

Cooperation and Community Responsibility:^{*}

A Folk Theorem for Repeated Matching Games with Names

(JOB MARKET PAPER)

Joyee Deb[†]

November 1, 2007

Abstract

When large communities transact with each other and players change rivals over time, players may not recognize each other or may have limited information about past play. Can players cooperate in such anonymous transactions? I analyze an infinitely repeated random matching game between two communities. Players' identities are unobservable and players only observe the outcomes of their own matches. Players may send an unverifiable message (a name) before playing each game. I show that for any such game, all feasible individually rational payoffs can be sustained in equilibrium if players are sufficiently patient. Cooperation is achieved not by the standard route of community enforcement or third-party punishments, but by "*community responsibility*". If a player deviates, her entire community is held responsible and punished by the victim.

^{*}I would like to thank my advisor, Johannes Hörner for his invaluable guidance and encouragement at every stage of this project. I also thank Eddie Dekel, Jeff Ely, Julio González-Díaz, Peter Klibanoff and Siddarth Madhav for comments and many helpful conversations.

[†]Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University;
j-deb@kellogg.northwestern.edu

1 Introduction

Would you lend a complete stranger \$10,000? How would you get your money back? Trusting people you don't know ... may sound like the height of foolishness. But a modern economy depends on exactly such impersonal exchange. Every day, people lend ... to strangers with every expectation that they'll be repaid. Vendors supply goods and services, trusting that they'll be compensated within a reasonable time. How does it all work?

From "Even Without Law, Contracts Have a Way of Being Enforced"

New York Times, October 10, 2002

Such impersonal exchange lies at the heart of this paper. In many situations communities of agents are involved in bilateral transactions with each other, and it may be reasonable to assume that agents do not recognize each other or have very limited information about each other's actions. In such situations, how does impersonal exchange take place? Can players achieve cooperative outcomes? This is the central question of this paper. Formally, I ask whether every feasible and individually rational payoff vector of a two-player game can be an equilibrium outcome in the infinitely repeated game between two communities, where players are anonymously randomly matched to one another in every period and players do not observe the complete history of past play.

It is well-known that when only two players interact repeatedly, any feasible and individually rational payoff can be sustained in equilibrium, provided players are sufficiently patient. Further, we know that the Folk Theorem extends (under appropriate conditions) to games with N players, with perfect monitoring, imperfect public monitoring and to certain games with private monitoring. Any feasible and individually rational payoff can be achieved using a mechanism of personal punishment. If a player deviates, her rival can credibly retaliate and punish her in the future. The threat of future punishment deters patient players from deviating. However, these results implicitly assume that players recognize their rivals, and so cooperation can be sustained through the threat of personalized punishments. In interactions in large communities where players meet each other infrequently and anonymously, personalized punishments are not possible. Players may change partners and may not know each other's true identities. So it is not possible for a victim to accurately punish the culprit. Can cooperation then not be sustained?

To examine this question, I consider an infinitely repeated stage-game played between two

communities. In every period, members of one community are randomly matched to members of the rival community.¹ Each player plays the stage-game with the opponent she is randomly matched to. Players cannot observe the entire pattern of play within the communities. I impose the strong informational restriction that players observe only the transactions they are personally engaged in. Further, they do not recognize each other. There is limited communication. I only allow players to introduce themselves (announce a name) before they play in each period. However, names are not verifiable, and the true identity of a player cannot be known through her announced name. Players cannot communicate in any other way within their community or communicate the identity of their past opponents. Within this setting of limited information, I examine what payoffs can be achieved in equilibrium.²

Achieving any feasible, individually rational payoff in equilibrium through only personalized punishments may be difficult as players are essentially anonymous. A form of punishment that has been used in similar settings is “community enforcement”. In community enforcement a player who deviates is punished not necessarily by the victim but by other players in the society who become aware of the deviation. For instance, in a prisoner’s dilemma (PD) game if a player faces a defection, she could punish any rival in the future by switching to defection forever. By starting to defect, she would spread the information that someone has defected. The defection action would spread (“contagion”) throughout the population, and cooperation would eventually break down. The credible threat of such a breakdown of cooperation can deter players from defecting in the first place. Earlier literature (e.g. Kandori (1992), Ellison (1994)) has shown that in a PD game, such community enforcement can be used to achieve efficiency. Why do players have the incentive to punish even when they know that they may not be matched to the original defector and may spread the contagion more quickly? In the PD, the maximum one-period gain from defecting is the same as the one-period loss from not defecting to slow down the contagion. Ellison (1994) establishes that the loss from starting the contagion is greater than the gain from slowing it down once it has started, even without any kind of communication. Consequently, it is possible that players fear a breakdown of cooperation enough that they will not deviate first, but do not fear so much that they are unwilling to spread the contagion once it has begun. However, in general, games may not have this feature and this contagious community

¹The assumption of two communities is not necessary. The results of this paper continue to hold if there just is one community of agents playing the repeated anonymous random matching game. See Section 3.5 for more on this.

²The result would extend to environments in which more information can be transmitted.

enforcement does not work. So far, little is known on how to attain cooperation in this setting with any game other than the PD.³

The main result I obtain is a possibility result - a Folk Theorem for infinitely repeated random matching games - which states that for *any* two-player game played between two communities, it is possible to sustain all feasible individually rational payoffs in a sequential equilibrium, provided players are sufficiently patient and can announce names. I establish the result constructively by identifying strategies that can attain any given feasible, individually rational payoff.⁴

Interestingly, in this paper, cooperation is sustained neither by personalized punishments nor by community enforcement. It is no longer the case that a deviator is punished by third-parties in her victim's community. On the contrary, if a player cheats she is punished only by her victim, but her entire community is held responsible for the deviation and everyone in her community is punished by her victim. I call this form of punishment "*community responsibility*".

The terminology is inspired by the community responsibility system, an institution prevalent in medieval Europe (Greif (2006)). Under the community responsibility system, if a member of any community defaulted or cheated, all members of her community were held legally liable for the default. The property of any member of the defaulter's community could be confiscated. The system internalized the cost of a default by each of their members on other members.⁵ This paper does not involve any exogenous enforcement institutions, but the equilibrium strategies used turn out to have a similar flavor in the sense that if a member of a community deviates in any transaction, the victim holds the entire community of the deviator liable for the deviation.

Further, in describing the utility of the community responsibility system in Europe, Greif (2006) makes the following observation:

Communal liability ... supported intercommunity impersonal exchange. Exchange did not require that the interacting merchants have knowledge about past conduct,

³In this paper, when I refer to sustaining cooperation or cooperative outcomes, I refer to any feasible and individually rational payoff that is not a static Nash equilibrium outcome.

⁴In establishing the main result, I focus on achieving identical payoffs within a community. See Remark 2 in Section 3.3 for a discussion of how this can be generalized.

⁵See Greif (2006). "Historical evidence ... supports the claim that the community responsibility system prevailed throughout Europe. ... In a charter granted to London in the early 1130s, King Henry I announced that 'all debtors to the citizens of London discharge these debts ... and if they refuse to pay ... then the citizens to whom the debts are due may take pledges either from the borough or from the village ... in which the debtor lives!'"

share expectations about trading in the future, have the ability to transmit information about a merchant's conduct to future trading partners, or know a priori the personal identity of each other.

This also captures the essence of why in the framework of this paper, community responsibility works in games where community enforcement does not. It is important to note that information transmission is critical to sustaining cooperation through community enforcement. Since dishonesty needs to be punished by third-parties who were not a part of the original dishonest transaction, information transmission is required to start the punishment. In Ellison (1994), when a player starts deviating after observing a deviation, she transmits the information to a third-party that a deviation has occurred. In a prisoner's dilemma, information about the deviation can be transmitted indirectly through the act of defection. For other games, transmission of some verifiable ("hard") information seems to be necessary. Kandori (1992) introduces local information processing, where information about past deviations is transmitted through labels which depend on a player's history of play. Takahashi (2007) allows verifiable first-order information. In this paper I obtain a Folk Theorem for general games without introducing any hard information in the model. Players are allowed to announce names, but names are not verifiable. There is no hard information. Community responsibility does not require third-parties to carry out punishments and consequently can be implemented even with these strong informational restrictions and little transmission of information.

What is community responsibility in the context of this paper and how does it work? Consider two communities with M players each. Players from the two communities are randomly matched into pairs to play the stage-game. We allow players to announce their names in each period before they play the stage-game. Think of each player as playing separate but identical games, one with each of the M names of her rival community. A player treats her interactions with each rival name separately and conditions play against any name on the history of play with that name. Play with each name proceeds in blocks of length T (i.e. T interactions). Players keep track of the stage of a block they are in separately for each possible name. Each player plays one of two strategies of the T -fold repeated stage-game within a block. At the beginning of each block, each player is indifferent between the two strategies, but one of the strategies ensures a low payoff for her opponent and the other a high payoff. So, in the first period of any block (called the plan period), each player mixes between the two strategies in a way to ensure that her opponent gets the target equilibrium payoff. The realized action in the

plan period serves as a coordination device and indicates the plan of play within that block. If a player plays certain actions, she is said to send a “good” (bad) plan and play in that block proceeds according to the strategy that is favorable (unfavorable) for the opponent. At the start of the next block, each player tailors her rival’s continuation payoff based on the actions played in the last block, by appropriately choosing the probability with which she sticks to or changes her chosen strategy. Players can control the continuation payoffs of their rivals by appropriately mixing between two strategies, irrespective of what their rival plays.

Since each player conditions play on the name announcement she hears, players may have incentives to misreport names. We construct strategies to prevent misreporting. For any pair of players, the second interaction in any block of length T is designated as the “signature period” and members of a pair play actions that serve as their “signatures”. The signatures for any pair of players are different pure actions based on the action realized in their first interaction. No player outside a pair can observe the action realized in their first interaction, and so no one can know what the appropriate signature action is. So, if a player outside a pair impersonates one of the players in this pair, she can end up playing the wrong signature in case it is a signature period, and her impersonation will get detected.

If a player observes an incorrect signature in a signature period, she knows that a deviation has occurred, though the identity of the deviator is unknown. She holds the deviator’s entire community responsible and punishes them all by switching to the bad plan with each of her rivals in their next plan period. Since every player is indifferent between her two strategies at the start of any block, she can switch to a strategy that is bad for all her opponents without affecting her own payoff in any way. Notice that punishments are carried out by the victim herself and not by third parties. However, innocent members of the deviator’s community are held responsible and punished. Indifference at the start of each block makes community responsibility a credible threat. For sufficiently patient players, this threat can deter impersonations and deviations from the equilibrium path.

At this point, it is worthwhile to ask the following question. If community responsibility prescribes punishing everybody in the community after a deviation and does not condition punishments on names, then why do we need names? It turns out that though the names are not used to personalize punishments off the equilibrium path, players need to use the names on the equilibrium path to tailor the continuation payoffs of their rivals.

In an extension, I show that the Folk Theorem extends to K -player games being played by

$K > 2$ communities, where players from each community are randomly matched in each period to form K -player groups to play the stage-game. The same idea of community responsibility is used to attain cooperation. Each community acts as the monitor of one other community. Say community 1 is the monitor of community 2. If any player in community 2 deviates, the player from community 1 whom she meets in that period holds the whole of community 2 responsible and punishes all members of community 2 that she meets in the future.

1.1 Related Literature

This paper is related to two independent streams of literature. First, it contributes to the literature that asks similar questions about cooperation and impersonal exchange. These questions have been asked earlier for the prisoner's dilemma in the framework of repeated random matching games. Second, this paper is related to recent literature on repeated games with imperfect private monitoring, because of methodological similarities.

1.1.1 Connection with Community Enforcement

Kandori (1992) is one of the early papers that studies community enforcement. Kandori studies the repeated prisoner's dilemma with anonymous random matching. Players only see the transactions they are personally involved in and there is no other form of information transmission. It is shown that if the loss from being cheated is large enough, and if players are sufficiently patient, efficiency can be achieved. Efficiency is achieved through contagion. If any player faces defection, she indiscriminately defects against any other player she meets in the future. Defection spreads like an epidemic and cooperation breaks down. The threat of such a breakdown in cooperation helps sustain cooperation on the equilibrium path.

Kandori (1992) also considers games beyond the prisoner's dilemma, but requires significantly more information in the model. He assumes the existence of a mechanism that assigns labels to players based on their history of play. Players who have deviated or seen a deviation can be distinguished from those who have not, by their labels. These labels naturally enable transmission of information and cooperation can be sustained through community enforcement.

Ellison (1994) generalizes Kandori's first result and shows that cooperation is possible in equilibrium for *any* prisoner's dilemma game using contagious strategies, with no information transmission. Contagious strategies however critically depend on the specific structure of the prisoner's dilemma (PD) - in particular on symmetry and on the existence of a Nash equilibrium

in strictly dominant strategies. As mentioned earlier, in the PD, the maximum short-term gain from deviating on the equilibrium path is the same as the short-term loss from not punishing when a player successfully slows down the contagion. Ellison (1994) proves that the future loss from deviating on equilibrium path is greater than the future gain from slowing down the contagion once it has started. Consequently, it is possible to make the short-term loss / gain from following equilibrium strategies to lie between the two future effects. This argument does not apply to a general game.⁶

In a recent paper, Takahashi (2007) again considers the repeated prisoner’s dilemma game with random matching but with a continuum of agents. Cooperation is sustained through community enforcement but by allowing first-order information. Players have available to them the complete history of past actions of their partner in each period.

A stark gap evident in the current literature is that little is known about games other than the prisoner’s dilemma. This paper tries to fill this gap. I consider general two-player games being played by communities of agents who are randomly matched in each period. I obtain the Folk Theorem for general games under the mild informational assumption that players can announce names, though announcements are not verifiable.

As discussed earlier, this paper is different from earlier work in that it does not use community enforcement, but introduces the alternate route of community responsibility. Since community enforcement involves third-party punishments, it requires information transmission. On the contrary, community responsibility requires only a victim (a player who directly observes a deviation) to punish. Consequently, community responsibility requires less information transmission and we can achieve the Folk Theorem without addition of any hard information.

This paper also goes beyond the current literature in considering repeated random matching games with more than two players. In an extension of the main model, the Folk Theorem is shown to also hold for K -player games played by K communities.

1.1.2 Connection with Repeated Games with Imperfect Private Monitoring

While this paper’s contribution is substantively related to community enforcement, the methodological content is closely related to recent advances in repeated games with imperfect private

⁶As Ellison (1994) points out, in general games this argument shows that a symmetric strategy profile (a, a) is an equilibrium outcome if the payoff from (a, a) strictly dominates the payoff of a static Nash equilibrium (s^*, s^*) and s^* is a best response to a (e.g. in games with a dominant strategy equilibrium).

monitoring.

Community responsibility depends on the fact that the player who detects a deviation is willing to punish the deviator's entire community. However, the detector may not have an incentive to punish if the punishment action either involves a short-term cost, or alters her own continuation payoff adversely. In my equilibrium construction, punishing is not costly in either of these ways. When a player has to punish, she is indifferent (in that period) between punishing and not punishing. Further, a player starts punishing only in periods when she is indifferent and is supposed to mix between all her actions on the equilibrium path. So even when a player punishes, her rival cannot know if her action is a punishment or not. So, the punishment action cannot change her continuation payoff. This indifference is important to the construction and is reminiscent of the equilibrium strategies used in the literature on imperfect private monitoring.

Ely-Hörner-Olszewski (2005) study belief-free equilibria in repeated games with imperfect private monitoring. The strategies have the special feature that in infinitely many periods, each player is indifferent between several actions that she can play. But her actions give different continuation payoffs to her opponent - some ensure a high payoff and others a low payoff. In equilibrium, each player chooses actions (mixes) based on her opponent's past play. She can choose an action that is favorable to reward her opponent or an unfavorable one to punish her.

Hörner-Olszewski (2006) generalize this idea with "*block strategies*" that are characterized by periodic indifference. Play proceeds in blocks of say T periods each. In each block of T periods, players use one of two strategies of the T -fold repeated game. The length T is chosen so that the average payoff of the four resulting strategy profiles surrounds the target payoff vector. For any player, one of the two strategies guarantees her opponent a continuation payoff higher than the target payoff, and the other guarantees her opponent a payoff strictly lower than the target payoff. So players are not indifferent over their opponent's choice of strategies. However, players can be made indifferent over their own two strategies at the start of each block, by appropriately choosing the probability of using these strategies in each block as a function of the play in the most recently elapsed block. In fact players are indifferent between these two strategies and weakly prefer them to all others. The target payoff is achieved by suitably choosing the probability with which each strategy is used in the initial block of the game.

In this paper, I build on the block strategies of Hörner-Olszewski (2006). As mentioned above, the block structure provides each player with infinitely many periods of indifference, which make the threat of punishments credible. However, in the random matching setting of

this paper, players need to use block strategies separately with each possible opponent they can be matched to. They have to track what stage of a block they are in separately for each rival. This is precisely where the names are used - to track the games with each opponent separately.

Random matching also poses other difficulties - the duration of a block (in calendar time) between any pair of players is now random. It is not clear that it is possible for a player to adjust her rival's continuation payoff by just appropriately mixing her two strategies in a block, as a function of the play in the most recently elapsed block. If a block takes a very large number of calendar time periods, the required adjustment in payoff may not be feasible. I show that it is actually possible to adjust continuation payoffs in a way that, in expectation, at every stage of a block, players are indifferent between their two strategies and prefer them to all others.

A novel feature of the construction in this paper is that it is possible to convert unverifiable information into hard information. The signature periods discussed above play exactly this role. Even though messages (names) are unverifiable, the signatures provide players a means to ensure that no one has an incentive to misreport their names - effectively converting the soft messages into hard information. Further, signatures enable this verification without enriching players' communication possibilities, but just through the actions available to players in the underlying game. This poses challenges as playing the right signature action has potential payoff consequences in the short-term, and continuation payoffs have to be specified to satisfy intertemporal incentives of players.

The rest of the paper is organized as follows. Section 2 presents the model. In Section 3, I establish the Folk Theorem and discuss its key features. In Section 4, I extend the result to $K > 2$ communities and multilateral matching. Section 5 concludes. Proofs are contained in the appendix.

2 Model and Notation

Players: The game is played by two communities of players. Each community I , $I \in \{1, 2\}$ comprises $M > 2$ players⁷, say $I := \{I_1, \dots, I_M\}$. To save notation, I will often denote a generic element of any community of players I by i .

Random Matching and Timing of Game: In each period $t \in \{1, 2, \dots\}$, players are randomly matched into pairs with each member l of Community 1 facing a member $l' := m_t(l)$

⁷See Section 3.4 for the case $M = 2$.

of Community 2. The matches are made independently and uniformly over time, i.e. for all histories, for all l, l' , $\Pr[l' = m_t(l)] = \frac{1}{M}$.⁸ After being matched, each member of a pair simultaneously announces a message (“*her name*”). Then, they play a two-player finite stage-game. The timing of the game is represented in Figure 1.




Figure 1: Timing of Events

Message Sets: Each community I has a set of messages $\mathcal{N}_I, I \in \{1, 2\}$. Let \mathcal{N}_I be the set of names of players in community I (i.e. $\mathcal{N}_I = \{I_1, \dots, I_M\}$).⁹ For any pair of matched players, the pair of announced messages (names) is denoted by $\nu \in \mathcal{N} := \mathcal{N}_1 \times \mathcal{N}_2$. For any I , let $\Delta(\mathcal{N}_I)$ denote the set of mixtures over messages in \mathcal{N}_I . Messages are not verifiable, in the sense that a player cannot verify if her rival is actually announcing her name. So, the true identity of a player cannot be known from her announced name. “*Truthful reporting*” by any player i means that player i announces name i . Any other announcement by player i is called “*misreporting*” or “*impersonating*”.

Stage-Game: The stage-game Γ has finite action sets $A_I, I \in \{1, 2\}$. Denote an action profile by $a \in A := A_1 \times A_2$. For each I , let $\Delta(A_I), I \in \{1, 2\}$ denote the set of mixtures of actions in A_I . Stage-game payoffs are given by a function $u : A \rightarrow \mathbb{R}^2$. Define \mathcal{F} to be the convex hull of the payoff profiles that can be achieved by pure action profiles in the stage-game. Formally, $\mathcal{F} := \text{conv}(\{(u(a) : a \in A)\})$. Let v_i^* denote the mixed action minmax value for any player i .

⁸Unlike in earlier literature, the result does not depend on the matching being uniform or independent over time. See Remark 1 in Section 3.3, for a discussion on how this assumption can be relaxed.

⁹An implicit assumption is that the sets of messages \mathcal{N}_I contain at least M distinct messages each. For instance, we can allow players to be silent by interpreting some message as silence. In the exposition, I use exactly M messages as this is the coarsest information that suffices. Also, M messages have the reasonable interpretation of player names.

For $i \in I$, $v_i^* := \min_{\alpha_{-i} \in \Delta(A_{-I})} \max_{a_i \in A_I} u_i(a_i, \alpha_{-i})$. Let \mathcal{F}^* denote the individually rational and feasible payoff set, i.e. $\mathcal{F}^* := \{v \in \mathcal{F} : v_i > v_i^* \forall i\}$. We consider games where \mathcal{F}^* has non-empty interior ($\text{Int } \mathcal{F}^* \neq \emptyset$).¹⁰ Let $\gamma := \max_{i, a, a'} \{|u_i(a) - u_i(a')|\}$.

All players have a common discount factor $\delta \in (0, 1)$. No public randomization device is assumed. All primitives of the model are common knowledge.

Information Assumptions: Players can observe only the transactions they are personally engaged in, i.e. each player knows the names that she encountered in the past and the action profiles played with each of these names. Since names are not verifiable, she does not know the true identity of the players she meets. She does not know what the other realized matches are and does not observe play between other pairs of players.

Histories, Strategies and Payoffs: We define histories and strategies as follows.

Definition 1 A *complete private t -period history* for a player i is given by $h_i^t := \{(\nu^1, a^1), \dots, (\nu^t, a^t)\}$, where (ν^τ, a^τ) , $\tau \in \{1, \dots, t\}$ represent the name profile and action profile observed by player i in period τ . The set of complete private t -period histories is given by $\mathcal{H}_i^t := (\mathcal{N} \times A)^t$. The set of all possible complete private histories for player i is $\mathcal{H}_i := \bigcup_{t=0}^{\infty} \mathcal{H}_i^t$ ($\mathcal{H}_i^0 := \emptyset$).

Definition 2 An *interim private t -period history* for player i is given by $k_i^t := \{(\nu^1, a^1), \dots, (\nu^{t-1}, a^{t-1}), \nu^t\}$ where ν^τ and a^τ , $\tau \in \{1, \dots, t\}$ represent respectively the name profile and action profile observed by player i in period τ . The set of interim private t -period histories is given by $\mathcal{K}_i^t := \mathcal{H}_i^{t-1} \times \mathcal{N}$. The set of all possible interim private histories for player i is $\mathcal{K}_i := \bigcup_{t=1}^{\infty} \mathcal{K}_i^t$.

Definition 3 A *strategy* for a player i in community $I \in \{1, 2\}$ is a mapping σ_i such that,

$$\text{for any } i \in I, \sigma_i : \mathcal{H}_i \cup \mathcal{K}_i \rightarrow \Delta(\mathcal{N}_I) \cup \Delta(A_I) \text{ such that } \begin{cases} \sigma_i(x) \in \Delta(\mathcal{N}_I) & \text{if } x \in \mathcal{H}_i, \\ \sigma_i(x) \in \Delta(A_I) & \text{if } x \in \mathcal{K}_i. \end{cases}$$

Σ_i is the set of i 's strategies. A strategy profile σ specifies strategies for all players (i.e. $\sigma \in \times_i \Sigma_i$).

¹⁰Observe that this restriction is not required in standard Folk Theorems for two-player games (e.g. Fudenberg and Maskin (1986)). It is however used in the literature on imperfect private monitoring (See Hörner-Olszewski (2006)). Note also that this restriction implies that $|A_i| \geq 2 \forall i$.

In some abuse of notation, for $k_i \in \mathcal{K}_i$ and $h_i \in \mathcal{H}_i$ we let $\sigma_i(a_i|k_i)$ and $\sigma_i(\nu_i|h_i)$ denote the probability with which i plays a_i and ν_i conditional on history k_i and h_i respectively, if she is using strategy σ_i . We denote equilibrium strategies by σ^* .

A player's payoff from a given strategy profile σ in the infinitely repeated random matching game is denoted by $U_i(\sigma)$. It is the normalized sum of discounted payoffs from the stage-games that the player plays in each period, i.e. $U_i(\sigma) := (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a_i^t, a_{-i}^t)$.

Beliefs: Given any strategy profile σ , after any private history, we can compute the beliefs that each player has over all the possible histories that are consistent with her observed private history. Denote such a system of beliefs by ξ .

Definition 4 *A strategy profile σ together with an associated system of beliefs ξ is said to be an **assessment**. The set of all assessments is denoted by Ψ .*

Solution Concept: The solution concept used here is sequential equilibrium. While sequential equilibrium (Kreps & Wilson (1982)) is formally defined for finite extensive form games, the notion can be extended naturally to this setting. Let Σ^0 denote the set of totally mixed strategies, i.e. $\Sigma^0 := \{\sigma : \forall i, \forall k_i \in \mathcal{K}_i, \forall a_i, \sigma_i(a_i|k_i) > 0 \text{ and } \forall i, \forall h_i \in \mathcal{H}_i, \forall \nu_i, \sigma_i(\nu_i|h_i) > 0\}$. In other words, strategy profiles in Σ^0 specify that in every period, players announce all the names with a strictly positive probability and play all feasible actions with strictly positive probability. If strategies belong to Σ^0 all possible histories are reached with positive probability. Players' beliefs can be computed using Bayes' Rule at all histories. Let Ψ^0 denote the set of all assessments (σ, ξ) such that $\sigma \in \Sigma^0$ and ξ is derived from σ using Bayes' Rule. We define sequential equilibrium in the following way.

Definition 5 *An assessment (σ^*, ξ^*) is said to constitute a **sequential equilibrium** if the assessment is*

(i) **sequentially rational**,

$$\begin{aligned} \forall i, \forall t, \forall h_i^t \in \mathcal{H}_i^t, \forall \sigma'_i, & \quad U_i(\sigma^*|h_i^t, \xi_i^*[h_i^t]) \geq U_i(\sigma'_i, \sigma_{-i}^*|h_i^t, \xi_i^*[h_i^t]), \\ \forall i, \forall t, \forall k_i^t \in \mathcal{K}_i^t, \forall \sigma'_i, & \quad U_i(\sigma^*|k_i^t, \xi_i^*[k_i^t]) \geq U_i(\sigma'_i, \sigma_{-i}^*|k_i^t, \xi_i^*[k_i^t]), \end{aligned}$$

and

(ii) **consistent** in the sense that there exists a sequence of assessments $\{\sigma^n, \xi^n\} \in \Psi^0$ such that for every player, and every interim and complete private history, the sequence converges to (σ^*, ξ^*) uniformly in t .

Later, we use the T -fold finitely repeated stage-game as well. To avoid confusing T -period strategies with the supergame strategies, we define the following.

Definition 6 *Consider the T -fold finitely repeated stage-game (ignoring the round of name announcements). Define an **action plan** to be a strategy of this finitely repeated game in the standard sense. Denote the set of all action plans by S_i^T .*

3 The Main Result

Theorem 1 *(Folk Theorem for Random Matching Games) Consider a finite two-player game and any $(v_1, v_2) \in \text{Int } \mathcal{F}^*$. There exists a sequential equilibrium that achieves payoffs (v_1, v_2) in the corresponding infinitely repeated random matching game with names with $2M$ players, if players are sufficiently patient.*

Before formally constructing the equilibrium, I first describe the overall structure.

3.1 Structure of Equilibrium

Each player plays M different but identical games, one with each of the M names in the rival community. Players report their names truthfully. So, on the equilibrium path, players really play separate games with each of the M possible opponents. They condition their play against any opponent only on their history of play against the same name.

3.1.1 T -period Blocks

Let $(v_1, v_2) \in \text{Int } \mathcal{F}^*$ be the target payoff profile. We will choose an appropriate positive integer T . Play between members of any pair of names then proceeds in blocks of T periods in which they meet. (Note that a block of length T for any pair of players comprises T interactions between them, and so typically takes more than T periods in calendar time.) In any block of T interactions, players use one of two action plans of the T -fold finitely repeated game. One of the action plans used by a player i ensures that rival name $-i$ cannot get on average more than v_{-i} , independently of what player $-i$ plays. The other action plan ensures that rival name $-i$ gets on average at least v_{-i} . In the initial period of a block (henceforth called “*plan period*”), each player randomizes between these two action plans in a way that ensures that the target payoff of her rival name is achieved in expectation. At the end of the block, by suitably choosing the

probability of sticking to or changing her action plan, each player tailors her rival's continuation payoff based on play in the last block. Conditional on truthful reporting of names, the form of strategies described above will be shown to constitute an equilibrium.

To ensure that players announce names truthfully in equilibrium, we need a device that enables players to detect impersonations and provides incentives to a detector to punish them.

3.1.2 Detecting Impersonations

I use a device called signatures to detect impersonations. In this paper, "detection" of an impersonation means that if a player impersonates, then with positive probability a player in the rival community will become aware in the current period or in the future that some deviation from equilibrium has occurred. This kind of detection along with appropriate incentives for the detector to punish impersonations will enable cooperation.

Every pair of players designates their second interaction in each block as the "*signature period*" and in this interaction, members of a pair play actions that serve as their "*signatures*". The signature depends on the action profile realized in the plan period of that block. Players use different pure actions depending on what action profile was realized in the plan period. No player outside the pair can observe the realized action in the plan period. Consequently, no one outside a pair knows what the correct signature for that pair is. When a player impersonates someone, her announced name could be in a signature period with the rival she is matched to. In this case, with positive probability the impersonator can play the wrong signature, and so get detected. When her rival observes the wrong signature, she knows that play is not on equilibrium path.

3.1.3 Community Responsibility

Now, if a player observes an incorrect signature in a signature period with any rival, she knows that someone has deviated. The nature of the deviation or the identity of the deviator is unknown - it is possible that her current rival reported her name truthfully but played the wrong signature or that she met an impersonator now or previously. She holds all the members of her rival community responsible for the deviation, and punishes them by switching to the bad action plan (with arbitrarily high probability) with each of her rivals in their next plan period. Note that she is indifferent between her two action plans at the start of any block. But the continuation payoffs her rivals get are different for these two action plans, with one plan

being strictly better than the other for her rivals. Consequently, she can punish the entire rival community without affecting her own payoff adversely.

3.2 Construction of Equilibrium Strategies

Consider any payoff profile $(v_1, v_2) \in \text{Int } \mathcal{F}^*$. Pick payoff profiles $w^{GG}, w^{GB}, w^{BG}, w^{BB}$ such that the following conditions hold.

1. $w_i^{GG} > v_i > w_i^{BB} \forall i \in \{1, 2\}$.
2. $w_1^{GB} > v_1 > w_1^{BG}$.
3. $w_2^{BG} > v_2 > w_2^{GB}$.

These inequalities imply that there exists \underline{v}_i and \bar{v}_i with $v_i^* < \underline{v}_i < v_i < \bar{v}_i$ such that the rectangle $[\underline{v}_1, \bar{v}_1] \times [\underline{v}_2, \bar{v}_2]$ is completely contained in the interior of $\text{conv}(\{w^{GG}, w^{GB}, w^{BG}, w^{BB}\})$ and further $\bar{v}_1 < \min\{w_1^{GG}, w_1^{GB}\}$, $\bar{v}_2 < \min\{w_2^{GG}, w_2^{BG}\}$, $\underline{v}_1 > \max\{w_1^{BB}, w_1^{BG}\}$ and $\underline{v}_2 > \max\{w_2^{BB}, w_2^{GB}\}$. See Figure 2 below for a pictorial representation.




Figure 2: Payoff Profiles

Clearly, there may not exist pure action profiles whose payoffs satisfy these relationships, but there exist correlated actions that achieve exactly these payoffs $w^{GG}, w^{GB}, w^{BG}, w^{BB}$. We can approximate these correlated actions using long enough sequences of different pure action profiles. In fact, we can find finite sequences of action profiles $\{a_1^{GG}, \dots, a_N^{GG}\}$, $\{a_1^{GB}, \dots, a_N^{GB}\}$, $\{a_1^{BG}, \dots, a_N^{BG}\}$, $\{a_1^{BB}, \dots, a_N^{BB}\}$ such that each vector w^{XY} , the average discounted payoff vector over the sequence $\{a_1^{XY}, \dots, a_N^{XY}\}$ satisfies the above relationships if δ is large enough.

Further, we can find $\epsilon \in (0, 1)$ small so that $v_i^* < (1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i < v_i < (1 - \epsilon)\bar{v}_i + \epsilon\underline{v}_i$. In the equilibrium construction that follows, when I refer to an action profile a^{XY} , I actually refer to the finite sequence of action profiles $\{a_1^{XY}, \dots, a_N^{XY}\}$ described above.

3.2.1 Defining Strategies at Complete Histories: Name Announcements

At any complete private history, players announce their names truthfully.

$$\forall i, \forall t, \forall h_i^t \in \mathcal{H}_i^t, \quad \sigma_i^*[h_i^t] = i.$$

3.2.2 Defining Strategies at Interim Histories: Actions

Partitioning of Histories: Now think of each player playing M separate games, one against each rival. Since players truthfully report names in equilibrium, players can condition play on the announced name.

Definition 7 *A pairwise game denoted by $\Gamma_{i,-i}$ is the “game” player i plays against name $-i$. Player i ’s private history of length t in this pairwise game is denoted by $\hat{h}_{i,-i}^t$ and comprises the last t interactions in the supergame for player i in which she faced name $-i$.*

Now, at any interim private history of the supergame, each player i partitions her history into M separate pairwise histories $\hat{h}_{i,-i}^t, -i \in \{1, \dots, M\}$ corresponding to each of her pairwise games $\Gamma_{i,-i}$. If her current rival name is j , she plays game $\Gamma_{i,j}$, i.e. for any interim history $k_i^t = \{(\nu^1, a^1), \dots, (\nu^{t-1}, a^{t-1}), \nu^t\}$, if $\nu_{-i}^t = j$, player i plays her pairwise game $\Gamma_{i,j}$.

Since equilibrium strategies prescribe truthful name announcement, a description of how $\Gamma_{i,-i}$ is played will complete the specification of strategies on the equilibrium path for the supergame.

Play of Game $\Gamma_{i,-i}$:

For ease of exposition, fix player i and a rival name $-i$. Play is specified in an identical manner for each rival name. For the rest of the section (since rival name $-i$ is fixed), to save on notation I denote player i ’s private histories $\hat{h}_{i,-i}^t$ in the pairwise game $\Gamma_{i,-i}$ by \hat{h}_i^t . Recall that a t -period history denoted by \hat{h}_i^t specifies the action profiles played in the last t periods of this game $\Gamma_{i,-i}$, and not in the last t calendar time periods.¹¹ Since in equilibrium, any history \hat{h}_i^t of $\Gamma_{i,-i}$ has

¹¹A period in $\Gamma_{i,-i}$ is really an interaction between player i and name $-i$. So, when I refer to $\Gamma_{i,-i}$, I use “interaction” and “period” interchangeably.

the same name profile in each period, we ignore the names while specifying how $\Gamma_{i,-i}$ is played on the equilibrium path.

The pairwise game $\Gamma_{i,-i}$ proceeds in blocks of T periods (Later we define T).

In the first period of every block (plan period), the action profile used by players i and $-i$ serves as a coordination device to determine play for the rest of the block. Partition the set of i 's actions into two non-empty subsets G_i and B_i . Let $\Delta(G_i)$ and $\Delta(B_i)$ denote the set of mixtures of actions in G_i and B_i respectively. If player i chooses an action from set G_i , she is said to send plan $P_i = G$. Otherwise she is said to send plan $P_i = B$.

Further, choose any four pure action profiles $g, b, x, y \in A$ such that $g_i \neq b_i \forall i \in \{1, 2\}$. Define a function $\psi : A \rightarrow \{g, b, x, y\}$ (the signatures) mapping one-period histories (or a pair of plans) to one of the action profiles as follows.

$$\psi(a) = \begin{cases} g & \text{if } a \in G_1 \times G_2, \\ b & \text{if } a \in B_1 \times B_2, \\ x & \text{if } a \in G_1 \times B_2, \\ y & \text{if } a \in B_1 \times G_2. \end{cases}$$

Suppose the observed plans are (P_1, P_2) . Define a set of action plans as follows.

$$\mathcal{S}_i := \left\{ s_i \in S_i^T : \forall \hat{h}_i^t = \left(a, \psi(a), (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), \dots, (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}) \right), a \in P_i \times G, \right. \\ \left. s_i[\hat{h}_i^1] = \psi_i([\hat{h}_i^1]) \text{ and } s_i[\hat{h}_i^t] = a_i^{P_2, P_1}, t \geq 2 \right\}.$$

Note that the set of action plans in \mathcal{S}_i restricts player i 's actions if her rival announced plan G . In particular, action plans in \mathcal{S}_i prescribe that player i use the correct signature and play $a_i^{P_2, P_1}$ if the announced plans were (P_1, P_2) . \mathcal{S}_i does not restrict the plan that player i can announce in the plan period or her play if her rival announced a B plan or her play after any deviations.

In equilibrium, in any T -period block of a pairwise game, players will choose action plans from \mathcal{S}_i . Players will use in fact one of two actions plans from \mathcal{S}_i , a favorable one which I denote by s_i^G and an unfavorable one which I denote by s_i^B . These are defined below.

Define partially a favorable action plan s_i^G such that

$$s_i^G[\emptyset] \in \Delta(G_i), \\ s_i^G[\hat{h}_i^1] = \psi_i([\hat{h}_i^1]), \text{ and} \\ \forall \hat{h}_i^t = \left(a, \psi(a), (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), \dots, (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}) \right), a \in P_i \times P_{-i}, t \geq 1, s_i^G[\hat{h}_i^t] = a_i^{P_2, P_1}.$$

Similarly, partially define an unfavorable action plan s_i^B such that

$$s_i^B[\emptyset] \in \Delta(B_i),$$

$$s_i^B[\hat{h}_i^1] = \psi_i([\hat{h}_i^1]),$$

$$\forall \hat{h}_i^t = \left(a, \psi(a), (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), \dots, (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}) \right), a \in P_i \times P_{-i}, t \geq 1, s_i^B[\hat{h}_i^t] = a_i^{P_2, P_1},$$

$$\forall t \geq r > 1, \text{ if } \hat{h}_i^r = \left(a, \psi(a), (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), \dots, (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), (a_i^{P_1, P_2}, a'_{-i}) \right), a \in P_i \times P_{-i}, a'_{-i} \neq a_{-i}^{P_2, P_1},$$

, then $s_i^B[\hat{h}_i^t] = \alpha_i^*$, and

$$\forall t > 2, \text{ if } \hat{h}_i^2 = (a, (\psi_i(a), a'_{-i})), a \in P_i \times P_{-i}, a'_{-i} \neq \psi_{-i}(a), \text{ then } s_i^B[\hat{h}_i^t] = \alpha_i^*.$$

Note that both action plans s_i^G and s_i^B belong to \mathcal{S}_i . s_i^G is an action plan in \mathcal{S}_i that prescribes sending a G plan at the start of a block. s_i^B prescribes sending plan B at the start of a block and minmaxing when i 's rival is the first to deviate from the plan proposed in the plan period. For any history not included in the definitions of s_i^G and s_i^B above, prescribe the actions arbitrarily.

Why do we call s_i^G favorable and s_i^B unfavorable? Suppose player 1 uses action plan s_1^G , her rival, player 2 gets a payoff strictly higher than \bar{v}_2 in each period, except possibly in the first two periods. This is because as long as player 1 plays s_1^G , the payoff to player 2 that is realized in any period except the first two is approximately w_2^{BG} or w_2^{GG} both of which are higher than \bar{v}_2 . Further, if player 1 plays s_1^B , player 2 gets a payoff strictly lower than \underline{v}_2 in all except at most two periods. In the plan period and in the first period where player 2 decides to deviate, she can potentially get a higher payoff. In all other periods, she receives w_2^{GB} , w_2^{BB} or v_2^* , all of which are strictly lower than \underline{v}_2 .

It is therefore possible to choose T large enough so that for some $\underline{\delta} < 1$, for all $\delta > \underline{\delta}$, i 's average payoff within the T -period block from any action plan $s_i \in \mathcal{S}_i$ against s_{-i}^G strictly exceeds \bar{v}_1 and her average payoff from using any action plan $s_i \in \mathcal{S}_i^T$ against s_{-i}^B is strictly below \underline{v}_1 . Assume from here on that $\delta > \underline{\delta}$.

Finally, I define two benchmark action plans which are used later to compute continuation payoffs for every possible history within a block. Define $r_i^G \in \mathcal{S}_i$ to be an action plan such that given any history \hat{h}_i^t , $r_i^G | \hat{h}_i^t$ gives the lowest payoffs against s_{-i}^G among all action plans in \mathcal{S}_i . Define $r_i^B \in \mathcal{S}_i^T$ to be an action plan such that given any history \hat{h}_i^t , $r_i^B | \hat{h}_i^t$ gives the highest payoffs against s_{-i}^B among all action plans in \mathcal{S}_i^T . Redefine \bar{v} and \underline{v} so that $\bar{v}_i := U_i^T(r_i^G, s_{-i}^G)$ and $\underline{v}_i := U_i^T(r_i^B, s_{-i}^B)$, where $U_i^T : \mathcal{S}_i^T \times \mathcal{S}_{-i}^T \rightarrow \mathbb{R}$ is the payoff function in the T -fold finitely

repeated game, where $U^T(\cdot)$ is the appropriately discounted and normalized sum of stage-game payoffs.

Now we are equipped to specify how player i plays her pairwise game $\Gamma_{i,-i}$. We call this i 's “*partial strategy*”.

Partial Strategies: Specification of Play in $\Gamma_{i,-i}$

- **Initial Period of $\Gamma_{i,-i}$:** In the first ever period when player i meets player $-i$, player i plays s_i^G with probability μ_0 and s_i^B with probability $(1 - \mu_0)$ where μ_0 solves

$$v_{-i} = \mu_0 \bar{v}_{-i} + (1 - \mu_0) \underline{v}_{-i}.$$

Note that since $(1 - \epsilon) \underline{v}_{-i} + \epsilon \bar{v}_{-i} < v_{-i} < \epsilon \underline{v}_{-i} + (1 - \epsilon) \bar{v}_{-i}$, we have $\mu_0, 1 - \mu_0 \geq \epsilon$.

- **Plan Period of a Non-Initial Block of $\Gamma_{i,-i}$:** If player i ever observed a deviation in a signature period of an earlier block in any pairwise game, she plays strategy s_i^B with probability $(1 - \beta^l)$ where l is the number of deviations she has seen so far (in the supergame) and $\beta > 0$ is small.

Otherwise, she plays strategy s_i^G with probability μ and s_i^B with probability $(1 - \mu)$ where the mixing probability μ is chosen to tailor player $-i$'s continuation payoff.

How are continuation payoffs determined? Continuation payoffs are specified in a way that makes each player indifferent between all action plans in S_i^T when her opponent plays s_{-i}^B and indifferent between all action plans in \mathcal{S}_i when her opponent plays s_{-i}^G . The average payoff from playing any action plan in S_i^T against the opponent's play of s_{-i}^B is adjusted to be exactly \underline{v}_i . Similarly, the average payoff from playing any action plan in \mathcal{S}_i against the opponent's play of s_{-i}^G is adjusted to be exactly \bar{v}_i . This is done as follows.

Let c denote the current calendar time period, and let $c(\tau)$, $\tau \in \{1, \dots, T\}$ denote the calendar time period of the τ^{th} stage of the most recently elapsed block in the pairwise game $\Gamma_{i,-i}$.

For any history \hat{h}_i^T observed (at calendar period c) by i in the most recently elapsed block, if s_i^B was played in the last block, define rewards $\omega_{-i}^B(\cdot)$ as

$$\omega_{-i}^B(\hat{h}_i^T) := \sum_{\tau=1}^T \pi_{\tau}^B$$

where,

$$\pi_{\tau}^B := \begin{cases} \frac{1}{\delta^{T+1-\tau}} \theta_{\tau}^B M^{T-\tau+1} & \text{if } c - c(\tau) = T + 1 - \tau \\ 0 & \text{otherwise,} \end{cases}$$

and θ_τ^B is the difference between $-i$'s continuation payoff in the last block from playing r_{-i}^B from period τ on and $-i$'s continuation payoff from playing the action observed by i at τ followed by reversion to r_{-i}^B from $(\tau + 1)$ on. Since r_{-i}^B gives i maximal payoffs, $\theta_\tau^B \geq 0$. Player i chooses $\mu \in (0, 1)$ to solve $\mu\bar{v}_{-i} + (1 - \mu)\underline{v}_{-i} = \underline{v}_{-i} + (1 - \delta)\omega_{-i}^B(\hat{h}_i^T)$. Since T is fixed, we can make $(1 - \delta)\omega_{-i}^B(\hat{h}_i^T)$ arbitrarily small, for large enough δ , and so the above continuation payoff will be feasible.

It is worthwhile to note how these rewards make player $-i$ indifferent between all action plans in S_{-i}^T when her opponent plays s_i^B . Suppose at some stage τ of a block, player $-i$ plays an action that gives her a payoff in the current period that is lower than that from playing r_{-i}^B . With probability $(\frac{1}{M})^{T+1-\tau}$ her next plan period with player i will be exactly $T + 1 - \tau$ calendar periods later, and in that case, she will receive a proportionately high reward $\theta_\tau^B M^{T+1-\tau}$. If her next plan period is not exactly $T + 1 - \tau$ periods later, she does not get compensated. However, in expectation, for any action that she may choose, the loss she will suffer today (compared to the benchmark action plan r_{-i}^B) is exactly compensated by the reward she will get in the future.

If s_i^G was played in the last block, specify punishments $\omega_{-i}^G(\cdot)$ as

$$\omega_{-i}^G(\hat{h}_i^T) := \sum_{\tau=1}^T \pi_\tau^G$$

where,

$$\pi_\tau^G := \begin{cases} \frac{1}{\delta^{T+1-\tau}} \min\{0, \theta_\tau^G\} M^{T-\tau+1} & \text{if } c - c(\tau) = T + 1 - \tau \\ 0 & \text{otherwise,} \end{cases}$$

and θ_τ^G is the difference between $-i$'s continuation payoff within the last block from playing r_{-i}^G from time τ on and $-i$'s continuation payoff from playing the action observed by i at period τ followed by reversion to r_{-i}^G from $\tau + 1$ on. Since r_{-i}^G gives $-i$ minimal payoffs, $\theta_\tau^G \leq 0$ for all actions are used by strategies in \mathcal{S}_{-i} .

Player i chooses $\mu \in (0, 1)$ to solve $\mu\bar{v}_{-i} + (1 - \mu)\underline{v}_{-i} = \bar{v}_{-i} + (1 - \delta)\omega_{-i}^G(\hat{h}_i^T)$. Again, since T is fixed, we can make $(1 - \delta)\omega_{-i}^G(\hat{h}_i^T)$ arbitrarily small, for large enough δ . We restrict attention to δ close enough to 1 so that

$$(1 - \delta)\omega_{-i}^B(\hat{h}_i^T) < \epsilon\underline{v}_{-i} + (1 - \epsilon)\bar{v}_{-i} - \underline{v}_{-i} \text{ and } (1 - \delta)\omega_{-i}^G(\hat{h}_i^T) > (1 - \epsilon)\underline{v}_{-i} + \epsilon\bar{v}_{-i} - \bar{v}_{-i}.$$

- **Signature Period and other Non-initial Periods:** Players use the designated signature $\psi(a)$ if a was the profile realized in the plan period of the block. For the rest of

the block, they play according to the announced plan (i.e. if the announced plans were (P_1, P_2) , then they play action profile a^{P_2, P_1}).

This completes the specification of strategies on the equilibrium path.

3.2.3 Beliefs of Players

At any private history, each player believes that in every period, she met the true owner of the name she encountered, and that no player has ever misrepresented her name.

3.3 Proof of Theorem 1

In this section, I show that the above strategies and beliefs constitute a sequential equilibrium. Here I prove sequential rationality of strategies on the equilibrium path. This is done in two steps. First, conditional on truthful reporting of names, the actions prescribed are shown to be optimal. Second, I show that it is incentive compatible to report one's name truthfully. The proof of sequential rationality off the equilibrium path and consistency of beliefs is relegated to the appendix.

As before, fix a player i and a rival $-i$. The partial strategy for player i in pairwise game $\Gamma_{i,-i}$ can be represented by an automaton that revises actions and states in every plan period of $\Gamma_{i,-i}$.

Set of States: The set of states of a player i is the set of continuation payoffs for her rival $-i$ and is the interval $[(1 - \epsilon)\underline{v}_{-i} + \epsilon\bar{v}_{-i}, \epsilon\underline{v}_{-i} + (1 - \epsilon)\bar{v}_{-i}]$.

Initial State: Player i 's initial state is the target payoff for her rival v_{-i} .

Decision Function: When player i is in state u , she uses strategy s_i^G with probability μ and s_i^B with probability $(1 - \mu)$ where μ solves

$$u = \mu [\epsilon\underline{v}_{-i} + (1 - \epsilon)\bar{v}_{-i}] + (1 - \mu) [(1 - \epsilon)\underline{v}_{-i} + \epsilon\bar{v}_{-i}].$$

Transition Function: For any history \hat{h}_i^T in the last T -period block for player i , if the action played was s_i^G then at the end of the block, the state transits to $\bar{v}_{-i} + (1 - \delta)\omega_{-i}^G(\hat{h}_i^T)$. If the realized action was s_i^B the new state is $\underline{v}_{-i} + (1 - \delta)\omega_{-i}^B(\hat{h}_i^T)$. Recall that for δ large enough, $(1 - \delta)\omega_{-i}^B(\hat{h}_i^T)$ and $(1 - \delta)\omega_{-i}^G(\hat{h}_i^T)$ can be made arbitrarily small, which ensures that the continuation payoff always lies within the interval $[(1 - \epsilon)\underline{v}_{-i} + \epsilon\bar{v}_{-i}, \epsilon\underline{v}_{-i} + (1 - \epsilon)\bar{v}_{-i}]$.

It can be easily seen that given i 's strategy, any strategy of player $-i$ whose restriction belongs to \mathcal{S}_{-i} is a best response. The average payoff within a block from playing r_{-i}^G against

s_i^G is exactly \bar{v}_{-i} , and that from playing r_{-i}^B against s_i^B is \underline{v}_{-i} . Moreover, the continuation payoffs are also \bar{v}_{-i} and \underline{v}_{-i} respectively. Any player's payoff is therefore $\mu_0 \bar{v}_{-i} + (1 - \mu_0) \underline{v}_{-i}$. Note also that each player is indifferent between all action plans in S_i^T when her opponent plays s_{-i}^B .

It remains to verify that players will truthfully report their names in equilibrium. First I show that if a player impersonates someone else in her community, irrespective of what action she chooses to play, she can get detected (i.e. with positive probability, someone in her rival community will become aware that some deviation has occurred). Then, the detector will punish the whole community of the impersonator. For sufficiently patient players, this threat is enough to deter impersonation.

At any calendar time t , define the state of play between any pair of players to be $k \in \{1, \dots, T\}$ where k is the stage of the current block they are playing in their pairwise game (e.g. for a plan period, $k = 1$). At time $(t + 1)$, they will either transit to state $k + 1$ with probability $\frac{1}{M}$, if they happen to meet again in the next calendar time period or remain in state k . Suppose at time t , player i_1 decides to impersonate i_2 . Player i_1 can form beliefs over the possible states of each of her rivals $j, j \in \{1, \dots, M\}$ with respect to i_2 , conditional on her own private history. Denote player i_1 's beliefs over the states of any pair of players by a vector (p_1, \dots, p_n) .

Fix a member j of the rival community, whom player i_1 can be matched to in the next period. Suppose player i_1 has met the sequence of names $\{j^1, \dots, j^{t-1}\}$. For any $t \geq 2$, her belief over states of j and i_2 is given by

$$\sum_{\tau=1}^{t-1} (1 - \mathbb{I}_{j=j^\tau}) \left(\frac{M-2}{M-1} \right)^{\sum_{i=1}^{\tau-1} (1 - \mathbb{I}_{j=j^i})} \frac{1}{M-1} (1, 0, \dots, 0) \prod_{k=\tau}^{t-1} [\mathbb{I}_{j=j^k} I + (1 - \mathbb{I}_{j=j^k}) H], \quad (1)$$

$$\text{where } H = \begin{bmatrix} \frac{M-2}{M-1} & \frac{1}{M-1} & 0 & 0 & \dots & 0 \\ 0 & \frac{M-2}{M-1} & \frac{1}{M-1} & 0 & \dots & 0 \\ \vdots & & & & & \\ \frac{1}{M-1} & 0 & 0 & 0 & \dots & \frac{M-2}{M-1} \end{bmatrix}$$

$$I \text{ is the } T \times T \text{ identity matrix, and } \mathbb{I}_{j=j^\tau} = \begin{cases} 1 & \text{if } j = j^\tau, \\ 0 & \text{otherwise.} \end{cases}$$

To see how we obtain the above expression, note that player i_1 knows that in periods when she met rival j , it is not possible that player i_2 also met j . Hence, she knows with certainty

that in these periods the state of play between players i_2 and j did not change. She believes that in other periods, the state of play would have changed according the transition matrix H . This gives the product term in the above expression. For any given calendar period τ , player i_1 can also use this information to compute the expected state of players i_2 and j conditioning on the event that i_2 and j met for the first time ever in period τ . For any τ , the probability that players i_2 and j met for the first time at period τ is given by $\left(\frac{M-2}{M-1}\right)^{\sum_{l=1}^{\tau-1} (1-\mathbb{I}_{j=j^l})} \frac{1}{M-1}$. Finally, player i_1 knows that the pair i_2 and j could not have met for the first time in a period that she met j herself, and so needs to condition only on such periods when she did not meet j .

Notice that the transition matrix H is irreducible and

$$\lim_{q \rightarrow \infty} (1, 0, \dots, 0) \cdot H^q = \left(\frac{1}{T}, \dots, \frac{1}{T}\right). \quad (2)$$

Further it can be easily shown that the following is true.

$$\forall q \geq 1, \quad [(1, 0, \dots, 0) \cdot H^q]_2 > 0, \quad (3)$$

where $[(1, 0, \dots, 0) \cdot H^q]_2$ represents the 2^{nd} component of $(1, 0, \dots, 0) \cdot H^q$

It follows from (2) and (3) that for any rival j whom i_1 has not met at least in one period, there exists a lower bound $\phi > 0$ such that the probability of j being in state 2 with i_2 is at least ϕ .

Now, when player i_1 announces name i_2 , she does not know which rival she will end up meeting that period. It follows that at $t \geq 2$, player i_1 assigns probability at least $\frac{\phi}{M(M-1)}$ to the event that the rival she meets is in state 2 with i_2 . (To see why, pick a rival j' whom i_1 did not meet in the first calendar time period ($t = 1$). With probability $\frac{1}{M}$, at time t , i_1 will meet this j' and with probability $\frac{1}{M-1}$ this j' would have met i_2 at $t = 1$ and period t could be their signature period.)

Consequently, if player i_1 announces her name to be i_2 , there is a minimal strictly positive probability $\epsilon^2 \frac{\phi}{M(M-1)}$ that her impersonation gets detected. This is because if the rival she meets is supposed to be in a signature period with i_2 , they should play one of the signatures g, b, x, y depending on the realized plan in their plan period. Since players mix with probability at least ϵ on both Plans G and B , player i_1 will play the wrong signature with probability at least ϵ^2 irrespective of the action she chooses. Player i_1 's rival will realize that some deviation has occurred, and she will switch to the bad plan B (almost certainly) with each of the players in i_1 's community in their next plan period.

Player i_1 will not misreport her name if her maximal potential gain from deviating is not greater than the minimal expected loss in continuation payoff from detection.

$$\text{Player } i_1 \text{'s maximal current gain from misreporting} = \left(1 - \frac{\delta}{\delta + M(1 - \delta)}\right) \gamma.$$

Note that because of the random matching process, the effective discount factor for any player in her pairwise games is not δ , but higher, i.e. $\frac{\delta}{\delta + M(1 - \delta)}$.

Player i_1 's minimal expected loss in continuation payoff from impersonation is given by

$$\text{Minimal loss from deviation} \geq \frac{\phi}{M(M-1)} \epsilon^2 (1 - \beta) \left(\frac{\delta}{\delta + M(1 - \delta)}\right)^T [v_i - ((1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i)].$$

To derive the above expression, observe that there is a minimal probability $\frac{\phi}{M(M-1)}$ that players j and i_2 are in a signature period. Conditional on this event, irrespective of the action i_1 plays, there is a minimal probability ϵ^2 that her deviation gets detected by her rival, j . Conditional on detection, player j will switch to playing the unfavorable strategy with probability $(1 - \beta)$ in the next plan period with i_1 . At best, i_1 and j 's plan period is $(T - 1)$ periods away, after which i_1 's payoff in her pairwise game with j will drop from the target payoff v_i to $(1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i$.

i_1 will not impersonate if her maximal current gain is outweighed by her loss in continuation payoff i.e. if the following inequality holds.

$$\left(1 - \frac{\delta}{\delta + M(1 - \delta)}\right) \gamma \leq \frac{\phi}{M(M-1)} \epsilon^2 (1 - \beta) \left(\frac{\delta}{\delta + M(1 - \delta)}\right)^T [v_i - ((1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i)].$$

For δ close enough to 1, this inequality is satisfied, and so misreporting one's name is not a profitable deviation at any $t \geq 2$.

Now consider incentives for truth-telling in the first period of the supergame. Suppose i_1 impersonates i_2 and meets rival j . In the next period, with probability $\frac{\epsilon^2}{M}$, i_2 will meet j and use the wrong signature, thus informing j that someone has deviated. By a similar argument as above, if δ is high enough, i_1 's potential current gain will be outweighed by the future loss in continuation payoff. \square

The interested reader may refer to the appendix for a formal proof of the consistency of beliefs and sequential rationality off the equilibrium path.

Remark 1 *General Matching Technologies:* A distinguishing feature of this result is that unlike earlier literature, it does not depend on the random matching being independent or

uniform. The assumption of uniform independent matching is made only for convenience. The construction continues to work for more general matching technologies. For instance it is enough to assume that for each player, the probability of being matched to each rival is strictly positive and the expected time until she meets each of her rivals again is bounded.

Remark 2 *Generalizable to Asymmetric Payoffs:* In this result, I restrict attention to the case where all members of a specific community get identical payoffs. With the same equilibrium strategies, it is possible to also achieve other asymmetric payoff profiles $(v_{i_1}, \dots, v_{i_M}, v_{j_1}, \dots, v_{j_M})$ with the property that for all possible pairs of rivals i and j , $(v_i, v_j) \in \text{int}(\mathcal{F}^*)$. Clearly, the feasibility of asymmetric payoff profiles does depend on the specifics of the matching process, in particular on the probability of meeting each rival.

Remark 3 *Asymmetric Discount Factors:* Unlike in earlier work (e.g. Ellison (1994)), the assumption of a common discount factor for all players is not necessary for the equilibrium construction of this paper.

3.4 Small Communities ($M = 2$)

An important feature that enables the above construction is that at any time, each player is uncertain about the states that the other players are in with respect to each other. This source of uncertainty ensures that if a player wants to impersonate somebody, she believes that she will get detected. This is no longer the case if we consider communities with just two members each. Since each player knows the sequence of names she has met, she knows the sequence of names her rivals have met (conditional on truthful revelation). So, each player knows with certainty which period of a block any pair of her rivals is in. Since the states of one's opponents' play are no longer random, the above construction does not apply.

In this section, I show that with some modification to the strategies, every feasible and individually rational payoff is still achievable.

3.4.1 Equilibrium Construction

As before, play proceeds in blocks of T interactions between any pair of players, but now each block starts with “*initiation periods*”. The first ever interaction between any two players is called their “*game initiation period*”. In this period, the players play a coordination game.

They each play two given actions (say a_1 and a_2 for player 1 and b_1 and b_2 for player 2) with equal probability. If the realized action profile is not (a_1, b_1) , the game is said to be initiated and players continue to play as described below. If the realized action profile is (a_1, b_1) , players replay the game initiation period. Once the pairwise game is initiated, it proceeds as before in blocks of T periods. Any new block of play also starts with similar initiation periods. In a block initiation period, players play as described above. If the realized profile is not (a_1, b_1) , they start playing their block action plans from the next period. Otherwise, they play the initiation period again. Once a block is initiated, play within the block proceeds exactly as in the earlier construction, i.e. players start the block with a plan period followed by a signature period and then play according to the announced plan of the block. Since the pairwise game after initiation is exactly the same as in the earlier construction, I omit a detailed description here.

The initiation periods ensure that no player can know precisely what state her rivals are in with respect to each other. In particular, no player knows whether a given period is a signature period for any pair of her rivals. Further, no player outside a pair can observe the action realized in the plan period, and so is unaware of the sequence of actions that is being played. Consequently, if anyone outside a pair tries to impersonate one of the members of the pair, she can end up playing the wrong action in case it is a signature period and thus get detected. If a deviation is detected, the detector punishes the entire rival community by switching to the unfavorable strategy with every rival in the next plan period. This threat is enough to deter deviation if players are sufficiently patient.

Since the construction is quite similar, the details of the proof are relegated to the appendix.

3.5 Cooperation within a Single Community

In many applications, it may be reasonable to assume that there is only one large community of players who interact repeatedly with each other, possibly in different roles. For example, consider a large community of traders over the internet, where people are repeatedly involved in a two-player game between a buyer and a seller. It is conceivable that no player is just a seller or just a buyer. Players switch roles in the trading relationship in each period, but each time play a trading game against another trader in the community. Can cooperation be sustained in this slightly altered environment?

It turns out that the same equilibrium construction works for a single community of agents. Any feasible and individually rational payoff can be sustained in equilibrium within a single

community of players in the same way, using the idea of community responsibility. To see how, consider a community of M players, being randomly matched in every period and playing a two-player stage-game. For ease of exposition, think of a two-player trading game played between a buyer and a seller. Suppose players are paired randomly each period, and a public randomization device determines the roles within each pair. (Say, players are designated buyers and sellers with equal probability).

Each player now plays one set of games as a buyer against $(M - 1)$ sellers and another set of games as a seller against $(M - 1)$ buyers. She tracks continuation payoffs separately for each possible opponent in exactly the same way as before. Now she treats the same name in a buyer role and a seller role separately. If a player detects a deviation as a seller (or buyer), she switches to a bad mood against all buyers (or sellers) at the earliest possible opportunity (i.e. at the start of a new T -period block with each opponent).

An interesting observation is that a single community actually facilitates detection of impersonations. If a player misreports her name, with positive probability she will meet the real owner of her reported name, and in this case her rival will know with certainty that an impersonation has occurred. This feature can be used to simplify the equilibrium strategies, and eliminate the need for special signature periods.

4 Community Responsibility with Multiple Communities

So far, we have analyzed the interaction between two communities of agents who repeatedly play a two-player game and shown that a Folk Theorem holds if players are sufficiently patient. This section establishes that the result generalizes to situations with random multilateral matching where $K > 2$ communities interact repeatedly. Agents from K different communities are randomly matched to form groups of K players each (called “*playgroups*”). Players first simultaneously introduce themselves, and then play a simultaneous move K -player stage-game. It turns out it is still possible to achieve any individually rational feasible interior payoff through community responsibility.

How does community responsibility work when there are multiple communities? In the two-player case, each player keeps track of her rival’s continuation payoff. Her own strategy is independent of her own continuation payoff, which is controlled by her rival. With K players, the challenge is that we need to ensure that each player can control the payoffs of all her rivals

simultaneously. This problem is resolved by making each community keep track of exactly one other community. The construction can be summarized as follows.

Every player tracks separately her play with every possible K player group she could be in. Play within any playgroup proceeds in blocks of T periods. Each community k acts as the *monitor* of one other community, say its *successor community* $k + 1$ (community K 's successor is community 1). At the beginning of each block, each player uses one of two continuation strategies. She is indifferent between them, but the strategy she chooses determines whether the continuation payoff of the player of her successor community in that playgroup is high or low. So, each player's payoff is tracked by her monitor in a playgroup. The monitor randomizes between her two strategies at the start of each block in a way to ensure that the target payoff of her successor is achieved. As before, conditional on truthful announcement of names, these types of strategies can be used to attain cooperative outcomes. As in the case of two communities, community responsibility is used to ensure truthful announcement of names. If any player deviates from the equilibrium strategies, she can be punished in two ways. First, the members of her specific playgroup can minmax her. Second, her monitor can hold her whole community responsible and punish the community by switching to the unfavorable strategy with all her playgroups at the start of the next block.

4.1 Model and Result

Multilateral Matching: There are K communities of agents with $M > 2$ members in each community I , $I \in \{1, \dots, K\}$. In each time period $t \in \{1, 2, \dots\}$, agents are randomly matched into groups of K members each, with one member from each community. Let \mathcal{G}_{-k} denote a group of $(K - 1)$ players with members from all except the k^{th} community. Let $m_t(\mathcal{G}_{-k})$ denote the member of the k^{th} community who is matched to the group \mathcal{G}_{-k} . Matches are made independently and uniformly over time, i.e. \forall histories, $\forall j \in$ community k , $\Pr[j = m_t(\mathcal{G}_{-k})] = \frac{1}{M}$. For any player i , the set of rivals she is matched with (say \mathcal{G}_{-i}) is said to constitute her *playgroup*. After being matched, players announce their names. However, names are not verifiable. Then, they play the K -player stage-game.

Stage-Game and Message Sets: As in the model with two communities, each community has a directory of names $\mathcal{N}_I : I \in \{1, \dots, K\}$ with M names each. A name profile of a playgroup is denoted by $\nu \in \mathcal{N} := \mathcal{N}_1 \times \dots \times \mathcal{N}_K$. Let $\Delta(\mathcal{N}_I)$ denote the set of mixtures of messages in \mathcal{N}_I . The stage-game Γ has finite action sets $A_I, I \in \{1, \dots, K\}$. Denote an action profile by

$a \in A := \prod_I A_I$. The set of mixtures of actions in A_I is denoted by $\Delta(A_I)$. Stage-game payoffs are given by a function $u : A \rightarrow \mathbb{R}^K$. Define \mathcal{F} to be the convex hull of the payoff profiles that can be achieved by pure action profiles in the stage-game. Formally, $\mathcal{F} := \text{conv}(\{u(a) : a \in A\})$. As before, denote the feasible and individually rational payoff set by $\mathcal{F}^* := \{v \in \mathcal{F} : v_i > v_i^* \forall i\}$, where v_i^* is the mixed action minmax value for player i . We consider games where \mathcal{F}^* has non-empty interior ($\text{Int}\mathcal{F}^* \neq \emptyset$). Let γ be also defined as before. All players have a common discount factor $\delta \in (0, 1)$.

Information Assumption: Players can observe only the transactions they are personally engaged in. So each player knows the names that she encountered in her playgroup in each period and the action profiles played in that playgroup. She does not know the true identity of her partners. She does not know the composition of other playgroups or how play proceeds in them.

The definitions of histories, strategies, action plans and sequential equilibrium can be easily extended to this setting in a way analogous to Section 2.

Theorem 2 (*Folk Theorem for Random Multilateral Matching Games*) *Consider a finite K -player game being played by $K > 2$ communities of M members each in a random matching setting. For any $(v_1, \dots, v_K) \in \text{Int}(\mathcal{F}^*)$, there exists a sequential equilibrium that achieves payoffs (v_1, \dots, v_K) in the infinitely repeated random matching game with names with KM players, if players are sufficiently patient.*

The equilibrium construction in the K -community case is similar to the two community case. So the formal specification of strategies and the proof of Theorem 2 are relegated to the appendix.

5 Conclusion

In games where large communities transact with each other, it is reasonable to assume that players change partners over time, they do not recognize each other or have very limited information about each other's actions. This paper investigates whether it is possible to achieve all individually rational and feasible payoffs in equilibrium in such anonymous transactions. To answer this question, I consider a repeated two-player game being played by two communities of agents. In every period, each player is randomly matched to another player from the rival community and the pair plays the two-player stage-game. Players do not recognize each other.

Further, they observe only the transactions they are personally involved in. I examine what payoffs can be sustained in equilibrium in this setting of limited information availability.

I obtain a strong possibility result by allowing players to announce unverifiable messages in every period. The main result is a Folk Theorem which states that for any two-player game played between two communities, it is possible to sustain all feasible individually rational payoffs in a sequential equilibrium, provided players are sufficiently patient. Though cooperation in anonymous random matching games has been studied before, little was known about games other than the prisoner's dilemma. This paper is an attempt to fill this gap in the literature.

Earlier literature has shown that though efficiency can be achieved in a repeated PD with no information transmission, with any other game, transmission of hard information seems necessary. Kandori (1992) assumes the existence of labels - players who have deviated or faced deviation can be distinguished from those who have not, by their labels. Takahashi (2007) assumes that players know the full history of past actions of her rival. To the best of my knowledge, this paper is the first to obtain a general Folk Theorem without adding any hard information in the model. Though players can announce names, it is unverifiable cheap talk.

An interesting feature of the strategies I use is that cooperation is not achieved by the customary community enforcement. In most settings with anonymous transactions, cooperation is sustained by implementing third-party sanctions. A player who deviates is punished by other people in the society, not necessarily by the victim. Here, cooperation is sustained by community responsibility. A player who deviates is punished only by the victim, but the victim holds the deviator's entire community responsible and punishes the whole community. It is this alternate form of punishment that allows us to obtain the Folk Theorem in a setting with such limited information.

An appealing feature of the equilibrium in this paper is that unlike earlier work, the construction applies to quite general matching technologies, and does not require uniform or independent matching. I also show that the Folk Theorem extends to a setting with multiple communities playing a K -player stage-game.

A question that remains unanswered in this paper is whether cooperation can be achieved in a general game with even less information than is used here. Can we obtain a Folk Theorem for general games without *any* transmission of information? If not, what is the minimal information transmission which will enable impersonal exchange between two large communities? This is the subject of future work.

6 Appendix

6.1 Sequential Equilibrium

Section 3.3 establishes optimality of strategies on the equilibrium path. Below, I prove sequential rationality off the equilibrium path and the consistency of beliefs. Strategies on the equilibrium path were specified in Section 3.2. Off-equilibrium strategies are defined as follows.

- $\forall i, \forall t, \forall h_i^t \in \mathcal{H}_i^t, \sigma_i^*[h_i^t] = i$.

In other words, after any complete private history including those in which they observed a deviation (own or other), players report their name truthfully,

- $\forall k_i^t = \{(\nu^1, a^1), \dots, (\nu^{t-1}, a^{t-1}), \nu^t\} \in \mathcal{K}_i^t$ with $\nu_i^t = i$ and $\nu_i^\tau \neq i$ for some τ , player i plays the partial strategy for pairwise game $\Gamma_{i,j}$ where $\nu_{-i}^t = j$.

In other words, at any t -period interim private history in which a player has misrepresented her name in at least one period, but has reported truthfully in the current period, she plays game $\Gamma_{i,-i}$ according to the partial strategy against the current rival name.

- $\forall k_i^t = \{(\nu^1, a^1), \dots, (\nu^{t-1}, a^{t-1}), \nu^t\} \in \mathcal{K}_i^t$ with $\nu_i^t \neq i$, $\sigma^*[k_i^t] = \operatorname{argmax}_{a_i \in A_i} U_i(a_i, \sigma_{-i}^* | \xi_i[k_i^t])$.

In other words, at any t -period interim private history in which a player has misrepresented her name in the current period, she plays the action that maximizes her expected utility given her beliefs and her rivals' equilibrium strategies.

- At any t -period interim private history in which a player has deviated by playing the wrong action, i.e. $\forall k_i^t = \{(\nu^1, a^1), \dots, (\nu^{t-1}, a^{t-1}), \nu^t\} \in \mathcal{K}_i^t$ with $a_i^\tau \neq \sigma_i^*[k_i^\tau]$ for some τ , $\sigma^*[k_i^t]$ prescribes the following.
 - If ν_{-i}^τ was in the unfavorable state (playing s_{-i}^B), player i should play her best response to the minmax strategy of her opponent for the rest of the block, and then revert to playing her partial strategy for her game $\Gamma_{i,-i}$ against this rival.
 - If ν_{-i}^τ was in the favorable state (playing s_{-i}^G), player i should continue playing s_{-i}^G for the rest of the block and revert to playing her partial strategy for her game $\Gamma_{i,-i}$ against this rival.

Optimality of Actions:

Lemma 1 *For any player i , misreporting ones name is not optimal after any history.*

Proof: Fix a player i . The proof of the Folk Theorem establishes optimality on the equilibrium path. So now consider any information set of player i reached off the equilibrium path, possibly after one or more deviations (impersonations or deviations in action) by player i herself or others. We compare i 's payoffs if she truthfully reports her name to her payoffs if she impersonates someone.

Consider the play between i and any rival name j who has observed d deviations so far. By misreporting and claiming to be i' , i can potentially get a short-term gain in the pairwise game with j .

$$\text{Maximal Gain} \leq \left(1 - \frac{\delta}{\delta + M(1 - \delta)}\right) \gamma.$$

However, by impersonating i' , player i increases the probability with which j will punish in case her deviation is detected. Player i 's minimal expected loss in continuation payoff from the deviation is given by the following.

$$\text{Minimal expected loss} \geq \frac{\phi}{M(M-1)} \epsilon^2 \left(\frac{\delta}{\delta + M(1 - \delta)}\right)^T (\beta^d - \beta^{d+1}) [v_i - ((1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i)].$$

To see how we obtain this expression, note that there is a minimal probability $\frac{\phi}{M(M-1)}$ that j and i' are supposed to be in a signature period. Conditional on this event, irrespective of what action i plays, there is a minimal probability ϵ^2 that her rival j will learn of a deviation. Conditional on detection, player j will switch to the unfavorable action plan with probability $(1 - \beta^{d+1})$ in the next plan period, instead of $(1 - \beta^d)$. At best, i and j 's plan period is $(T - 1)$ periods away, after which i 's payoff in her pairwise game with j will drop from the target payoff v_i to $(1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i$. (As before, in the pairwise game between i and j , the effective discount factor is not δ but higher, i.e. $\frac{\delta}{\delta + M(1 - \delta)}$.)

So, player i will not misreport her name if the maximal gain from deviating is outweighed by the minimal expected loss in continuation payoff, i.e. if the following inequality holds.

$$\gamma \left(1 - \frac{\delta}{\delta + M(1 - \delta)}\right) \leq \frac{\phi}{M(M-1)} \epsilon^2 \left(\frac{\delta}{\delta + M(1 - \delta)}\right)^T \beta^d (1 - \beta) [v_i - ((1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i)].$$

It can be seen that the above inequality holds for sufficiently large δ . Hence, at any information set off the equilibrium path, i does not find it profitable to misreport her name. \square

This establishes that the strategies are optimal, since conditional on truthful reporting of names, it is optimal to play the specified actions.

Consistency of Beliefs:

For any player i , perturb the strategies as follows. (Fix $\eta > 0$ small.)

- At any t -period complete private history, player i announces her name truthfully with probability $(1 - \frac{\eta^2}{e^t})$ and announces an incorrect name with complementary probability (randomizing uniformly between other possible names).
- At any interim t -period private history, player i plays the equilibrium action with probability $(1 - \eta^{\frac{1}{2^t}})$. She plays other actions with complementary probability (randomizing uniformly across the other possible actions).

Now, consider any t -period complete private history of player i . We will show that she believes with probability 1 that there have been no impersonations in the past.

Any observed history is consistent with the sequence of events that there have been no impersonations but only deviations in action. Consider the sequence of events of no impersonations and t deviations in action. If this sequence is consistent with the observed history, the probability that player i assigns to this sequence of events is given by

$$\prod_{s=1}^t \left(1 - \frac{\eta^2}{e^s}\right) \eta^{\frac{1}{2^s}}.$$

Since $\sum_{n=1}^k \frac{1}{2^n}$ is bounded above by 1, it follows that the probability of any number of deviations in action is bounded below by $\eta(1 - \eta)$. Hence any sequence of events with no name deviations and some action deviations will be assigned probability that is greater than

$$\eta(1 - \eta) \prod_{s=1}^t \left(1 - \frac{\eta^2}{e^s}\right).$$

Further, we can show that the above expression is bounded below by a constant κ uniformly in t . To see how, note that

$$\begin{aligned} \eta(1 - \eta) \prod_{s=1}^t \left(1 - \frac{\eta^2}{e^s}\right) &\geq \eta(1 - \eta) \prod_{s=1}^t \left(1 - \frac{1}{e^s}\right) \\ &\geq \eta(1 - \eta) \prod_{s=1}^{\infty} \left(1 - \frac{1}{e^s}\right) \end{aligned}$$

We know that the series $\sum_{s=1}^{\infty} \frac{1}{e^s}$ converges, which implies that the infinite product $\prod_{s=1}^{\infty} (1 - \frac{1}{e^s})$ converges.¹² Since the infinite product converges, there exists a constant κ such that

$$\forall t, \eta(1 - \eta) \prod_{s=1}^t \left(1 - \frac{\eta^2}{e^s}\right) \geq \eta(1 - \eta)\kappa.$$

¹²This follows from the result that for $u_n \in [0, 1)$, $\prod_{n=1}^{\infty} (1 - u_n) > 0 \iff \sum_{n=1}^{\infty} u_n < \infty$. (See Rudin: *Real and Complex Analysis*)

Now we analyze sequences of events which are consistent with the observed history and which involve at least one impersonation.

Consider sequences with only one impersonation. The probability of this set of events is given by

$$p(1) = \sum_{r=1}^t \frac{\eta^2}{e^r} \prod_{q \neq r} \left(1 - \frac{\eta^2}{e^q}\right).$$

The probability of the set of events with exactly two impersonations is given by

$$p(2) = \sum_{\tau=1}^t \frac{\eta^2}{e^\tau} \left[\sum_{r>\tau} \frac{\eta^2}{e^r} \prod_{q \neq r, q \neq \tau} \left(1 - \frac{\eta^2}{e^q}\right) \right].$$

Similarly for sequences of events with l impersonations, we have

$$p(l) = \sum_{\tau_1=1}^t \frac{\eta^2}{e^{\tau_1}} \sum_{\tau_2>\tau_1} \frac{\eta^2}{e^{\tau_2}} \cdots \sum_{\tau_{l-1}>\tau_{l-2}} \frac{\eta^2}{e^{\tau_{l-1}}} \sum_{\tau_l>\tau_{l-1}} \frac{\eta^2}{e^{\tau_l}} \prod_{q \neq \tau_i, i \in \{1, \dots, l\}} \left(1 - \frac{\eta^2}{e^q}\right).$$

Hence the probability of the sequences of events that are consistent with the observed history and involve any impersonations is given by $P := \sum_{l=1}^t P(l)$. Collecting terms differently (in powers of e), we can see that for any t ,

$$P \leq \sum_{m=1}^t \eta^2 \frac{1}{e^m} \sum_i^{\sqrt{2m}} m^i \tag{4}$$

$$\begin{aligned} &\leq \sum_{m=1}^{\infty} \eta^2 \frac{1}{e^m} \sum_i^{\sqrt{2m}} m^i \\ &= \eta^2 \sum_{m=1}^{\infty} \frac{1}{e^m} \frac{m(-1 + m\sqrt{2m})}{(-1 + \sqrt{m})(1 + \sqrt{m})}. \end{aligned} \tag{5}$$

The first inequality follows from two observations. First, any term with a given power of e , say e^m , can belong to a sequence of events with at most $\sqrt{2m}$ impersonations. Second, if there i impersonations in m periods, there are less than m^i ways in which this can occur.

The series $\sum a_m$ in expression (5) is convergent. Denote the limit by Λ . Convergence follows from the observation that

$$\lim_{m \rightarrow \infty} \frac{a_{m+1}}{a_m} = \frac{1}{e} < 1.$$

Hence, for any t , $P < \eta^2 \Lambda$.

Given any observed history h_i^t of player i , by Bayes' Rule, the probability i assigns to a consistent sequence of events with no impersonations is given by

$$\frac{\Pr(\text{Consistent events with no impersonations})}{\Pr(\text{All consistent events})}$$

$$\geq \frac{\eta(1-\eta)\kappa}{\eta(1-\eta)\kappa + \eta^2\Lambda}.$$

As $\eta \rightarrow 0$, the above expression approaches 1 uniformly for all t . In other words, as perturbations vanish, after any history player i believes that with probability 1 there were no impersonations in the past. \square

6.2 Proof of Folk Theorem for Small Communities ($M = 2$)

Consider any payoff profile $(v_1, v_2) \in \text{Int } \mathcal{F}^*$. We proceed just as in the equilibrium construction of Theorem 1. Pick payoff profiles $w^{GG}, w^{GB}, w^{BG}, w^{BB}$ such that the following conditions hold

1. $w_i^{GG} > v_i > w_i^{BB} \quad \forall i \in \{1, 2\}$.
2. $w_1^{GB} > v_1 > w_1^{BG}$.
3. $w_2^{BG} > v_2 > w_2^{GB}$.

These inequalities imply that there exists \underline{v}_i and \bar{v}_i with $v_i^* < \underline{v}_i < v_i < \bar{v}_i$ such that the rectangle $[\underline{v}_1, \bar{v}_1] \times [\underline{v}_2, \bar{v}_2]$ is completely contained in the interior of $\text{conv}(\{w^{GG}, w^{GB}, w^{BG}, w^{BB}\})$ and further $\bar{v}_1 < \min\{w_1^{GG}, w_1^{GB}\}$, $\bar{v}_2 < \min\{w_1^{GG}, w_1^{BG}\}$, $\underline{v}_1 > \max\{w_1^{BB}, w_1^{BG}\}$ and $\underline{v}_2 > \max\{w_1^{BB}, w_1^{GB}\}$.

We can find finite sequences of action profiles $\{a_1^{GG}, \dots, a_N^{GG}\}$, $\{a_1^{GB}, \dots, a_N^{GB}\}$, $\{a_1^{BG}, \dots, a_N^{BG}\}$, $\{a_1^{BB}, \dots, a_N^{BB}\}$ such that each vector w^{XY} , the average discounted payoff vector over the sequence $\{a_1^{XY}, \dots, a_N^{XY}\}$ satisfies the above relationships if δ is large enough.

Further, we can find $\epsilon \in (0, 1)$ small so that $v_i^* < (1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i < v_i < (1 - \epsilon)\bar{v}_i + \epsilon\underline{v}_i$. In what follows, when we refer to an action profile a^{XY} , we actually refer to the finite sequence of action profiles $\{a_1^{XY}, \dots, a_N^{XY}\}$ described above.

6.2.1 Defining Strategies at Complete Histories: Name Announcements

At complete private histories, players report names truthfully, (i.e. $\forall i, \forall t, \forall h_i^t \in \mathcal{H}_i^t, \sigma_i^*[h_i^t] = i$).

6.2.2 Defining Strategies at Interim Histories: Actions

Partitioning of Histories:

At any interim private history, each player i partitions her history into M separate histories corresponding to each of her pairwise games $\Gamma_{i,-i}$. If her current rival name is j , she plays game $\Gamma_{i,j}$. Since equilibrium strategies prescribe truthful name announcement, a description of $\Gamma_{i,j}$

will complete the specification of strategies on the equilibrium path for the supergame.

Play of Game $\Gamma_{i,-i}$:

Fix player i and a name $-i$ in i 's rival community. Play is specified in an identical manner for each possible rival name. As before, we denote player i 's history in this pairwise game by \hat{h}_i^t . The game $\Gamma_{i,-i}$ between i and $-i$ proceeds in blocks of T interactions, but with each block starting with "initiation periods".

Initiation Periods of Game $\Gamma_{i,-i}$: The first ever interaction between two player i and $-i$ is called the "*game initiation period*". In this period, player 1 (from community 1) plays two given actions (say a_1 and a_2) with equal probability and player 2 (from community 2) plays two actions (say b_1 and b_2) with equal probability. If the realized action profile is not (a_1, b_1) , the game is said to be initiated and players continue to play as described below. If the realized action profile is (a_1, b_1) , players replay the game initiation period. Once the game is initiated, the game proceeds in blocks of T interactions. Any non-initial block of play also starts with similar initiation periods. In a block initiation period, players play as described above. If the realized profile is not (a_1, b_1) , they start playing their block action plans from the next period. Otherwise, they play the initiation period again.

T -period Blocks in $\Gamma_{i,j}$: Once a block is initiated, players use block action plans just like in the construction with $M > 2$ players. In the first period (plan period) of a block, players i and $-i$ take actions which inform each other about the plan of play for the rest of the block. Partition the set of i 's actions into two non-empty subsets G_i and B_i . If player i chooses an action from set G_i , she is said to send plan $P_i = G$. Otherwise she is said to send plan $P_i = B$.

Further, choose any four pure action profiles $g, b, x, y \in A$ such that $g_i \neq b_i \forall i \in \{1, 2\}$. Define the signature function $\psi : A \rightarrow \{g, b, x, y\}$ mapping one-period histories to one of the action profiles as follows.

$$\psi(a) = \begin{cases} g & \text{if } a \in G_1 \times G_2, \\ b & \text{if } a \in B_1 \times B_2, \\ x & \text{if } a \in G_1 \times B_2, \\ y & \text{if } a \in B_1 \times G_2. \end{cases}$$

Suppose the observed plans are (P_1, P_2) .

Define a set of action plans of the standard T -period finitely repeated stage-game as follows.

$$\mathcal{S}_i := \left\{ s_i \in S_i^T : \forall \hat{h}_i^t = \left(a, \psi(a), (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), \dots, (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}) \right), a \in P_i \times G, \right. \\ \left. s_i[\hat{h}_i^1] = \psi_i([\hat{h}_i^1]) \text{ and } s_i[\hat{h}_i^t] = a_i^{P_2, P_1}, t \geq 2 \right\}.$$

As before, in equilibrium, players will use actions plans from the above set. Each player uses one of two actions plans s_i^G and s_i^B , just as before.

Define partially a favorable action plan s_i^G such that

$$s_i^G[\emptyset] \in \Delta(G_i),$$

$$s_i^G[\hat{h}_i^1] = \psi_i([\hat{h}_i^1]), \text{ and}$$

$$\forall \hat{h}_i^t = \left(a, \psi(a), (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), \dots, (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}) \right), a \in P_i \times P_{-i}, t \geq 1, s_i^G[\hat{h}_i^t] = a_i^{P_2, P_1}.$$

Similarly, partially define an unfavorable action plan s_i^B such that

$$s_i^B[\emptyset] \in \Delta(B_i),$$

$$s_i^B[\hat{h}_i^1] = \psi_i([\hat{h}_i^1]),$$

$$\forall \hat{h}_i^t = \left(a, \psi(a), (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), \dots, (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}) \right), a \in P_i \times P_{-i}, t \geq 1, s_i^B[\hat{h}_i^t] = a_i^{P_2, P_1},$$

$$\forall t \geq r > 1, \forall \hat{h}_i^t \text{ after } \hat{h}_i^r = \left(a, \psi(a), (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), \dots, (a_i^{P_2, P_1}, a_{-i}^{P_2, P_1}), (a_i^{P_1, P_2}, a'_{-i}) \right),$$

$$a \in P_i \times P_{-i}, a'_{-i} \neq a_{-i}^{P_2, P_1}, \quad s_i^B[\hat{h}_i^t] = \alpha_i^*, \text{ and}$$

$$\forall \hat{h}_i^t \text{ after } \hat{h}_i^2 = (a, (\psi_i(a), a'_{-i})), a \in P_i \times P_{-i}, a'_{-i} \neq \psi_{-i}(a), t > 2, \quad s_i^B[\hat{h}_i^t] = \alpha_i^*.$$

As before, it is possible to choose T large enough so that for some $\underline{\delta} < 1, \forall \delta > \underline{\delta}$, i 's average payoff within the block from any action plan $s_i \in \mathcal{S}_i$ against s_{-i}^G strictly exceeds \bar{v}_1 and her average payoff from using any action plan $s_i \in \mathcal{S}_i^T$ against s_{-i}^B is strictly below \underline{v}_1 . Assume from here on that $\delta > \underline{\delta}$.

Define the two benchmark action plans used to compute continuation payoffs. Let $r_i^G \in \mathcal{S}_i$ be an action plan such that given any history \hat{h}_i^t , $r_i^G|\hat{h}_i^t$ gives the lowest payoffs against s_{-i}^G among all action plans in \mathcal{S}_i . Define $r_i^B \in \mathcal{S}_i^T$ to be an action plan such that given any history \hat{h}_i^t , $r_i^B|\hat{h}_i^t$ gives the highest payoffs against s_{-i}^B among all action plans in \mathcal{S}_i^T . Redefine \bar{v} and \underline{v} so that $U_i(r_i^G, s_{-i}^G) = \bar{v}_i$ and $U_i(r_i^B, s_{-i}^B) = \underline{v}_i$.

Partial Strategies: Specification of Play in $\Gamma_{i,-i}$

The following describes how player i plays in the game $\Gamma_{i,-i}$. We call this i 's ‘‘partial strategy’’.

- **Game Initiation Period:** Player i plays actions a_1 and a_2 and Player $-i$ plays actions b_1 and b_2 with equal probability.

- **Period following Game Initiation Period:** If the realized action profile is not (a_1, b_1) , the game is said to be initiated and players continue to play as described below. If the realized action profile is (a_1, b_1) , players replay the initiation period in their next meeting.
- **First Plan Period of $\Gamma_{i,-i}$:** In the first ever period that player i meets player $-i$ after their game is initiated, player i mixes between s_i^G and s_i^B in the following way.

- If the first plan period of game $\Gamma_{i,-i}$ occurs in the calendar period immediately following the first initiation period of the game, and action profile a was realized in the initiation period, then player i plays s_i^G with probability μ_0 and s_i^B with probability $(1 - \mu_0)$ where μ_0 solves

$$v_{-i} + \frac{1 - \delta}{\delta} \frac{8}{3} \rho(a) = \mu_0 \bar{v}_{-i} + (1 - \mu_0) \underline{v}_{-i},$$

where ρ is the difference in player $-i$'s payoff from the action profile (a_1, b_1) and the profile a .

- Otherwise, player i plays s_i^G with probability μ_0 and s_i^B with probability $(1 - \mu_0)$, where μ_0 solves

$$v_{-i} = \mu_0 \bar{v}_{-i} + (1 - \mu_0) \underline{v}_{-i},$$

For discount factor δ close enough to 1, the payoffs v_{-i} and $v_{-i} + \frac{1-\delta}{\delta} 4\rho$ both lie in the interval $[(1 - \epsilon) \underline{v}_{-i} + \epsilon \bar{v}_{-i}, \epsilon \underline{v}_{-i} + (1 - \epsilon) \bar{v}_{-i}]$. Henceforth, assume that δ is large enough. Further, in both the above cases, $\mu_0, 1 - \mu_0 \geq \epsilon$.

- **Block Initiation Period:** In the initiation period of a non-initial block, player i plays actions a_1 and a_2 and Player $-i$ plays actions b_1 and b_2 with equal probability.
- **Period following Block Initiation Period:** If the realized action profile in the last interaction was not (a_1, b_1) , the next block is said to be initiated and players continue to play as described below. If the realized action profile is (a_1, b_1) , players replay the initiation period.
- **Plan Period of a Non-Initial Block of $\Gamma_{i,-i}$:** If player i ever observed a deviation in a signature period of an earlier block, she plays strategy s_i^B with probability $(1 - \beta^l)$ where l is the number of deviations she has seen so far and $\beta > 0$ is small.

Otherwise, she plays strategy s_i^G with probability μ and s_i^B with probability $(1 - \mu)$ where the mixing probability μ is used to tailor player $-i$'s continuation payoff, as shown below. Let c be the current calendar time period, and $c(\tau)$, $\tau \in \{1, \dots, T\}$ denote the calendar

time period of the τ^{th} period of the most recently elapsed block. For any history \hat{h}_i^T observed (at calendar period c) by i in the most recently elapsed block, if s_i^B was played in the last block, we define rewards $\omega_{-i}^B(\cdot)$ as

$$\omega_{-i}^B(\hat{h}_i^T) := \sum_{\tau=1}^T \pi_{\tau}^B$$

where

$$\pi_{\tau}^B = \begin{cases} \frac{1}{\delta^{T+2-\tau}} \theta_{\tau}^B \frac{4}{3} 2^{T+2-\tau} + \frac{1}{\delta} \frac{8}{3} \rho^B(a) & \text{if } c - c(\tau) = T + 2 - \tau \\ 0 & \text{otherwise.} \end{cases}$$

θ_{τ}^B is the difference between $-i$'s continuation payoff in the last block from playing r_{-i}^B from time τ on and $-i$'s continuation payoff from playing the action observed by i at period τ followed by reversion to r_{-i}^B from $(\tau + 1)$ on, and $\rho^B(a)$ is the difference between the maximum possible one-period payoff in the stage-game and player $-i$'s payoff from profile a . Since r_{-i}^B gives i maximal payoffs, $\theta_{\tau}^B \geq 0$. Also by definition, $\rho^B(a) \geq 0$.

Player i chooses $\mu \in (0, 1)$ to solve $\mu \bar{v}_{-i} + (1 - \mu) \underline{v}_{-i} = \underline{v}_{-i} + (1 - \delta) \omega_{-i}^B(\hat{h}_i^T)$.

If s_i^G was played in the last block, we specify punishments $\omega_{-i}^G(\cdot)$ as

$$\omega_{-i}^G(\hat{h}_i^T) := \sum_{\tau=1}^T \pi_{\tau}^G$$

where,

$$\pi_{\tau}^G = \begin{cases} \frac{1}{\delta^{T+2-\tau}} \min\{0, \theta_{\tau}^G\} \frac{4}{3} 2^{T+2-\tau} + \frac{1}{\delta} \frac{8}{3} \rho^G(a) & \text{if } c - c(\tau) = T + 1 - \tau \\ 0 & \text{otherwise,} \end{cases}$$

θ_{τ}^G is the difference between $-i$'s continuation payoff within the last block from playing r_{-i}^G from time τ on and $-i$'s continuation payoff from playing the action observed by i at period τ followed by reversion to r_{-i}^G from $\tau + 1$ on and $\rho^G(a)$ is the difference between the minimum possible one-period payoff in the stage-game and player $-i$'s payoff from profile a . Since r_{-i}^G gives $-i$ minimal payoffs, $\theta_{\tau}^G \leq 0$ for all actions are used by strategies in \mathcal{S}_{-i} . By definition, $\rho^G(a) \geq 0$.

Player i chooses $\mu \in (0, 1)$ to solve $\mu \bar{v}_{-i} + (1 - \mu) \underline{v}_{-i} = \bar{v}_{-i} + (1 - \delta) \omega_{-i}^G(\hat{h}_i^T)$.

Note that since T is fixed, we can make $(1 - \delta) \omega_{-i}^G(\hat{h}_i^T)$ and $(1 - \delta) \omega_{-i}^B(\hat{h}_i^T)$ arbitrarily small, for large enough δ . We restrict attention to δ close enough to 1 so that

$$(1 - \delta) \omega_{-i}^B(\hat{h}_i^T) < \epsilon \underline{v}_{-i} + (1 - \epsilon) \bar{v}_{-i} - \underline{v}_{-i} \text{ and } (1 - \delta) \omega_{-i}^G(\hat{h}_i^T) > (1 - \epsilon) \underline{v}_{-i} + \epsilon \bar{v}_{-i} - \bar{v}_{-i}.$$

For such δ , the continuation payoff at every period always lies within the interval $[(1 - \epsilon)\underline{v}_{-i} + \epsilon\bar{v}_{-i}, \epsilon\underline{v}_{-i} + (1 - \epsilon)\bar{v}_{-i}]$.

- **Signature Period and other Non-initial Periods:** Players use the designated signature $\psi(a)$ if a was the profile realized in the plan period of the block. For the rest of the block, they play according to the announced plan.

This completes the specification of strategies on the equilibrium path.

6.2.3 Beliefs of Players

At any private history, each player believes that in every period, she met the true owners of the names she encountered, and that no player ever misrepresented her name.

6.2.4 Proof of Equilibrium

First we show that conditional on truthful reporting of names, these strategies constitute an equilibrium.

Note that any player i is indifferent across her actions in the initiation period of a game against any rival $-i$. This is because any gain that player i can get over her payoff from profile a in the initiation period will be wiped out in expectation. With probability $\frac{3}{8}$, she expects to meet player $-i$ again in the next calendar time period and initiate the game. In this case, player $-i$ will adjust her continuation payoff to exactly offset any gain or loss she made in the initiation period.

Once the game is initiated, the strategies of any pair of players can be represented by an automaton which revises actions and states in every plan period. The following describes the automaton for any player $-i$.

Set of states: The set of states of a player $-i$ is the set of continuation payoffs for her rival i and is the interval $[(1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i, \epsilon\underline{v}_i + (1 - \epsilon)\bar{v}_i]$.

Initial State: Player $-i$'s initial state is the target payoff for her rival v_i .

Decision Function: When $-i$ is in state u , she uses s_{-i}^G with probability μ and s_{-i}^B with probability $(1 - \mu)$ where μ solves $u = \mu[\epsilon\underline{v}_i + (1 - \epsilon)\bar{v}_i] + (1 - \mu)[(1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i]$

Transition Function: For any history \hat{h}_{-i}^T for player $-i$, if the realized action plan is s_{-i}^G then at the end of the block, the state transits to $\bar{v}_i + (1 - \delta)\omega_i^G(\hat{h}_{-i}^T)$. If the realized action plan is s_{-i}^B the new state is $\underline{v}_i + (1 - \delta)\omega_i^B(\hat{h}_{-i}^T)$.

It can be easily seen that given $-i$'s action plan, any action plan of player i whose restriction belongs to \mathcal{S}_i is a best response. The average payoff within a block from playing r_i^G against s_{-i}^G is exactly \bar{v}_i , and that from playing r_i^B against s_{-i}^B is \underline{v}_i . Moreover, the continuation payoffs are also \bar{v}_i and \underline{v}_i respectively. Any player's payoff is therefore $\mu_0 \bar{v}_i + (1 - \mu_0) \underline{v}_i$.

Note that each player is indifferent between all action plans in S_i^T when her rival plays s_{-i}^B . At any stage τ of a block, player i believes that with probability $\frac{3}{4}(\frac{1}{2})^{T+2-\tau}$, her next plan period with $-i$ is exactly $(T + 2 - \tau)$ periods away, and in that case, for any action she chooses now she will receive a proportionately high reward $\frac{4}{3}\theta_\tau^B 2^{T+2-\tau}$. In expectation, any loss she suffers today is exactly compensated for in the future. Similarly, in an initiation period of any block, player i believes that with probability $\frac{3}{8}$ that she will initiate the block in the next calendar time period, and again for any action that she chooses now, she gets a proportionate reward / punishment.

It remains to check if players will truthfully report their names. At any calendar time t , define the state of play between any pair of players to be $k \in \{0, 1, \dots, T\}$, where k is the stage of the current block they are in (with $k=0$ for the initiation period). Suppose at period t , player i_1 impersonates i_2 and meets rival j . Player i_1 can form beliefs over the possible states that each of her rivals j_1 and j_2 are in with respect to player i_2 , conditional on her own private history. Based on her own history, i_1 knows how many times her rivals have met. Suppose player i_1 knows that player i_2 has met rival j_1 J_1 times and met the other rival J_2 times. Player i_1 has a belief over the possible states that j_1 and i_2 are in. Represent a player's beliefs by a vector (p_0, \dots, p_T) .

For any $t \geq 2$, player i_1 's belief over the states of j_1 and i_2 is given by:

$$(1, 0, \dots, 0) \cdot H^{J_1}, \text{ where } H = \begin{bmatrix} \frac{1}{4} & \frac{3}{4} & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & 0 & \dots & 0 \\ \vdots & & & & & & \\ 0 & 0 & 0 & 0 & 0 & \dots & 1 \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 \end{bmatrix}$$

To obtain the above expression, note that for any pair of players, conditional on meeting, if they are in stage $k = 0$, they transit to state 1 with probability $\frac{3}{4}$ and stay in the same state with probability $\frac{1}{4}$. Otherwise, in every meeting, they move to the next state. The transition matrix

H^{J_1} is irreducible, and the limiting distribution is

$$\lim_{q \rightarrow \infty} (1, 0, \dots, 0) \cdot H^q = \left(\frac{4}{3T+4}, \frac{3}{3T+4}, \frac{3}{3T+4}, \dots, \frac{3}{3T+4} \right).$$

Further, it can be easily shown that

$\forall q \geq 3$, $[(1, 0, \dots, 0) \cdot H^q]_3 > 0$ where $[(1, 0, \dots, 0) \cdot H^q]_3$ is the 3rd component of $(1, 0, \dots, 0) \cdot H^q$

It follows that for any rival j whom player i_1 has not met in at least three periods in the past, there is a lower bound $\phi > 0$ such that the probability of j being in the signature period with player i_2 is at least ϕ . Now, when i_1 announces the name i_2 , she does not know which rival she will end up meeting. However, for any $t \geq 5$, player i_1 must assign probability at least ϕ to the event that her rival is supposed to be in a signature period with i_2 . This is because at any $t \geq 5$ there is at least one rival whom i_1 has not met for three periods in the past. Consequently, if she impersonates, there is a minimal strictly positive probability $\phi\epsilon^2$ that her lie gets detected. i_1 will not impersonate i_2 if her maximal gain is outweighed by the minimal expected loss from deviation.

$$\text{Player } i_1 \text{'s maximal current gain from impersonation} = \left(1 - \frac{\delta}{\delta + 2(1 - \delta)} \right) \gamma.$$

Her expected loss in continuation payoff is given by the following expression.

$$\text{Minimal loss from deviation} \geq \phi\epsilon^2(1 - \beta) \left(\frac{\delta}{\delta + 2(1 - \delta)} \right)^T [v_i - ((1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i)].$$

So player i_1 will not impersonate if the following inequality holds.

$$\left(1 - \frac{\delta}{\delta + 2(1 - \delta)} \right) \gamma \leq \phi\epsilon^2(1 - \beta) \left(\frac{\delta}{\delta + 2(1 - \delta)} \right)^T [v_i - ((1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i)].$$

For δ close enough to 1, this inequality is satisfied and misreporting one's name is not a profitable deviation.

Now consider incentives for truth-telling at $t \leq 4$. Suppose player i_1 wants to impersonate player i_2 at $t = 1$. She believes that with probability $\frac{3}{4}$ the game will get initiated in the current period and with probability $\frac{1}{4}$ the rival she meets now (say player j) will meet the true i_2 in the next two calendar time periods. In this case, irrespective of what player i_2 plays at $t = 3$, with probability ϵ , player j will become aware that a deviation occurred. In other words, at $t = 1$, player i_1 believes that with probability $\frac{3\epsilon}{16}$ her deviation will be detected at $t = 3$, and one of her rivals will switch to her unfavorable strategy forever. By a similar argument as above, if δ is high enough, player i_1 's potential current gain from impersonation will be outweighed by the long-term loss in continuation payoff. Similar arguments apply for $t = 2, 3, 4$. \square

6.3 Proof of Folk Theorem for Multilateral Matching

This section contains the formal equilibrium construction for the case of multiple communities.

6.3.1 Structure of Equilibrium

In equilibrium, players all report their names truthfully. Each player plays the prescribed equilibrium strategies separately against each possible playgroup that she can be matched to. On the equilibrium path, players condition play with a particular playgroup only on the history of play vis-à-vis that group of names. It is as if each player is playing separate but identical games with M^{K-1} different playgroups.

T -period Blocks: For any target payoff profile $(v_1, \dots, v_K) \in \text{int}(\mathcal{F}^*)$, we choose an appropriate positive integer T . Play between members of any group of K players proceeds in blocks of T periods. In a block each player i uses one of two action plans of the T -period finitely repeated game. One of the action plans used by a player i ensures that player $(i + 1)$ in that playgroup cannot get more than v_{i+1} , the target payoff for $i + 1$. The other action plan ensures that $i + 1$ gets at least v_{i+1} . We call i the *monitor* of her *successor* $(i + 1)$. (Player M monitors player 1.) In the plan period of a block, each player randomizes between the two action plans so as to achieve the target payoff of her successor in this playgroup. The action profile played in the plan period acts as a coordination device that informs the players of the plan of play for the rest of the block for this group. At the next plan period, each player's continuation payoff is again adjusted by her monitor based on the action profiles played in the last block with that playgroup. Conditional on players reporting their names truthfully we show that the above form of strategies constitute an equilibrium. Impersonations are detected and punished in a similar way as before.

Detecting Impersonations: The second period of a block is designated as the signature period and all players play actions that serve as their signatures. The signature used depends on the action profile realized in the plan period of the block. No player outside the specific K -player group can observe the action in the plan period. Consequently, if anyone outside the playgroup tries to impersonate one of the members, she can end up playing the wrong signature in case it is a signature period, and so get detected.

Community Responsibility: If a player sees an incorrect action or signature, she knows that someone has deviated, though the identity of the deviator or the nature of the deviation is

unknown. (In fact every player in the playgroup knows that a deviation has occurred.) The deviator's entire community can be punished by the relevant monitor. The monitor just switches to the bad action plan with every playgroup in their next plan period. Since every player is indifferent between her two action plans at the start of any block, the relevant monitor can punish her successor's entire community without adversely affecting her own payoff.

6.3.2 Preliminaries

Consider any payoff profile $(v_1, \dots, v_K) \in \text{Int}(\mathcal{F}^*)$. There exist 2^K payoff profiles w^P such that the following conditions hold.

1. $w_i^P > v_i$ if $P_i = G$.
2. $w_i^P < v_i$ if $P_i = B$.

These conditions imply that there exists \underline{v}_i and \bar{v}_i with $v_i^* < \underline{v}_i < v_i < \bar{v}_i$ such that the rectangle $[\underline{v}_1, \bar{v}_1] \times \dots \times [\underline{v}_K, \bar{v}_K]$ is contained in the interior of $\text{conv}(\{w^P : P = (P_1, \dots, P_K), P_i \in \{G, B\}\})$ and further, for all i , $\bar{v}_i < \min\{w_i^P : P_i = G\}$ and $\underline{v}_i > \max\{w_i^P : P_i = B\}$.

Now we can choose finite sequences of pure action profiles $\{a_1^P, \dots, a_N^P\}$, with $P = (P_1, \dots, P_K)$, $P_i \in \{G, B\}$, so that the vectors w^P , the payoffs (average discounted) from the sequence of action profiles $\{a_n^P\}_{n=1}^N$ for any plan profile P satisfy the above relationships. As before, choose $\epsilon \in (0, 1)$ small so that $v_i^* < (1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i < v_i < (1 - \epsilon)\bar{v}_i + \epsilon\underline{v}_i$.

Henceforth, when we refer to an action profile a^P , we actually refer to the finite sequence of action profiles $\{a_1^P, \dots, a_N^P\}$.

6.3.3 Name Announcements at Complete Histories

After any complete history (and the null history), players report their names truthfully.

6.3.4 Actions at Interim Histories

Partitioning of Histories:

At any interim private history, each player i partitions her history into M^{K-1} separate histories corresponding to different games (denoted by $\Gamma_{i, \mathcal{G}_{-i}}$) with each possible playgroup \mathcal{G}_{-i} . If her current playgroup's name profile is \mathcal{G}_{-i} , she plays game $\Gamma_{i, \mathcal{G}_{-i}}$. Fix a player i and a playgroup \mathcal{G}_{-i} . Below, I describe how game $\Gamma_{i, \mathcal{G}_{-i}}$ is played. Let \hat{h}_i^t denote a t -period history in the game $\Gamma_{i, \mathcal{G}_{-i}}$. It specifies the action profiles played in the last t interactions of i with the playgroup \mathcal{G}_{-i} .

Play of Game $\Gamma_{i, \mathcal{G}_{-i}}$:

The game $\Gamma_{i, \mathcal{G}_{-i}}$ between i and playgroup \mathcal{G}_{-i} proceeds in blocks of T periods. In the first period (the plan period) of a block, players take actions which inform their rivals about the plan of play for the rest of the block. Partition the set of player i 's actions into two non-empty subsets G_i and B_i . If player i chooses an action from set G_i , she is said to send plan $P_i = G$. Otherwise she is said to send plan $P_i = B$.

Further, choose any two pure action profiles $g, b \in A$ such that $g_i \neq b_i \forall i \in \{1, \dots, K\}$. Define the signature function $\psi : A \rightarrow A$ mapping one-period histories to action profiles such that,

$$\psi(a) = \begin{cases} g & \text{if } a_i \in G_i \forall i, \\ b & \text{if } a_i \in B_i \forall i. \end{cases}$$

Define $\psi(\cdot)$ arbitrarily otherwise. Suppose the observed plans are (P_1, \dots, P_K) . Let $\tilde{P} = (P_K, P_1, \dots, P_{K-1})$.

Define a set of action plans of a T -period finitely repeated game as follows.

$$\mathcal{S}_i := \left\{ s_i \in S_i^T : \forall \hat{h}_i^t = \left(a, \psi(a), a^{\tilde{P}}, \dots, a^{\tilde{P}} \right), a_{i-1} \in G_{i-1}, s_i[\hat{h}_i^1] = \psi([\hat{h}_i^1]) \text{ and } s_i[\hat{h}_i^t] = a_i^{\tilde{P}} \forall t \geq 1 \right\}.$$

\mathcal{S}_i includes action plans that prescribe playing the correct signature and playing according to the plan announced in the plan period if ones monitor announced a favorable plan G , and everyone in the playgroup used the correct signature and played as per the plan so far. In equilibrium, players use action plans from the above set. Within a block, they use one of two plans s_i^G and s_i^B which are defined below.

Define partially a favorable action plan s_i^G such that

$$\begin{aligned} s_i^G[\emptyset] &\in \Delta(G_i), \\ s_i^G[\hat{h}_i^1] &= \psi_i([\hat{h}_i^1]), \text{ and} \\ \forall \hat{h}_i^t &= \left(a, \psi(a), a^{\tilde{P}}, \dots, a^{\tilde{P}} \right), t \geq 1, s_i^G[\hat{h}_i^t] = a_i^{\tilde{P}}. \end{aligned}$$

We partially define an unfavorable action plan s_i^B such that

$$\begin{aligned} s_i^B[\emptyset] &\in \Delta(B_i), \\ s_i^B[\hat{h}_i^1] &= \psi_i([\hat{h}_i^1]), \\ \forall \hat{h}_i^t &= \left(a, \psi_i(a), a^{\tilde{P}}, \dots, a^{\tilde{P}} \right), t \geq 1, s_i^B[\hat{h}_i^t] = a_i^{\tilde{P}}, \end{aligned}$$

$\forall \hat{h}_i^t$ after $\hat{h}_i^r = (a, \psi(a), a^{\tilde{P}}, \dots, a^{\tilde{P}}, \dots, a^{\tilde{P}}, a')$, with $j : a'_j \neq a_j^{\tilde{P}}, a'_k = a_j^{\tilde{P}} \forall k \neq j, t \geq r > 1$,

$$s_i^B[\hat{h}_i^t] = \alpha_{ji}^*, \quad \text{where } \alpha_{ji}^* \text{ is } i\text{'s action in action profile } \alpha_j^* \text{ which minmaxes player } j, \text{ and}$$

$\forall \hat{h}_i^t$ after $\hat{h}_i^2 = (a, a')$, with $j : a'_j \neq \psi_j(a), a'_k = \psi_k(a) \forall k \neq j, t > 2$,

$$s_i^B[\hat{h}_i^t] = \alpha_{ji}^*, \text{ where } \alpha_{ji}^* \text{ is } i\text{'s action in action profile } \alpha_j^* \text{ which minmaxes player } j.$$

For any history not included in the definitions of s_i^G and s_i^B above, prescribe the actions arbitrarily. Given a plan profile \tilde{P} , these strategies specify $\psi(a)$ and $a^{\tilde{P}}$ until the first unilateral deviation. (In case of simultaneous deviations, these strategies also specify $\psi(a)$ and $a^{\tilde{P}}$.) If a player j unilaterally deviates, then strategy s_i^B specifies that other players in her playgroup minmax her.

Notice that if player i 's monitor ($i-1$) uses strategy s_{i-1}^G , i gets a payoff strictly more than \bar{v}_i in each period, except possibly the first two periods. Further, if i 's monitor plays s_{i-1}^B , player i gets a payoff strictly lower than \underline{v}_i in all except at most two periods. It is therefore possible to choose T large enough so that for some $\underline{\delta} < 1, \forall \delta > \underline{\delta}$, i 's average payoff within the block from any strategy $s_i \in \mathcal{S}_i$ against s_{-i}^G strictly exceeds \bar{v}_i and her average payoff from using any strategy $s_i \in \mathcal{S}_i^T$ against s_{-i}^B is strictly below \underline{v}_i .

Now we define two benchmark action plans which are used to compute continuation payoffs. For any $s_j \in \{s_j^G, s_j^B\}$ define $r_{i+1}^G \in \mathcal{S}_i$ to be an action plan such that given any history \hat{h}_{i+1}^t , $r_{i+1}^G | \hat{h}_{i+1}^t$ gives player $i+1$ the lowest payoffs against s_i^G and s_j for $j \neq i, i+1$ among all action plans in \mathcal{S}_{i+1} . Define $r_{i+1}^B \in \mathcal{S}_i^T$ to be an action plan such that given any history \hat{h}_{i+1}^t , $r_{i+1}^B | \hat{h}_{i+1}^t$ gives the highest payoffs against s_i^B and s_j for $j \neq i, i+1$ among all action plans in \mathcal{S}_{i+1}^T . Redefine \bar{v} and \underline{v} so that $U_{i+1}(r_{i+1}^G, s_i^G) = \bar{v}_{i+1}$ and $U_{i+1}(r_{i+1}^B, s_i^B) = \underline{v}_{i+1}$.

In other words, \bar{v}_i is the lowest payoff player i can get if she uses an action plan in \mathcal{S}_i and her monitor plays her favorable action plan, while \underline{v}_i represents the highest payoff that player i can get irrespective of what she plays when her monitor plays her unfavorable plan.

Partial Strategies: Specifying Play in $\Gamma_{i, \mathcal{G}_{-i}}$

Players play the following strategies in the pairwise games $\Gamma_{i, \mathcal{G}_{-i}}$.

- Players always report their names truthfully.
- Each player plays the following strategies separately against each possible playgroup that she could be in.

- **Initial Period of $\Gamma_{i, \mathcal{G}_{-i}}$:** Player i plays s_i^G with probability μ_0 and s_i^B with probability $(1 - \mu_0)$ where μ_0 solves $v_{i+1} = \mu_0 \bar{v}_{i+1} + (1 - \mu_0) \underline{v}_{i+1}$. Note that since $(1 - \epsilon) \underline{v}_i + \epsilon \bar{v}_i < v_i < \epsilon \underline{v}_i + (1 - \epsilon) \bar{v}_i \forall i$, we will have $\mu_0, 1 - \mu_0 \geq \epsilon$.
- **Plan Period of a Non-Initial Block:** If player i ever observed a deviation in the signature period of an earlier block with any playgroup, she plays s_i^B with probability $(1 - \beta^l)$, where l is the number of deviations she has seen so far and $\beta > 0$ is small. Otherwise, she plays s_i^G with probability μ and s_i^B with probability $(1 - \mu)$ where the mixing probability μ is used to tailor $(i + 1)$'s continuation payoff. For any history \hat{h}_i^T observed (at calendar time c) by i in the last block, specify $(i + 1)$'s continuation payoff as follows. Let c denote the current calendar time period, and let $c(t), t \in \{1, \dots, T\}$ denote the calendar time period of the t^{th} period of the most recently elapsed block.

If s_i^B was played in the last block, we specify the reward $\omega_{i+1}^B(\cdot)$ as

$$\omega_{i+1}^B(\hat{h}_i^T) := \sum_{\tau=1}^T \pi_{\tau}^B$$

where,

$$\pi_{\tau}^B = \begin{cases} \frac{1}{\delta^{T+1-\tau}} \theta_{\tau}^B M^{(K-1)(T+1-\tau)} & \text{if } c - c(\tau) = T + 1 - \tau \\ 0 & \text{otherwise,} \end{cases}$$

and θ_t^B is the difference between $(i + 1)$'s continuation payoff within the last block from playing r_{i+1}^B from time t on and $(i + 1)$'s continuation payoff from playing the action observed by i at period t as in history h_i^t followed by reversion to r_{i+1}^B from $t + 1$ on. Notice that $\theta_t^B \geq 0$. If s_i^B was played in the last block, player i chooses $\mu \in (0, 1)$ to solve $\mu \bar{v}_{i+1} + (1 - \mu) \underline{v}_{i+1} = \underline{v}_{i+1} + (1 - \delta) \omega_{i+1}^B(\hat{h}_i^T)$.

If s_i^G was played in the last block, we specify punishments $\omega_{i+1}^G(\cdot)$ as

$$\omega_{i+1}^G(\hat{h}_i^T) := \sum_{\tau=1}^T \pi_{\tau}^G$$

where,

$$\pi_{\tau}^G = \begin{cases} \frac{1}{\delta^{T+1-\tau}} \min\{0, \theta_{\tau}^B\} M^{(K-1)(T+1-\tau)} & \text{if } c - c(\tau) = T + 1 - \tau \\ 0 & \text{otherwise,} \end{cases}$$

and θ_t^G is the difference between $(i + 1)$'s continuation payoff within the last block

from playing r_{i+1}^G from time t on and $(i+1)$'s continuation payoff from playing the action observed by i at period t as in history \hat{h}_i^t followed by reversion to r_{i+1}^G from $t+1$ on. Note that $\theta_i^G \leq 0$ for all actions that are used by strategies in \mathcal{S}_{i+1} . If s_i^G was played in the last block, player i chooses $\mu \in (0, 1)$ to solve $\mu\bar{v}_{i+1} + (1-\mu)\underline{v}_{i+1} = \bar{v}_{i+1} + (1-\delta)\omega_{i+1}^G(\hat{h}_{-i}^T)$.

We restrict attention to δ close enough to 1 so that

$$(1-\delta)\omega_{i+1}^B(\hat{h}_i^T) < \epsilon\underline{v}_{i+1} + (1-\epsilon)\bar{v}_{i+1} - \underline{v}_{i+1} \text{ and}$$

$$(1-\delta)\omega_{i+1}^G(\hat{h}_i^T) > (1-\epsilon)\underline{v}_{i+1} + \epsilon\bar{v}_{i+1} - \bar{v}_{i+1}.$$

Then, continuation payoffs lie within the interval $[(1-\epsilon)\underline{v}_{i+1} + \epsilon\bar{v}_{i+1}, \epsilon\underline{v}_{i+1} + (1-\epsilon)\bar{v}_{i+1}]$.

- **Signature Periods and other Non-initial Periods:** In signature periods, players use the designated signature $\psi_i(a)$ if a was the profile realized in the plan period. For the rest of the block, they play as per the announced plan.

6.3.5 Beliefs of Players

After every history, players believe that in every period so far, they met the true owners of the names they encountered.

6.3.6 Proof of Theorem 2

Here, we prove optimality on the equilibrium path. Since the proof for consistency of beliefs and sequential rationality off the equilibrium path are identical to the two community case, these proofs are omitted. First we show that conditional on truthful reporting of names, these strategies constitute an equilibrium.

Fix a player i and a rival playgroup \mathcal{G}_{-i} . The partial strategy for player i in her game $\Gamma_{i, \mathcal{G}_{-i}}$ can be represented by an automaton that revises actions and states in every plan period. The following describes the automaton for any player i .

Set of States: The set of states of a player i in a game with a particular playgroup is the set of continuation payoffs for her successor $i+1$ in that playgroup and is the interval $[(1-\epsilon)\underline{v}_i + \epsilon\bar{v}_i, \epsilon\underline{v}_i + (1-\epsilon)\bar{v}_i]$.

Initial State: Player i 's initial state is the target payoff for her successor v_{i+1} .

Decision Function: When i is in state u , she uses action plan s_i^G with probability μ and s_i^B with probability $(1-\mu)$ where μ solves $u = \mu [\epsilon \underline{v}_{i+1} + (1-\epsilon) \bar{v}_{i+1}] + (1-\mu) [(1-\epsilon) \underline{v}_{i+1} + \epsilon \bar{v}_{i+1}]$

Transition Function: For any history \hat{h}_i^T in the last T -period block for player i , if the realized action plan is s_i^G then at the end of the block, the state transits to $\bar{v}_{i+1} + (1-\delta)\omega_{i+1}^G(\hat{h}_i^T)$. If the realized action is s_i^B the new state is $\underline{v}_{i+1} + (1-\delta)\omega_{i+1}^B(\hat{h}_i^T)$.

It can be easily seen that given i 's strategy, any strategy of player $i+1$ whose restriction belongs to \mathcal{S}_{i+1} is a best response. The average payoff within a block from playing r_{i+1}^G against s_i^G is exactly \bar{v}_{i+1} , and that from playing r_{i+1}^B against s_i^B is \underline{v}_i . Moreover, the continuation payoffs are also \bar{v}_{i+1} and \underline{v}_{i+1} respectively. Any player's payoff is therefore $\mu_0 \bar{v}_i + (1-\mu_0) \underline{v}_i$.

Further, as in the case of two communities, each player is indifferent between all possible action plans when her monitor plays the unfavorable action plan. At any stage τ of a block, she believes that with probability $(\frac{1}{M^{K-1}})^{T+1-\tau}$ her next plan period with this playgroup is exactly $T+1-\tau$ calendar time periods away, and in that case, for any action she chooses now she will receive a proportionate reward $\theta_\tau^B M^{(K-1)(T+1-\tau)}$. This makes her indifferent across all action plans in expectation.

It remains to verify that players will truthfully report their names in equilibrium. We show below that if a player impersonates someone else in her community, irrespective of the action she plays, there is a positive probability that her playgroup will become aware that a deviation has occurred. Further, if a deviation is detected, her monitor will punish her whole community (which includes her in particular). For sufficiently patient players this threat is enough to deter impersonation.

At any calendar time t , define the state of play between any player i and any rival playgroup \mathcal{G}_{-i} to be $k \in \{1, \dots, T\}$ where k is the period of the current block they are playing in. At time $(t+1)$, they will either transit to state $k+1$ with probability $\frac{1}{M^{K-1}}$ (if i happens to meet the same playgroup again in the next calendar time period) or remain in state k .

Suppose at time t player i_1 decides to impersonate i_2 . Conditional on her private history, i_1 can form beliefs over the possible states that each of her possible playgroups is in with respect to i_2 . Suppose i_1 has met the sequence of playgroups $\{\mathcal{G}_{-i}^1, \dots, \mathcal{G}_{-i}^{t-1}\}$. She knows that the playgroup she meets in any period remains in the same state with i_2 in that period. Fix any playgroup \mathcal{G}_{-i} whom i_1 can be matched to. Player i_1 has a belief over the possible states \mathcal{G}_{-i} is in with respect to i_2 . Represent i_1 's beliefs over the states by a vector (p_1, \dots, p_n) .

For any $t \geq 2$, her belief over states of \mathcal{G}_{-i} and i_2 is given by

$$\sum_{\tau=1}^{t-1} \left(1 - \mathbb{I}_{\mathcal{G}_{-i} = \mathcal{G}_{-i}^\tau}\right) \left(\frac{M-2}{M-1}\right)^{\sum_{l=1}^{\tau-1} \left(1 - \mathbb{I}_{\mathcal{G}_{-i} = \mathcal{G}_{-i}^l}\right)} \frac{1}{M-1} (1, 0, \dots, 0) \prod_{k=\tau}^{t-1} [\mathbb{I}_{j=j^k} I + (1 - \mathbb{I}_{j=j^k}) H], \quad (6)$$

$$\text{where } H = \begin{bmatrix} \frac{M-2}{M-1} & \frac{1}{M-1} & 0 & 0 & \dots & 0 \\ 0 & \frac{M-2}{M-1} & \frac{1}{M-1} & 0 & \dots & 0 \\ \vdots & & & & & \\ \frac{1}{M-1} & 0 & 0 & 0 & \dots & \frac{M-2}{M-1} \end{bmatrix}$$

$$I \text{ is the } T \times T \text{ identity matrix, and } \mathbb{I}_{\mathcal{G}_{-i} = \mathcal{G}_{-i}^\tau} = \begin{cases} 1 & \text{if } \mathcal{G}_{-i} = \mathcal{G}_{-i}^\tau, \\ 0 & \text{otherwise.} \end{cases}$$

To derive the above expression, note that player i_1 knows that in periods when she met playgroup \mathcal{G}_{-i} it is not possible that i_2 met the same playgroup. Hence in these periods, the state of play between i_2 and \mathcal{G}_{-i} did not change. In other periods the state changed according to the transition matrix H . This leads to the last product term. Now for any calendar period τ , player i_1 can use this information to compute the state of play between i_2 and \mathcal{G}_{-i} conditioning on the event that they met for the first time ever in period τ . For any τ , the probability that i_2 and \mathcal{G}_{-i} met for the first time at period τ is given by $\left(\frac{M-2}{M-1}\right)^{\sum_{l=1}^{\tau-1} \left(1 - \mathbb{I}_{\mathcal{G}_{-i} = \mathcal{G}_{-i}^l}\right)} \frac{1}{M-1}$. Finally player i_1 knows that i_2 and \mathcal{G}_{-i} could not have met for the first time in a period when she herself met playgroup \mathcal{G}_{-i} , and so does not need to condition on such periods.

Notice that the initial state $(1, 0, \dots, 0)$ and H form an irreducible Markov chain with

$$\lim_{q \rightarrow \infty} (1, 0, \dots, 0) \cdot H^q = \left(\frac{1}{T}, \dots, \frac{1}{T}\right). \quad (7)$$

Further it can be easily shown that the following is true.

$$\forall q \geq 1, \quad [(1, 0, \dots, 0) \cdot H^q]_2 > 0, \quad (8)$$

where $[(1, 0, \dots, 0) \cdot H^q]_2$ represents the 2^{nd} component of $(1, 0, \dots, 0) \cdot H^q$.

It follows from (7) and (8) that for any playgroup \mathcal{G}_{-i} whom i_1 has not met at least in one period, there exists a lower bound $\phi > 0$ such that the probability of \mathcal{G}_{-i} being in state 2 with i_2 is at least ϕ .

Now, when i_1 announces name i_2 , she does not know which playgroup she will end up meeting that period. It follows that at $t \geq 2$, player i_1 assigns probability at least $\frac{\phi}{M^{K-1}(M-1)}$ to the event that the rival she meets is in state 2 with i_2 . (To see why, pick a playgroup \mathcal{G}'_{-i} whom i_1 did not meet in the first calendar time period ($t = 1$). With probability $\frac{1}{M^{K-1}}$, at time t , i_1 will meet this \mathcal{G}'_{-i} and with probability $\frac{1}{M-1}$ this \mathcal{G}'_{-i} would have met i_2 at $t = 1$ and period t could be their signature period.)

Consequently, if player i_1 impersonates i_2 , there is a strictly positive probability $\epsilon^K \frac{\phi}{M^{K-1}(M-1)}$ that the impersonation will get detected. This is because if the playgroup she meets is supposed to be in a signature period with i_2 , they should play one of the actions profiles g, b, x, y depending on the realized plan in their plan period. Since players mix with probability at least ϵ on both Plans G and B , with probability at least ϵ^K , i_1 will play the wrong action irrespective of what action she chooses. Her playgroup will be informed of a deviation, and her monitor will switch to the bad plan B with all playgroups in the next respective plan period.

i_1 will not impersonate any other player if her maximal potential gain from deviating is not greater than the minimal expected loss in continuation payoff from detection.¹³

$$i_1\text{'s maximal current gain from misreporting} = \left(1 - \frac{\delta}{\delta + M^{K-1}(1 - \delta)}\right) \gamma.$$

$$\text{Loss in continuation payoff} \geq \frac{\phi}{M^{K-1}(M-1)} \epsilon^K (1-\beta) \left(\frac{\delta}{\delta + M^{K-1}(1 - \delta)}\right)^T [v_i - ((1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i)].$$

To derive the expected loss in continuation payoff, note that there is a minimal probability $\frac{\phi}{M^{K-1}(M-1)}$ that i_2 and playgroup \mathcal{G}_{-i} are in a signature period. Conditional on this event, irrespective of the action played, there is a minimal probability ϵ^K that player i_1 's deviation is detected by playgroup \mathcal{G}_{-i} . Conditional on detection, the relevant monitor will switch to the unfavorable strategy with probability $(1 - \beta)$ in the next plan period with i_1 . At best, this plan period is $T - 1$ periods away, after which player i_1 's payoff will drop from v_1 to $(1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i$. i_1 will not impersonate if the following inequality holds.

$$\left(1 - \frac{\delta}{\delta + M^{K-1}(1 - \delta)}\right) \gamma \leq \frac{\phi}{M^{K-1}(M-1)} \epsilon^K (1-\beta) \left(\frac{\delta}{\delta + M^{K-1}(1 - \delta)}\right)^T [v_i - ((1 - \epsilon)\underline{v}_i + \epsilon\bar{v}_i)].$$

For δ close enough to 1, this inequality is satisfied, and so misreporting one's name is not a profitable deviation. Now consider incentives for truth-telling in the first period of the supergame.

¹³As before, because of the random matching process, the effective discount factor for any player in her pairwise game is not δ , but $\frac{\delta}{\delta + M^{K-1}(1 - \delta)}$.

Suppose i_1 impersonates i_2 at $t = 1$ and meets playgroup \mathcal{G}_{-i} . In the next period, with probability $\frac{\epsilon^K}{M^{k-1}}$, i_2 will meet the same playgroup \mathcal{G}_{-i} and use the wrong signature, thus informing \mathcal{G}_{-i} that someone has deviated. By a similar argument as above, if δ is high enough, i_1 's potential current gain will be outweighed by the future loss in continuation payoff caused by her monitor's punishment. \square

References

- [1] Avinash Dixit. On modes of economic governance. *Econometrica*, 71(2):449–481, 2003.
- [2] Avinash Dixit. Trade expansion and contract enforcement. *Journal of Political Economy*, 111(6):1293–1317, 2003.
- [3] Glenn Ellison. Cooperation in the prisoner’s dilemma with anonymous random matching. *Review of Economic Studies*, 61(3):567–88, July 1994.
- [4] Jeffrey C. Ely, Johannes Hörner, and Wojciech Olszewski. Belief-free equilibria in repeated games. *Econometrica*, 73(2):377–415, 03 2005.
- [5] Jeffrey C. Ely and Juuso Välimäki. A robust folk theorem for the prisoner’s dilemma. *Journal of Economic Theory*, 102(1):84–105, January 2002.
- [6] Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, 1986.
- [7] Avner Greif. Contract enforceability and economic institutions in early trade: the maghribi traders’ coalition. *American Economic Review*, 83(3):525–48, June 1993.
- [8] Avner Greif. *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press, 2006.
- [9] Joseph E Harrington. Cooperation in a one-shot prisoner’s dilemma. *Games and Economic Behavior*, 8:364–377, 1995.
- [10] Johannes Hörner and Wojciech Olszewski. The folk theorem for games with private almost-perfect monitoring. *Econometrica*, 74(6):1499–1544, 2006.
- [11] Michihiro Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59(1):63–80, January 1992.
- [12] Daniel B Klein. Promise keeping in the great society: A model of credit information sharing. *Economics and Politics*, 4:117–136, 1992.
- [13] Masahiro Okuno-Fujiwara and Andrew Postlewaite. Social norms and random matching games. *Games and Economic Behavior*, 9(1):79–109, April 1995.
- [14] Barry R. Weingast Paul R. Milgrom, Douglas C. North. The role of institutions in the revival of trade: The law merchant, private judges, and the champagne fairs. *Economics and Politics*, 2(1):1–23, 1990.

- [15] Robert W. Rosenthal. Sequences of games with varying opponents. *Econometrica*, 47(6):1353–66, November 1979.
- [16] Walter Rudin. *Real and Complex Analysis*. McGraw Hill Publishing Company, 1987.
- [17] Satoru Takahashi. Community enforcement when players observe partners' past play. 2007.