

Communication with Tokens in Repeated Games on Networks*

Alexander Wolitzky

Stanford University

February 1, 2013

Abstract

A key obstacle to coordination and cooperation in many networked environments is that behavior in each bilateral relationship is not observable to individuals outside that relationship: that is, information is *local*. This paper investigates when players can use communication to replicate any outcome that would have been sustainable were this information public. A benchmark result is that if only cheap talk communication is possible then public information can only be replicated if the network is 2-connected: that is, if no player can prevent the flow of information to another. In contrast, the main result is that public information can always be replicated if in addition to cheap talk the players have access to undifferentiated tokens that can be freely transferred among neighbors (which bear some resemblance to certain models of fiat money). Necessary conditions are provided for such tokens to expand the equilibrium payoff set, relative to what would be achievable without explicit communication or with cheap talk communication only.

*For helpful comments, I thank Daron Acemoglu, Nageeb Ali, Susan Athey, Boaz Barak, Gabe Carroll, Sylvain Chassang, Matt Elliott, Glenn Ellison, Ben Golub, Adam Kalai, Chris Phelan, David Rahman, Ilya Segal, Andy Skrzypacz, Joel Sobel, Juuso Toikka, Rakesh Vohra, Randy Wright, and seminar audiences at MIT, Microsoft Research, the Minneapolis Federal Reserve, UCLA, the 2012 Tel Aviv Economic Theory Conference in Honor of Jerry Green, and the 2012 Chicago Federal Reserve Workshop on Money, Banking, Payments, and Finance.

1 Introduction

Consider three people—1, 2, and 3—arranged on a line: 1 and 2 have a relationship, and 2 and 3 have a relationship, but 1 and 3 do not. In this situation, 1 and 3 might hope to keep 2 on good behavior by threatening “community enforcement”: if 2 cheats 1, then 3 cheats 2. But if 2 cheats 1, how does 3 find out? She doesn’t have a relationship with 1, and 2 clearly can’t be trusted to tell her. So the group has a problem.

In this example, the obstacle to sustaining cooperation is that information about individuals’ past behavior in a bilateral relationship is *local*: it is common knowledge within the relationship, but is not observable to outsiders. In addition, letting the players communicate locally does not enable them to sustain certain outcomes—like “cooperation”—that would have been sustainable if this information were public to all players. In the language of this paper, local communication does not *replicate* public information in these examples.¹

The goal of this paper is to compare two communication technologies—cheap talk and physical tokens—in terms of their ability to replicate public information. I study a fairly general model of repeated games on networks, in which monitoring is public within relationships but nonexistent across relationships (*locally public monitoring*), and say that a given communication technology *replicates* public information if it enables the players to sustain any payoff vector that would have been sustainable if monitoring were public to all players. A benchmark result is that if only cheap talk communication with one’s neighbors is available then public information can be replicated only if the network is 2-connected (i.e., it remains connected after any node is removed).² The intuition is simple: under 2-connectedness, there are always at least two independent paths through which a piece of information can reach each player, so no single player can prevent information from reaching another (i.e.,

¹This informational impediment to community enforcement is potentially widely relevant, as community enforcement on networks is thought to be a key feature of the economics of risk-sharing (Bloch, Genicot, and Ray, 2008; Ambrus, Möbius, and Szeidl, 2010), favor-trading (Karlan, Möbius, Rosenblat, and Szeidl, 2009; Jackson, Rodriguez-Barraquer, and Tan, 2011), and trade without external enforcement (Milgrom, North, and Weingast, 1990; Dixit, 2003; Greif, 2006).

²This is related to a result of Renault and Tomala (1998). The precise connection is discussed below.

there are no “information gatekeepers”), and therefore the players can coordinate as well as if all information were public. Conversely, if the network is not 2-connected then there are games in which the local nature of monitoring imposes costs on the players, as the above example shows.

Indeed, this example suggests that to overcome local monitoring players must be allowed not only to talk but also to exchange some form of *evidence*. In this paper, I focus on a very specific form of evidence: the players are endowed with undifferentiated, divisible tokens that they can freely transfer to their neighbors. The crucial difference between tokens and talk is that one player cannot send another more tokens than she has, while a player can always send any cheap talk message. In particular, “talk” messages can be manipulated arbitrarily, while “token” messages can only be manipulated downward.

The main result of the paper is that public information can always be replicated with tokens: that is, tokens allows players to overcome the “information gatekeeper” problem associated with non-2-connected networks. The main idea is to initially endow “leaf players”—like 1 and 3 in the examples—with tokens, and to endow “non-leaf players”—like 2—with none. Non-leaf players must then obtain tokens from leaf players in order to convince others that they have behaved well, which disciplines their behavior. In particular, non-leaf players are prevented from cheating some leaf players while concealing this information from others.³ The result is presented in quite a general setting, however, which necessitates the use of somewhat complicated sequences of tokens transfers to ensure that non-leaf players cannot misrepresent their information.

I then apply this result to study when tokens are *essential*, in that the equilibrium payoff set is strictly larger with tokens than with no communication, or *strongly essential*, in that the equilibrium payoff set is strictly larger with tokens than with cheap talk alone.⁴ I show that a sufficient condition for tokens to be essential is that the network contains a “nice” subnetwork, which is a subtree in which every bilateral relationship has a product structure (Fudenberg and Levine, 1994) and in which there is some payoff vector that can be sustained in equilibrium with public monitoring that cannot be sustained in a “locally

³However, the result holds for any strictly positive vector of initial token endowments, so it is not actually necessary that non-leaf players start with no tokens.

⁴The terminology here is borrowed from the literature on monetary theory (e.g., Lagos and Wright, 2008).

public equilibrium” (a generalization of perfect public equilibrium) with private monitoring. In many games, the condition that the network contains a nice subnetwork reduces to the condition that it contains a subtree of size at least three.

I study tokens rather than some other form of evidence for two reasons. First, tokens are intuitively a fairly minimal form of evidence. Allowing more sophisticated forms of evidence, like tokens that are tagged with different colors, or letters with signatures that cannot be forged, would only make the positive results of this paper easier to prove. Conversely, the main result fails if—contrary to my assumptions—initial token endowments are uncertain or tokens are indivisible. Second, tokens are inspired by the “tangible useless objects” (Wallace, 2001) used to model fiat money in the literature on the microfoundations of money (Kiyotaki and Wright, 1989, 1993).⁵ Unlike that literature, this paper is not in any way intended to provide a theory of how money is used in reality. However, examining the limits of what agents can achieve by transferring abstract tokens in arbitrarily complicated ways may be informative about what restrictions on agents’ information or behavior may be useful in monetary models. I briefly discuss this possibility in the conclusion.

The paper proceeds as follows: Section 2 relates the paper to the literatures on repeated games, networks, and the microfoundations of money. Section 3 presents the model. Section 4 contains benchmark results on replicating public information with cheap talk. Section 5 presents the main result on replicating public information with tokens. Section 6 presents examples showing that various conditions for the main result cannot be dispensed with. Section 7 shows how the main result can be applied to show when tokens are essential in a broad class of games. Section 8 concludes. Omitted proofs are contained in the Appendix.

2 Related Literature

The seminal paper on community enforcement in repeated games is Kandori (1992), who shows that cooperation is sustainable in the repeated prisoner’s dilemma with anonymous random matching with a simple form of hard evidence: exogenously determined labels, such as “guilty” or “innocent” (cf Ellison, 1994; Okuno-Fujiwara and Postlewaite, 1995).

⁵For example, this paper is closely related to Kocherlakota (1998, 2002), as discussed below.

However, most of the subsequent literature on community enforcement has not considered hard evidence.⁶ There is also a literature on the folk theorem in general private monitoring repeated games with communication, dating back to Compte (1998) and Kandori and Matsushima (1998). The main differences between my paper and this literature is that I restrict attention to repeated games on networks and compare the equilibrium payoff set with different communication technologies for a fixed discount factor.⁷

There is also a rapidly growing literature on repeated games on networks, which often emphasizes the role of communication. The folk theorems of Ben-Porath and Kahneman (1996) and Renault and Tomala (1998) are related to the benchmark results of Section 4 and are discussed there.⁸ Most of the rest of the literature studies more specific games. For example, Ahn and Suominen (2001) and Balmaceda and Escobar (2011) study how local communication among buyers can dissuade a seller from providing a low-quality good, and Lippert and Spagnolo (2011) and Ali and Miller (2012) study how local communication can help sustain cooperation in a repeated prisoner’s dilemma.

Also related is the large computer science-based literature on secure information transmission in networks. See Linial (1994) for a survey aimed at game theorists and see Tomala (2011), Renou and Tomala (2011), and Renault, Renou, and Tomala (2012) for recent contributions by game theorists. More broadly, Koessler and Forges (2008) survey the literature on multistage communication with certifiable information, and Forges (2009) surveys the literature on implementing communication equilibrium outcomes with private communication. In particular, the latter paper discusses how communication equilibrium outcomes may be implemented using private authentication keys or sealed envelopes (Ben-Porath, 1998; Krishna, 2007; Izmalkov, Lepinski, and Micali, 2011). However, to the best of my knowledge, no papers in this literature consider communication technologies resembling physical tokens. From this perspective, one interpretation of the results of this paper is that they show that undifferentiated tokens can often substitute for private authentication keys or sealed

⁶See Fujiwara-Greve, Okuno-Fujiwara, and Suzuki (2012) for a recent exception.

⁷McLean, Obara, and Postlewaite (2012) investigate when players in private monitoring repeated games are willing to publicly report their observations. Their results rest on players being “informationally small,” which is not the case in my model.

⁸A recent series of papers by Laclau (2012a, 2012b, 2012c) provides additional folk theorems for repeated games on networks with various communication technologies. None of these technologies resemble physical tokens.

envelopes in facilitating secure information transmission in networks.⁹

Finally, this paper relates to the large literature on the microfoundations of money. Much of this literature is concerned with the informational role of money—often modeled as undifferentiated physical tokens—albeit in models that are very different from mine. In particular, I provide sufficient conditions for tokens to be essential in games with a finite, non-anonymous population of players interacting on a fixed network, relative to what could be achieved with cheap talk alone, when tokens may be used in arbitrarily complicated ways. In contrast, most of the monetary theory literature considers games with a continuum of anonymous players interacting at random, does not compare money with cheap talk, and focuses on simple exchanges of money for goods; for example, this is the setting in Kiyotaki and Wright (1993).¹⁰ A natural question here is why models with non-anonymous agents have any relevance for monetary theory, given that the fact that money is used “anonymously” is sometimes taken as one of its defining characteristics (e.g., Ostroy and Starr, 1974). While this is a hard question to answer a priori, the fact that money is often repeatedly exchanged in non-anonymous, long-run relationships (risk-sharing, inter-bank lending, etc.) raises the possibility that models with anonymous agents may not tell the whole story.

A prominent paper on monetary theory that shares with mine the goal of comparing physical tokens with other information technologies is Kocherlakota (1998). Kocherlakota’s main result is that money is often inessential in the presence of a form of public information.¹¹ In contrast, I show that tokens are often essential when only private information is avail-

⁹Relative to this literature, tokens are a way of making a player’s message set depend on the past messages she has sent and received. If a player’s message set could be made to depend on past messages in an arbitrary way, ensuring truthful information transmission would be trivial: simply specify that a player must pass on all messages she receives. Thus, the advantage of tokens per se is that they are a natural and easily interpretable way of introducing a dependence of message sets on past messages.

¹⁰There are some exceptions, however. Araujo (2004) adapts the arguments of Kandori and Ellison to show that money is essential in sufficiently large finite games with anonymous random matching. Aliprantis, Camera, and Puzzello (2007) present a model with an infinite but non-anonymous population where money is essential even though players occasionally meet in centralized markets. Kocherlakota and Wallace (1998) show that money is essential with a continuum of players and random matching in the presence of sufficiently unreliable public monitoring of individual actions. Corbae, Temzelides, and Wright (2003) investigate the essentiality of money in a model with directed matching that in some cases resembles trade on a network.

¹¹More precisely, Kocherlakota’s notion of *memory* is perfect information about one’s partners’ past play, their partners’ past play, and so on. The idea that a primary role of money is replicating public information (“memory,” “record-keeping”) goes back at least at Starr (1972), Ostroy (1973), and Ostroy and Starr (1974).

able. Kocherlakota also gives an example in which (indivisible) money can replicate public information and an example in which it cannot. In my model, tokens can replicate public information quite generally; the primary reason for this difference is that players in my model are non-anonymous, which makes tokens—and repeated game effects more generally—much more powerful than in Kocherlakota’s model.

In a follow-up paper, Kocherlakota (2002) allows money to be infinitely divisible and shows that undifferentiated money can then replicate public information if money holdings are observable (the “one-money theorem”), and that tagged money (e.g., red money and black money) can replicate public information even if—as in my model—players can conceal money (the “two-money theorem”).¹² In contrast, in my model plain, undifferentiated tokens can replicate public information even if tokens are concealable.¹³ The key difference is again that my players are non-anonymous and can make more complicated sequences of transfers.¹⁴

3 Model

This section describes the repeated game without communication and the notion of replicating public information. I add cheap talk to the model in Section 4 and add tokens in Section 5.

Players: There is a finite set of players $N = \{1, \dots, n\}$ arranged on an undirected and connected network $L \subseteq P_2(N)$, the set of 2-element subsets of N , where $\{i, j\} \in L$ denotes a link between players i and j .¹⁵ The network is fixed over time and players “know” the entire network. The network will determine the structure of players’ actions, payoffs, information, and—in subsequent sections—communication. In particular, it will become clear that the assumption that L is connected is essentially without loss of generality, as the fact that players only “interact” with their neighbors implies that if L is not connected one

¹²A very similar idea appears in Townsend (1987). See also Townsend (1980) for a canonical monetary theory model emphasizing “spatial separation” of agents.

¹³I also show that in my model infinite divisibility of tokens can be replaced by targeted disbursements of tokens from a “central planner.”

¹⁴However, even with tagged tokens my results would not follow from Kocherlakota’s. Among other differences, Kocherlakota’s model involves “trading mechanisms,” while there are no mechanisms or contracts in my model (i.e., all transfers of tokens are completely voluntary).

¹⁵Links are denoted with braces rather than parentheses to emphasize that $\{i, j\}$ and $\{j, i\}$ refer to the same link.

can replicate the analysis on each connected component of L . Let $N_i = \{j : \{i, j\} \in L\}$ be the set of player i 's neighbors, and let $d(i, j)$ be the distance (shortest path length) between players i and j .

Stage game: Player i 's stage-game action set is $A_i = \prod_{j \in N_i} A_{i,j}$, where the $A_{i,j}$ are arbitrary finite sets interpreted as player i 's possible actions toward player j . There is a set of signal profiles $Z = \prod_{\{i,j\} \in L} Z_{i,j}$, where the $Z_{i,j} = Z_{j,i}$ are arbitrary finite sets interpreted as the signals that can be generated by the interaction between players i and j . It is assumed that the signal $z_{i,j}$ is “locally public,” in that it is identically equal to $z_{j,i}$ but is completely uninformative about any other $z_{i',j'}$. In particular, there are probability distributions $\pi_{i,j}(\cdot | a_{i,j}, a_{j,i}) = \pi_{j,i}(\cdot | a_{j,i}, a_{i,j})$ such that the probability of signal $z_{i,j}$ conditional on action pair $(a_{i,j}, a_{j,i})$ is $\pi_{i,j}(z_{i,j} | a_{i,j}, a_{j,i})$, independent of the signal realizations for other pairs of players, so that the probability of signal profile $z = (z_{i,j})_{\{i,j\} \in L}$ given action profile $a = (a_i)_{i \in N}$ is given by $\pi(z | a) = \prod_{\{i,j\} \in L} \pi_{i,j}(z_{i,j} | a_{i,j}, a_{j,i})$.¹⁶ Hence, Player i 's stage-game expected payoff is $u_i(a) = \sum_{j \in N_i} \sum_{z_{i,j} \in Z_{i,j}} \pi_{i,j}(z_{i,j} | a_{i,j}, a_{j,i}) u_{i,j}^*(z_{i,j}, a_{i,j})$, where $u_{i,j}^* : Z_{i,j} \times A_{i,j} \rightarrow \mathbb{R}$ gives player i 's realized payoff from her interaction with player j . To save on notation, let $u_{i,j}(a_{i,j}, a_{j,i}) = \sum_{z_{i,j} \in Z_{i,j}} \pi_{i,j}(z_{i,j} | a_{i,j}, a_{j,i}) u_{i,j}^*(z_{i,j}, a_{i,j})$, and note that $u_i(a) = \sum_{j \in N_i} u_{i,j}(a_{i,j}, a_{j,i})$. Thus, $u_{i,j} : A_{i,j} \times A_{j,i} \rightarrow \mathbb{R}$ gives player i 's expected payoff from her interaction with player j . For $\{i, j\} \in L$, I will refer to the two-player game $(A_{i,j}, A_{j,i}, Z_{i,j}, \pi_{i,j}, u_{i,j}, u_{j,i})$, which captures the direct relationship between i and j , as the (i, j) -game.

I assume throughout the paper that each (i, j) -game has a mutual-minmax Nash equilibrium. That is, I assume that every mixed action set $\Delta(A_{i,j})$ contains an element $\alpha_{i,j}^*$ such that the mixed action profile $\alpha^* = (\alpha_i^*)_{i \in N} = \left((\alpha_{i,j}^*)_{j \in N_i} \right)_{i \in N}$ is a stage-game Nash equilibrium and

$$u_{i,j}(\alpha_{i,j}^*, \alpha_{j,i}^*) = \min_{\alpha_{j,i} \in \Delta(A_{j,i})} \max_{\alpha_{i,j} \in \Delta(A_{i,j})} u_{i,j}(\alpha_{i,j}, \alpha_{j,i}) \text{ for all } \{i, j\} \in L.$$

This assumption ensures that the worst possible punishments can be delivered “link by link,” and thus do not require punishers to coordinate. It is needed for my results, because

¹⁶Recall that $\{i, j\} = \{j, i\}$, so there are the same number of terms in this product as there are links in the network.

generally an outsider will be able to tell when a deviation occurs in the relationship between two players but will not be able to tell which one of them deviated.

Repeated game: The players play a repeated game in discrete time. At the beginning of period $t \in \{0, 1, \dots\}$, each player i chooses an action $a_{i,t} \in A_i$. The signal z_t is then drawn from $\pi(\cdot|a)$, payoffs are realized, and player i observes $(z_{i,j,t})_{j \in N_i}$.¹⁷ Letting $h_{i,t} = (a_{i,t}, (z_{i,j,t})_{j \in N_i})$, player i 's time t history is $h_i^t = (h_{i,\tau})_{\tau=0}^{t-1}$ for $t \geq 1$, and every player has trivial initial history $h_i^0 = h^0$; in addition, let $h_{i,j,t} = (a_{i,j,t}, z_{i,j,t})$ so that player i 's time t (i, j) -game history is $h_{i,j}^t = (h_{i,j,\tau})_{\tau=0}^{t-1}$. Letting H_i^t be the set of player i 's time- t histories, a behavior strategy of player i 's is a map $\sigma_i : H_i^t \rightarrow \Delta(A_i)$, and player i 's behavior strategy in the (i, j) -game, $\sigma_{i,j} : H_i^t \rightarrow \Delta(A_{i,j})$, is given by projecting $\sigma_i(h_i^t)$ onto $\Delta(A_{i,j})$. Players have common discount factor $\delta \in (0, 1)$. Denote the resulting repeated game by Γ_{PRI} , where the subscript PRI emphasizes that signal $z_{i,j}$ is private to the pair of players $\{i, j\}$ (though it is locally public between i and j).

Solution concept: The solution concept for the benchmark results concerning cheap talk communication is sequential equilibrium (SE).¹⁸ For the main result concerning communication with tokens, it will be important that tokens are infinitely divisible. This makes action spaces infinite, which necessitates using perfect Bayesian equilibrium (PBE) for the main result. I define an appropriate version of PBE for this model in Section 5.

Replicating public information: Let Γ_{PUB} be the game in which the entire signal z is public. That is, Γ_{PUB} is derived from Γ_{PRI} by letting $h_{i,t}$ equal $(a_{i,t}, z_t)$ rather than $(a_{i,t}, (z_{i,j,t})_{j \in N_i})$. Let E_{PUB} be the set of SE payoffs of game Γ_{PUB} , and let \tilde{E}_{PUB} be the set of PBE payoffs of game Γ_{PUB} . Below, I will define games Γ_{PRI}^{PUBCT} , Γ_{PRI}^{PRICT} , and Γ_{PRI}^{TOK} by adding public cheap talk, private (i.e., local) cheap talk, or tokens to the game Γ_{PRI} . The solution concept for Γ_{PRI}^{PUBCT} and Γ_{PRI}^{PRICT} is SE, and the corresponding SE payoff sets are denoted E_{PRI}^{PUBCT} and E_{PRI}^{PRICT} , respectively. The solution concept for Γ_{PRI}^{TOK} is PBE, and the corresponding PBE payoff set is denoted \tilde{E}_{PRI}^{TOK} . I will say that public cheap talk (resp., private cheap talk, tokens) *can replicate public information* if $E_{PRI}^{PUBCT} \supseteq E_{PUB}$

¹⁷Thus, player i observes her own payoff.

¹⁸As usual, sequential equilibrium in repeated games with finite information sets in every period is defined by putting the product topology on the space of beliefs.

(resp., $E_{PRI}^{PRICT} \supseteq E_{PUB}$, $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PUB}$).¹⁹ Informally, communication replicates public information if any payoff vector that can be attained in equilibrium when all local information is made public can also be attained with communication.

4 Replicating Public Information with Cheap Talk

This section establishes benchmark results on when cheap talk may be used to replicate public information. I start with the fairly trivial benchmark of public cheap talk and then move to the more significant benchmark of private cheap talk. The results of this section are broadly similar to results in the literature and I do not view them as primary contributions of this paper; rather, they give a point of departure for the main analysis of Sections 5 through 7.

4.1 Public Cheap Talk

A game with public cheap talk $\Gamma_{PRI}^{PUBCT}(Y)$ is derived by augmenting the game Γ_{PRI} with a finite message set $Y = (Y_1, \dots, Y_n)$ such that after players observe their private signals they simultaneously send public messages $y_i \in Y_i$. Formally, letting $h_{i,t} = (a_{i,t}, (z_{i,j,t})_{j \in N_i}, (y_{j,t})_{j \in N})$, there are now two kinds of histories for every time t , denoted $h_i^{t-} = (h_{i,\tau})_{\tau=0}^{t-1}$ (called *action histories*) and $h_i^{t+} = ((h_{i,\tau})_{\tau=0}^{t-1}, a_{i,t}, (z_{i,j,t})_{j \in N_i})$ (called *communication histories*), and a strategy maps action histories to $\Delta(A_i)$ and maps communication histories to $\Delta(Y_i)$.²⁰ Let $E_{PRI}^{PUBCT}(Y)$ be the SE payoff set of $\Gamma_{PRI}^{PUBCT}(Y)$, and let $E_{PRI}^{PUBCT} = \bigcup_Y E_{PRI}^{PUBCT}(Y)$, where the union is taken over all finite sets Y .

The first benchmark result is that public cheap talk can always replicate public information. The proof is very simple: Take a SE strategy profile σ^{PUB} in the public monitoring game Γ_{PUB} . Specify that after every round of play the players publicly report what signals they observe and then play according to σ^{PUB} , taking the reported signals as the true ones. If any reports disagree, play the mutual-minmax profile α^* forever. This strategy profile

¹⁹There are games for which these inclusions are strict. I omit the proof of this fact, since it is not used in the paper.

²⁰I continue to denote generic histories by h_i^t . That is, h_i^t may denote either an action history or a communication history.

yields the same payoffs as σ^{PUB} , and it can be shown to be a SE profile.²¹

Theorem A *Public cheap talk can replicate public information (i.e., $E_{PRI}^{PUBCT} \supseteq E_{PUB}$).*

In addition, the proof of Theorem A shows that there is a single finite message set Y such that $E_{PRI}^{PUBCT}(Y) \supseteq E_{PUB}$. The same will be true of the later results on communication through private cheap talk or tokens. Thus, in all cases there is no need to tailor the message set (or the initial endowment of tokens) to the desired payoff vector.

Theorem A is related to Theorem 1 of Ben-Porath and Kahneman (1996), which establishes the folk theorem for repeated games with public communication where each player is perfectly observed by at least two others. Here, it is enough that i and j observe the same $z_{i,j}$, because there is a mutual-minmax Nash equilibrium. Also, Theorem A is not a folk theorem but rather a result about replicating public information for fixed δ .

4.2 Private Cheap Talk

In this paper, “private cheap talk” means communication along the links of the network. That is, with private cheap talk players can communicate directly with their neighbors but not with other players. However, players can communicate indirectly with players to whom they are not linked by passing information from one link to another. This requires multiple rounds of communication after every round of play, which I allow. Indeed, in any communication round a player may learn something that she would like to pass on. To accommodate this, I allow for infinitely many rounds of communication after each round of play.²²

A game with private cheap talk $\Gamma_{PRI}^{PRICT}(Y)$ is derived by augmenting the game Γ_{PRI} with a finite message set $Y = \left((Y_{i,j})_{j \in N_i} \right)_{i \in N}$ such that after players observe their private signals

²¹This argument clearly relies heavily on the assumption that signal $z_{i,j}$ is locally public between i and j . If signals were not even locally public, then one would be in the setting of general repeated games with private monitoring, and public information could not be replicated even with public communication. However, it may be the case that if signals are “almost” locally public then public communication can “almost” replicate public information. I do not pursue this question here.

²²An alternative to having infinitely many rounds of communication would be having a finite but unbounded number of rounds, where communication continues only as long as some player keeps talking. However, in this alternative model it is not clear how to interpret the assumption that the players know when everyone is done talking and it is time to move to the next period. In any case, only finitely many rounds of communication are needed on-path.

they have infinitely many opportunities to simultaneously send private messages $y_{i,j}^k \in Y_{i,j}$ to their neighbors, where the subscript denotes a message from i to j and the superscript $k \in \mathbb{N}$ denotes the number of the communication round. Formally, the stage game is now a long cheap talk game as modeled by Aumann and Hart (2003). That is, letting $h_{i,t} = \left(a_{i,t}, (z_{i,j,t})_{j \in N_i}, (y_{i,j,t}^k, y_{j,i,t}^k)_{j \in N_i, k \in \mathbb{N}} \right)$, there are now infinitely many kinds of histories for every period t , denoted $h_i^{t-} = (h_{i,\tau})_{\tau=0}^{t-1}$ (action histories), $h_i^{t,0} = \left((h_{i,\tau})_{\tau=0}^{t-1}, a_{i,t}, (z_{i,j,t})_{j \in N_i} \right)$, and $h_i^{t,k} = \left((h_{i,\tau})_{\tau=0}^{t-1}, a_{i,t}, (z_{i,j,t})_{j \in N_i}, (y_{i,j,t}^{k'}, y_{j,i,t}^{k'})_{j \in N_i, k' \in \{1, \dots, k\}} \right)$, for $k \in \mathbb{N}$ (communication histories). A strategy σ_i is now a measurable function that maps action histories h_i^{t-} to $\Delta(A_i)$ and maps communication histories $h_i^{t,k}$ to $\Delta\left(\prod_{j \in N_i} Y_{i,j}\right)$. Let $E_{PRI}^{PRICT}(Y)$ be the SE payoff set of $\Gamma_{PRI}^{PRICT}(Y)$, and let $E_{PRI}^{PRICT} = \bigcup_Y E_{PRI}^{PRICT}(Y)$.²³

The second benchmark result is that private cheap talk can replicate public information for all games if and only if the network L is 2-connected. Recall that a network is 2-connected if there are at least two independent paths (i.e., two paths with disjoint sets of internal nodes) between every pair of nodes. The main idea is again quite simple: Start with a SE profile σ^{PUB} in game Γ_{PUB} . Specify that after every round of play there are multiple rounds of communication in which players report both the signals they have observed directly and the signals that have been reported to them in earlier rounds, until all signals have been reported to all players. The players then play according to σ^{PUB} , taking the reported signals as the true ones. If a player sends or receives an inconsistent report, she then reports that there has been a deviation, and the news of the deviation spreads throughout the network and leads all players to play the mutual-minmax profile α^* . The assumption that the network is 2-connected implies that no player can deceive another about the signals: if a player i lies about a signal to one of her neighbors, the neighbor will eventually receive a conflicting report via a path that does not include i , and will then revert to α^* .

Conversely, if the network L is not 2-connected then there are three players—call them 1, 2, and 3—such that 1 and 2 are linked, 2 and 3 are linked, and the unique path from 1 to 3 is the one through 2. It is not difficult to find specifications of the (1, 2) and (2, 3) games such

²³Aumann and Hart prove that in a single long cheap talk game the induced mapping from strategies to payoffs is measurable, so that the game is well-defined. Their proof immediately extends to the current repeated game model with finite players, actions, signals, and messages. It also immediately extends to the model of Section 5, where introducing divisible tokens makes the message sets countably infinite.

that, when all other (i, j) games are taken to be trivial games with $u_{i,j}(a_{i,j}, a_{j,i}) = 0$ for all $(a_{i,j}, a_{j,i}) \in A_{i,j} \times A_{j,i}$, it follows that private cheap talk cannot replicate public information. In particular, it suffices to take the (1, 2) and (2, 3) games to be the asymmetric prisoner's dilemmas described in the example below (to which I return later in the paper).

Theorem B *Private cheap talk can replicate public information (i.e., $E_{PRI}^{PRICT} \supseteq E_{PUB}$) if the network L is 2-connected.*

Conversely, if the network L is not 2-connected, then there exists a game (A, Z, π, u, δ) for which private cheap talk cannot replicate public information (i.e., $E_{PUB} \setminus E_{PRI}^{PRICT} \neq \emptyset$).

The first part of Theorem B is related to Theorem 2.6 of Renault and Tomala (1998), which gives a Nash folk theorem for repeated games with a 2-connected monitoring network without explicit communication. Theorem B avoids some complications that emerge in their paper by allowing explicit communication and assuming a mutual-minmax Nash equilibrium (though Theorem B is for sequential equilibrium rather than Nash). Moreover, the results differ in that Theorem B is about replicating public information for fixed δ .

4.2.1 Example: Asymmetric Prisoner's Dilemma on a Line

There are three players on a line and each relationship is a prisoner's dilemma with "locally perfect" monitoring (players 1 and 3 take female pronouns; player 2 takes male pronouns). Formally, $L = \{\{1, 2\}, \{2, 3\}\}$, $A_{i,j} = \{C, D\}$ for $\{i, j\} \in L$, $Z_{i,j} = A_{i,j} \times A_{j,i}$, $\pi_{i,j}((a_{i,j}, a_{j,i}) | a_{i,j}, a_{j,i}) = 1$, and the payoff matrix in the (1, 2) game is

$$\begin{array}{cc} & C & D \\ C & 1, 1 & -l_1, 1 + g_{2,1} \\ D & 1 + g_1, -l_{2,1} & 0, 0 \end{array}$$

while the payoff matrix in (2, 3) game is

$$\begin{array}{cc} & C & D \\ C & 1, 1 & -l_3, 1 + g_{2,3} \\ D & 1 + g_3, -l_{2,3} & 0, 0 \end{array}$$

where in both matrices player 2 is the column player (so 1 is the row player in the first matrix and 3 is the row player in the second). Assume that for each matrix the sum of the players' payoffs is maximized at outcome (C, C) .²⁴ In addition, assume the following.

$$g_1 \leq \frac{\delta}{1-\delta}, g_3 \leq \frac{\delta}{1-\delta}, g_{2,1} > \frac{\delta}{1-\delta}, g_{2,3} < \frac{\delta}{1-\delta}, g_{2,1} + g_{2,3} \leq 2\frac{\delta}{1-\delta}.$$

The following result shows that private cheap talk may fail to replicate public information when the network is not 2-connected. The intuition is that with public monitoring player 2 can be made to cooperate in the $(1, 2)$ game by specifying that defection in the $(1, 2)$ game leads to permanent defection in both games, but with private cheap talk player 2 can defect in the $(1, 2)$ game while concealing this deviation from player 3 and continuing to cooperate in the $(2, 3)$ game.

Proposition 1 *In this example of an asymmetric prisoner's dilemma on a line, private cheap talk cannot replicate public information.*

Proof. I show that the payoff vector $(1, 2, 1)$ is an element of E_{PUB} but not E_{PRI}^{PRICT} , which proves the result.

The payoff vector $(1, 2, 1)$ can be attained only if the outcome in both games is (C, C) in every period. To see that this is possible in E_{PUB} , consider the *multilateral grim trigger* profile when players play C (in both games, in the case of player 2) if the outcome in both games has always been (C, C) , and play D otherwise. Then player 1 has no profitable deviation under the assumption $g_1 \leq \frac{\delta}{1-\delta}$, player 3 has no profitable deviation under the assumption $g_3 \leq \frac{\delta}{1-\delta}$, and player 2 has no profitable deviation under the assumption $g_{2,1} + g_{2,3} \leq 2\frac{\delta}{1-\delta}$ (this last observation follows because player 2's is most tempted to simultaneously deviate to D in both games, as a deviation in either game leads to (D, D) forever in both).

Now suppose toward a contradiction that for some message set Y there exists in $\Gamma_{PRI}^{PRICT}(Y)$ a SE profile σ in which the outcome in both games is (C, C) in every period. Replace $\sigma_{2,3}$ with a strategy $\tilde{\sigma}_{2,3}$ that for each action history h_2^{t-} depends only on $\left(z_{2,3,\tau}, (y_{2,3,\tau}^k, y_{3,2,\tau}^k)_{k \in \mathbb{N}}\right)_{\tau=0}^{t-1}$ but has the same marginals over $A_{2,3}$ conditional on $\left(z_{2,3,\tau}, (y_{2,3,\tau}^k, y_{3,2,\tau}^k)_{k \in \mathbb{N}}\right)_{\tau=0}^{t-1}$ as does $\sigma_{2,3}$,

²⁴That is, assume that $g_{2,1} - l_1, g_1 - l_{2,1}, g_{2,3} - l_3$, and $g_3 - l_{2,3}$ are all less than 1.

and similarly for communication histories. Then, when player 3 plays σ_3 , the distribution of outcomes in the (2, 3) game when player 2 plays (2, 3) game strategy $\tilde{\sigma}_{2,3}$ is the same as when he plays (2, 3) game strategy $\sigma_{2,3}$, which is to say that the outcome is (C, C) in every period. Hence, if player 2 deviates to always playing D in the (1, 2) game and playing $\tilde{\sigma}_{2,3}$ in the (2, 3) game, his payoff is $(1 - \delta)(1 + g_{2,1}) + \delta(0) + 1$, which is greater than his equilibrium payoff of 2 under the assumption $g_{2,1} > \frac{\delta}{1-\delta}$. So there can be no such SE. ■

5 Replicating Public Information with Tokens

I now turn to the main part of the analysis, where players have access to tokens in addition to private cheap talk.

A game with tokens Γ_{PRI}^{TOK} is similar to a game with private cheap talk, except that in addition to sending cheap talk messages players can also transfer quantities of undifferentiated, infinitely divisible tokens to each other.²⁵ The difference between cheap talk and tokens is that a player can send any cheap talk message she wants, but can only send tokens that she is currently holding: for example, any player can say “message number 5,” but only a player with at least 5 tokens can make a 5 token transfer. Thus, tokens are a natural way of making a player’s “message set” depend on the past messages she has sent or received, which turns out to allow players to replicate public information.

Formally, a game with tokens $\Gamma_{PRI}^{TOK}(Y, m^0)$ is derived from the game with private cheap talk $\Gamma_{PRI}^{PRICT}(Y)$ by specifying an initial endowment of tokens $m^0 = (m_1^0, \dots, m_n^0)$, with $m_i^0 \in \mathbb{Q}_+$ for all $i \in N$ (where \mathbb{Q}_+ denotes the non-negative rationals), and allowing players to transfer tokens concurrently with their messages. That is, at every history in $\Gamma_{PRI}^{PRICT}(Y)$ where player i chooses a message $y_{i,j} \in Y_{i,j}$ to send to player j , she now chooses a pair $(y_{i,j}, m_{i,j}) \in Y_{i,j} \times \mathbb{R}_+$ to send to player j , subject to the constraint that $\sum_{j \in N_i} m_{i,j} \leq m_i$, where m_i is player i ’s current token holding, and the vector of token holding is then updated

²⁵It would be essentially equivalent to let players transfer only tokens and not also cheap talk messages: simply replace the cheap talk messages with transfers of tiny amount of tokens that nonetheless convey a large amount of information. The current approach of allowing both cheap talk and tokens is slightly easier to exposit.

to

$$m_i^t = m_i + \sum_{j \in N_i} (m_{j,i} - m_{i,j}).^{26}$$

A strategy is *feasible* if it satisfies $\sum_{j \in N_i} m_{i,j} \leq m_i$ at every communication history $h_i^{t,k}$. I also now allow players to send messages and transfers concurrently with actions (i.e., at an action history h_i^{t-} , player i now chooses a triple $(a_i, (y_{i,j})_{j \in N_i}, (m_{i,j})_{j \in N_i})$).²⁷

Assuming that tokens are infinitely divisible makes action spaces infinite and necessitates using a version of PBE rather than SE (the motivation for divisible tokens, as well as alternative models, are discussed below). There is no off-the-shelf version of PBE that seems appropriate in this model. For example, consider 4 players on a line: $L = \{\{1, 2\}, \{2, 3\}, \{3, 4\}\}$. On the one hand, assuming that player 1 does not update her belief about player 3's private history after observing a 0-probability move by player 2 seems too strong, as this move may have been a response to a deviation by player 3. On the other hand, letting player 2 update his beliefs about player 3's private history after observing a 0-probability move by player 1 seems implausible, as 1's play can only affect 3 via 2. To take care of these concerns, I introduce an extension of weak perfect Bayesian equilibrium that captures the idea that information can only flow along the links of the network.²⁸ Let $L \setminus \{i\}$ denote the network L with node i removed, and let C_j^i denote the component of $L \setminus \{i\}$ containing j .²⁹ Let $\mu_i(h_j^t | h_i^t)$ denote player i 's belief that player j 's private history is h_j^t when player i 's private history is h_i^t .

Definition 1 *A network weak perfect Bayesian equilibrium (PBE) is a weak perfect Bayesian equilibrium satisfying the following two additional properties:*

1. *If $j \in N_i$ and $j' \notin C_j^i$, then $\mu_i(h_{j'}^t | h_i^t)$ does not depend on $h_{i,j}^t$.*

²⁶The point of allowing players to transfer only rational quantities of tokens is to ensure that strategy spaces remain countable.

²⁷This modification plays a "technical" role in the proof of the main result, discussed in footnote 54. It is not needed either if monitoring in all (i, j) games is (locally) perfect or if monitoring in all (i, j) games has full support.

²⁸Recall that a weak perfect Bayesian equilibrium is an assessment (σ, μ) such that σ_i is sequentially rational given beliefs μ_i about the vector of private histories $(h_j^t)_{j=1}^n$ and μ_i is updated according to Bayes' rule whenever possible.

²⁹That is, C_j^i is the set of players $j' \in L$ such that there is a path in L from j to j' that does not contain i .

2. If $j \in N_i$, then for all $j' \in N$, $\mu_i \left(h_{j'}^{t,k} | h_i^{t,k} \right)$ does not depend on $y_{i,j}^{t,k'}$ for $k' > k - d(j, j')$, and does not depend on $z_{i,j}^t$ if $k < d(j, j')$.

The first requirement says that information does not “jump over” player i . The second requirement says that information does not propagate faster than one link per round of communication. In particular, these requirements rule out “grim trigger” beliefs, where a player who observes a single deviation believes that all of her opponents immediately revert to α^* .³⁰

The following is the main result of the paper. Here, $\tilde{E}_{PRI}^{TOK} = \bigcup_{(Y, m^0)} \tilde{E}_{PRI}^{TOK}(Y, m^0)$, where $\tilde{E}_{PRI}^{TOK}(Y, m^0)$ is the (network weak) PBE payoff set in $\Gamma_{PRI}^{TOK}(Y, m^0)$.

Theorem 1 *Tokens can replicate public information (i.e., $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PUB}$).*

Theorem 1 shows that in networked environments tokens enable players to sustain any payoff vector that would be sustainable were all information public, even in non-2-connected networks in which the presence of information gatekeepers rules out this possibility when only cheap talk is available. The fact that tokens can replicate public information even when the network is not 2-connected will form the basis of the later results on when tokens are essential.

To see the overall approach of the proof, suppose for simplicity that L is a tree (so that, in particular, L is not 2-connected). Let σ^{PUB} be a PBE strategy profile in game Γ_{PUB} . Initially, endow each of the “leaf players” in L (i.e., players with only one neighbor) with a large number of tokens, and endow “non-leaf players” with none. Have players initially play as in σ^{PUB} . After each round of play, first have players repeatedly report their signals to each other as in the model with private cheap talk: this is called the “reporting subphase” in

³⁰While sequential equilibrium is traditionally defined only for finite games (Kreps and Wilson, 1982), the definition extends immediately to games with countably infinite action spaces. Thus, one could technically apply sequential equilibrium in the current model. However, in games with countably infinite action spaces sequential equilibrium imposes strong restrictions that go well beyond Kreps and Wilson’s original motivation. For example, in community enforcement games it is often convenient to specify that players believe that deviations in period t are much more likely than deviations in period $t - 1$, so that 0-probability moves are always interpreted as deviations rather than as responses to earlier deviations. But this is impossible when action spaces are countably infinite, as there must be some actions that are vanishingly unlikely to occur as deviations in period t . This difficulty and others make it extremely difficult to work with sequential equilibrium in the current model.

the proof. Then have players use tokens to check that no players have misreported signals: this is called the “confirmation subphase” and is described below. In the next round, have players play according to σ^{PUB} , taking the reported signals as the true ones.

As in Theorems A and B, if players can be induced to report their signals truthfully, the above construction yields an equilibrium with the same payoffs as σ^{PUB} . Thus, the key insight behind Theorem 1 is that the confirmation subphase can be constructed so as to ensure that no player can mislead another about the value of any signal. The construction is as follows: Assign a natural number q to every possible vector of signals z . At the beginning of the period t confirmation subphase, each leaf player i thinks that the true vector of period t signals is some \hat{z}^i (where all the \hat{z}^i 's are the same on path). The confirmation subphase starts with some leaf player—say, player 1—sending q_1 tokens down the (unique) path toward another leaf player—call him player 2—where q_1 is the number assigned to \hat{z}^1 . The non-leaf players on this path then pass these tokens on to player 2. When player 2 receives the tokens, he checks whether the number of tokens received equals the number q_2 assigned to \hat{z}^2 . If it does, he adds an additional q_2 tokens to the transfer he received, and passes this new larger “pot” of tokens on to the next leaf player, player 3. This process continues until each leaf player gets the chance to add tokens to the pot, and the pot is then returned to player 1. Finally, if this returned pot is of the size she expected, player 1 then sends an additional large transfer down the path to each leaf player in turn, each of whom returns this transfer to player 1.³¹

In the proof itself, every number q_i in the preceding paragraph is replaced by $q_i/2^t$. That is, the size of the transfers made in the confirmation subphase halve each period. As discussed below, this trick serves to keep leaf players from running out of tokens.

While it is hard to give a complete intuition for why this construction works without just

³¹The point of these last transfers may be seen in the example of three players on a line (where player 2 is the middle (non-leaf) player, contrary to the numbering scheme used in the proof): Suppose player 2 reports outcome 10 to player 1 and outcome 5 to player 3. Then player 1 begins the confirmation phase by transferring 10 tokens to player 2. Player 2 is now supposed to pass the 10 tokens on to player 3, but suppose he instead passes only 5 tokens to her. Player 3 then sees the report of outcome 5 confirmed, and passes 5 tokens back to player 2. Now player 1 expects to receive a transfer of 20 tokens from player 2, but player 2 only has 15 tokens, so no matter how much player 2 sends her she will learn that there has been a deviation. But—crucially—player 3 will *not* learn this: she will continue to think the outcome is 5. Adding the final transfers from player 1 to the other leaf players fixes this problem, as now all leaf players learn about any deviation that player 1 discovers.

proving the theorem, it is worth noting the following three important facts. First, if some player fails to pass on the correct number of tokens toward the next leaf player, then player 1 does not get back a pot of the correct size, and thus does not make the final confirmation transfers to the other leaf players (who can then infer that a deviation occurred). Second, if any two players disagree about the vector of period t signals, then player 1 again does not get back a pot of the correct size, and since signals are locally public at least one player will disagree with any misreported signal. Finally, the transfers can be constructed so that a player can never save enough tokens in an earlier round to mislead another player in a later round.³²

To illustrate the construction, consider the asymmetric prisoner’s dilemma on a line of Section 4.2.1. Assign a number between 1 and 16 to each of the 16 possible stage game outcomes, $((C, C), (C, C)), ((C, C), (C, D)), \dots, ((D, D), (D, D))$. For concreteness, suppose the number 1 is assigned to outcome $((C, C), (C, C))$ (the desired outcome, which as we have seen cannot be sustained in Γ_{PRI}^{PRICT}). Initially, endow players 1 and 3 with a large number of tokens (the number used in the proof would be 128 here, although since tokens are infinitely divisible one could also have normalized this number to 1), and endow player 2 with none.

On-path, play in period t proceeds as follows: Players cooperate, yielding outcome $((C, C), (C, C))$. In the reporting subphase, players truthfully report their observations to their neighbors; in particular, player 2 tells player 1 that the outcome in the $(2, 3)$ game was (C, C) and tells player 3 that the outcome in the $(1, 2)$ game was (C, C) . At this point, players 1 and 3 both believe that the overall outcome was $((C, C), (C, C))$, but this has not yet been “confirmed.” The confirmation subphase starts with player 1 sending $1/2^t$ tokens (i.e., $q_1/2^t$ tokens, recalling that 1 is the number assigned to $((C, C), (C, C))$) to player 2 (the next player on the path from 1 to 3). In the next round, player 2 then sends these $1/2^t$ tokens to player 3. Player 3 now notes that $1/2^t$ equals $q_3/2^t$ (as she also believes

³²Another issue in the proof of Theorem 1 is that explicitly describing off-path play is intractable. This is because when player i sees player j deviate, player i may wish to conceal this information if passing it on would lead the other players to punish her (and, if the network is not 2-connected, she may be able to do this). Therefore, I specify that when player i observes a deviation by player j , she only minmaxes player j himself as well as the other players in C_j^i . The rest of off-path play is specified implicitly by a procedure described in the Appendix.

that the outcome was $((C, C), (C, C))$, and she indicates this by sending a total of $2/2^t$ tokens back to player 2, who then sends these $2/2^t$ tokens to player 1. Player 1 notes that the extra $1/2^t$ tokens contributed by player 3 matches her beliefs that the outcome was $((C, C), (C, C))$. She therefore sends a large “confirmation” transfer to player 2 (in the proof, this transfer would consist of $128 - 64/2^t$ tokens), who then sends these tokens on to player 3. Finally, player 3 sends these tokens back to player 2, who then sends them back to player 1, completing the confirmation subphase.

In contrast, suppose player 2 deviates to D in the $(1, 2)$ -game in period t , yielding outcome $((C, D), (C, C))$. Then in the reporting subphase, player 2 may still report to player 3 that the outcome was $((C, C), (C, C))$ (e.g., this is what he would do if the players tried to sustain $((C, C), (C, C))$ in Γ_{PRI}^{PRICT}). But this misreport will be detected in the confirmation subphase as follows. Since player 1 observes an off-path action by player 2 at an on-path history, she punishes player 2 by *both* playing D forever *and* never again sending him tokens. Player 3 now expects to receive $1/2^t$ tokens from player 2, but player 2 has no tokens to send her (recall that he started period t with no tokens, as he returned all tokens to player 1 at the end of period $t - 1$ —of course, the proof must also verify that he could not have profitably deviated by retaining tokens in period $t - 1$). So when no tokens arrive, this constitutes an off-path (lack of a) transfer from player 2 at an on-path history, and player 3 also punishes player 2 by playing D forever and never again sending him tokens. Hence, player 2 is punished by both players 1 and 3 for deviating in the $(1, 2)$ -game, and as in Γ_{PRI}^{PUBCT} this deters the deviation.

I conclude this section with two remarks on Theorem 1. First, the choice of initial token endowment m^0 is not crucial. As a consequence, even in setting where initial endowments are exogenously determined (rather than being a “choice variable,” as I have assumed), public information can still be replicated for a wide range of initial endowments; for example, Theorem 1 goes through whenever all players start with a positive number of tokens. The idea is that if any non-leaf players are endowed with tokens, they can be induced to transfer all of their tokens to player 1 at the beginning of the game. Formally, one can show the following result.³³

³³A *spanning tree* is a connected subnetwork with no cycles that contains all the nodes in the original

Proposition 2 *Suppose the initial token endowment m^0 is exogenously given. If there exists a spanning tree $L' \subseteq L$ such that $m_i^0 > 0$ for all leaf players i in L' , then $\tilde{E}_{PRI}^{TOK}(Y, m^0) \supseteq \tilde{E}_{PUB}$, where Y is the message set from the proof of Theorem 1.*

However, Proposition 2 does continue to assume that the initial token endowment m^0 is common knowledge. Proposition 4 below shows that Theorem 1 may fail with uncertain endowments.

Second, Theorem 1 relies on the assumption that tokens are infinitely divisible. This assumption serves two roles in the proof. First, it lets one ensure that leaf players never run out of tokens. This could potentially be addressed by instead “rebalancing” token holdings between rounds, although doing this is not trivial. Second—and more importantly—it allows the size of the final confirmation transfer in each period to increase over time. This ensures that a player who deviates by saving some tokens in one period cannot use them to mimic a later confirmation transfer. Both of these roles of infinite divisibility could instead be filled by simply disbursing more tokens to the leaf players every period, if this were allowed (contrary to my assumptions). For example, Theorem 1 would go through if tokens are indivisible but $|Z|$ tokens are disbursed from the “planner” to each leaf player in every period. In contrast, Proposition 5 below shows that Theorem 1 may fail with indivisible tokens if such disbursements are not allowed.

6 Examples

This section presents examples showing that the assumptions that the network and the initial token endowment are commonly known and that tokens are infinitely divisible cannot be completely dispensed with.

6.1 Unknown Network

Consider the following model: There are three players. At the beginning of the game, Nature flips two independent fair coins to determine whether players 1 and 2 linked and

network.

whether players 2 and 3 are linked, respectively. Only players 1 and 2 observe whether they become linked, and similarly for players 2 and 3. Thus, the network is stochastic and is realized once-and-for-all at the start of the game, and there is common knowledge of the ex ante distribution over networks but not of the realization.

If players 1 and 2 are linked, they play the following (1, 2)-game; if players 2 and 3 are linked, they play the following (2, 3)-game (player 2 is always the column player).

		(1, 2)-game		(2, 3)-game	
		A	B	X	Y
A	1, -1	0, 0	X	3, 3	0, 0
B	0, 0	0, 0	Y	0, 0	1, 1

Thus, in the (1, 2)-game player 2 has the chance to transfer a util to player 1, and the (2, 3)-game is a coordination game. Assume $\delta \geq \frac{1}{2}$.

Take the public information benchmark here, Γ_{PUB} , to be as in the main model, with the modification that all players observe the realized network at the beginning of the game. To give the players a chance to replicate this benchmark with private information, introduce a round of communication after the network is realized but before the first action phase. In this model, a small extension of Theorem A shows that public cheap talk can replicate public information (in particular, the players can be induced to truthfully report the realized network by specifying Nash reversion in case of disagreement). However, the following result shows that tokens cannot replicate public information here.

Proposition 3 *In this example with an unknown network, tokens cannot replicate public information.*

Proof. I show that payoff vector $(\frac{1}{4}, \frac{5}{4}, \frac{3}{2})$ is in \tilde{E}_{PUB} but not \tilde{E}_{PRI}^{TOK} .

For Γ_{PUB} , consider the following strategy profile.

- If realized network is $\{\{1, 2\}, \{2, 3\}\}$, 2 and 3 play X in the (2, 3)-game if 2 has always played A in the (1, 2)-game, and otherwise play Y. 2 plays A in the (1, 2)-game if he has always played in A in the (1, 2)-game, and otherwise plays B.

- If realized network is $\{\{1, 2\}\}$, 2 always plays B .
- If realized network is $\{\{2, 3\}\}$, 2 and 3 always play X .
- If realized network is $\{\emptyset\}$, there is nothing to play.

This is a PBE under the assumption $\delta \geq \frac{1}{2}$ (as this is the condition that ensures that it is not profitable for 2 to deviate to B in the $(1, 2)$ -game when the realized network is $\{\{1, 2\}, \{2, 3\}\}$), and it yields payoff vector $(\frac{1}{4}, \frac{5}{4}, \frac{3}{2})$ (as each possible network is realized with probability $\frac{1}{4}$). So $(\frac{1}{4}, \frac{5}{4}, \frac{3}{2}) \in \tilde{E}_{PUB}$.

Now suppose toward a contradiction that $(\frac{1}{4}, \frac{5}{4}, \frac{3}{2}) \in \tilde{E}_{PRI}^{TOK}$. Note that if the realized network is $\{\{1, 2\}\}$ then 2 always plays B . Hence, for 1 to get payoff $\frac{1}{4}$, 2 must play A with probability 1 if the realized network is $\{\{1, 2\}, \{2, 3\}\}$. In addition, for 3 to get payoff $\frac{3}{2}$, the outcome in every period of the $(2, 3)$ -game must be (X, X) with probability 1 if the realized network is $\{\{2, 3\}\}$. However, any $(2, 3)$ -game strategy that is feasible for 2 when the realized network is $\{\{2, 3\}\}$ is also feasible for 2 when the realized network is $\{\{1, 2\}, \{2, 3\}\}$, as 2 has the option of never passing tokens to 1 (and of ignoring any tokens he might receive from 1). Let $\sigma_{2,3}^{\{\{2,3\}\}}$ be player 2's equilibrium $(2, 3)$ -game strategy when the realized network is $\{\{2, 3\}\}$. Then when the realized network is $\{\{1, 2\}, \{2, 3\}\}$, it is feasible for 2 to deviate to always playing B in the $(1, 2)$ -game while playing $\sigma_{2,3}^{\{\{2,3\}\}}$ in the $(2, 3)$ -game, and this deviation yields payoff $\frac{3}{2}$ (as the distribution of outcomes in the $(2, 3)$ -game depends only on players 2 and 3's strategies in the $(2, 3)$ -game, which after this deviation are the same as they are in equilibrium when the realized network is $\{\{2, 3\}\}$), which is greater than his equilibrium payoff of $\frac{5}{4}$. Hence, $(\frac{1}{4}, \frac{5}{4}, \frac{3}{2}) \notin \tilde{E}_{PRI}^{TOK}$. ■

6.2 Unknown Initial Endowment of Tokens

Consider the asymmetric prisoner's dilemma on a line of Section 4.2.1. Suppose players 1 and 3 start with m tokens each, while player 2 starts with 0 tokens with probability $\frac{1}{2}$ and starts with m tokens with probability $\frac{1}{2}$, where only he knows which event obtains. Denote this stochastic token endowment by \tilde{m} . Note that the network in this example, $L = \{\{1, 2\}, \{2, 3\}\}$, is itself a tree and the leaf players 1 and 3 always start with a positive

number of tokens, so Proposition 2 shows that $\tilde{E}_{PRI}^{TOK}(Y, m^0) \supseteq \tilde{E}_{PUB}$ when m^0 is taken to be either the deterministic endowment where player 2 starts with 0 tokens or the deterministic endowment where player 2 starts with m tokens. In contrast, the following result shows that public information cannot be replicated with the assumed stochastic endowment.

Proposition 4 *In this example with an unknown initial endowment of tokens, public information cannot be replicated with the assumed stochastic endowment (i.e., $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI}^{TOK}(Y, \tilde{m}) \neq \emptyset$, where $\tilde{E}_{PRI}^{TOK}(Y, \tilde{m})$ is the PBE payoff set with stochastic endowment \tilde{m}).*

Proof. I show that payoff vector $(1, 2, 1)$ is in \tilde{E}_{PUB} but not $\tilde{E}_{PRI}^{TOK}(Y, \tilde{m})$. That $(1, 2, 1) \in \tilde{E}_{PUB}$ was already proved in the proof of Proposition 1.

Suppose toward a contradiction that $(1, 2, 1) \in \tilde{E}_{PRI}^{TOK}(Y, \tilde{m})$. Then in every period the outcome in both the $(1, 2)$ -game and the $(2, 3)$ -game must be (C, C) with probability 1, for both possible initial endowments. However, if a $(2, 3)$ -game strategy $\sigma_{2,3}^0$ is feasible for 2 when his realized endowment is 0 and he plays his equilibrium $(1, 2)$ -game strategy, then strategy $\sigma_{2,3}^0$ is also feasible for 2 when his realized endowment is m and he plays any $(1, 2)$ -game strategy that never involves passing tokens to 1, as in every period his token holding is at least as great in the second case as in the first.³⁴ Therefore, when 2's realized endowment is m , it is feasible for him to deviate to always playing D in the $(1, 2)$ -game while playing $\sigma_{2,3}^0$ in the $(2, 3)$ -game (and never passing tokens to 1), and this deviation yields payoff at least $(1 - \delta)(1 + g_{2,1}) + \delta(0) + 1$, which is greater than his equilibrium payoff of 2 under the maintained assumption $g_{2,1} > \frac{\delta}{1-\delta}$. Hence, $(1, 2, 1) \notin \tilde{E}_{PRI}^{TOK}(Y, \tilde{m})$. ■

6.3 Indivisible Tokens

In this subsection only, assume that players can only transfer integer quantities of tokens (i.e., tokens are *indivisible*). Consider the game given by $L = \{\{1, 2\}, \{2, 3\}\}$ with $(1, 2)$ -game and $(2, 3)$ -game as in Section 6.1 (i.e., the game is exactly as in Section 6.1 but with the network known to be $\{\{1, 2\}, \{2, 3\}\}$). The following result shows that indivisible tokens cannot replicate public information in this example.

³⁴In particular, in the second case his token holding is m plus his net transfer from 3, while in the first case his token holding is his net transfer from 1 plus his net transfer from 3, and his net transfer from 1 cannot exceed 1's initial endowment m .

Proposition 5 *In this example, indivisible tokens cannot replicate public information.*

Proof. I show that with indivisible tokens payoff vector $(1, 2, 3)$ is in \tilde{E}_{PUB} but not \tilde{E}_{PRI}^{TOK} .

For Γ_{PUB} , the grim trigger strategy profile, “2 and 3 play X in the $(2, 3)$ -game if 2 has always played A in the $(1, 2)$ -game, and otherwise play Y ; 2 plays A in the $(1, 2)$ -game if he has always played in A in the $(1, 2)$ -game, and otherwise plays B ,” is a PBE under the assumption $\delta \geq \frac{1}{2}$ and yields payoff $(1, 2, 3)$.

Suppose toward a contradiction that $(1, 2, 3) \in \tilde{E}_{PRI}^{TOK}$. Then for some (Y, m^0) there exists a PBE in $\Gamma_{PRI}^{TOK}(Y, m^0)$ in which the outcome is $((A), (X, X))$ in every period. Note that player 3’s token holding is measurable with respect to h_2^t , as it simply equals m_3^0 plus the net transfer of tokens from player 2 to player 3. Let h_2^t be a history such that player 3’s token holding is maximal over all on-path h_2^t (this exists because tokens are indivisible and finite in number).

I claim that player 2 has a profitable deviation at h_2^t . Note that for every subsequent on-path history h_2^τ , the net token transfer from player 2 to player 3 between histories h_2^t and h_2^τ is non-positive, as otherwise player 3’s token holding would be greater at h_2^τ than at h_2^t . Hence, player 2’s equilibrium $(2, 3)$ -game continuation strategy is feasible for him regardless of his continuation strategy against player 1, so long as he does not transfer tokens to player 1. Therefore, it is a profitable deviation for player 2 to play B in every subsequent period in the $(1, 2)$ -game, never again transfer tokens to player 1, and continue to play his equilibrium continuation strategy against player 3. Hence, $(1, 2, 3) \notin \tilde{E}_{PRI}^{TOK}$. ■

7 From Replication to Essentiality

A final set of results shows how Theorem 1 can be used to show that tokens are essential—in that the PBE payoff set is larger with tokens than without them—in a broad class of games.

I use the following definition.³⁵

Definition 2 *Tokens are essential if $\tilde{E}_{PRI}^{TOK} \supsetneq \tilde{E}_{PRI}$.*

³⁵In this section, I consider PBE rather than SE in Γ_{PRI} and Γ_{PRI}^{PRICT} , to facilitate comparison with Γ_{PRI}^{TOK} . The PBE concept here is as in Section 5, with the exception that the second additional requirement imposed there is not applicable in Γ_{PRI} . Consistent with the notation in the rest of the paper, \tilde{E}_{PRI} and \tilde{E}_{PRI}^{PRICT} are the PBE payoff sets in Γ_{PRI} and Γ_{PRI}^{PRICT} , respectively.

Tokens are strongly essential if $\tilde{E}_{PRI}^{TOK} \supsetneq \tilde{E}_{PRI}^{PRICT}$.

The latter property is indeed stronger because $\tilde{E}_{PRI}^{PRICT}(Y) \supseteq \tilde{E}_{PRI}$ for every message set Y , as messages can always be ignored.

How can one tell whether tokens are essential in a particular game? Recall that \tilde{E}_{PRI} , \tilde{E}_{PRI}^{PRICT} , and \tilde{E}_{PRI}^{TOK} are PBE payoff sets in private monitoring games for fixed discount factors. Such sets are usually impossible to characterize. However, this section shows that essentiality can often be verified by building on Theorem 1.

A first observation is that $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PRI}$ and $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PRI}^{PRICT}$ are trivially true: any PBE in Γ_{PRI} or Γ_{PRI}^{PRICT} can be turned into a payoff-equivalent PBE in Γ_{PRI}^{TOK} by specifying that players never make transfers and ignore transfers if they are made (in particular, $\tilde{E}_{PRI}^{TOK}(Y, m^0) \supseteq \tilde{E}_{PRI}$ and $\tilde{E}_{PRI}^{TOK}(Y, m^0) \supseteq \tilde{E}_{PRI}^{PRICT}(Y)$ for any (Y, m^0)). Combining this observation with Theorem 1 yields the following corollary.

Corollary 1 *Tokens are essential if $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI} \neq \emptyset$. Tokens are strongly essential if $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI}^{PRICT} \neq \emptyset$.*

Proof. By Theorem 1, $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PUB}$. So $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI} \neq \emptyset$ (resp., $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI}^{PRICT} \neq \emptyset$) implies that $\tilde{E}_{PRI}^{TOK} \setminus \tilde{E}_{PRI} \neq \emptyset$ (resp., $\tilde{E}_{PRI}^{TOK} \setminus \tilde{E}_{PRI}^{PRICT} \neq \emptyset$). The observation that $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PRI}$ completes the proof for “essential,” and the observation that $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PRI}^{PRICT}$ completes the proof for “strongly essential.” ■

Combining Corollary 1 and Proposition 1 shows that tokens are strongly essential in the asymmetric prisoner’s dilemma of Section 4.2.1. However, in general it can be hard to know when $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI} \neq \emptyset$ or $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI}^{PRICT} \neq \emptyset$. Fortunately, one can often establish essentiality while restricting attention to the following much more tractable class of strategies.

Definition 3 *A locally public strategy σ_i is a strategy in Γ_{PRI} where $\sigma_{i,j}$ depends only on $(z_{i,j,\tau})_{\tau=0}^{t-1}$, for all $j \in N_i$. A locally public equilibrium (LPE) in Γ_{PRI} is a PBE in Γ_{PRI} in locally public strategies. Denote the LPE payoff set in Γ_{PRI} by E_{PRI}^{LPE} .*

A local cheap talk strategy σ_i is a strategy in Γ_{PRI}^{PRICT} where $\sigma_{i,j}$ depends only on $\left\{ z_{i,j,\tau}, (y_{i,j,\tau}^k, y_{j,i,\tau}^k)_{k \in \mathbb{N}} \right\}_{\tau=0}^{t-1}$, for all $j \in N_i$. A local cheap talk equilibrium (LCTE) in Γ_{PRI}^{PRICT} is a PBE in Γ_{PRI}^{PRICT} in local cheap talk strategies. Denote the LCTE payoff set in Γ_{PRI}^{PRICT} by E_{PRI}^{LCTE} .

Thus, a locally public strategy is one where player i conditions her play in her relationship with player j only on the history of locally public signals between i and j , and a local cheap talk strategy is one where player i conditions her play in her relationship with player j (including the messages she sends to j) only on the history of locally public signals and cheap talk between i and j . Locally public equilibrium is the natural analog of the standard perfect public equilibrium (PPE) in repeated games with imperfect public monitoring, and local cheap talk equilibrium is the natural analog of PPE when players can send messages only about mutually public information.³⁶ Note that with local cheap talk strategies players have very little to talk about, since they do not condition their messages on information that the receiver does not already have. In particular, one can show that the set of LCTE payoffs is simply the set of LPE payoffs in the auxiliary game where each pair of players is given access to a public randomizing device.

I now show that the condition that $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI} \neq \emptyset$ (resp., $\tilde{E}_{PUB} \setminus \tilde{E}_{PRI}^{PRICT} \neq \emptyset$) in Corollary 1 may be replaced with something like $\tilde{E}_{PUB} \setminus E_{PRI}^{LPE} \neq \emptyset$ (resp., $\tilde{E}_{PUB} \setminus E_{PRI}^{LCTE} \neq \emptyset$). To do this, I introduce the notion of a “nice” subnetwork.

For any subnetwork $M \subseteq L$, let $\tilde{E}|_M$ be the PBE payoff set in the game where M is the original network, or equivalently the PBE payoff set in the game where all links $\{i, j\} \notin M$ are deleted (so that $\tilde{E}_{PRI}|_M$ is the PBE payoff set in this game with private monitoring, $\tilde{E}_{PUB}|_M$ is the PBE payoff set in this game with public monitoring, etc.). For future reference, the game where M is the original network will be denoted $\Gamma|_M$. Finally, for any set X , let $\text{co}(X)$ denote the convex hull of X . I now introduce a key definition.³⁷

Definition 4 *A subnetwork $M \subseteq L$ is nice if it has the following three properties.*

1. *M is a subtree of L . That is, for any two players $i, j \in M$, there is a unique path from i to j in L , and every node in this path is contained in M .³⁸*

³⁶In particular, it is straightforward to verify that PBE and SE strategy profiles coincide once one restricts to locally public or local cheap talk strategies, so Definition 3 could equivalently have been stated with reference to SE.

³⁷I slightly abuse notation here by letting M stand for both a subnetwork of L and the set of nodes in that subnetwork.

³⁸This is stronger than the condition that M is itself a tree: it is not enough that there is a unique path from i to j in M .

2. For all $\{i, j\} \in M$, the (i, j) -game has a product structure. That is, $Z_{i,j} = Z_{i,j}^i \times Z_{i,j}^j$ and $\pi_{i,j}(z_{i,j}|a_{i,j}, a_{j,i}) = \pi_{i,j}^i(z_{i,j}^i|a_{i,j}) \pi_{i,j}^j(z_{i,j}^j|a_{j,i})$.

3. $\tilde{E}_{PUB}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$.

In addition, M is truly nice if the last condition can be strengthened to $\tilde{E}_{PUB}|_M \setminus \text{co}(E_{PRI}^{LCTE}|_M) \neq \emptyset$.

The following theorem is the key tool for determining when tokens are essential.

Theorem 2 *Tokens are essential if L contains a nice subnetwork. Tokens are strongly essential if L contains a truly nice subnetwork.*

For example, if L is a tree, all (i, j) -games in L have a product structure, and $\tilde{E}_{PUB} \setminus \text{co}(E_{PRI}^{LPE}) \neq \emptyset$, then Theorem 2 says that tokens are essential.³⁹ However, Theorem 2 is much more general than this because L itself need not be nice. In particular, the condition that L is a tree is very strong, but the condition that L contains a subtree is trivial. However, not any subtree will do: in particular, if every (i, j) -game in M has a product structure, then $\tilde{E}_{PUB}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$ can hold only if M contains at least three players. Conversely, in many games—including the “trading favors” example discussed below—it is possible to show that any subtree of size at least three is truly nice, and conclude that tokens are strongly essential whenever L contains a subtree of size at least three. This condition provides a very simple method of verifying essentiality in these applications even though characterizing \tilde{E}_{PRI} , \tilde{E}_{PRI}^{PRICT} , and \tilde{E}_{PRI}^{TOK} remains intractable.

The intuition for Theorem 2 is as follows: If M is a tree and all (i, j) -games in M have a product structure, then in the game where M is the original network it is without loss of generality to restrict attention to LPE. If in addition M is a subtree of L , then the equilibrium payoff set on L equals the sum of the equilibrium payoff set on M and the equilibrium payoff set on $L \setminus M$ (in Γ_{PRI}). So if tokens expands the equilibrium payoff set on M while restricting attention to LPE, then it also expands the (unrestricted) equilibrium payoff set on L .⁴⁰

³⁹Technically, one can show that $E_{PUB} \setminus E_{PRI}^{LPE} \neq \emptyset$ is sufficient in this case.

⁴⁰It may be seen from the proof of Theorem 2 that—consistent with this intuition—tokens only circulate among players in M and the expansion in the equilibrium payoff set on L comes entirely from expanding the payoff set available to players in M .

A previous version of this paper considered some leading classes of games in which the results of this section may be used to show that tokens are essential. In particular, it is shown there that Theorem 1 may be adapted to cover continuous time “trading favors” games (Möbius, 2000; Hauser and Hopenhayn, 2010), and that in such games every subtree of size at least three is truly nice, implying that tokens are essential if the network contains a subtree of size at least three.

8 Conclusion

This paper has compared cheap talk and divisible, undifferentiated, physical tokens as means of replicating public information in repeated games on networks. The main result is that public information can always be replicated when tokens are available—in contrast, it can only be replicated when the network is 2-connected if tokens are unavailable. In addition, the tokens considered in this paper are “close” to the minimal communication technology needed for this result, in that the result may fail if the initial endowment of tokens is unknown or if tokens are indivisible.

In addition, the main result on replicating public information leads to a simple sufficient condition for tokens to expand the equilibrium payoff set: tokens are essential in this sense if the network contains a nice subnetwork (that is, a subtree on which replicating public information may be shown to be valuable while restricting attention to locally public equilibria). In many games, this condition reduces to the property that the network contains a subtree of size at least three—a simple and easily verifiable condition.

The physical tokens studied in this paper bear a strong technological resemblance to the “tangible useless objects” (Wallace, 2001) used to model fiat money in monetary theory, while the way they are used in the proof of the main result bears no resemblance to the way money is used in reality. This suggests that an important direction for future research is to consider limits on players’ information or “rationality” that might make simpler and more realistic ways of using tokens constrained optimal. For example, it might be useful to study models where players are “more anonymous” than in this paper but “less anonymous” than in standard continuum agent–random matching models of money, or models where players

use maxmin optimal strategies or other boundedly rational rules in the face of uncertainty about the distribution of tokens.

Appendix

Proof of Theorem A

Let $Y_i = \prod_{j \in N_i} Z_{i,j}$. I show that $E_{PRI}^{PUBCT}(Y) \supseteq E_{PUB}$.

Let σ^{PUB} be a SE strategy profile in game Γ_{PUB} . Define a strategy profile σ^{PRI} in game $\Gamma_{PRI}^{PUBCT}(Y)$ as follows:

- Initially, play as in σ^{PUB} (i.e., $\sigma_i^{PRI}(h^{0-}) = \sigma_i^{PUB}(h^{0-})$).
- Always report all signals truthfully (i.e., $\sigma_i^{PRI}(h^{t+}) = (z_{i,j,t})_{j \in N_i}$ for all h^{t+}).
- Let $\hat{z}_{i,j,t}$ be the (i,j) coordinate of $y_{i,t}$. At history h_i^{t-} , if $\hat{z}_{i',j',\tau} = \hat{z}_{j',i',\tau}$ for all $\{i',j'\} \in L$ and all $\tau < t$, then let $\sigma_i^{PRI}(h_i^{t-}) = \sigma_i^{PUB}(\hat{h}_i^t)$, where $\hat{h}_i^t = \left(a_{i,\tau}, (\hat{z}_{i',j',\tau})_{\{i',j'\} \in L} \right)_{\tau=0}^{t-1}$.⁴¹ If instead $\hat{z}_{i',j',\tau} \neq \hat{z}_{j',i',\tau}$ for some $\{i',j'\} \in L$ and $\tau < t$, then let $\sigma_i^{PRI}(h_i^{t-}) = \alpha_i^*$.

Clearly, σ^{PRI} yields the same payoff vector as σ^{PUB} . It remains only to show that σ^{PRI} is a SE profile. To see this, note first that player i does not have a profitable deviation at an action history h_i^{t-} with $\hat{z}_{i',j',\tau} = \hat{z}_{j',i',\tau}$ for all $\{i',j'\} \in L$ and all $\tau < t$. This follows because her continuation payoff from playing any action a_i at history h_i^{t-} under σ^{PRI} is the same as her continuation payoff from playing a_i at history \hat{h}_i^t under σ^{PUB} , and σ^{PUB} is a SE. Second, player i does not have a profitable deviation at an action history h_i^{t-} with $\hat{z}_{i',j',\tau} \neq \hat{z}_{j',i',\tau}$ for some $\{i,j\} \in L$ and $\tau < t$, because starting from such a history her opponents play α_{-i}^* forever and α_i^* is a best response to α_{-i}^* . Finally, player i does not have a profitable deviation at a communication history h^{t+} . In the case where $\hat{z}_{i,j,\tau} = \hat{z}_{j,i,\tau}$ for all $\{i,j\} \in L$ and all $\tau < t$, this follows because her continuation payoff when she conforms to σ^{PRI} equals her expected continuation payoff under σ^{PUB} conditional on reaching history \hat{h}_i^t , playing

⁴¹To clarify the notation here, note that the reported signals $\hat{z}_{i,j,t}$ and $\hat{z}_{j,i,t}$ are not indentially equal (unlike the true signals $z_{i,j,t}$ and $z_{j,i,t}$), so in general the vector $(\hat{z}_{i,j,t})_{\{i,j\} \in L}$ is not well-defined (recalling that $\{i,j\} = \{j,i\}$ by definition). But this vector is well-defined whenever $\hat{z}_{i,j,t} = \hat{z}_{j,i,t}$ for all $\{i,j\} \in L$.

$a_{i,t}$, and observing signals $(z_{i,j,t})_{j \in N_i}$, while her continuation payoff when she deviates equals $u_i(\alpha^*)$, which is weakly less. In the case where $\hat{z}_{i,j,\tau} \neq \hat{z}_{j,i,\tau}$ for some $\{i, j\} \in L$ and $\tau < t$, it follows because her continuation payoff equals $u_i(\alpha^*)$ whether she conforms or deviates. Hence, player i does not have a profitable deviation at any history, so σ^{PRI} is a SE profile.

Proof of Theorem B

For the converse, let the (1, 2) and (2, 3) games be as in Section 4.2.1, and let all other (i, j) games be trivial games with $u_i(a_{i,j}, a_{j,i}) = 0$ for all $(a_{i,j}, a_{j,i}) \in A_{i,j} \times A_{j,i}$. Arguing as in the proof of Proposition 1 now implies that payoff vector $(1, 2, 1, 0, \dots, 0) \in E_{PUB} \setminus E_{PRI}^{PRICT}$.

For the main part of the theorem, I first introduce one unorthodox piece of terminology. Throughout the Appendix, say that an (action or communication) history h_i^t is *on-path* under strategy profile σ if it is reached with positive probability under σ or if there exists another history \tilde{h}_i^t that differs from h_i^t only in player i 's past actions $(a_{i,\tau})_{\tau=0}^t$ such that \tilde{h}_i^t is reached with positive probability under σ . A history is *off-path* otherwise.⁴²

Let $Y_{i,j} = \prod_{\{i',j'\} \in L} (Z_{i',j'} \cup \{0_{i',j'}\}) \cup \{alert\}$, where *alert* and $0_{i',j'}$ are arbitrary disjoint messages not contained in any $Z_{i',j'}$. If a message $y_{i,j}$ is not *alert* and the $\{i', j'\}$ coordinate of $y_{i,j}$ is an element of $Z_{i',j'}$ (rather than $0_{i',j'}$), then I refer to the $\{i', j'\}$ coordinate of $y_{i,j}$ as an $\{i', j'\}$ *report*.⁴³ I show that $E_{PRI}^{PRICT}(Y) \supseteq E_{PUB}$.

Let σ^{PUB} be a SE strategy profile in game Γ_{PUB} . I construct a strategy profile σ^{PRI} in game $\Gamma_{PRI}^{PRICT}(Y)$ which will be shown to be a SE profile with the same payoff vector as σ^{PUB} . I first describe play at action histories, then describe play at on-path communication histories, and finally describe play at off-path communication histories.

Action Histories: Initially, play as in σ^{PUB} (i.e., $\sigma_i^{PRI}(h^{0-}) = \sigma_i^{PUB}(h^{0-})$). At subsequent on-path action histories, $\sigma_i^{PRI}(h_i^{t-}) = \sigma_i^{PUB}(\hat{h}_i^t)$, where $\hat{h}_i^t = \left(a_{i,\tau}, (\hat{z}_{i',j',\tau})_{\{i',j'\} \in L} \right)_{\tau=0}^{t-1}$ and $\hat{z}_{i',j',\tau}$ is the $\{i', j'\}$ report player i received in the period τ communication phase.

If player i received conflicting $\{i', j'\}$ reports in some period $\tau < t$, or did not receive

⁴²The point of this terminology is that if player i “trembles” at an action history but nonetheless an on-path signal is generated, then player i will want to “forget” about the deviation. By calling the resulting history “on-path,” it will be possible to insist that player i plays her mutual-minmax action α_i^* at all “off-path” histories, which is convenient for constructing equilibria.

⁴³In contrast, $0_{i',j'}$ may be interpreted as a null report meaning “no report of $z_{i',j'}$ ”.

an $\{i', j'\}$ report in some period $\tau < t$, then h_i^{t-} is an off-path history (as will become clear from the description of the communication phase below). At off-path action histories, $\sigma_i^{PRI}(h_i^{t-}) = \alpha_i^*$.

On-Path Communication Histories: In round 1, each player i sends message $\left((z_{i,j,t})_{j \in N_i}, (0_{i',j',t})_{\{i',j'\} \neq \{i,j \in N_i\}} \right)$ to every player $j \in N_i$. In subsequent rounds, if all $\{i', j'\}$ reports that player i has sent or received so far equal $\hat{z}_{i',j',t}$, then player i sends every $j \in N_i$ the message with $\{i', j'\}$ report $\hat{z}_{i',j',t}$ for those $\{i', j'\}$ for which she has received a report and with $\{i', j'\}$ coordinate $0_{i',j',t}$ for those $\{i', j'\}$ for which she has not yet received a report. Consequently, if player i has sent or received conflicting $\{i', j'\}$ reports for some $\{i', j'\}$, or has sent or received *alert*, then her history is off-path.

Off-Path Communication Histories: Send *alert* to all $j \in N_i$.

Note that if all players follows σ^{PRI} , then for every player $i \in N$ and every on-path action history h_i^{t-} , $\sigma_i^{PRI}(h_i^{t-}) = \sigma_i^{PUB}(\hat{h}_i^t)$, where $\hat{h}_i^t = \left(a_{i,\tau}, (z_{i',j',\tau})_{\{i',j'\} \in L} \right)_{\tau=0}^{t-1}$. Therefore, σ^{PRI} yields the same payoff vector as σ^{PUB} . It remains to show that σ^{PRI} is a SE profile.⁴⁴

I first claim that if any player i deviates from σ^{PRI} at any communication history $h_i^{t,k}$, then every player $j \neq i$ plays α_j^* in all subsequent periods.

The first step in proving the claim is showing that if player i deviates from σ^{PRI} at any communication history $h_i^{t,k}$, then some other player reaches an off-path history during the period t communication phase. This is clearly true if player i deviates by sending *alert*, as *alert* is never sent on-path. It is also true if player i deviates by sending (to some $j \in N_i$) a message with $\{i', j'\}$ coordinate $0_{i',j',t}$ rather than sending an $\{i', j'\}$ report, or by sending an $\{i', j'\}$ report rather than $0_{i',j',t}$, as player j “knows” at what rounds player i sends an $\{i', j'\}$ report on-path.⁴⁵ The only remaining possibility is that player i deviates by sending an $\{i', j'\}$ report $\hat{z}_{i',j',\tau} \neq z_{i',j',\tau}$ to some $j \in N_i$. Assume without loss of generality that $i' \neq i$. Let (i', j_1, \dots, j_l, j) be a path from i' to j that does not include i , which exists by 2-connectedness. Then in round 1 player i' sends $\{i', j'\}$ report $z_{i',j',\tau}$ to player j_1 , and by

⁴⁴As in the proof of Theorem A, I will show that (σ^{PRI}, μ) is a SE for any consistent belief system μ .

⁴⁵In particular, player i sends an $\{i', j'\}$ report at round k if and only if $k \leq \min\{d(i, i'), d(i, j')\} + 1$.

induction in round $l' + 1$ player $j_{l'}$ either sends $\{i', j'\}$ reports $z_{i',j',\tau}$ to player $j_{l'+1}$ or sends *alert* to player $j_{l'+1}$. In either case, player j receives either $\{i', j'\}$ report $z_{i',j',\tau}$ or *alert* in round $l + 1$, so $h_j^{t, \max\{k+1, l+1\}}$ is an off-path history.

The second—and final—step in proving the claim is showing that if all players except possibly i conform to σ^{PRI} and some player $j \neq i$ reaches an off-path history during the period t communication phase, then every player $i' \neq i$ plays $\alpha_{i'}^*$ in all subsequent periods. To see this, note that at an off-path history reached by player j during the period t communication phase (call it $h_j^{t,k}$), player j sends *alert* to all of his neighbors. By induction, each player $i' \neq i$ receives *alert* in round $k + d$, where d is the length of the shortest path between j and i' that does not include i (which exists by 2-connectedness). Hence, every player $i' \neq i$ reaches an off-path history during the period t communication phase. Therefore, every subsequent action history is off-path for all $i' \neq i$, so all $i' \neq i$ play $\alpha_{i'}^*$ in all subsequent periods.

It follows from the claim that no player has a profitable deviation at an on-path history: First, at any on-path action history h_i^{t-} , player i 's continuation payoff from playing any action a_i is the same as her continuation payoff from playing a_i at history \hat{h}_i^t under σ^{PUB} , and σ^{PUB} is a SE. Second, at any on-path communication history $h_i^{t,k}$, player i 's continuation payoff from conforming to σ^{PRI} equals her continuation payoff under σ^{PUB} conditional on reaching history \hat{h}_i^t , playing $a_{i,t}$, and observing some subset of the period t signals $(z_{i,j,t})_{\{i,j\} \in L}$, while her continuation payoff from deviating equals $u_i(\alpha^*)$, which is weakly less.

Finally, I argue that no player has a profitable deviation at an off-path history. The key observation is that if player i is at an off-path history then regardless of her future play all of her opponents will play α^* in every subsequent period. This is immediate from the claim if player i is the only player that has deviated from σ^{PRI} and player i has deviated at a communication history. If player i deviated from σ^{PRI} at an action history and an off-path signal $z_{i,j}$ was generated, then player j is at an off-path history.⁴⁶ Similarly, if some player $j \neq i$ has deviated from σ^{PRI} , then that player is at an off-path history. In either of these cases, the second paragraph of the proof of the claim implies that all players $i' \neq i$ play α^* in every subsequent period. Therefore, if i conforms to σ^{PRI} her continuation payoff is $u_i(\alpha^*)$,

⁴⁶If player i deviated at an action history and an on-path signal was generated, then player i 's resulting history is classified as on-path.

while if she deviates her continuation payoff is weakly less.

Proof of Theorem 1

Let $Y_{i,j} = \prod_{\{i',j'\} \in L} (Z_{i',j'} \cup \{0_{i',j'}\}) \cup \{alert\}$, as in the proof of Theorem B. To define m^0 , first let L' be an arbitrary spanning tree of L (i.e., a connected subnetwork with no cycles that contains all the nodes in L), and let $N'_i \subseteq N_i$ be the set of player i 's neighbors in L' . Renumber the players such that the *leaf players* in L' (i.e., the players with only one neighbor in L') are numbered $1, 2, \dots, n'$.⁴⁷ Now define m^0 by letting $m_i^0 = 4n'|Z|$ for all $i \in \{1, \dots, n'\}$ and $m_i^0 = 0$ for all $i \in \{n'+1, \dots, n\}$ (in particular, only leaf players start with tokens). In addition, number the elements of Z from 1 to $|Z|$. I show that $\tilde{E}_{PRI}^{TOK}(Y, m^0) \supseteq \tilde{E}_{PUB}$.

Let σ^{PUB} be a PBE strategy profile in Γ_{PUB} . I construct a profile σ^{PRI} in $\Gamma_{PRI}^{TOK}(Y, m^0)$ which will be shown to be an PBE profile with the same payoffs as σ^{PUB} . I first describe play at on-path action phase histories, then describe play at on-path communication phase histories (which are now broken into a “reporting subphase” followed by a “confirmation subphase”), and finally describe off-path play and beliefs.

Actions (On-Path): Initially, play as in σ^{PUB} (i.e., $\sigma_i^{PRI}(h^{0-}) = \sigma_i^{PUB}(h^{0-})$). In subsequent periods, $\sigma_i^{PRI}(h_i^{t-}) = \sigma_i^{PUB}(\hat{h}_i^t)$, where $\hat{h}_i^t = \left(a_{i,\tau}, (\hat{z}_{i',j',\tau})_{\{i',j'\} \in L} \right)_{\tau=0}^{t-1}$ and $\hat{z}_{i',j',\tau}$ is the $\{i', j'\}$ report player i received in the period τ reporting subphase. If player i received conflicting $\{i', j'\}$ reports or did not receive an $\{i', j'\}$ report in some period $\tau < t$, then h_i^{t-} is an off-path history (as will become clear from the description of the reporting subphase below) and $\sigma_i^{PRI}(h_i^{t-})$ is therefore given by the description of off-path play below.

Reporting Subphase (On-Path): The reporting subphase consists of the first $n-1$ rounds of the communication phase, during which the players report all signals of which they have been informed to their neighbors in L' and do not make transfers.⁴⁸ In particular,

⁴⁷That is, the set of leaf players is $\{i : |N'_i| = 1\}$, not to be confused with $\{i : |N_i| = 1\}$. The set of players $\{i : |N_i| = 1\}$ also plays a role in the proof, but I reserve the terminology “leaf players” for $\{i : |N'_i| = 1\}$.

⁴⁸Throughout, *transfer* means transfer of tokens.

player i sends message $(0_{i',j'})_{\{i',j'\} \in L}$ to every player $j \in N_i \setminus N'_i$ in every round of the reporting subphase. In round 1, player i sends message $\left((z_{i,j',t})_{j' \in N_i}, (0_{i',j',t})_{\{i',j'\} \neq \{i,j \in N_i\}} \right)$ to every player $j \in N'_i$. In rounds 1 through $n - 1$, player i sends every $j \in N'_i$ the message with $\{i', j'\}$ report $\hat{z}_{i',j',t}$ if all $\{i', j'\}$ reports she has sent or received in earlier rounds equal $\hat{z}_{i',j',t}$, and with $\{i', j'\}$ coordinate $0_{i',j'}$ if she has not yet received an $\{i', j'\}$ report. (Note that if all players conform then they all learn all of the true signals in the course of the reporting subphase.)

Confirmation Subphase (On-Path): The confirmation subphase consists of all but the first $n - 1$ communication rounds. In every round of the confirmation subphase, every player i sends message $(0_{i',j'})_{\{i',j'\} \in L}$ to all $j \in N_i$, and in addition one player transfers tokens to one of her neighbors (until a certain round is reached after which no tokens are transferred). I now describe the details of these transfers for the time t confirmation subphase.⁴⁹ In what follows, let $\hat{z}^i = (\hat{z}_{i',j'}^i)_{\{i',j'\} \in L}$ be the vector of $\{i', j'\}$ reports received by player i in the time- t reporting subphase (noting that if a confirmation subphase history of player i 's is on-path then player i must have received consistent (i.e., non-conflicting) $\{i', j'\}$ reports for all $\{i', j'\} \in L$ in the reporting subphase), and let q_i be the number between 1 and $|Z|$ assigned to \hat{z}^i .⁵⁰ In addition, let $p_{i,j}$ denote the (unique) path from player i to player j in L' , and let $p_{i,j}^l$ denote the l^{th} player in this path, for $l \in \{1, \dots, d(i, j) + 1\}$.⁵¹ Finally denote a transfer of x tokens by $\$x$.

- Round $n + \sum_{j=1}^{i-1} d(j, j+1) + l - 1$, $i \in \{1, \dots, n' - 1\}$, $l \in \{1, \dots, d(i, i+1)\}$: Player $p_{i,i+1}^l$ sends $\$iq_i/2^t$ to player $p_{i,i+1}^{l+1}$.
- Round $n + \sum_{j=1}^{n'-1} d(j, j+1) + l - 1$, $l \in \{1, \dots, d(n', 1)\}$: Player $p_{n',1}^l$ sends $\$n'q_{n'}/2^t$ to player $p_{n',1}^{l+1}$.
- Round $n + \sum_{j=1}^{n'-1} d(j, j+1) + d(n', 1) + 2 \sum_{j=2}^{i-1} d(1, j) + l - 1$, $i \in \{2, \dots, n'\}$, $l \in \{1, \dots, d(1, i)\}$: Player $p_{1,i}^l$ sends $\$(4 - 1/2^{t-1})n'|Z|$ to player $p_{1,i}^{l+1}$.

⁴⁹The description given here is concise and complete but perhaps difficult to read. See Section 5 for a verbal description of on-path play in the confirmation subphase.

⁵⁰Note that on-path $q_i = q_j$ for all i, j , but if some player has deviated then it may be the case that players i and j are both at on-path histories and yet $q_i \neq q_j$.

⁵¹Note that $p_{i,j} \neq p_{j,i}$. In particular, $p_{i,j}^l = p_{j,i}^{d(i,j)-l+2}$. For example, $p_{i,j}^1 = i = p_{j,i}^{d(i,j)+1}$.

- Round $n + \sum_{j=1}^{n'-1} d(j, j+1) + d(n', 1) + 2 \sum_{j=2}^{i-1} d(1, j) + d(1, i) + l - 1$, $i \in \{2, \dots, n'\}$, $l \in \{1, \dots, d(1, i)\}$: Player $p_{i,1}^l$ sends $\$(4 - 1/2^{l-1})n'|Z|$ to player $p_{i,1}^{l+1}$.
- Round $k \geq n + \sum_{j=1}^{n'-1} d(j, j+1) + d(n', 1) + 2 \sum_{j=2}^{n'} d(1, j)$: No transfers are made.

Off-Path Play and Beliefs: For players i and $j \in N_i$, say that player i *punishes* player j at history h_i^t if player i plays $\alpha_{i,j}^*$ at all subsequent action histories, sends *alert* to player j at all subsequent communication histories, and never again transfers tokens to player j .⁵² I first specify the following aspects of off-path play:

1. If player i receives *alert* from player j at any history h_i^t (on or off-path), then i punishes every player $j' \in N_i \cap C_j^i$.
2. If player i satisfies $|N_i| = 1$ and i sends an off-path signal $z_{i,j}$, message $y_{i,j}$, or transfer $z_{i,j}$ to j at an on-path history h_i^t (i.e., a signal, message, or transfer that i never sends to j at h_i^t under the specification of on-path play), then i punishes j .⁵³
3. If player i receives a transfer $m_{j,i} > 0$ from a player $j \notin N'_i$ at any history h_i^t (on or off-path), then i punishes every player $j' \in N_i \cap C_j^i$.
4. Player i never sends a transfer $m_{i,j} > 0$ to a player $j \notin N'_i$.

Off-path beliefs and the remaining aspects of off-path play are jointly defined by the following recursive procedure, which partitions histories according to their “number of steps off-path.” As will become evident, the beliefs generated by this procedure satisfy the property that a player at a d step off-path history believes that with probability 1 all other players are at histories that are at most d steps off-path.

- Classify history h_i^t as 0 *steps off-path* if it is on-path. Thus, we have already specified play and beliefs at 0 step off-path histories.

⁵²Note that, by definition, if player i punishes player j at history h_i^t and history $h_i^{t'}$ is a successor of h_i^t then i punishes j at $h_i^{t'}$.

⁵³To be precise, say that the $\{i, j\}$ signal $z_{i,j}$ is both “sent” from i to j and “received” by i from j (there is no such ambiguity for reports or transfers).

- Say that a signal, message, or transfer received by player i from player j (resp., sent from player i to player j) at a d step off-path history h_i^t is $d + 1$ steps off-path if it is never received by i from j (resp., sent from i to j) at h_i^t under i 's beliefs $\mu_i(\cdot|h_i^t)$ and the specification of play at $\{0, \dots, d\}$ step off-path histories. Classify the resulting history $h_i^{t'}$ as $d + 1$ steps off-path.
- Specify that if player i receives a $d + 1$ step off-path signal, message, or transfer from j at a d step off-path history h_i^t , then i punishes every player $j' \in N_i \cap C_j^i$.
- Specify that if player i receives a $d+1$ step off-path signal, message or transfer from player j at a d step off-path history h_i^t , she believes that every player $j' \in C_j^i \setminus \{j\}$ received *alert* from every player in $N_{j'} \cap C_j^i$ at history $h_{j'}^t$.⁵⁴ The remaining aspects of i 's beliefs about players $j' \in C_j^i$ are arbitrary, subject to the requirement that with probability 1 history $h_{j'}^t$ was at most d steps off-path. Beliefs about players in $N \setminus C_j^i$ are determined by i 's beliefs at $\{0, \dots, d\}$ step off-path histories and the assumption (necessary for PBE) that they do not depend on $h_{i,j}$. In particular, after receiving a $d+1$ step off-path message from player j , player i remains certain that players in $N \setminus C_j^i$ are at histories at most d steps off-path.
- If player i sends a $d + 1$ step off-path signal, message or transfer to player j at a d step off-path communication history, then her beliefs about all players are determined by her beliefs at $\{0, \dots, d\}$ step off-path histories and the specification of play at $\{0, \dots, d\}$ step off-path histories.
- Observe that if player i sends or receives a $d+1$ step off-path message or transfer at a d step off-path communication history, she now faces a distribution of opposing (i, j) -game action plans for all $j \in N_i$ determined by her beliefs, the specification of play at $\{0, \dots, d\}$ step off-path histories, and the fact that any player j at a d' step off-path history h_j^t punishes i if he receives a $d' + 1$ step off-path signal, message, or transfer from i , for all $d' \leq d$.⁵⁵ Specify that player i 's continuation

⁵⁴If h_i^t is an action history, this would not be possible if players were not allowed to send messages concurrently with actions. However, this contingency cannot arise if monitoring in all (i, j) games has full support, and if monitoring in all (i, j) games is perfect then one could specify that player i believes that every player $j' \in C_j^i \setminus \{j\}$ observed an off-path action rather than receiving *alert*.

⁵⁵Note that the distribution of opposing (i, j) -action plans but not opposing strategies is specified, as we

play at histories consistent with this distribution of opposing action plans is (Nash) optimal.⁵⁶

- Classify history $h_i^{t'}$ as $d + 1$ steps off-path if it reached with positive probability following a d step off-path history h_i^t given the above continuation play and beliefs. Thus, we have specified play and beliefs at $d + 1$ step off-path histories.

This completes the description of off-path play and beliefs, and thus completes the description of σ^{PRI} .

It is clear that σ^{PRI} yields the same payoffs as σ^{PUB} . I now show that σ^{PRI} is a PBE.

A preliminary observation, which I will use repeatedly, is that a leaf player i never transfer tokens at an off-path history h_i^t if she has conformed to σ^{PRI} in the past. To see this, note that at any off-path history h_i^t where i has conformed to σ^{PRI} in the past, i has received an off-path signal, message, or transfer from some player j at an earlier on-path history, and hence i punishes every player $j' \in N_i \cap C_j^i$ at h_i^t . Since L' spans L and i is a leaf player, $C_j^i = N \setminus \{i\}$. Hence, i never again transfers tokens to any player.

I now prove a key lemma, which says that if player i deviates from σ^{PRI} then each of her neighbors either minmaxes her or plays as if she had conformed to σ^{PRI} .

Lemma 1 *For every pair of players i and $j \in N_i$, every strategy σ_i , and every action history h_j^{t+1-} reached under strategy profile $(\sigma_i, \sigma_{-i}^{PRI})$, $\sigma_{j,i}^{PRI}(h_j^{t+1-}) \in \left\{ \alpha_{j,i}^*, \sigma_{j,i}^{PUB} \left(\hat{h}_j^{t+1} \right) \right\}$, where $\hat{h}_j^{t+1} = \left(a_{j,\tau}, (z_{i',j',\tau})_{\{i',j'\} \in L} \right)_{\tau=0}^t$.*

Proof. Suppose, toward a contradiction, that $\sigma_{j,i}^{PRI}(h_j^{t+1-}) \notin \left\{ \alpha_{j,i}^*, \sigma_{j,i}^{PUB} \left(\hat{h}_j^{t+1} \right) \right\}$ for some $j \in N_i$. Note that if j ever received an off-path signal, transfer, or message, then the player j' from whom he received it must be in C_i^j (since only i deviates from σ^{PRI}), so j plays $\alpha_{j,i}^*$,

have not yet specified player j 's play towards his other neighbors after a deviation by i , nor have we specified his play toward i after deviations by his other neighbors. However, these aspects of player j 's strategy are irrelevant for computing player i 's optimal continuation play.

⁵⁶One might worry that player i could fail to have a “best response” here because $m_{i,j}$ can take on infinitely many values. However, $m_{i,j}$ takes on only finitely many values on-path, and we have specified that playing any off-path $m_{i,j}$ leads every $j' \in N_i \cap C_j^i$ to punish i . In addition, player i 's continuation payoff against players $j' \notin C_j^i$ is non-decreasing in her token holding, as more continuation strategies against players $j' \notin C_j^i$ are feasible when she holds more tokens. Hence, any off-path $m_{i,j}$ is “weakly dominated” by $m_{i,j} = 0$, so in effect player i need only choose among the finitely many on-path values on $m_{i,j}$ and $m_{i,j} = 0$.

(as j punishes every player in $N_i \cap C_{j'}^j$, and $j' \in C_i^j$ implies $C_{j'}^j = C_i^j$). Hence, history h_j^{t+1-} must be on-path, and there must be a period $t' \leq t$ such that in the period t' communication phase j received a consistent vector of $\{i', j'\}$ reports that does not equal $(z_{i', j', t'})_{\{i', j'\} \in L}$. I consider three cases, deriving a contradiction in each:

Case 1: Player i is a leaf player. Since only i deviates from σ^{PRI} , if player j is at an on-path history with incorrect reports then it must be that some player $j' \in N_i \setminus \{j\}$ received an off-path signal, transfer, or message from player i at an on-path history $h_{j'}^{t'}$ with $t' \leq t$ (note that it is not possible that j' received an incorrect but on-path report from i , because the fact that j is i 's only neighbor in L' implies that all reports received by j' from i are off-path). Then j' sends *alert* to all players in $N_{j'} \cap C_i^{j'}$, which because i is a leaf player includes player $p_{j', j}^2$. By induction, all players in $p_{j', j}$, including player j , receive *alert* during the period t' communication phase. This contradicts the hypothesis that h_j^{t+1-} is on-path.

Case 2: Player i is not a leaf player, and history $h_1^{t', 0}$ is off-path. Let $t_0 \leq t'$ be the first time τ such that history $h_1^{\tau, 0}$ is off-path. I will show that player j does not receive the $\$(4 - 1/2^{t_0-1})n'|Z|$ transfer in period t_0 , which contradicts hypothesis that history h_j^{t+1-} is on-path.

I first claim that no non-leaf player has any tokens at the beginning of period τ for all $\tau \leq t_0 - 1$. The proof is by induction on τ . The claim is immediate for $\tau = 0$. Suppose it is true for some $\tau \leq t_0 - 2$. Then if player 1's first transfer in the period τ communication phase equals q_1 and some non-leaf player does not fully pass on one or more of the transfers he receives in the period τ communication phase, then player 1 does not receive either the $\$n'q_1/2^\tau$ transfer or the $\$(4 - 1/2^{\tau-1})n'|Z|$ transfer in period τ (since all leaf players are following σ^{PRI} and no non-leaf player has any tokens at the beginning of period τ). But then history $h_1^{\tau+1}$ would be off-path, and since $\tau + 1 < t_0$ this would contradict the definition of t_0 . Hence, it must be that no non-leaf player has any tokens at the beginning of period $\tau + 1$. The claim follows by induction.

Next, I claim that the joint token holdings of all non-leaf players at the beginning of period t_0 is at most $\$(4 - 1/2^{t_0-2})n'|Z|$. To see this, suppose that the non-leaf players collectively try to maximize their joint token holdings in the period $t_0 - 1$ communication phase. Note that every token that the non-leaf players do not pass on to a leaf player out

of any on-path transfer they receive reduces the size of the next on-path transfer sent by a leaf player by more than one token, and that leaf players do not send transfers at off-path histories. So the joint token holdings of the non-leaf players is maximized when they pass on all on-path transfers except the last one, which is of size $\$(4 - 1/2^{t_0-2}) n' |Z|$.

Now if player j receives a transfer of size $\$(4 - 1/2^{t_0-1}) n' |Z|$ in period t_0 it must be that the joint token holdings of the non-leaf players (including player j if he is a non-leaf player) reaches $\$(4 - 1/2^{t_0-1}) n' |Z|$ at some point during period t_0 . However, it can be seen that the joint token holdings of the non-leaf players at any point in period t_0 is no more than $\$(4 - 1/2^{t_0-2}) n' |Z| + (n' - 1) |Z| / 2^{t_0}$, since they start the period with at most $\$(4 - 1/2^{t_0-2}) n' |Z|$ and can obtain at most $\$(n' - 1) |Z| / 2^{t_0}$ more in the course of the communication phase (by sending $\$|Z| / 2^{t_0}$ to player 2 in the appropriate round and eventually receiving $\$n' |Z| / 2^{t_0}$ from player n'). Finally,

$$(4 - 1/2^{t_0-2}) n' |Z| + (n' - 1) |Z| / 2^{t_0} < (4 - 1/2^{t_0-1}) n' |Z|.$$

Hence, player j does not receive the $\$(4 - 1/2^{t_0-1}) n' |Z|$ transfer in period t_0 .

Case 3: Player i is not a leaf player, and history $h_1^{t',0}$ is on-path. If player 1 does not receive a consistent vector of reports in the period t' reporting subphase, then the argument is as in Case 2. So suppose that she does, and denote this vector by $(\hat{z}_{i',j'})_{\{i',j'\} \in L}$. Note that it is not possible for all players other than i to have the same consistent—but incorrect—vector of reports at the start of the period t' confirmation phase, as if $\hat{z}_{i',j'} \neq z_{i',j'}$ then players i' and j' cannot have consistent vector $(\hat{z}_{i',j'})_{\{i',j'\} \in L}$. So there is some player $i' \notin \{1, i\}$ who at the start of confirmation phase is either off-path or is on-path with consistent vector $(\tilde{z}_{i',j'})_{\{i',j'\} \in L} \neq (\hat{z}_{i',j'})_{\{i',j'\} \in L}$. Let q be the number assigned to $(\hat{z}_{i',j'})_{\{i',j'\} \in L}$, and let $q' \neq q$ be the number assigned to $(\tilde{z}_{i',j'})_{\{i',j'\} \in L}$.

Consider two cases:

1. $q < q'$: Let κ be the communication round where player i' first receives a transfer on-path. I first claim that player i' punishes every player in $N_{i'} \cap C_i^{i'}$ at round $\kappa + 1$. To see this, first note that no non-leaf player begins period t' with any tokens, by the same argument as in Case 2 (because $h_1^{t',0}$ is on-path). Hence, for no joint strategy of

the non-leaf players is their joint token holding at round κ greater than $q/2^{t'}$ times the number of leaf players who send transfers prior to round κ . Since i' receives a transfer of $q'/2^{t'}$ times this number under σ^{PRI} , any transfer she receives at round κ is off-path. In addition, since only i deviates from σ^{PRI} , the first player j' from whom i receives an on-path signal, transfer, or message must be in $C_i^{i'}$. Hence, player i' punishes every player in $N_{i'} \cap C_{j'}^{i'} = N_{i'} \cap C_i^{i'}$ at round $\kappa + 1$.

I now consider two subcases, and show that in each one player 1 does not receive her expected $\$n'q/2^{t'}$ transfer in period t' . First, suppose i' lies on the path from 1 to i in L' . Let $l \neq i$ be a leaf player such that i lies on path from 1 to l in L' (which exists since i is not a leaf player). Then neither i nor l receive a transfer in period t' , because i' punishes every player in $N_{i'} \cap C_i^{i'}$ at round $\kappa + 1$ and all players except i only transfer tokens along the links of L' (even off-path). Hence, l never sends a transfer in period t' (as leaf players do not send transfers off-path), no leaf player sends a transfer that is more than $\$q/2^{t'}$ greater than the transfer she receives, and non-leaf players begin period t' with no tokens, so player 1 does not receive the $\$n'q/2^{t'}$ transfer. Next, suppose i' does not lie on the path from 1 to i in L' . Let l' be a leaf player (possibly equal to i') such that i' lies on path from 1 to l' in L' . Then any transfer sent by l' will not reach either i or 1, because i' punishes every player in $N_{i'} \cap C_i^{i'}$ at round $\kappa + 1$ and all players except i only transfer tokens along the links of L' . So again player 1 does not receive the $\$n'q/2^{t'}$ transfer.

Now since player 1 does not receive the $\$n'q/2^{t'}$ transfer, she does not send the $\$(4 - 1/2^{t'-1})n'|Z|$ transfer. Finally, as argued in Case 2, the non-leaf players can collectively obtain no more than $\$(n' - 1)|Z|/2^{t'} < \$(4 - 1/2^{t'-1})n'|Z|$ in the course of the period t' confirmation phase, so it follows that player j does not receive the $\$(4 - 1/2^{t'-1})n'|Z|$ transfer. This contradicts the hypothesis that history $h_j^{t'+1-}$ is on-path.

2. $q > q'$: If player i' receives an off-path transfer at round κ (as well as if she punishes every player in $N_{i'} \cap C_i^{i'}$ at round $\kappa + 1$ due to an earlier deviation), the argument is as in the $q < q'$ case. The remaining case is where in round κ player i' receives a transfer

equal to $q'/2^{t'}$ times the number of leaf players who send transfers prior to round κ . Let l be a leaf player (possibly equal to i') such that i' lies on the path from 1 to l in L' . Then the first transfer l receives in the period t' confirmation phase is at most $q'/2^{t'}$ times the number of leaf players who send transfers prior to this round. Hence, l then sends a transfer that is at most $q'/2^{t'}$ greater than the transfer she received. It follows that player 1 does not receive her expected $\$n'q'/2^{t'}$ transfer in period t' , because non-leaf players begin period t' with no tokens, no leaf player sends a transfer that is more than $\$q/2^{t'}$ greater than the transfer she receives, and some leaf player (player l) sends a transfer that is only at most $\$q'/2^{t'}$ greater than the transfer she receives. This yields a contradiction as in the $q < q'$ case.

■

Lemma 1 is not quite enough to rule out on-path deviations. The following lemma will also be needed.

Lemma 2 *Suppose that under strategy profile $(\sigma_i, \sigma_{-i}^{PRI})$ an off-path action history h_j^{t-} is reached for some $j \in N_i$. Then $\sigma_{j',i}^{PRI}(h_{j'}^{t'-}) = \alpha_{j',i}^*$ for all $t' > t$ and all $j' \in N_i$.*

Proof. It suffices to show that if h_j^{t-} is off-path for some $j \in N_i$ then the next action history $h_{j'}^{t+1-}$ is off-path for all $j' \in N_i$. For if $h_{j'}^{t+1-}$ is off-path then the first off-path signal, transfer, or message received by player j' must have come from a player $j'' \in C_i^{j'} = C_{j''}^{j'}$, and hence $\sigma_{j',i}^{PRI}(h_{j'}^{t'-}) = \alpha_{j',i}^*$ for all $t' > t$.

I now show that $h_{j'}^{t+1-}$ is off-path for all $j' \in N_i$, considering three cases.

Case 1: Player i is a leaf player. Since i is a leaf player, any incorrect report she sends is off-path, as on-path she only sends reports to her neighbor $j' \in N_i'$ and j' observes $z_{i,j'}$. Hence, since h_j^{t-} is off-path, player i must have sent an off-path signal, transfer or message to some player $j' \in N_i$ at some time $t' < t$. The same argument as in Case 1 of the proof of Lemma 1 now implies that every player $j'' \in N_i$ receives *alert* in the period t' communication phase. So $h_{j''}^{t+1-}$ is off-path for all $j'' \in N_i$.

Case 2: Player i is not a leaf player and history $h_1^{t_0,0}$ is off-path. The same argument as in Case 2 of the proof of Lemma 1 implies that no player receives the $\$(4 - 1/2^{t_0-1})n'|Z|$ transfer in the first period t_0 at which $h_1^{t_0,0}$ is off-path. So $h_{j'}^{t+1-}$ is off-path for all $j' \in N_i$.

Case 3: Player i is not a leaf player and history $h_1^{t,0}$ is on-path. Since only i deviates from σ^{PRI} , the first player j' from whom player j received an off-path signal, transfer, or message must lie in C_i^j , as only i deviates from σ^{PRI} . Hence, player j punishes every player in $N_{i'} \cap C_{j'}^j = N_{i'} \cap C_i^j$ at history $h_j^{t,0}$. The same argument as in the first subcase of Case 3 of the proof of Lemma 1 now implies that no player receives the $\$(4 - 1/2^{t-1})n'|Z|$ transfer in period t . So $h_{j''}^{t+1-}$ is off-path for all $j'' \in N_i$. ■

Together, Lemmas 1 and 2 imply that there are no profitable deviations at on-path histories, as follows. It is clear that there are no profitable deviations at on-path action histories, as playing any action a_i at an on-path action history h_i^{t-} under σ^{PRI} yields the same continuation payoff as does playing action a_i at history \hat{h}_i^t under σ^{PUB} . Now suppose, toward a contradiction, that player i has a profitable deviation at an on-path period t communication history. By Lemmas 1 and 2, such a deviation must lead some of i 's neighbors to start minmaxing i in period $t+1$ and lead the rest of them to play $\sigma_{j,i}^{PUB}(\hat{h}_j^{t+1})$ in period $t+1$ and then start minmaxing i in period $t+2$. Such a deviation is weakly worse for i than conforming to σ^{PRI} in the period t communication phase, deviating to her myopic best response in the period $t+1$ action phase, and playing α_i^* from period $t+2$ on, since the latter deviation yields a weakly higher payoff in period $t+1$ (as best-responding to an arbitrary mixed action gives a weakly higher payoff than best-responding to the minmax mixed action) and the same payoff in all subsequent periods. But the latter deviation is not profitable, since there are no profitable deviations at on-path action histories, so the proposed deviation cannot be profitable, either.

Finally, I argue that there are no profitable deviations at off-path histories. Start with a lemma.

Lemma 3 *If the specification of off-path play requires that player i punishes player j at history h_i^t , then player i believes that every player $j' \in N_i \cap C_j^i$ punishes player i at history $h_{j'}^{t'}$, where $h_{j'}^{t'}$ is the history immediately following $h_{j'}^t$ (i.e., $h_{j'}^{t'} = h_{j'}^{t,k+1}$ if $h_i^t = h_i^{t,k}$; $h_{j'}^{t'} = h_{j'}^{t,0}$ if $h_i^t = h_i^{t-}$).*

Proof. Player i is only required to punish j at off-path histories. I consider each of the different ways in which i may reach an off-path history.

First, i may receive an off-path signal, message, or transfer from a player $j' \in N_i \cap C_j^i$ at an on-path history h_i^τ . If $|C_j^i| = 1$, then $j' = j$, so $|N_j| = 1$ and j sent an off-path signal, message, or transfer to i at on-path history h_j^τ (as if $N_j = \{i\}$ and h_i^τ is on-path then h_j^τ must be on-path as well), so j punishes i . If $|C_j^i| \neq 1$, then if $j \neq j'$ then i believes that j received *alert* from a player $j'' \in N_i \cap C_{j'}^i$ at history h_j^τ . Hence, i believes that j punishes every player in $N_j \cap C_{j''}^i$, which includes i , at h_i^τ . Alternatively, if $|C_j^i| \neq 1$ and $j = j'$, then i believes that some player $j'' \in N_j$ (with $j'' \neq i$) received alert from j at on-path history h_j^τ , and hence that j will receive *alert* from j'' at history $h_{j''}^\tau$. Hence, i believes that j punishes i at history $h_{j''}^\tau$.

Second, i may have sent or received an off-path signal, message, or transfer to/from a player outside of $N_i \cap C_j^i$ at an on-path history h_i^τ . Then if i is required to punish j it is because i subsequently (1) received *alert* from a player $j' \in N_i \cap C_j^i$, (2) received a transfer from a player $j' \in N_i \cap C_j^i \setminus N_i'$, or (3) received a $d + 1$ step off-path signal, message, or transfer from a player $j' \in N_i \cap C_j^i$ at a d step off-path history. Since transfers are never sent along links outside of L under σ^{PRI} , (2) also represents a $d + 1$ step off-path transfer. So, since $j \in N_i \cap C_{j'}^i$, both (2) and (3) lead i to believe that j received *alert* from a player $j'' \in N_i \cap C_{j'}^i$ at history h_j^τ . Hence, i believes that j punishes every player in $N_j \cap C_{j''}^i$, which includes i . For (1), if this *alert* represents a $d + 1$ step off-path message, the same argument applies. If not, then it must be that some player $j' \in N_i \cap C_j^i$ received an off path signal, message, or transfer from i . In this case, i believes that j received *alert* from a player $j'' \in C_{j'}^i$, and hence punishes every player in $N_j \cap C_{j''}^i$, which includes i .

Finally, i may have sent an off-path signal, message, or transfer to a player $j' \in N_i \cap C_j^i$ at an on-path history h_i^τ . If $|N_i| = 1$, then since h_i^τ is on-path i believes that h_j^τ is on-path, and in addition i believes that this signal, message, or transfer is never received by j from i at h_j^τ under σ^{PRI} (as in this case h_i^τ is measurable with respect to h_j^τ , so any signal, message, or transfer that is never sent from i to j at h_i^τ is also never received by j from i at h_j^τ). Hence, i believes that j punishes i . If instead $|N_i| \neq 1$, then i is required to punish j only if i subsequently received *alert* from a player $j' \in N_i \cap C_j^i$, received a transfer from a player $j' \in N_i \cap C_j^i \setminus N_i'$, or received a $d + 1$ step off-path signal, message, or transfer from a player $j' \in N_i \cap C_j^i$ at a d step off-path history, in which case the same argument as in the preceding

paragraph applies. ■

Note that the only path in L from a player in $N_i \cap C_j^i$ to a player in $N_i \setminus C_j^i$ is the one through i , so if player i 's continuation strategy against players $j' \in N_i \cap C_j^i$ maximizes her (i, j') -game continuation payoff for all $j' \in N_i \cap C_j^i$ as well as the transfer she receives from every player $j' \in N_i \cap C_j^i$ in every round then it maximizes her payoff overall (for any fixed continuation strategy against players $j' \in N_i \setminus C_j^i$). By Lemma 3, at any off-path history h_i^t where $\sigma_{i,j}$ is specified, player i punishes every player $j' \in N_i \cap C_j^i$ and player i believes that every player $j' \in N_i \cap C_j^i$ punishes player i at history h_i^t regardless of i 's strategy. Therefore, every player $j' \in N_i \cap C_j^i$ plays $\alpha_{j',i}^*$ and does not transfer tokens to player i at all future histories. Hence, it is optimal for player i to play $\alpha_{i,j'}^*$, send *alert*, and not transfer tokens to every player $j' \in N_i \cap C_j^i$ at all future histories. Finally, transferring $m_{i,j} > 0$ to a player $j \notin N_i'$ leads all $j' \in N_i \cap C_j^i$ to punish player i , so it is optimal for player i to never make such a transfer. It follows that player i does not have a profitable deviation at any off-path history, completing the proof.

Proof of Proposition 1

I sketch the necessary modification of the proof of Theorem 1, omitting the details.

Let L' be such a spanning tree, and renumber the leaf players in L' by $1, \dots, n'$, as in the proof of Theorem 1. Let $\varepsilon = \min_{i \in \{1, \dots, n'\}} m_i^0$. Add a new “redistribution subphase” to the start of the period 0 communication phase. In it, all non-leaf players pass all their tokens to player 1. Let $x = \sum_{i=n'+1}^n m_i^0$ be the joint initial token holding of the non-leaf players, so that player 1 receives $\$x$ in the redistribution subphase. The rest of the strategy profile is as in the proof of Theorem 1, except that throughout $\$q_i/2^t$ is replaced with $\$\frac{q_i}{2^t} \left(\frac{\varepsilon}{4n'|Z|} \right)$ and $\$(4 - 1/2^{t-1})n'|Z|$ is replaced with $\$x + (1 - 1/2^{t+1})\varepsilon$, reflecting the fact that players $2, \dots, n'$ now end the redistribution subphase with as little as $\$\varepsilon$ rather than $\$4n'|Z|$ and player 1 ends the redistribution subphase with as little as $\$x + \varepsilon$ rather than $\$4n'|Z|$.

The proof that this is a SE profile is a minor extension of the proof of Theorem 1. Intuitively, the facts that non-leaf players end the redistribution subphase with no tokens and that the “confirmation transfer” $\$x + (1 - 1/2^{t+1})\varepsilon$ is greater than $\$x$ and increases each period imply that no player can mislead another about the history of signals.

Proof of Theorem 2

I first prove the result for “essential,” and then describe how it must be modified for “strongly essential.”

I start by introducing the notion of an M -local public equilibrium (M -LPE), where M is an arbitrary subnetwork of L . This is defined to be a PBE in Γ_{PRI} in which $\sigma_{i,j}(h_i^t)$ depends only on $(z_{i,j,\tau})_{\tau=0}^{t-1}$ for all $\{i,j\} \in M$, and $\sigma_{i,j}(h_i^t)$ depends only on $\left((a_{i,j',\tau}, z_{i,j',\tau})_{\{i,j'\} \in L \setminus M}\right)_{\tau=0}^{t-1}$ for all $i \in M$ and $j \notin M$. That is, a M -LPE is a PBE in which players in M condition their play in a relationship with another player in M only on past play in that relationship, and condition their play in a relationship with a player outside M only on past play with players outside M . Denote the M -LPE payoff set in Γ_{PRI} by \tilde{E}_{PRI}^{MLPE} .

For the rest of the proof, assume that M is a nice subnetwork of L .

First, I claim that $\tilde{E}_{PRI} = \tilde{E}_{PRI}^{MLPE}$. The argument adapts the proof of Theorem 5.2 of Fudenberg and Levine (1994), which shows that the SE payoff set and PPE payoff set coincide in repeated games with imperfect public monitoring and a product structure. In particular, fix a PBE σ in Γ_{PRI} , any let $\{i,j\} \in M$. Because M is a subtree of L , player i 's beliefs at history h_i^t about player j 's private history depend only on $(a_{i,j,\tau}, z_{i,j,\tau})_{\tau=0}^{t-1}$; this follows from the first additional requirement in the definition of PBE. Given this, the fact that Γ_{PRI} has a product structure implies that player i 's beliefs about player j 's private history depend only on $(z_{i,j,\tau})_{\tau=0}^{t-1}$, by Bayes rule. Now replace $\sigma_{i,j}$ with a strategy that depends only on $(z_{i,j,\tau})_{\tau=0}^{t-1}$ but has the same marginals over $A_{i,j}$ conditional on $(z_{i,j,\tau})_{\tau=0}^{t-1}$ as does $\sigma_{i,j}$. Do this for every $\{i,j\} \in M$. In addition, again because M is a subtree of L , for any $\{i,j\} \in L$ with $i \in M$ and $j \notin M$, player i 's beliefs at history h_i^t about player j 's private history depends only on $\left((a_{i,j',\tau}, z_{i,j',\tau})_{\{i,j'\} \in L \setminus M}\right)_{\tau=0}^{t-1}$. For any such i,j , replace $\sigma_{i,j}$ with a strategy that depends only on $\left((a_{i,j',\tau}, z_{i,j',\tau})_{\{i,j'\} \in L \setminus M}\right)_{\tau=0}^{t-1}$ but has the same marginals over $A_{i,j}$ conditional on $\left((a_{i,j',\tau}, z_{i,j',\tau})_{\{i,j'\} \in L \setminus M}\right)_{\tau=0}^{t-1}$ as does $\sigma_{i,j}$. Then the resulting strategy profile (after both kinds of replacements) is a M -LPE with the same payoffs as σ ; this is because for every pure strategy of any player i , she faces the same distribution over outcomes whether her opponents follow the original strategy profile or the modified strategy profile.

Second, I claim that $\tilde{E}_{PRI}^{MLPE} = E_{PRI}^{LPE}|_M + \tilde{E}_{PRI}|_{L \setminus M}$.⁵⁷ To see this, given a LPE σ' in $\Gamma_{PRI}|_M$ and a PBE σ'' in $\Gamma_{PRI}|_{L \setminus M}$, define a strategy profile σ in Γ_{PRI} by letting $\sigma_{i,j}(h_i^t) = \sigma'_{i,j}((z_{i,j,\tau})_{\tau=0}^{t-1})$ if $\{i,j\} \in M$ and $\sigma_{i,j}(h_i^t) = \sigma''_{i,j}\left(\left((a_{i,j',\tau}, z_{i,j',\tau})_{\{i,j'\} \in L \setminus M}\right)_{\tau=0}^{t-1}\right)$ if $\{i,j\} \in L \setminus M$. Then it is straightforward to check that σ is a M -LPE in Γ_{PRI} and that payoffs under σ are the sum of payoffs under σ' and σ'' , so $\tilde{E}_{PRI}^{MLPE} \supseteq E_{PRI}^{LPE}|_M + \tilde{E}_{PRI}|_{L \setminus M}$. Similarly, given a M -LPE in Γ_{PRI} , σ , define strategy profiles σ' in $\Gamma_{PRI}|_M$ and σ'' in $\Gamma_{PRI}|_{L \setminus M}$ by $\sigma'_{i,j}(h_i^t) = \sigma_{i,j}(h_i^t)$ for all $\{i,j\} \in M$ and $\sigma''_{i,j}(h_i^t) = \sigma_{i,j}(h_i^t)$ for all $\{i,j\} \notin M$. Then σ' is a LPE in $\Gamma_{PRI}|_M$, σ'' is a PBE in $\Gamma_{PRI}|_{L \setminus M}$, and payoffs under σ are the sum of payoffs under σ' and σ'' , so $\tilde{E}_{PRI}^{MLPE} \subseteq E_{PRI}^{LPE}|_M + \tilde{E}_{PRI}|_{L \setminus M}$. Combining the inclusions yields $\tilde{E}_{PRI}^{MLPE} = E_{PRI}^{LPE}|_M + \tilde{E}_{PRI}|_{L \setminus M}$.

Third, I claim that $\tilde{E}_{PRI}^{TOK}|_M \supseteq E_{PRI}^{LPE}|_M$ and $\tilde{E}_{PRI}^{TOK}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$. The inclusion follows because $\tilde{E}_{PRI}^{TOK}|_M \supseteq \tilde{E}_{PRI}|_M$ (by the observation preceding Corollary 1) and $\tilde{E}_{PRI}|_M \supseteq E_{PRI}^{LPE}|_M$ (because LPE refines PBE). The inequality follows because $\tilde{E}_{PRI}^{TOK}|_M \supseteq \tilde{E}_{PUB}|_M$ (by Theorem 1) and $\tilde{E}_{PUB}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$ (because M is nice).

Finally, I claim that $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PRI}^{TOK}|_M + \tilde{E}_{PRI}|_{L \setminus M}$. To see this, for any message set and vector of initial token holdings (\tilde{Y}, \tilde{m}^0) in $\Gamma_{PRI}|_M$, define message set and initial token holdings (Y, m^0) in Γ_{PRI} by $Y_{i,j} = \tilde{Y}_{i,j}$ if $\{i,j\} \in M$, $Y_{i,j} = \emptyset$ if $\{i,j\} \notin M$, $m_i^0 = \tilde{m}_i^0$ if $i \in M$, and $m_i^0 = 0$ if $i \notin M$. Then $\tilde{E}_{PRI}^{TOK}(Y, m^0) \supseteq \tilde{E}_{PRI}^{TOK}(\tilde{Y}, \tilde{m}^0)|_M + \tilde{E}_{PRI}|_{L \setminus M}$, as given a PBE σ' in $\tilde{E}_{PRI}^{TOK}(\tilde{Y}, \tilde{m}^0)|_M$ and a PBE σ'' in $\tilde{E}_{PRI}|_{L \setminus M}$ one can construct a PBE σ in $\tilde{E}_{PRI}^{TOK}(Y, m^0)$ with payoffs equal to the sum of payoffs under σ' and σ'' by letting $\sigma_{i,j}(h_i^t) = \sigma'_{i,j}(\tilde{h}_i^t)$ if $\{i,j\} \in M$, where \tilde{h}_i^t is derived from h_i^t by deleting actions, signals, messages, and transfers along links $\{i,j\} \notin M$, and letting $\sigma_{i,j}(h_i^t) = \sigma''_{i,j}(\hat{h}_i^t)$ if $\{i,j\} \notin M$, where \hat{h}_i^t is derived from h_i^t by deleting actions and signals along links $\{i,j\} \in M$ and deleting all messages and transfers.

Combining the four claims, one has

$$\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PRI}^{TOK}|_M + \tilde{E}_{PRI}|_{L \setminus M} \supseteq E_{PRI}^{LPE}|_M + \tilde{E}_{PRI}|_{L \setminus M} = \tilde{E}_{PRI}^{MLPE} = \tilde{E}_{PRI},$$

where the strict inclusion uses the fact that, for any sets X , X' , and W , if $X \supseteq X'$ and

⁵⁷The notation here is that for sets $A, B \subseteq \mathbb{R}^n$, $A + B = \{a + b : a \in A, b \in B\}$.

$X \setminus \text{co}(X') \neq \emptyset$ then $X + W \supsetneq X' + W$ (as can be seen from a separating hyperplane argument). Therefore, $\tilde{E}_{PRI}^{TOK} \supsetneq \tilde{E}_{PRI}$.

The proof for “strongly essential” is almost identical. In place of a M -local public equilibrium, define a M -local cheap talk equilibrium to be a PBE in Γ_{PRI}^{PRICT} in which players in M condition their play (including messages) in a relationship with another player in M only on past play in that relationship, and condition their play in a relationship with a player outside M only on past play with players outside M . Let \tilde{E}_{PRI}^{MLCTE} be the set of M -local cheap talk equilibrium payoffs in Γ_{PRI}^{PRICT} . Then $\tilde{E}_{PRI}^{PRICT} = \tilde{E}_{PRI}^{MLCTE}$ by the same argument as for $\tilde{E}_{PRI} = \tilde{E}_{PRI}^{MLPE}$, with the addition that strategies about which message to send may also need to be replaced by M -local cheap talk strategies with the same marginals. Next, $\tilde{E}_{PRI}^{MLCTE} = E_{PRI}^{LCTE}|_M + \tilde{E}_{PRI}|_{L \setminus M}$ by the same argument as for $\tilde{E}_{PRI}^{MLPE} = E_{PRI}^{LPE}|_M + \tilde{E}_{PRI}|_{L \setminus M}$, and $\tilde{E}_{PRI}^{TOK}|_M \supseteq E_{PRI}^{LCTE}|_M$ and $\tilde{E}_{PRI}^{TOK}|_M \setminus \text{co}(E_{PRI}^{LCTE}|_M) \neq \emptyset$ by the same argument as for $\tilde{E}_{PRI}^{TOK}|_M \supseteq E_{PRI}^{LPE}|_M$ and $\tilde{E}_{PRI}^{TOK}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$ (where the statement that $\tilde{E}_{PUB}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$ is strengthened to $\tilde{E}_{PUB}|_M \setminus \text{co}(E_{PRI}^{LCTE}|_M) \neq \emptyset$, which is possible when M is truly nice). Combining these inclusions with $\tilde{E}_{PRI}^{TOK} \supseteq \tilde{E}_{PRI}^{TOK}|_M + \tilde{E}_{PRI}|_{L \setminus M}$ as in the “essential” case yields $\tilde{E}_{PRI}^{TOK} \supsetneq \tilde{E}_{PRI}^{PRICT}$.

References

- [1] Ahn, I. and M. Suominen (2001), “Word-Of-Mouth Communication and Community Enforcement,” *International Economic Review*, 42, 399-415.
- [2] Ali, S.N. and D. Miller (2012), “Ostracism,” mimeo.
- [3] Aliprantis, C.D., G. Camera, and D. Puzzello (2007), “Anonymous Markets and Monetary Trading,” *Journal of Monetary Economics*, 54, 1905-1928.
- [4] Ambrus, A., M. Möbius, and A. Szeidl (2010), “Consumption Risk-Sharing in Social Networks,” mimeo.
- [5] Araujo, L. (2004), “Social Norms and Money,” *Journal of Monetary Economics*, 51, 241-256.
- [6] Aumann, R.J. and S. Hart (2003), “Long Cheap Talk,” *Econometrica*, 71, 1619-1660.
- [7] Balmaceda F. and J. Escobar (2011), “Self Governance in Social Networks of Information Transmission,” mimeo.

- [8] Ben-Porath, E. and Kahneman, M. (1996), "Communication in Repeated Games with Private Monitoring," *Journal of Economic Theory*, 70, 281-297.
- [9] Ben-Porath, E. (1998), "Correlation without Mediation: Expanding the Set of Equilibrium Outcomes by Cheap Pre-Play Procedures," *Journal of Economic Theory*, 80, 108-122.
- [10] Bloch, F., G. Genicot, and D. Ray (2008), "Informal Insurance in Social Networks," *Journal of Economic Theory*, 143, 36-58.
- [11] Compte, O. (1998), "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66, 597-626.
- [12] Corbae, D., T. Temzelides, and R. Wright (2003), "Directed Matching and Monetary Exchange," *Econometrica*, 71, 731-756.
- [13] Dixit, A. (2003), "Trade Expansion and Contract Enforcement," *Journal of Political Economy*, 111, 1293-1317.
- [14] Ellison, G. (1994), "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61, 567-588.
- [15] Forges, F. (2009), "Correlated Equilibrium and Communication in Games," *Encyclopedia of Complexity and Systems Science*, edited by R. Meyers, Springer.
- [16] Fudenberg, D. and D.K. Levine (1994), "Efficiency and Observability with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62, 103-135.
- [17] Fujiwara-Greve, T., M. Okuno-Fujiwara, and N. Suzuki (2012), "Voluntarily Separable Prisoner's Dilemma with Reference Letters," *Games and Economic Behavior*, 74, 504-516.
- [18] Greif, A. (2006), *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*, Cambridge: Cambridge University Press.
- [19] Hauser, C. and H. Hopenhayn (2010), "Trading Favors: Optimal Exchange and Forgiveness," mimeo.
- [20] Izmalkov, S., S. Micali, and M. Lepinski (2011), "Perfect Implementation," *Games and Economic Behavior*, 71, 121-140.
- [21] Jackson, M.O., T. Rodriguez-Barraquer, and X. Tan (2011), "Social Capital and Social Quilts: Network Patterns of Favor Exchange," *American Economic Review*, 102, 1857-1897.
- [22] Kandori, M. (1992), "Social Norms and Community Enforcement," *Review of Economic Studies*, 59, 63-80.
- [23] Kandori, M. and H. Matsushima (1998), "Private Observation, Communication and Collusion," *Econometrica*, 66, 627-652.

- [24] Karlan, D., M. Möbius, T. Rosenblat, and A. Szeidl (2009), “Trust and Social Collateral,” *Quarterly Journal of Economics*, 124, 1307-1361.
- [25] Koessler, P. and F. Forges (2008), “Multistage Communication with and without Verifiable Types,” *International Game Theory Review*, 10, 145-164.
- [26] Kiyotaki, N. and R. Wright (1989), “On Money as a Medium of Exchange,” *Journal of Political Economy*, 97, 927-954.
- [27] Kiyotaki, N. and R. Wright (1993), “A Search-Theoretic Approach to Monetary Economics,” *American Economic Review*, 83, 63-77.
- [28] Kocherlakota, N.R. (1998), “Money is Memory,” *Journal of Economic Theory*, 81, 232-251.
- [29] Kocherlakota, N.R. (2002), “The Two-Money Theorem,” *International Economic Review*, 43, 333-346.
- [30] Kocherlakota, N. and N. Wallace (1998), “Incomplete Record-Keeping and Optimal Payment Arrangements,” *Journal of Economic Theory*, 81, 272-289.
- [31] Kreps, D.M., and R. Wilson (1982) “Sequential Equilibria,” *Econometrica*, 50, 863-894.
- [32] Krishna, R.V. (2007), “Communication in Games of Incomplete Information,” *Mathematics of Operations Research*, 10, 117-153.
- [33] Laclau, M. (2012a), “A Folk Theorem for Repeated Games Played on a Network,” *Games and Economic Behavior*, 76, 711-737.
- [34] Laclau, M. (2012b), “Communication in Repeated Network Games with Imperfect Monitoring,” mimeo.
- [35] Laclau, M. (2012c), “Repeated Games with Local Monitoring and Private Communication,” mimeo.
- [36] Lagos, R. and R. Wright (2008), “When is Money Essential? A Comment on Aliprantis, Camera, and Puzzello,” mimeo.
- [37] Linial, N. (1994), “Games Computers Play: Game-Theoretic Aspects of Computing,” in R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory with Economic Applications*, Volume II, 1340-1395, North-Holland: Amsterdam.
- [38] Lippert, S. and G. Spagnolo (2011), “Networks of Relations and Word-of-Mouth Communication,” *Games and Economic Behavior*, 72, 202-217.
- [39] McLean, R., I. Obara, and A. Postlewaite (2012), “Informational Smallness and Private Monitoring,” mimeo.
- [40] Milgrom, P.R., D.C. North, and B.R. Weingast (1990), “The Role of Institutions in the Revival of Trade: the Law Merchant, Private Judges, and the Champagne Fairs,” *Economics and Politics*, 2, 1-23.

- [41] Möbius, M.M. (2001), “Trading Favors,” mimeo.
- [42] Okuno-Fujiwara, M. and A. Postlewaite (1995), “Social Norms and Random Matching Games,” *Games and Economic Behavior*, 9, 79-109.
- [43] Ostroy, J.M. (1973), “The Informational Efficiency of Monetary Exchange,” *American Economic Review*, 63, 597-610.
- [44] Ostroy, J.M. and R.M. Starr (1974), “Money and the Decentralization of Exchange,” *Econometrica*, 42, 1093-1113.
- [45] Renault, J., L. Renou, and T. Tomala (2012), “Secure Message Transmission on Directed Networks,” mimeo.
- [46] Renault, J. and T. Tomala (1998), “Repeated Proximity Games,” *International Journal of Game Theory*, 27, 539-559.
- [47] Renou, L. and T. Tomala (2011), “Mechanism Design and Communication Networks,” *Theoretical Economics*, forthcoming.
- [48] Starr, R.M. (1972), “The Structure of Exchange in Barter and Monetary Economies,” *Quarterly Journal of Economics*, 86, 290-302.
- [49] Tomala, T. (2011), “Fault Reporting in Partially Known Networks and Folk Theorems,” *Operations Research*, 59, 754-763.
- [50] Townsend, R.M. (1980), “Models of Money with Spatially Separated Agents,” in *Models of Monetary Economics*, edited by J.H. Kareken and N. Wallace, Minneapolis: Federal Reserve Bank of Minneapolis.
- [51] Townsend, R. (1987), “Economic Organization with Limited Communication,” *American Economic Review*, 77, 954-970.
- [52] Wallace, N. (2001), “Whither Monetary Economics?” *International Economic Review*, 42, 847-869.