, pp. 9-30.

[17] Galbraith, J. R., Designing Matrix Organizations that Actually Work: How IBM, Procter & Gamble and Others Design for Success, 2009, Jossey-Bass Business & Management.

[18] Gerardi, D., Unmediated Communication in Games with Complete and Incomplete Information, 2004, Journal of Economic Theory, 114, pp. 104-131.

[19] Harris, M., and A. Raviv, Organization Design, 2002, Management Science, 48, pp.

[20] Ivanov, M., Communication via a Strategic Mediator, 2010, Journal of Economic Theory 145, pp. 869–884.

[21] Laclau, M., Communication in Repeated Network Games With Imperfect Monitoring, 2014, Games and Economic Behavior, 87, pp. 136-160.

[22] Laclau, M., A Folk Theorem for Repeated Games Played on a Network, 2012, Games and Economic Behavior, 76 , pp. 711-737.

[23] Linial, N., Game-Theoretic Aspects of Computing, 1994, Handbook of Game Theory with Economic Applications, Aumann, R.J. and Hart, S. (eds), 2 , pp. 1339–1395.

[24] Migrow, D., Designing Communication Hierarchies, 2021, Journal of Economic Theory 198, pp. 1-30.

[25] Myerson, R.B., Multistage Games with Communication, 1986, Econometrica, 54, pp. 323-358.

[26] Myerson, R.B., Game Theory: Analysis of Conflicts, 1991, Harvard University Press.

[27] Renault, J., and T. Tomala, Repeated Proximity Games, 1998, International Journal of Game Theory, 27, pp. 93-109.

[28] Renault, J., and T. Tomala, Learning the State of Nature in Repeated Game with Incomplete Information and Signals, 2004, Games and Economic Behavior, 47, pp. 124-156.

[29] Renault, J., and T. Tomala, Probabilistic Reliability and Privacy of Communication Using Multicast in General Neighbor Networks, 2008, Journal of Cryptology, 21, pp. 250-279.

[30] Renault, J., and L. Renou, and T. Tomala, Secure Message Transmission on Directed Networks, 2014, Games and Economic Behavior, 2014, 85, pp. 1-18.

[31] Renou, L., and T. Tomala, Mechanism Design and Communication Networks, 2012, Theoretical Economics, 7, pp. 489-533.

[32] Rivera, T., Incentives and The Structure of Communication, 2018, Journal of Economic Theory, 175, pp. 201-247.

[33] Schröter, A., Distribution of Decision Power in Matrix Organizations: A Qualitative Survey, 2014, Thesis, University of Gloucester.

[34] Srivastava, S. M., A course on Borel sets, 2008, Vol. 180, Springer Science & Business Media.

[35] Tomala, T., Fault Reporting in Partially Known Networks and Folk Theorems, 2011, Operations Research, 59, pp. 754-763.

[36] Wolitzky, A., Communication with Tokens in Repeated Games on Networks, 2015, Theoretical Economics, 10, pp. 67-101.

MARIE LACLAU, HEC PARIS AND GREGHEC-CNRS, 1 RUE DE LA LIBÉRATION, 78351 JOUY-EN-JOSAS, FRANCE

Email address: laclau(at)hec.fr

LUDOVIC RENOU, QUEEN MARY UNIVERSITY OF LONDON AND CEPR, MILES END, E1 4NS, LONDON, UK

Email address: lrenou.econ(at)gmail.com

XAVIER VENEL, LUISS GUIDO CARLI UNIVERSITY, 32 VIALE ROMANIA, 00197 ROME, ITALY

Email address: xvenel(at)luiss.it