

Communication with Tokens in Repeated Games on Networks*

Alexander Wolitzky

Stanford University

September 29, 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). Sufficient 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, Johannes Hörner, Matt Jackson, Adam Kalai, Chris Phelan, David Rahman, Ilya Segal, Andy Skrzypacz, Joel Sobel, Juuso Toikka, Rakesh Vohra, Randy Wright, and seminar audiences at Bocconi, Cambridge, Central European University, MIT, Microsoft Research, the Minneapolis Federal Reserve, Toulouse, 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

Community enforcement of norms in networked environments 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). A crucial issue for sustaining community enforcement is information sharing: agents cannot rely on community enforcement if they cannot effectively share information about each other’s behavior. While many potential obstacles to information sharing exist, an especially salient one for economists is incentives: a groups of agents cannot effectively share information if they do not have individual incentives to do so.

To see the problem, consider three agents—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. Suppose that 1 and 3 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 does not have a relationship with 1, and 2 clearly cannot be trusted to tell her. So community enforcement fails.

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 via cheap talk does not enable them to sustain certain outcomes that would have been sustainable if this information were public to all players. In the language of this paper, cheap talk does not *replicate* public information in this example.

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 if and only if the network is 2-connected (i.e., it remains connected after any node is removed).¹ 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., there are no “information gatekeepers”), and therefore the players can coordinate as well as if all information were public.

Conversely, in non-2-connected networks, players can replicate public information only if they can not only talk but also exchange some form of hard evidence. In this paper, I focus on a very specific form of evidence: players are endowed with undifferentiated 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. Thus, 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 above example—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. For example, non-leaf players are prevented from cheating some leaf players while concealing this information from others. 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. In the conclusion, I discuss the possibility of using tokens in simpler ways in more special settings.

I then apply this result to study when tokens are *essential*, in that the equilibrium payoff set is strictly larger with tokens than without.² I show that a sufficient condition for tokens

¹This is related to a result of Renault and Tomala (1998). The precise connection is discussed below.

²The terminology here is borrowed from the literature on monetary theory (e.g., Lagos and Wright, 2008).

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 public equilibrium” with private monitoring. In many games, the network contains a nice subnetwork if and only if it contains a subtree of size at least three—a slightly stronger condition than not being 2-connected.

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 unforgeable signatures, 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; Kocherlakota 1998, 2002). 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.

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 gives a benchmark result on replicating public information with cheap talk. Section 5 presents the main result on replicating public information with tokens. Section 6 shows how the main result can be applied to show when tokens are essential in a broad class of games. Section 7 concludes. Appendix A presents examples showing that various conditions for the main result cannot be dispensed with. Omitted proofs are contained in Appendix B.

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). Most of the subsequent literature on community enforcement has not considered hard evidence. There is also a literature on the folk theorem in private monitoring repeated games with communication (Compte, 1998; Kandori and Matsushima, 1998). In this literature, the folk theorems of Ben-Porath and Kahneman (1996) and Renault and Tomala (1998) are closely related to the benchmark result of Section 4; the main difference is that I compare the equilibrium payoff set with different communication technologies for a fixed discount factor.³

Also related is the large computer science-based literature on secure information transmission in networks. Livial (1994) gives a survey aimed at 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. 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 sometimes substitute for private authentication keys or sealed envelopes in facilitating secure information transmission in networks.⁴

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

³McLean, 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.

⁴Relative 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. 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.

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).⁵ 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.

The papers on monetary theory most closely related to mine are Kocherlakota (1998) and Kocherlakota (2002). Kocherlakota (1998) shows that money is often inessential when information is public, and also gives an example in which money cannot replicate public information.⁶ Kocherlakota (2002) shows that this example relies on the assumption that money is indivisible, and shows that infinitely divisible money can replicate public information if money holdings are observable, 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.⁷ Thus, a key difference between my model and Kocherlakota (2002) is that I allow only undifferentiated tokens. On the other hand, unlike Kocherlakota I also allow multiple rounds of transfers in every period, so that in every period each player can (indirectly) receive tokens from all other players. Hence, my main result shows that indirect transfers can substitute for tagged money as a means of replicating public information, even when these

⁵There 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.

⁶More 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). See Zhu and Maenner (2012) for a recent contribution.

⁷A very similar idea appears in Townsend (1987). See also Townsend (1980) for a canonical monetary theory model emphasizing “spatial separation” of agents.

transfers must be made through strategic third parties.⁸

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 $N = \{1, \dots, n\}$, with $n \geq 3$, 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 . 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. 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 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 $z_{i',j'}$. That is, 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 $a =$

⁸There are other differences between my model and Kocherlakota’s. First, Kocherlakota’s model involves “trading mechanisms,” while there are no mechanisms or contracts in my model. However, Kocherlakota’s use of trading mechanisms is actually quite similar to my use of the assumption that every bilateral game has a mutual-minmax Nash equilibrium, in that the two devices simplify off-path play in similar ways. Second, the role of the divisibility of tokens is very different in the two models. In Kocherlakota’s model, the key role of divisibility is allowing precise communication about arbitrarily long histories of play. In my model, players can communicate every period, so precise communication does not require an infinitely large message space. Instead, divisibility is used to scale transfers over time to prevent players from running out of tokens. I thank an anonymous referee for bringing these points to my attention.

$(a_i)_{i \in N}$ is given by $\pi(z|a) = \prod_{\{i,j\} \in L} \pi_{i,j}(z_{i,j}|a_{i,j}, a_{j,i})$.⁹ 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 . 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: 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 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 $(z_{i,j,t})_{j \in N_i}$.¹⁰ 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).

⁹Links are denoted with braces rather than parentheses to emphasize that $\{i, j\}$ and $\{j, i\}$ refer to the same link. Thus, there are the same number of terms in this product as there are links in the network.

¹⁰Thus, player i observes her own payoff.

Solution concept: For the main results 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) rather than sequential equilibrium.^{11,12} Unfortunately, there is no off-the-shelf version of PBE that seems appropriate in this model. For example, consider three players on a line: $L = \{\{1, 2\}, \{2, 3\}\}$. On the one hand, assuming that player 1 does not update her belief about player 3's history after observing an unexpected move by player 2 seems too restrictive, 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 history after observing an unexpected move by player 1 seems too permissive, as 1's play can only affect 3 via 2. In light of these issues, I use the following extension of weak perfect Bayesian equilibrium.¹³ 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 .¹⁴ 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 with the property that if $j \in N_i$ and $j' \notin C_j^i$, then $\mu_i(h_{j'}^t | h_i^t)$ does not depend on $h_{i,j}^t$.*

Note that this definition rules out “grim beliefs,” where whenever a player observes a deviation she believes that all of her opponents also observed a deviation.

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

¹¹Action spaces will remain countable, so there is no problem in defining sequential equilibrium. However, in dynamic games with countably infinite action spaces sequential equilibrium imposes unusually strong restrictions. 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.

¹²In earlier versions of the paper, the solution concept for the benchmark result concerning cheap talk communication (where action spaces are finite) was sequential equilibrium. For the sake of consistency, all results in the current version are stated for PBE.

¹³Recall 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.

¹⁴In other words, 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 .

$(a_{i,t}, (z_{i,j,t})_{j \in N_i})$. Let E_{PUB} be the set of PBE payoffs of game Γ_{PUB} . Below, I will define games Γ_{PRI}^{CT} and Γ_{PRI}^{TOK} by adding cheap talk and tokens to the game Γ_{PRI} , and will denote the corresponding PBE payoff sets by E_{PRI}^{CT} and E_{PRI}^{TOK} . I will say that cheap talk (resp., tokens) *can replicate public information* if $E_{PRI}^{CT} \supseteq E_{PUB}$ (resp., $E_{PRI}^{TOK} \supseteq E_{PUB}$). 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 a benchmark result on when cheap talk may be used to replicate public information. It is broadly similar to results in the literature and is intended primarily as a point of departure for the main analysis of Sections 5 and 6.

Throughout the paper, players can communicate directly with their neighbors only. However, players can communicate with non-neighbors indirectly by passing information from one link to another. This requires multiple rounds of communication after every round of play. 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.¹⁵

A game with cheap talk $\Gamma_{PRI}^{CT}(Y)$ 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 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 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 infinitely many kinds of histories for every period t , denoted $h_i^{t-} = (h_{i,\tau})_{\tau=0}^{t-1}$ (called *action histories*), $h_i^{t,0} = \left((h_{i,\tau})_{\tau=0}^{t-1}, a_{i,t}, (z_{i,j,t})_{j \in N_i} \right)$,

¹⁵An alternative would be allowing 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.

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 $k \in \mathbb{N}$ (called *communication histories*). A strategy σ_i is a measurable function that maps action histories h_i^{t-} to $\Delta(A_i)$ and maps communication histories $h_i^{t,k}$ to $\Delta\left((Y_{i,j})_{j \in N_i}\right)$.¹⁶ Let $E_{PRI}^{CT}(Y)$ be the PBE payoff set of $\Gamma_{PRI}^{CT}(Y)$, and let $E_{PRI}^{CT} = \bigcup_Y E_{PRI}^{CT}(Y)$, where the union is taken over all finite sets Y .¹⁷

The benchmark result is that 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 simple: Start with a PBE 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 mislead 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^* .¹⁸

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 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

¹⁶I continue to denote generic histories by h_i^t . That is, h_i^t may denote either an action history or a communication history.

¹⁷Aumann 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.

¹⁸This 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 more powerful communication technologies like sealed letters or public broadcasts. However, it may be the case that if signals are “almost” locally public then it is possible to “almost” replicate public information. I do not pursue this question here.

all $(a_{i,j}, a_{j,i}) \in A_{i,j} \times A_{j,i}$, cheap talk cannot replicate public information. For example, it suffices to take the (1, 2) and (2, 3) games to be the asymmetric prisoner’s dilemmas described in the example below (which I return to later in the paper).

Theorem A *Cheap talk can replicate public information (i.e., $E_{PRI}^{CT} \supseteq E_{PUB}$) 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 cheap talk cannot replicate public information (i.e., $E_{PUB} \setminus E_{PRI}^{CT} \neq \emptyset$).*

The first part of Theorem A 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 A avoids some complications that emerge in their paper by allowing explicit communication and assuming a mutual-minmax Nash equilibrium (although Theorem A is for sequential equilibrium rather than Nash). Also, Theorem A is not a folk theorem but rather a result about replicating public information for fixed δ .¹⁹ An earlier version of this paper also showed that public information can *always* be replicated if *public* cheap talk is available; that result bears a similar relation to the folk theorem of Ben-Porath and Kahneman (1996) as Theorem A does to Renault and Tomala’s result.

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

	C	D
C	1, 1	$-l_1, 1 + g_{2,1}$
D	$1 + g_1, -l_{2,1}$	0, 0

¹⁹In some games, communication becomes superfluous as discounting vanishes, in which case the results of this paper are only relevant for lower discount factors. In other games, communication expands the equilibrium payoff set even in the limit (either because the folk theorem fails in the “local” games or because pooling incentives across local games expands the set of individually rational payoffs).

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) .²⁰ 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 cheap talk may fail to replicate public information when the network is not 2-connected.

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

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

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

²⁰That 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 PBE. ■

5 Replicating Public Information with Tokens

I now turn to the main part of the analysis, where in addition to sending cheap talk messages players can transfer quantities of undifferentiated, infinitely divisible tokens. 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.

Formally, a game with tokens $\Gamma_{PRI}^{TOK}(Y, m^0)$ is derived from the game with cheap talk $\Gamma_{PRI}^{CT}(Y)$ 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.²¹ That is, at every history in $\Gamma_{PRI}^{CT}(Y)$ where player i chooses a message $y_{i,j} \in Y_{i,j}$ to send to player j , she now chooses a pair $(y_{i,j}, m_{i,j}) \in Y_{i,j} \times \mathbb{Q}_+$ 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 to

$$m'_i = m_i + \sum_{j \in N_i} (m_{j,i} - m_{i,j}).$$

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 $(a_i, (y_{i,j})_{j \in N_i}, (m_{i,j})_{j \in N_i})$).²² Let

²¹The point of allowing players to transfer only rational quantities of tokens is to ensure that strategy spaces remain countable, so that the game is well-defined.

²²This modification plays a “technical” role in the proof of the main result, discussed in footnote 41. It is not needed if either monitoring in all (i, j) games is (locally) perfect or monitoring in all (i, j) games has full support.

$E_{PRI}^{TOK} = \bigcup_{(Y, m^0)} E_{PRI}^{TOK}(Y, m^0)$, where $E_{PRI}^{TOK}(Y, m^0)$ is the PBE payoff set in $\Gamma_{PRI}^{TOK}(Y, m^0)$.

The following is the main result of the paper.

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

Theorem 1 shows that in networked environments tokens enable players to sustain any payoff vector that would be sustainable were all information public. 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.

The key feature of tokens that makes Theorem 1 possible is that tokens enable the secure communication of non-exclusive information. Consider the following one-shot game. Initially, each player observes several signals (i.e., realizations of arbitrary random variables), with each signal being observed by at least two players. Players then engage in unboundedly many rounds of cheap talk and token transfers with their neighbors. A key fact is that secure communication is possible in this game, in that there exists a strategy profile σ such that if the players conform to σ , they all learn all the signals, while if any one player deviates from σ , then each player either learns all the signals or reaches an off-path history.²³

The construction in the one-shot game—which is a simpler version of the construction in the proof of Theorem 1—is as follows. Let L' be a spanning tree of L (i.e., a connected subnetwork of L with no cycles that still contains all n players). Number the “leaf players” in L' (i.e., players with only one neighbor) from 1 to $n' \leq n$. Endow each of the leaf players in L' with a large number of tokens, and endow non-leaf players with none. Assign a natural number q to every possible signal profile z . Have players repeatedly report their signals to each other as in the model with cheap talk. If at the end of this “reporting subphase,” player i has received a consistent vector of reported signals \hat{z}^i (which is the case on-path), let q_i be the corresponding natural number (at off-path histories, players immediately and permanently stop sending reports and tokens). The “confirmation subphase” then begins with player 1 sending q_1 tokens down the (unique) path in L' toward player 2. The non-leaf players on this path then pass these tokens on to player 2. When player 2 receives the

²³This notion of secure communication is weaker than what is possible with access to sealed envelopes or a mediator. Indeed, tokens are weaker than these technologies: with three players on a line, 1 and 3 cannot use tokens to correlate their play without being observed by 2.

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, 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. Observe that player 1 sends these final transfers only if the “pot” contains $q_1 n'$ tokens when it is returned to her: otherwise, her history at the round where she is supposed to send the final transfers is off-path, and players do not transfer tokens at off-path histories.

Under this strategy profile, no single deviator can mislead another player about any signal. Suppose that at the end of the reporting subphase $q_1 \neq q_i$ for some $i \in N$ (if all the q_i 's are the same, then all players have learned the true signals, as each signal is observed by at least two players). I claim that when the pot of tokens is returned to player 1 it will contain fewer than $q_1 n'$ tokens, so that player 1 will not send her final transfers and consequently all players will eventually reach off-path histories (when they do not receive their final transfers). To see this, first note that no leaf player ever adds more than q_1 tokens to the pot, and that player i adds q_i tokens to the pot only if the pot contains $q_i \times (\text{number of leaf players before } i)$ tokens when it reaches her. Hence, if $q_1 > q_i$, then the pot contains at most $q_1 (n' - 1) + q_i < q_1 n'$ tokens when it is returned to player 1, while if $q_1 < q_i$ then the pot contains at most $q_1 (n' - 1) < q_1 n'$ tokens when it is returned to player 1 (as in this case the pot contains at most $q_1 \times (\text{number of leaf players before } i)$ tokens when it reaches player i , which is less than $q_i \times (\text{number of leaf players before } i)$).

This basic construction illustrates the key reason why tokens let players replicate public information. The full construction required for Theorem 1 is more complicated in two main respects.

1. The strategies in the basic construction are “rational” (in that no one deviator can mislead another player), while the strategies in the full construction must be sequentially rational. This makes specifying off-path play delicate, which in turn makes checking on-path incentives more difficult. For example, it is not always sequentially rational for players to stop sending reports and tokens off-path.

2. The basic construction is for a one-shot game, and new issues arise in the repeated game. For example, players' intertemporal budget of tokens must be tracked, so that players do not run out of tokens or have incentives to secretly store tokens from period to period. A trick that helps here is scaling down the transfers each period, taking advantage of the assumption that tokens are infinitely divisible.

To illustrate the construction, consider the asymmetric prisoner's dilemma on a line of Section 4. 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}^{CT}). Initially, endow players 1 and 3 with a large number of tokens, 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 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 final transfer to player 2, 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}^{CT}). But this misreport will be detected in the confirmation subphase

as follows. Since player 1 observes a deviation by player 2, 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, 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, which deters the deviation.

I conclude this section with three 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, 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.

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 $E_{PRI}^{TOK}(Y, m^0) \supseteq E_{PUB}$, where Y is the message set from the proof of Theorem 1.*

However, Proposition 2 continues to assume that the initial token endowment m^0 is common knowledge. Proposition 4 in Appendix A shows that Theorem 1 may fail with uncertain endowments.

Second, Theorem 1 relies on the assumption that tokens are infinitely divisible. This allows transfers to be scaled down over time, which prevents leaf players from running out of tokens and helps ensure 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 in Appendix A shows that Theorem

1 may fail with indivisible tokens if such disbursements are not allowed.

Third, as long as tokens are infinitely divisible it would be essentially equivalent to let players exchange only tokens and not also cheap talk messages. Here is a sketch of the required modifications to the construction in Theorem 1. Assign a natural number q to every partial signal profile $(z_{i,j})_{\{i,j\} \in L', L' \subseteq L}$ (rather than only to complete profiles $(z_{i,j})_{\{i,j\} \in L}$). In the period t reporting subphase, replace cheap talk reports of signal profiles with transfers of $\varepsilon q/2^t$ tokens, where q is the corresponding natural number and $\varepsilon > 0$ is a constant. Finally, in the confirmation subphase, adjust the final transfers so that non-leaf players relinquish the additional tokens they received in the reporting subphase. For sufficiently small ε , token holdings after every history will be close enough to those in the construction with cheap talk for the proof of Theorem 1 to go through.

6 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.

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

Tokens are strongly essential if $E_{PRI}^{TOK} \supsetneq E_{PRI}^{CT}$.

The latter property is indeed stronger because $E_{PRI}^{CT}(Y) \supseteq E_{PRI}$ 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 E_{PRI} , E_{PRI}^{CT} , and 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 $E_{PRI}^{TOK} \supseteq E_{PRI}$ and $E_{PRI}^{TOK} \supseteq E_{PRI}^{CT}$ are trivially true: any PBE in Γ_{PRI} or Γ_{PRI}^{CT} 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, $E_{PRI}^{TOK}(Y, m^0) \supseteq E_{PRI}$ and $E_{PRI}^{TOK}(Y, m^0) \supseteq E_{PRI}^{CT}(Y)$ for any (Y, m^0)). Combining this observation with Theorem 1 yields the following corollary.

Corollary 1 *Tokens are essential if $E_{PUB} \setminus E_{PRI} \neq \emptyset$. Tokens are strongly essential if $E_{PUB} \setminus E_{PRI}^{CT} \neq \emptyset$.*

Proof. By Theorem 1, $E_{PRI}^{TOK} \supseteq E_{PUB}$. So $E_{PUB} \setminus E_{PRI} \neq \emptyset$ (resp., $E_{PUB} \setminus E_{PRI}^{CT} \neq \emptyset$) implies that $E_{PRI}^{TOK} \setminus E_{PRI} \neq \emptyset$ (resp., $E_{PRI}^{TOK} \setminus E_{PRI}^{CT} \neq \emptyset$). Observing that $E_{PRI}^{TOK} \supseteq E_{PRI}$ (resp., $E_{PRI}^{TOK} \supseteq E_{PRI}^{CT}$) completes the proof. ■

Combining Corollary 1 and Proposition 1 shows that tokens are strongly essential in the asymmetric prisoner’s dilemma of Section 4. However, in general it can be hard to know when $E_{PUB} \setminus E_{PRI} \neq \emptyset$ or $E_{PUB} \setminus E_{PRI}^{CT} \neq \emptyset$. 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}^{CT} 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}^{CT} is a PBE in Γ_{PRI}^{CT} in local cheap talk strategies. Denote the LCTE payoff set in Γ_{PRI}^{CT} 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 .²⁴

I now show that the condition that $E_{PUB} \setminus E_{PRI} \neq \emptyset$ (resp., $E_{PUB} \setminus E_{PRI}^{CT} \neq \emptyset$) in Corollary 1 may be replaced with something like $E_{PUB} \setminus E_{PRI}^{LPE} \neq \emptyset$ (resp., $E_{PUB} \setminus E_{PRI}^{LCTE} \neq \emptyset$). To do this, I introduce the notion of a “nice” subnetwork (intuitively, a “subtree” on which tokens can be shown to expand the equilibrium payoff set while restricting attention to LPE).

²⁴With 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 can access a public randomizing device.

For any subnetwork $M \subseteq L$, let $E|_M$ be the PBE payoff set in the game where M is the original network, or equivalently in the game where all links $\{i, j\} \notin M$ are deleted (so that $E_{PRI}|_M$ is the PBE payoff set in this game with private monitoring, $E_{PUB}|_M$ is the PBE payoff set with public monitoring, etc.). The game where M is the original network will be denoted $\Gamma|_M$. For any set X , let $\text{co}(X)$ denote the convex hull of X .

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 .²⁵²⁶*
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. *$E_{PUB}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$.*

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

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 $E_{PUB} \setminus \text{co}(E_{PRI}^{LPE}) \neq \emptyset$, then Theorem 2 says that tokens are essential.²⁷ However, Theorem 2 is much more general than this because L itself need not be nice. For example, in many games 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 games even though characterizing E_{PRI} , E_{PRI}^{CT} , and E_{PRI}^{TOK} remains intractable.

²⁵This 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 .

²⁶I slightly abuse notation here by letting M stand for both a subnetwork of L and the set of nodes in that subnetwork.

²⁷Technically, one can show that $E_{PUB} \setminus E_{PRI}^{LPE} \neq \emptyset$ is sufficient in this case.

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 they also expand the (unrestricted) equilibrium payoff set on L .²⁸

A previous version of the 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.

7 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, public information can only be replicated when the network is 2-connected if tokens are unavailable. 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.

²⁸It 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 .

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 investigate whether simpler ways of using tokens are sufficient to replicate public information (or merely expand the equilibrium payoff set) in more special environments. For example, in the asymmetric prisoner’s dilemma example of Section 4, the players can sustain cooperation simply by having player 1 give player 2 a token whenever player 2 cooperates with her, and having player 3 cooperate with player 2 only if she first receives a token from player 2—an arrangement in which a token resembles a dollar bill used to “buy” cooperation. An alternative approach is to consider limits on players’ information or “rationality” that might make simpler and more realistic ways of using tokens constrained optimal. For instance, 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 A: Examples

This appendix 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.

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 are linked and 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
			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. It may be shown that in this model 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 E_{PUB} but not 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 $(\frac{1}{4}, \frac{5}{4}, \frac{3}{2}) \in E_{PUB}$.

Now suppose toward a contradiction that $(\frac{1}{4}, \frac{5}{4}, \frac{3}{2}) \in E_{PRI}^{TOK}$. 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 always play A 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 always be (X, X) 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, $(\frac{1}{4}, \frac{5}{4}, \frac{3}{2}) \notin E_{PRI}^{TOK}$. ■

Unknown Initial Endowment of Tokens

Consider the game given by $L = \{\{1, 2\}, \{2, 3\}\}$ with $(1, 2)$ -game and $(2, 3)$ -game as in the previous example (i.e., the game is exactly as in the previous example but with the network known to be $\{\{1, 2\}, \{2, 3\}\}$). 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 $E_{PRI}^{TOK}(Y, m^0) \supseteq E_{PUB}$ 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., $E_{PUB} \setminus E_{PRI}^{TOK}(Y, \tilde{m}) \neq \emptyset$, where $E_{PRI}^{TOK}(Y, \tilde{m})$ is the PBE payoff set with stochastic endowment \tilde{m}).*

Proof. I show that payoff vector $(1, 2, 3)$ is in E_{PUB} but not $E_{PRI}^{TOK}(Y, \tilde{m})$.

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 $(1, 2, 3) \in E_{PRI}^{TOK}(Y, \tilde{m})$. Then there is a PBE in which the outcome in every period is $(A, (X, X))$, 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.²⁹ Therefore, when 2’s realized endowment is m , it is feasible for him to deviate to always playing B 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 $3 > 2$. Hence, $(1, 2, 3) \notin E_{PRI}^{TOK}(Y, \tilde{m})$. ■

Indivisible Tokens

In this subsection only, assume that players can only transfer integer quantities of tokens (i.e., tokens are *indivisible*). Consider the same game as in the previous example (but with a deterministic initial endowment). The following result shows that indivisible tokens cannot replicate public information in this example.

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

²⁹In 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 of m .

Proof. I show that with indivisible tokens payoff vector $(1, 2, 3)$ is in E_{PUB} but not E_{PRI}^{TOK} . That $(1, 2, 3) \in E_{PUB}$ was already proved in the proof of Proposition 4.

Suppose toward a contradiction that $(1, 2, 3) \in E_{PRI}^{TOK}$. 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, $(1, 2, 3) \notin E_{PRI}^{TOK}$. ■

Appendix B: Omitted Proofs

Proof of Theorem A

For the converse, let the $(1, 2)$ and $(2, 3)$ games be as in Section 4, 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 $(1, 2, 1, 0, \dots, 0) \in E_{PUB} \setminus E_{PRI}^{CT}$.

For the main part of the theorem, I first introduce one non-standard 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.³⁰

³⁰The 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

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*.³¹ I show that $E_{PRI}^{CT}(Y) \supseteq E_{PUB}$.

Let σ^{PUB} be a PBE strategy profile in game Γ_{PUB} . I construct a strategy profile σ^{PRI} in game $\Gamma_{PRI}^{CT}(Y)$ which will be shown to be a PBE 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.³² If player i received conflicting $\{i',j'\}$ reports in some period $\tau < t$, 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 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 $j \in N_i$.

“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.

³¹In contrast, $0_{i',j'}$ may be interpreted as a null report meaning “no report of $z_{i',j'}$ ”.

³²To 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 $\{i,j\} \in L$.

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 PBE profile.

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.³³ 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 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

³³In particular, player i sends an $\{i', j'\}$ report at round k if and only if $k \geq \min\{d(i, i'), d(i, j')\} + 1$.

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 PBE. 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.³⁴ 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 $u_i(\alpha^*)$, 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 A. Let L' be an arbitrary spanning tree of 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' are numbered $1, 2, \dots, n'$.³⁵ 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 $E_{PRI}^{TOK}(Y, m^0) \supseteq E_{PUB}$.

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

³⁴If 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.

³⁵The set of leaf players in L' 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 $\{i : |N'_i| = 1\}$.

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.³⁶ Specifically, 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.³⁷ 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 $\{i',j'\}$ reports for all $\{i',j'\} \in L$ in the reporting subphase), and let q_i be the number

³⁶Throughout, *transfer* means transfer of tokens.

³⁷The 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.

between 1 and $|Z|$ assigned to \hat{z}^i . 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 $l \in \{1, \dots, d(i,j) + 1\}$.³⁸ 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 $p_{1,i}^{l+1}$.
- 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 $\$(4 - 1/2^{t-1})n'|Z|$ 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 .³⁹ 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 .⁴⁰

³⁸Note that $p_{i,j} \neq p_{j,i}$. In particular, $p_{i,j}^l = p_{j,i}^{d(i,j)-l+2}$. For example, $p_{i,j}^1 = i = p_{j,i}^{d(i,j)+1}$.

³⁹Note 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 $h_i^{t'}$.

⁴⁰To 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).

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.”

- Classify history h_i^t as 0 *steps off-path* if it is on-path. Thus, play and beliefs at 0 step off-path histories have already been specified
- 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 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 $h_{j'}^t$ (and therefore punishes i if $j' \in N_i$).⁴¹ The remaining aspects of i 's beliefs about players $j' \in C_j^i$ are arbitrary. 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

⁴¹If 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*.

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 $d' \leq d$.⁴² Specify that player i 's continuation play at histories consistent with this distribution of opposing action plans is (Nash) optimal.⁴³
- Classify history h_i^t as $d + 1$ steps off-path if it is reached with positive probability following a $d + 1$ 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.

An important preliminary observation is that a leaf player i never transfers tokens at an off-path history h_i^t if she conforms to σ^{PRI} . To see this, note that at any off-path history h_i^t where i has conformed to σ^{PRI} , 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.

⁴²Note that the distribution of opposing (i, j) -action plans but not opposing strategies is specified, as we 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.

⁴³One 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 of $m_{i,j}$ and $m_{i,j} = 0$.

The following key lemma 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 $\alpha_{j,i}^*$ (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 the 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 is $\$q_1$ 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'). But

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

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 $h_1^{t_0,0}$ 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 $h_{j'}^\tau$. 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 $h_{j'}^\tau$, and hence that j will receive *alert* from j'' at history $h_{j'}^\tau$. Hence, i believes that j punishes i at history $h_{j'}^t$.

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 $h_{j'}^\tau$. 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 PBE 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 signal profile.

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 E_{PRI}^{MLPE} .

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

First, I claim that $E_{PRI} = E_{PRI}^{MLPE}$. The argument adapts the proof of Theorem 5.2 of Fudenberg and Levine (1994), which shows that the sequential equilibrium payoff set and perfect public equilibrium payoff set coincide in repeated games with imperfect public monitoring and a product structure. 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 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 $E_{PRI}^{MLPE} = E_{PRI}^{LPE}|_M + E_{PRI}|_{L \setminus M}$.⁴⁴ 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 $E_{PRI}^{MLPE} \supseteq E_{PRI}^{LPE}|_M + E_{PRI}|_{L \setminus M}$. 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 $E_{PRI}^{MLPE} \subseteq E_{PRI}^{LPE}|_M + E_{PRI}|_{L \setminus M}$. Combining the inclusions yields $E_{PRI}^{MLPE} = E_{PRI}^{LPE}|_M + E_{PRI}|_{L \setminus M}$.

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

Finally, I claim that $E_{PRI}^{TOK} \supseteq E_{PRI}^{TOK}|_M + E_{PRI}|_{L \setminus M}$. 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 $E_{PRI}^{TOK}(Y, m^0) \supseteq E_{PRI}^{TOK}(\tilde{Y}, \tilde{m}^0)|_M + E_{PRI}|_{L \setminus M}$, as given a PBE σ' in $E_{PRI}^{TOK}(\tilde{Y}, \tilde{m}^0)|_M$ and a PBE σ'' in $E_{PRI}|_{L \setminus M}$ one can construct a

⁴⁴The notation here is that for sets $A, B \subseteq \mathbb{R}^n$, $A + B = \{a + b : a \in A, b \in B\}$.

PBE σ in $E_{PRI}^{TOK}(Y, m^0)$ 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

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

where the strict inclusion uses the fact that, for any sets X , X' , and W , if $X \supseteq X'$ and $X \setminus \text{co}(X') \neq \emptyset$ then $X + W \supsetneq X' + W$ (as can be seen from a separating hyperplane argument). Therefore, $E_{PRI}^{TOK} \supsetneq E_{PRI}$.

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}^{CT} 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 E_{PRI}^{MLCTE} be the set of M -local cheap talk equilibrium payoffs in Γ_{PRI}^{CT} . Then $E_{PRI}^{CT} = E_{PRI}^{MLCTE}$ by the same argument as for $E_{PRI} = E_{PRI}^{MLPE}$, 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, $E_{PRI}^{MLCTE} = E_{PRI}^{LCTE}|_M + E_{PRI}|_{L \setminus M}$ by the same argument as for $E_{PRI}^{MLPE} = E_{PRI}^{LPE}|_M + E_{PRI}|_{L \setminus M}$, and $E_{PRI}^{TOK}|_M \supseteq E_{PRI}^{LCTE}|_M$ and $E_{PRI}^{TOK}|_M \setminus \text{co}(E_{PRI}^{LCTE}|_M) \neq \emptyset$ by the same argument as for $E_{PRI}^{TOK}|_M \supseteq E_{PRI}^{LPE}|_M$ and $E_{PRI}^{TOK}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$ (where the statement that $E_{PUB}|_M \setminus \text{co}(E_{PRI}^{LPE}|_M) \neq \emptyset$ is strengthened to $E_{PUB}|_M \setminus \text{co}(E_{PRI}^{LCTE}|_M) \neq \emptyset$, which is possible when M is truly nice). Combining these inclusions with $E_{PRI}^{TOK} \supseteq E_{PRI}^{TOK}|_M + E_{PRI}|_{L \setminus M}$ as in the “essential” case yields $E_{PRI}^{TOK} \supsetneq E_{PRI}^{CT}$.

References

- [1] Aliprantis, C.D., G. Camera, and D. Puzzello (2007), “Anonymous Markets and Monetary Trading,” *Journal of Monetary Economics*, 54, 1905-1928.

- [2] Ambrus, A., M. Möbius, and A. Szeidl (2010), "Consumption Risk-Sharing in Social Networks," mimeo.
- [3] Araujo, L. (2004), "Social Norms and Money," *Journal of Monetary Economics*, 51, 241-256.
- [4] Aumann, R.J. and S. Hart (2003), "Long Cheap Talk," *Econometrica*, 71, 1619-1660.
- [5] Ben-Porath, E. and Kahneman, M. (1996), "Communication in Repeated Games with Private Monitoring," *Journal of Economic Theory*, 70, 281-297.
- [6] 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.
- [7] Bloch, F., G. Genicot, and D. Ray (2008), "Informal Insurance in Social Networks," *Journal of Economic Theory*, 143, 36-58.
- [8] Compte, O. (1998), "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66, 597-626.
- [9] Corbae, D., T. Temzelides, and R. Wright (2003), "Directed Matching and Monetary Exchange," *Econometrica*, 71, 731-756.
- [10] Dixit, A. (2003), "Trade Expansion and Contract Enforcement," *Journal of Political Economy*, 111, 1293-1317.
- [11] Ellison, G. (1994), "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *Review of Economic Studies*, 61, 567-588.
- [12] Forges, F. (2009), "Correlated Equilibrium and Communication in Games," *Encyclopedia of Complexity and Systems Science*, edited by R. Meyers, Springer.
- [13] Fudenberg, D. and D.K. Levine (1994), "Efficiency and Observability with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62, 103-135.
- [14] Greif, A. (2006), *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*, Cambridge: Cambridge University Press.
- [15] Hauser, C. and H. Hopenhayn (2010), "Trading Favors: Optimal Exchange and Forgiveness," mimeo.
- [16] Izmalkov, S., S. Micali, and M. Lepinski (2011), "Perfect Implementation," *Games and Economic Behavior*, 71, 121-140.
- [17] 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.

- [18] Kandori, M. (1992), “Social Norms and Community Enforcement,” *Review of Economic Studies*, 59, 63-80.
- [19] Kandori, M. and H. Matsushima (1998), “Private Observation, Communication and Collusion,” *Econometrica*, 66, 627-652.
- [20] Karlan, D., M. Möbius, T. Rosenblat, and A. Szeidl (2009), “Trust and Social Collateral,” *Quarterly Journal of Economics*, 124, 1307-1361.
- [21] Kiyotaki, N. and R. Wright (1989), “On Money as a Medium of Exchange,” *Journal of Political Economy*, 97, 927-954.
- [22] Kiyotaki, N. and R. Wright (1993), “A Search-Theoretic Approach to Monetary Economics,” *American Economic Review*, 83, 63-77.
- [23] Kocherlakota, N.R. (1998), “Money is Memory,” *Journal of Economic Theory*, 81, 232-251.
- [24] Kocherlakota, N.R. (2002), “The Two-Money Theorem,” *International Economic Review*, 43, 333-346.
- [25] Kocherlakota, N. and N. Wallace (1998), “Incomplete Record-Keeping and Optimal Payment Arrangements,” *Journal of Economic Theory*, 81, 272-289.
- [26] Koessler, P. and F. Forges (2008), “Multistage Communication with and without Verifiable Types,” *International Game Theory Review*, 10, 145-164.
- [27] Kreps, D.M., and R. Wilson (1982) “Sequential Equilibria,” *Econometrica*, 50, 863-894.
- [28] Krishna, R.V. (2007), “Communication in Games of Incomplete Information,” *Mathematics of Operations Research*, 10, 117-153.
- [29] Lagos, R. and R. Wright (2008), “When is Money Essential? A Comment on Aliprantis, Camera, and Puzzello,” mimeo.
- [30] Linal, 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.
- [31] McLean, R., I. Obara, and A. Postlewaite (2012), “Informational Smallness and Private Monitoring,” mimeo.
- [32] 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.
- [33] Möbius, M.M. (2001), “Trading Favors,” mimeo.
- [34] Okuno-Fujiwara, M. and A. Postlewaite (1995), “Social Norms and Random Matching Games,” *Games and Economic Behavior*, 9, 79-109.

- [35] Ostroy, J.M. (1973), "The Informational Efficiency of Monetary Exchange," *American Economic Review*, 63, 597-610.
- [36] Ostroy, J.M. and R.M. Starr (1974), "Money and the Decentralization of Exchange," *Econometrica*, 42, 1093-1113.
- [37] Renault, J. and T. Tomala (1998), "Repeated Proximity Games," *International Journal of Game Theory*, 27, 539-559.
- [38] Starr, R.M. (1972), "The Structure of Exchange in Barter and Monetary Economies," *Quarterly Journal of Economics*, 86, 290-302.
- [39] 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.
- [40] Townsend, R. (1987), "Economic Organization with Limited Communication," *American Economic Review*, 77, 954-970.
- [41] Wallace, N. (2001), "Whither Monetary Economics?" *International Economic Review*, 42, 847-869.
- [42] Zhu, T. and E. Maenner (2012), "Pairwise Monitoring, Information Revelation, and Message Trading," mimeo.