NORTHWESTERN UNIVERSITY

Essays in Economic Theory

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Managerial Economics and Strategy

By

Joyee Deb

EVANSTON, ILLINOIS

June 2008

© Copyright by Joyee Deb 2008

All Rights Reserved

Abstract

Essays in Economic Theory

Joyee Deb

This dissertation comprises three chapters, each focusing on a different question in economic theory. The first two chapters focus on repeated games and reputations, while the third is about large games. In “Cooperation and Community Responsibility”, I study whether cooperation can be sustained between communities where members interact repeatedly but with different people in the community. When communities are large 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, where payers' identities are unobservable and players only observe 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. In “Observability and Sorting in a Market for Names” I study the value of reputations. I ask whether firm names can be tradeable assets when changes in name ownership are observable? Earlier literature suggests that non-observability is critical to tradeable names. Yet, casual empiricism suggests that shifts in name ownership are often observable. I show how firm names can be traded under full observability. In equilibrium, even when consumers see a reputed name being divested they continue trusting it and so, these names are tradeable.

I also demonstrate an appealing "sorting" property. Competent firms separate themselves by buying valuable names, or only incompetent firms use worthless names. In "Large Games of Limited Individual Impact" (co-authored with Ehud Kalai), I study robustness properties of equilibria. Bayesian Nash equilibria that are not ex-post stable are a poor modeling tool for many applications. Earlier literature showed that Bayesian equilibria are ex-post stable in games with a large number of anonymous players, with finite types and actions and continuous payoff functions. These assumptions limit the applicability of the results in important games like market games, location games etc. We identify a broad class of large games that satisfies ex-post stability, without requiring finiteness or anonymity. We show that one regularity condition on payoff functions (a version of Lipschitz continuity) can guarantee ex-post stability.

Acknowledgments

I am indebted to my advisor Johannes Hörner. I learnt the most from him, and have much to thank him for - the countless hours I spent in his office, the many stimulating academic discussions, his (often tiresome!) perfectionism and attention to detail, his excellent work ethic (though I cursed it on Sunday mornings!), his dedicated advising, and his incredible involvement and support at every stage. This dissertation would certainly not exist, but for him!

I am very grateful to Ehud Kalai for his constant attention and guidance. Our weekly Thursday meetings formed an integral part of my life at grad school and they taught me a lot - about large games, economics, academics and life in general! I thank Peter Klibanoff, who was extremely encouraging and has been a constant sounding board for me from my first day at MEDS - answering questions that I didn't know whom else to ask. I am grateful to Jeff Ely who was an excellent resource for me, always ready to think about and discuss problems that I was working on, and surprise me with his insight. I thank Peter Eső for his encouragement and enthusiasm whenever I told him about my work.

There is a formidable list of smart people whom I met in the faculty at MEDS and Economics from whom I learnt a lot - in the classroom, at seminars, in hallway conversations, over coffee, at TG's, and I thank them all. I thank Julio González-Díaz for being an excellent colleague - I thoroughly enjoyed working with him, engaging in academic discussions, and am really grateful for his careful editing of the first chapter of this dissertation. I also thank all my fellow grad students and colleagues - in particular, Siddarth Madhav, Anirudhha Gupta and Sourav Bhattacharya - for patiently sitting through many unpolished presentations and editing unfinished work. I thank my teachers at St. Stephen's College - in particular Geetha Venkataraman - who played an important role in my intellectual development.

I could not have asked for a better set of classmates and friends - Alexei, Anandha, Ani, Basit, Frances, Mustafa, Neha, Nenad, Priscilla, Rinki, Ron, Serra, Siddarth, Sourav, Susan, Tapas - and thank them for making the years of grad school among the most memorable in my life. I will always cherish the late nights (homework or otherwise!), birthday celebrations, innumerable coffee breaks and lunches, hours of academic and (more often!) non-academic debate, adventure sports, board games, and what have you!

I thank my quirky but wonderful family. Words are not enough to thank Piyush who moved half way across the world, and has just always been there for me. I would not have been in grad school but for him. I thank Arjun and Yuliya for their love and relentless confidence in me. Finally, this dissertation is dedicated to my parents - Ma who taught me to live, laugh and love, and Baba who was the kind of intellectual I dream to be.

To my parents.

For love and inspiration.

Contents

List of Figures

CHAPTER 1

Cooperation and Community Responsibility: A Folk Theorem for Repeated Matching Games with Names

1.1. 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]

[1] The 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 1.3.5 for more on this.
[2] The result would extend to environments in which more information can be transmitted.

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 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

[3]In 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.

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, 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.

[4] In establishing the main result, I focus on achieving identical payoffs within a community. See Remark 2 in Section 1.3.3 for a discussion of how this can be generalized.

[5] See 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'."

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 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 deters impersonations and deviations from equilibrium.

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 with $K > 2$ communities, where players from each community are matched 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. Say community 1 is the monitor of community 2. If a player in community 2 deviates, the player from community 1 whom she meets holds community 2 responsible and punishes all members of community 2 that she meets in the future.

1.1.1. 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.

1.1.1.1. 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

[6]As 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).

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.

1.1.1.2. 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 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 1.2 presents the model. In Section 1.3, I establish the Folk Theorem and discuss its key features. In Section 1.4, I extend the result to $K > 2$ communities and multilateral matching. Section 1.5 concludes. Proofs are contained in the appendix.

1.2. Model and Notation

Players: The game is played by two communities of players. Each community I, $I \in \{1, 2\}$ comprises $M > 2$ players[7], say $I := \{I_1, \ldots, I_M\}$. 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 $t \in \{1, 2, \ldots\}$, players are randomly matched into pairs with each member l of Community 1 facing a member $l' := m_t(l)$ 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.

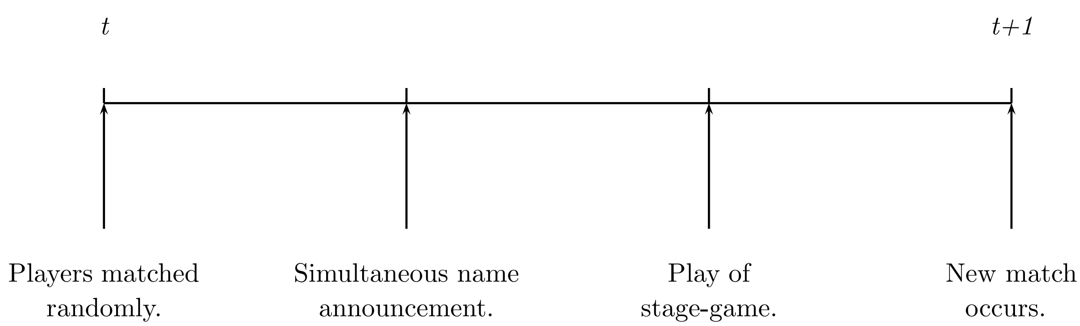

Figure 1.1. Timing of Events

Message Sets: Each community I has a set of messages $\mathscr{N}_I, I \in \{1, 2\}$. Let $\mathscr{N}_I$ be the set of names of players in community I (i.e. $\mathscr{N}_I = \{I_1, \ldots, I_M\}$).[9] For any pair of matched players, the pair of announced messages (names) is denoted by $\nu \in \mathscr{N} := \mathscr{N}_1 \times \mathscr{N}_2$. For any I, let $\Delta(\mathscr{N}_I)$ denote the set of mixtures over messages in $\mathscr{N}_I$. Messages are not verifiable, in the sense that

[7] See Section 1.3.4 for the case $M = 2$.

[8] Unlike in earlier literature, the result does not depend on the matching being uniform or independent over time. See Remark 1 in Section 1.3.3, for a discussion on how this assumption can be relaxed.

[9] An implicit assumption is that the sets of messages $\mathscr{N}_I$ 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.

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 $u : A \to \mathbb{R}^2$. Define $\mathscr{F}$ to be the convex hull of the payoff profiles that can be achieved by pure action profiles in the stage-game. Formally, $\mathscr{F} := \text{conv}\left(\{(u(a) : a \in A\}\right)$. Let v_i^* denote the mixed action minmax value for any player i. 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 $\mathscr{F}^*$ denote the individually rational and feasible payoff set, i.e. $\mathscr{F}^* := \{v \in \mathscr{F} : v_i > v_i^* \ \forall i\}$. We consider games where $\mathscr{F}^*$ has non-empty interior ($\text{Int}\, \mathscr{F}^* \neq \emptyset$).[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* $h_i^t :=$ $\{(\nu^1, a^1), \ldots,$

$(\nu^t, a^t)\}$*, where* (ν^τ, a^τ)*,* $\tau \in \{1, \ldots, t\}$ *represent the name profile and action profile observed by*

[10] Observe 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 $|A_i| \geq 2 \ \forall i$.

player i in period τ. The set of complete private t-period histories is given by $\mathscr{H}_i^t := (\mathscr{N} \times A)^t$. The set of all possible complete private histories for player i is $\mathscr{H}_i := \bigcup_{t=0}^{\infty} \mathscr{H}_i^t$ ($\mathscr{H}_i^0 := \emptyset$).

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

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

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

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

In some abuse of notation, for $k_i \in \mathscr{K}_i$ and $h_i \in \mathscr{H}_i$ 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. $\Sigma^0 := \{\sigma : \forall i, \forall k_i \in \mathscr{K}_i, \forall a_i, \sigma_i(a_i|k_i) > 0 \text{ and } \forall i, \forall h_i \in \mathscr{H}_i, \forall \nu_i, \sigma_i(\nu_i|h_i) > 0\}$. 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**,*

$$\forall i, \forall t, \forall h_i^t \in \mathscr{H}_i^t, \forall \sigma_i', \qquad U_i(\sigma^*|h_i^t, \xi_i^*[h_i^t]) \geq U_i(\sigma_i', \sigma_{-i}^*|h_i^t, \xi_i^*[h_i^t]),$$
$$\forall i, \forall t, \forall k_i^t \in \mathscr{K}_i^t, \forall \sigma_i', \qquad U_i(\sigma^*|k_i^t, \xi_i^*[k_i^t]) \geq U_i(\sigma_i', \sigma_{-i}^*|k_i^t, \xi_i^*[k_i^t]),$$

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.*

1.3. The Main Result

Theorem 1. ***(Folk Theorem for Random Matching Games)*** *Consider a finite two-player game and any $(v_1, v_2) \in \operatorname{Int} \mathscr{F}^*$. 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.

1.3.1. 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.

1.3.1.1. T-period Blocks. Let $(v_1, v_2) \in \operatorname{Int} \mathscr{F}^*$ 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, we need a device that enables players to detect impersonations and provides incentives to a detector to punish them.

1.3.1.2. 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.

1.3.1.3. 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.

1.3.2. Construction of Equilibrium Strategies

Consider any payoff profile $(v_1, v_2) \in \text{Int}\,\mathscr{F}^*$. 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 $\bar{v}_1 < \min\{w_1^{GG}, w_1^{GB}\}$, $\bar{v}_2 < \min\{w_1^{GG}, w_1^{BG}\}$, $\underline{v}_1 > \max\{w_1^{BB}, w_1^{BG}\}$ and $\underline{v}_2 > \max\{w_1^{BB}, w_1^{GB}\}$. See Figure 2 below for a pictorial representation.

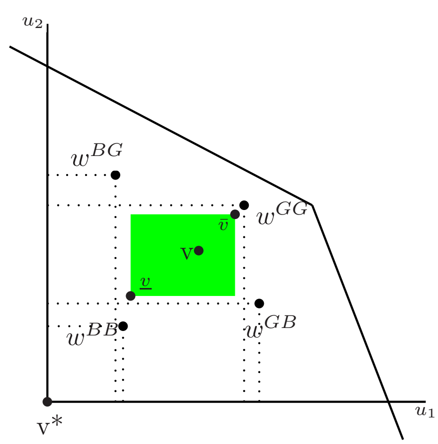

Figure 1.2. Payoff Profiles

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.

1.3.2.1. Defining Strategies at Complete Histories: Name Announcements. At any complete private history, players announce their names truthfully.

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

1.3.2.2. 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 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 $\hat{h}^t_{i,-i}$ in the pairwise game $\Gamma_{i,-i}$ by $\hat{h}^t_i$. Recall that a t-period history denoted by $\hat{h}^t_i$ 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 $\hat{h}^t_i$ of $\Gamma_{i,-i}$ has the same name profile in each period, we ignore names when specifying how $\Gamma_{i,-i}$ is played.

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 $\psi : A \to \{g, b, x, y\}$ (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.

$$\mathscr{S}_i := \Big\{ s_i \in S^T_i : \forall \hat{h}^t_i = \Big(a, \psi(a), (a^{P_2,P_1}_i, a^{P_2,P_1}_{-i}), \ldots, (a^{P_2,P_1}_i, a^{P_2,P_1}_{-i}) \Big), a \in P_i \times G,$$

$$s_i[\hat{h}^1_i] = \psi_i([\hat{h}^1_i]) \text{ and } s_i[\hat{h}^t_i] = a^{P_2,P_1}_i, t \geq 2 \Big\}.$$

Note that the set of action plans in $\mathscr{S}_i$ restricts player i's actions if her rival announced plan G. In particular, action plans in $\mathscr{S}_i$ prescribe that player i use the correct signature and play a^{P_2,P_1}_i

[11] A 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.

if the announced plans were (P_1, P_2). $\mathscr{S}_i$ 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 choose action plans from $\mathscr{S}_i$. Players will use in fact one of two actions plans from $\mathscr{S}_i$, 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

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

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

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

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

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

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

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

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

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

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

Note that both action plans s_i^G and s_i^B belong to $\mathscr{S}_i$. s_i^G is an action plan in $\mathscr{S}_i$ 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 $s_i \in \mathscr{S}_i$ against s_{-i}^G strictly exceeds $\bar{v}_1$ and her average payoff from using any action plan $s_i \in S_i^T$ 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 $r_i^G \in \mathscr{S}_i$ 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 $\mathscr{S}_i$. Define $r_i^B \in S_i^T$ 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 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 $U_i^T : S_i^T \times S_{-i}^T \to \mathbb{R}$ 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 $(1-\mu)$ 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 $\mathscr{S}_i$ 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 $\mathscr{S}_i$ 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 $c(\tau)$, $\tau \in \{1, \ldots, T\}$ 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

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

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

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

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

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

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 $\mathscr{S}_{-i}$.

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

$$(1-\delta)\omega^B_{-i}(\hat{h}^T_i) < \epsilon\underline{v}_{-i} + (1-\epsilon)\bar{v}_{-i} - \underline{v}_{-i} \text{ and } (1-\delta)\omega^G_{-i}(\hat{h}^T_i) > (1-\epsilon)\underline{v}_{-i} + \epsilon\bar{v}_{-i} - \bar{v}_{-i}.$$

- **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.

1.3.2.3. Beliefs of Players. At any private history, each player believes that in every period, she met the true owners of the names she met, and that no player has ever misreported her name.

1.3.3. 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^G_i with probability μ and s^B_i

with probability $(1-\mu)$ where μ solves

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

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 $\mathscr{S}_{-i}$ 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 $k \in \{1,\ldots,T\}$ 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 $j, j \in \{1, \ldots, M\}$ 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 $(p_1, \ldots, p_n)$.

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 $\{j^1, \ldots, j^{t-1}\}$. For any $t \geq 2$, her belief over states of j and i_2 is given by

$$\sum_{\tau=1}^{t-1} (1 - \mathbb{I}_{j=j^\tau}) \left(\frac{M-2}{M-1}\right)^{\sum_{l=1}^{\tau-1}\left(1-\mathbb{I}_{j=j^l}\right)} \frac{1}{M-1}(1, 0, \ldots, 0) \prod_{k=\tau}^{t-1} \left[\mathbb{I}_{j=j^k} I + (1 - \mathbb{I}_{j=j^k}) H\right], \tag{1.1}$$

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

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

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 $\left(\frac{M-2}{M-1}\right)^{\sum_{l=1}^{\tau-1}\left(1-\mathbb{I}_{j=j^l}\right)} \frac{1}{M-1}$. 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

$$\lim_{q\to\infty}(1,0,\ldots,0)\cdot H^q = \left(\frac{1}{T},\ldots,\frac{1}{T}\right). \tag{1.2}$$

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

$$\forall q \geq 1, \qquad [(1,0,\ldots,0)\cdot H^q]_2 > 0, \tag{1.3}$$

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

It follows from (1.2) and (1.3) 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

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

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.

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

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. □

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 asymmetric payoff profiles* $(v_{i_1}, \ldots, v_{i_M}, v_{j_1}, \ldots, v_{j_M})$ *with the property that for all possible pairs of rivals* i *and* j, $(v_i, v_j) \in int(\mathscr{F}^*)$. *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.*

1.3.4. 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 players 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.

1.3.4.1. 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.

1.3.5. 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.

1.4. 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 for sufficiently patient players. This section establishes that the result generalizes to situations with random multilateral matching where $K > 2$ communities interact. 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.

1.4.1. Model and Result

Multilateral Matching: There are K communities of agents with $M > 2$ members in each community I, $I \in \{1, \ldots, K\}$. In each period $t \in \{1, 2, \ldots\}$, agents are randomly matched into groups of K members each, with one member from each community. Let $\mathscr{G}_{-k}$ denote a group of $(K-1)$ players with members from all except the k^{th} community. Let $m_t(\mathscr{G}_{-k})$ denote the member of the k^{th} community who is matched to the group $\mathscr{G}_{-k}$. Matching is independent and uniform, i.e. $\forall$ histories, $\forall j \in$ community $k, \Pr[j = m_t(G_{-k})] = \frac{1}{M}$. For any player i, the set of rivals she is matched with (say $\mathscr{G}_{-i}$) is said to constitute her *playgroup*. After being matched, players announce their names. 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 $\mathscr{N}_I : I \in \{1, \ldots, K\}$ with M names each. A name profile of a playgroup is denoted by $\nu \in \mathscr{N} := \mathscr{N}_1 \times \ldots \times \mathscr{N}_K$. Let $\Delta(\mathscr{N}_I)$ denote the set of mixtures of messages in $\mathscr{N}_I$. The stage-game Γ has finite action sets $A_I, I \in \{1, \ldots, K\}$. 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 $u : A \to \mathbb{R}^K$. Define $\mathscr{F}$ to be the convex hull of the payoff profiles that can be achieved by pure action profiles in the stage-game. Formally, $\mathscr{F} := \text{conv}\left(\{u(a) : a \in A\}\right)$. Denote the feasible and individually rational payoff set by $\mathscr{F}^* := \{v \in \mathscr{F} : v_i > v_i^* \ \forall i\}$, where v_i^* is the mixed action minmax value for player i. We consider games where $\mathscr{F}^*$ has non-empty interior ($Int\mathscr{F}^* \neq \emptyset$). Let γ be defined as before. All players have 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 1.2.

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 $(v_1, \ldots, v_K) \in Int(\mathscr{F}^*)$, there exists a sequential equilibrium that achieves payoffs $(v_1, \ldots, v_K)$ 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 construction and proof of Theorem 2 are relegated to the appendix.

1.5. 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.

CHAPTER 2

Observability and Sorting in a Market for Names

2.1. Introduction

While describing the recent acquisition of IBM's ThinkPad name by computer manufacturer Lenovo, the New York Times wrote:[1]

> When Lenovo, the Chinese personal computer maker, bought I.B.M.'s personal computing business for USD 1.75 billion in December 2004... Lenovo executives assumed rightly that the I.B.M. brand would still resonate in the United States market and serve to assuage the worries of existing and prospective customers about the I.B.M. ThinkPad line of laptops. Lenovo also realized there would be concern among American customers about buying from a China-based company they had never heard of.

This situation highlights two important phenomena related to firm reputations.

First, firms may publicly buy and sell their names like other valuable tradeable assets. As in the case of Lenovo and ThinkPad, the sale of a well-established name may be public because it is covered widely by the business press. Changes of firm ownership may be publicly known because disclosure is mandated by law. Even when not mandated by law, we see that a new owner may choose to make it known - haven't we seen local restaurants announce "Under New Management"?

Second, the market for firm names can exhibit a sorting property, in the sense that well-established names are usually bought by good firms. Consider the example of 'Waterman' a

[1]See "Quickly Erasing I and B and M" by Glenn Rifkin & Jenna Smith, *New York Times*, April 12, 2006

famous premium brand of pens which changed ownership multiple times. Each time the brand was sold, it was bought by a well-established firm, once by Gillette in 1992 and more recently by Rubbermaid. IBM was bought by Lenovo the largest PC manufacturer in China. When Nabisco sold its well-known 'Shredded Wheat' cereal brand, the potential buyers were trusted companies, Kraft and General Mills. In fact it must be such a sorting property that enables consumers to trust a name even after they know it has been sold. In the Lenovo example, even after the ThinkPad name was sold, it appears that consumers continued to trust and buy it. They must have believed that a firm capable of buying a name such as ThinkPad was likely to be "good", and would continue to provide the same quality of products and services.[2]

The existing theory on firm reputations does not explain either of these two phenomena. On the contrary, the standard models (e.g. Kreps 1990, Tadelis 1999) have two opposite features. First, non-observability of name trading is a key assumption. The main result in this literature is that names are traded in all equilibria, but this result relies critically on the non-observability assumption which is the source of value for names. Consumers believe that the current owner is responsible for the good name or record of the firm, and good past record generates expectation of good future performance. Good names become tradeable because a firm can secretly buy a good name and create expectation of good performance and earn higher revenues.

The second feature of current models is the absence of sorting equilibria. A key result is that the trading equilibria are all pooling equilibria. It is not possible for good firms to separate themselves from the bad firms by buying valuable names. There are always some bad firms using valuable names in equilibrium.

This brings us to the two main questions addressed in this paper. First, can we do away with the assumption of non-observability and develop a theory to explain why firm names are valuable even when clients can observe changes in name ownership? Second, under what conditions and how can the market for firm names separate good firms from bad ones?

[2]See "Brands Still Easier to Buy than Create" by Kenneth N. Gilpin, *New York Times*, September 14, 1992.

I consider an infinite horizon economy with generations of firms and consumers interacting in each period. Both firms and consumers live for only one period. Consumers are homogeneous and are on the long side of the market. Firms are of two types - competent and incompetent. Competent firms can choose to work hard or be lazy. Working hard is costly but likely to result in good quality products or services. Laziness always results in bad quality. Incompetent firms are always lazy and so incapable of producing good quality. Consumers buy a product or service from the firm for which they pay upfront. At the time of purchase, consumers do not see the type of the firm or the quality of the product. They only see the name of the firm, and must pay a wage based on the name.

Firms choose to appear in the market under different names which they buy in a competitive market for names. The basic intuition of the model can be understood using two types of names. So, for most part of the paper, I focus on this case. Firms can choose to enter with new names (N) or successful names (S). Entering with an N-name is cheaper than entering with an S-name. After collecting the wages, competent firms can choose to work hard or be lazy. At the end of the period, each firm's reputation or name[3] changes based on the quality of products it has provided. Names evolve according to a fixed transition rule (potentially random) which specifies a firm's reputation at the end of the period, based on its original reputation and the quality of products provided. Before retiring, a firm can sell its reputation to a new entrant.

In this environment, the existence of a market for names affects incentives of firms in two ways. First, it influences the effort choice of the firm. The continuation payoff from selling a valuable name can make firms work hard to produce good quality. Also, the market may give firms incentives to buy one name rather than another one. For instance, a firm may choose to buy a costly name because consumers pay higher wages for it, or because a costly name is more likely to remain good and gives a higher continuation payoff.

[3]For the rest of the paper, I use "firm reputation" and "firm name" interchangeably.

It is worth noting here why this setting is particularly well-suited to analyze the value of firm names when name trading is observable. In a departure from the literature which always considers overlapping generations of firms, I consider a model where firms and consumers are all short-lived. Therefore it is common knowledge that name ownership changes every period. When a consumer meets a firm, she knows that the firm bought its name before entering.

With full observability of name changes, I examine the existence of equilibria in which the market for names both causes competent firms to work hard and sorts firms according to their type. I define a class of equilibria called sorting high-effort equilibria (SHE) where at least one type of firm has a strict incentive to not use one of the two names, and competent firms always work hard. In the main result of the paper, I characterize necessary and sufficient conditions for the existence of such sorting high-effort equilibria (SHE). I show that SHE exist provided the cost of effort is low enough for competent firms.

Two kinds of sorting arise in equilibrium. It is possible for competent firms to separate themselves by being the only ones buying the valuable successful names. I call these situations "Trust S-Names equilibria". In these equilibria, when consumers see a successful name, they trust it to be a competent firm and pay the corresponding high price. The higher wages for an S-name provide firms incentives to buy these names. Higher continuation payoffs from an S reputation - provided the same reputation is maintained - give competent firms the incentive to work hard, and guarantee that incompetent firms do not find it worthwhile to buy these names. Note that there is effort exertion by competent firms in equilibrium even though they are sorted out from the incompetent firms. In other words, here, unlike in earlier models, moral hazard suffices to provide incentives to exert effort, even in the absence of adverse selection.

The second type of sorting that arises is termed "Mistrust N-Names". Here, incompetent firms give themselves away by being the only firms using the cheap names. Consumers treat cheap names with mistrust and pay them corresponding low wages. Competent firms force this situation to arise by always buying S-names before entering the market. Note that both types of sorting allow firms to still pool on one name. In "Trust S-Names equilibria", S-names separate competent firms from incompetent ones. But some competent firms may still pool with incompetent firms on N-names. Similarly, in "Mistrust N-Names" equilibria, firms separate with N-names but may pool on S-names. In fact it will be shown that, generically, pooling on one name is required for the existence of high-effort equilibria in this model.

The type of sorting high-effort equilibria that arises depends on the transition rule posited. I show that for deterministic transition rules, the only type of sorting sustainable in a market with two names is the "Mistrust N-Names" type. Equilibria with deterministic rules have the appealing feature that they remain equilibria even in richer information structures, for instance when consumers observe not just a name but the full history of outcomes of a particular name. With random transition rules, both types of sorting can arise in equilibrium, as long as the cost of hard work is low enough. I characterize the transition rules that give rise to each of the two types of sorting.[4]

I examine the relationship between observability and separation. It turns out that under observability, the requirement for separation in equilibrium is not a restrictive one. Relaxing the requirement for sorting does not extend the range of parameter values under which high-effort equilibria exist.

Finally in an extension of the model, I study the welfare implications of observability of name trades. Since I restrict attention to high-effort equilibria, the total surplus of firms and consumers is constant across regimes. However, observability does affect consumer and firm surplus. I present examples where observability is irrelevant, observability makes consumers better off and worse off.

[4]Sorting high-effort equilibria that arise with random transition rules are perhaps less appealing because they fail to remain equilibria if consumers actually observed the full history of outcomes of a firm name. However, I conjecture that a two-state market with random transitions may be equivalently represented by a richer market with more states and deterministic transitions.

The rest of the paper is structured as follows: Section 2 describes the basic model. Section 3 presents two simple examples where firm ownership changes are observable and sorting high-effort equilibria exist. In Section 4, I consider a market with more general transition rules and characterize the necessary and sufficient conditions under which sorting high-effort equilibria exist. Section 5 presents an extension which allows us to compare regimes with and without observability. Section 6 discusses the relationship between the results in this paper and those in related literature. Section 7 concludes. Proofs of most results are in the appendix.

2.2. The Basic Model

There is a continuum of firms of unit measure. Firms live for one period only. In each period, they meet clients who are also short-lived. Consumers are homogeneous and are on the long side of the market. The following stage game is played by firms and consumers each period. Assume that the same play has occurred forever into the past. At the beginning of each period, the firm enters the market. At the time of entry, it has to choose a name for itself. After entry, the firm meets one client who pays it an upfront wage for its service. All a client sees at the time of purchase is the name of that particular firm. So, the wage paid depends only on the observed name. After collecting its wage, the firm makes an action choice. It has two choices: work hard (H) or be lazy (L). Being lazy is costless while working hard involves a cost $c > 0$. There are two possible outcomes that can arise from the action chosen - good (G) or bad (B). The probability of a good outcome given hard work is $(1 - \rho)$ with $\rho \in (0, 1)$. If the firm is lazy, a bad outcome occurs with probability 1. Firms are of two types, competent (C) or incompetent (I). A proportion $\phi \in (0, 1)$ is competent. A competent firm can choose to work hard or be lazy. An incompetent firm is incapable of working hard. After the firm takes its action, the outcome occurs, and the firm's name changes based on a pre-determined transition rule. The firm can sell its (changed) name before retiring. Figure 1 describes the sequence of events in a period.

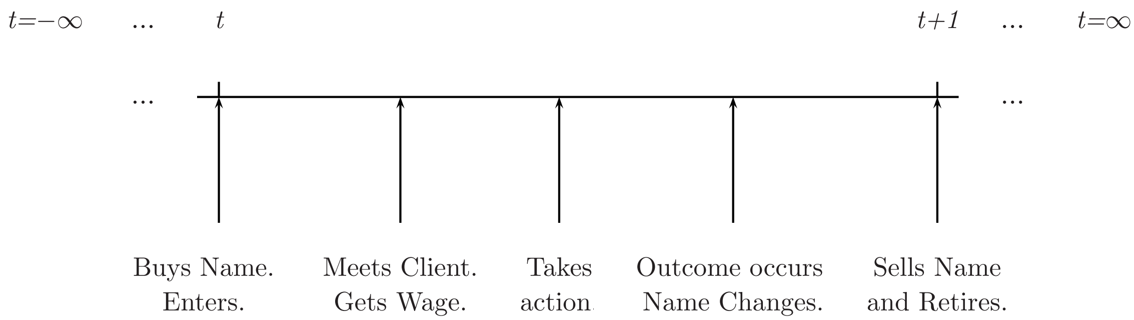

Figure 2.1. Timing of Events within a Period

2.2.1. The Market for Firm Names

In the market for firm names, there are two kinds of names, S and N. By convention, I call the cheaper name N and the more expensive name S. The market for names is competitive and names trade at prices V_S and V_N respectively. An entrant can buy (and an exiting firm can sell) a name of its choice from a retiring firm (to an entrant) at these prices. Without loss of generality, normalize the price of the cheaper name V_N to 0. I also assume that there is no shortage of the cheaper name. Think of entering with an N-name to be equivalent to entering with no reputation at all. The market is characterized by a transition rule. Define a non-random transition rule as follows.

Definition 8. *A* ***non-random transition rule*** *is a function* $f : \{S, N\} \times \{G, B\} \rightarrow \{S, N\}$.

A transition rule can be thought of as a device that describes how a name evolves based on the performance of its owner. As mentioned earlier, at the end of the period, the firm sells its name if possible, to a new entrant and then retires. Later, I generalize the analysis by allowing random transitions rules, which specify not a deterministic transition but a distribution over future names based on the original name and the outcome.

The market for names affects the incentives of firms in two ways. First, it influences the effort choice of the competent type. The expected continuation payoff from selling a valuable name can make competent firms exert effort, even though they are short-lived. Second, the market may give firms incentives to buy one name rather than another. For instance, a firm may choose to buy a costly 'good' name because consumers pay higher wages for 'good' names, or because the transition rule is such that a 'good' name gives a higher continuation payoff.

Notice that changes in name ownership are trivially observable in this setting. When consumers see a name, they know that the firm bought this name at the time of entry.

2.2.2. Payoffs

Firms have a discount factor of δ. The net payoff to a firm consists of the wages it receives plus the discounted proceeds from selling its name less the price it pays to buy its name, less the cost of effort. Clients get utility 0 from a bad outcome and utility 1 from a good one. Since clients can observe only the firm name and cannot observe the outcome at the time of payment, they pay firms wages equal to their expected utility conditional on the observed name, given the firms' strategies. Denote the wages conditional on the name by w_S and w_N. Since clients make no real decision in the game, they are not explicitly modeled as players in what follows.

2.2.3. Definition of Equilibrium

In this paper, we consider simple Markovian equilibria. An incompetent firm's strategy, denoted by μ_S, specifies the probability with which it chooses an S-name. A competent firm's strategy, denoted by (σ_S, e_S, e_N), specifies the probability with which it chooses an S-name and a probability of working hard conditional on each name.

Definition 9. *A **steady-state equilibrium** consists of strategies of firms and a price of an S-name, V_S such that*

1. *The strategies are optimal for the firms (given the wages), and*
2. *Demand equals supply in the market for S-names at price V_S.*

We do not require a market clearing condition for N-names which are in unlimited supply. We want to find equilibria in which the market for names solves the moral hazard problem and sorts competent and incompetent firms. So we define a sorting high effort equilibrium as follows.

Definition 10. *A **sorting high-effort equilibrium (SHE)** is a steady-state equilibrium in which*

1. *There exists at least one name that is chosen by one type of firm and not by the other,*
2. *Each firm strictly prefers a name it chooses to one that it does not choose,*
3. *Competent firms strictly prefer to work hard on the equilibrium path.*

This definition requires partial sorting. It turns out that the set of parameters for which SHE exist with full sorting is non-generic. In this model, we cannot apply the standard repeated game arguments to derive the values of names that can sustain high effort. Here, the continuation payoff or the values of names cannot be chosen arbitrarily but must satisfy the market clearing conditions in the market for names.

2.3. Example of SHE

In this section, I discuss two examples of sorting high-effort equilibria. These examples demonstrate different reputational effects. In the first, incompetent firms give themselves away by being the only firms entering with cheap N-names. I call these SHE "Mistrust Cheap Names" or "Mistrust N-Names" equilibria.

Definition 11. *A **"Mistrust N-Names"** or "Mistrust Cheap Names" equilibrium is a SHE where the incompetent type is the only one using cheap N-names. (S-names are used by both types)*

In these equilibria, S-names are used by both kinds of firms. So incompetent firms are actually indifferent between N and S-names. On the other hand, competent firms who can work hard can get a better continuation payoff with an S-name, and so strictly prefer to use S-names. As a result, clients treat cheap names with mistrust and pay low (zero) wages to a firm with a cheap name. I show that in a market with two names and non-random transition rules, this is the only kind of sorting that can be sustained. In the second example, I present a "Trust Expensive Names" equilibrium, where only copetent firms buy expensive names.

Definition 12. *A* ***"Trust S-Names"*** *or "Trust Expensive Names" equilibrium is a SHE where the competent type is the only one using the valuable S-names. (N-names are used by both types)*

To demonstrate this effect, we consider a market with three kinds of firm names. (These equilibria do not exist with only two names.) There is an expensive name that only competent firms buy. The other names are used by both types of firms. In equilibrium, clients know this and when they see this expensive name, they pay the highest possible wages. Since incompetent firms are not capable of getting good outcomes, they do not find it worthwhile to pay the high price of the expensive name.

2.3.1. "Mistrust Cheap Names" Equilibria: An Example

Consider the transition rules represented by the two automata in Figure 2. We will see that with these transition rules, "Mistrust Cheap Names" equilibria can be sustained. In equilibrium, incompetent firms will reveal their type by being the only firms with cheap N-names.

To see why, observe that under both these rules, conditional on a bad outcome N-names have strictly better continuation payoffs than S-names. So, I-firms are willing to pay for S-names, a price V_S which is the difference between the increase in wages from using an S-name and the decrease in the continuation payoff from using it. At this price, C-firms choose to buy S-names and work hard. Though they can be lazy, use N-names and still earn a higher continuation payoff,

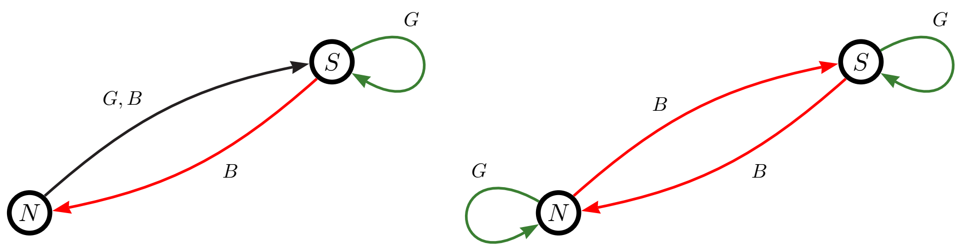

Figure 2.2. Transitions for "Mistrust Cheap Names" equilibria

they choose not to do so, as the increase in continuation payoffs from using an N-name does not compensate for the decrease in wages that they suffer from using it.

So, competent firms signal their competence by buying S-names and then work hard to get a high continuation payoff. Incompetent firms either use a costless name and get a high continuation payoff, or they impersonate competent firms and earn high wages but lose out on continuation payoff. Proposition 1 formalizes this intuition.

Lemma 1. *SHE must have C types using only S-names.*

This observation enables us to prove that the above example is the only SHE sustainable with non-random transition rules and two names. First, I introduce some notation. Any non-random transition rule f can be represented by a vector of zero's and one's (f_1, f_2, f_3, f_4), where

$$f(S,G) = \begin{cases} S & \text{if } f_1{=}1 \\ N & \text{otherwise} \end{cases} \qquad f(S,B) = \begin{cases} S & \text{if } f_2{=}1 \\ N & \text{otherwise} \end{cases}$$

$$f(N,G) = \begin{cases} S & \text{if } f_3{=}1 \\ N & \text{otherwise} \end{cases} \qquad f(N,B) = \begin{cases} S & \text{if } f_4{=}1 \\ N & \text{otherwise.} \end{cases}$$

First consider equilibria in which C-firms use both S and N-names. For C-firm to be indifferent between the two names, the following must be true:

$$-V_S + w_S - c + \delta(1-\rho)f_1 V_S + \delta\rho f_2 V_S = w_N - c + \delta(1-\rho)f_3 V_S + \delta\rho f_4 V_S. \tag{2.1}$$

For the C-type to prefer working hard to being lazy on the equilibrium path, we need

$$c < \delta(1-\rho)(f_1 - f_2)V_S \qquad \text{and} \qquad c < \delta(1-\rho)(f_3 - f_4)V_S. \tag{2.2}$$

For (2.2) to hold, both S and N-names must transition to an S-name after a good outcome and to an N-name after a bad one. (i.e. $f_1 = 1$, $f_2 = 0$, $f_3 = 1$ and $f_4 = 0$). But by (2.1), this implies

$$V_S = w_S - w_N. \tag{2.3}$$

A sorting equilibrium requires that I-firms play S or N but not both. So, one of the following must be true:

$$-V_S + w_S + \delta f_2 V_S < (>) w_N + \delta f_4 V_S.$$

Since $f_2 = f_4 = 0$, this implies that $V_S > (<) w_S - w_N$, contradicting (2.3). Next, consider sorting equilibria in which the C-type uses only N-names. Then, the I-type must use S-names. So the following must be true:

$$V_S \leq w_S - w_N + \delta V_S(f_2 - f_4).$$

But, in this equilibrium, $w_S = 0 \leq w_N$. This implies that $V_S \leq 0$, which is not possible. If $V_S \leq 0$, this would destroy incentives for C-firms to work hard in any state. This concludes the proof.

The intuition for the second part is that since only I-types buy S-names, there is no benefit in wages from an S-name. I-types buy S-names only because an S-name gives a better (weakly) continuation payoff than an N-name in case of a bad outcome. Therefore, at best an I-firm gets a continuation payoff of δV_S, but must pay V_S for it. This will leave the I-firm with a non-positive net payoff. So it cannot be that incompetent firms buy S-names. This contradicts the hypothesis. So, in equilibrium C-types can play only S-names.

Proposition 1. *Sorting high-effort equilibria exist in the market with non-random transitions if and only if*

$$\phi \leq \frac{1}{1+\rho} \text{ and } c < \frac{2\delta\phi(1-\rho)^2}{(1+\delta)(1+\phi-\phi\rho)}.$$

The equilibria can be characterized as follows:

1. Competent firms buy only S-names.

2. Incompetent firms buy S-names with probability $\mu_S = \frac{1}{2} - \frac{\rho\phi}{2(1-\phi)}$.

3. Firms earn wages that equal the expected utility to the consumer conditional on the firm's name, $w_S = \frac{\phi(1-\rho)}{\phi+(1-\phi)\mu_S}$ and $w_N = 0$.

4. S-names trade at price $V_S = \frac{w_S}{1+\delta}$.

The interested reader may refer to the appendix for the proof of Proposition 1. Notice that if the proportion of competent types is very high (i.e. for $\phi > \frac{1}{1+\rho}$), there will be a shortage of S-names in the market. There will be a high demand for S-names but there will be insufficient incompetent firms creating S-names. So the market for S-names cannot clear.

When $\phi < \frac{1}{1+\rho}$, the equilibria characterized above are partially sorting equilibria in that S-names are bought by both types of firms, but only I-firms use N-names. At the knife-edge case $\phi = \frac{1}{1+\rho}$, full separation can be sustained, where C-firms use S-names and I-firms use N-names.

An important feature of this equilibrium is that it remains an equilibrium even if consumers could observe not only the name of the firm but also the complete history of outcomes of the name. This is an appealing property because it implies that the existence of the SHE above is not dependent on the implementation of any specific transition rule. (Refer to Section 4.4.1 for further discussion of this property.)

2.3.2. "Trust Expensive Names" Equilibria: An Example

Proposition 1 establishes that the only SHE that exist with two names and non-random transition rules are of the "Mistrust Cheap Names" type. The next example therefore uses a richer market with three possible firm names to demonstrate "Trust Expensive Names" equilibria under non-random transition rules. There is a valuable firm name that is bought only by competent firms. The other names are bought by both types of firms. Competent firms work hard in equilibrium.

2.3.2.1. A Richer Market for Names. Consider a richer market for names with three names: S_1, S_2 and S_3. Without loss of generality, let S_3 be the cheapest name with its price normalized to 0. Denote the wages conditional on the name by w_1, w_2 and w_3. I restrict attention again to non-random transition rules which are functions $f : \{S_1, S_2, S_3\} \times \{G, B\} \to \{S_1, S_2, S_3\}$. An incompetent firm's strategy, denoted by (μ_1, μ_2, μ_3), specifies the probability with which it chooses S_1, S_2 and S_3-names respectively. A competent firm's strategy, denoted by $(\sigma_1, \sigma_2, \sigma_3, e_1, e_2, e_3)$, specifies the probabilities with which it chooses S_1, S_2 and S_3-names respectively, and the probability of working hard conditional on each name. The definitions of equilibrium and sorting high-effort equilibrium can be extended to this richer environment in a natural way.

Definition 13. *A **steady-state equilibrium** for a given transition rule consists of strategies of firms and prices of names V_1 and V_2 such that*

1. *The strategies are optimal for the firm (given the transition rule and the wages), and*
2. *Demand equals supply in the markets for S_1 and S_2-names at prices V_1 and V_2 respectively.*

Definition 14. *A **sorting high-effort equilibrium (SHE)** is a steady-state equilibrium in which*

1. *There exists at least one name that is chosen by one type of firm and not by the other,*
2. *Each firm strictly prefers the names it chooses to those it does not choose, and*
3. *Competent firms choose to work hard on the equilibrium path.*

Note that I assume that firms work hard on and off the equilibrium path. This is just a convenient assumption. If this restriction were dropped, SHE would exist for a larger set of parameters.

We can extend the definition of "Trust Expensive Names" equilibria in a similar way.

Definition 15. *A* ***"Trust Expensive Names"*** *equilibrium is a SHE in which there is some valuable name that is bought only by competent firms. (The other names may be bought by both types.)*

2.3.2.2. Example. Consider transition rules given by the automaton in Figure 3 below.

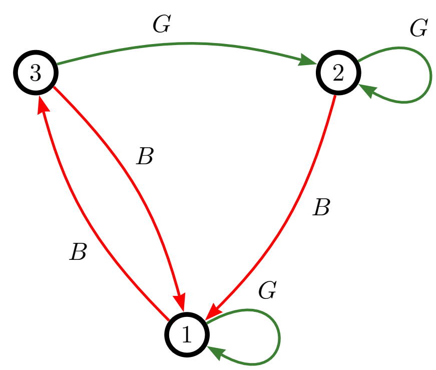

Figure 2.3. Transition for "Trust Expensive Names" Equilibria

We show that with these transition rules, "Trust Expensive Names" equilibria exist. In these equilibria, competent firms are the only ones buying S_1-names. S_2 and S_3-names are bought by both types. Competent firms work hard irrespective of the name they enter with. There are other transition rules which also give rise to SHE. Refer to the appendix for a description of all SHE sustainable under non-random transition rules in this richer market with three names. The necessary and sufficient conditions for such an equilibrium are as follows:

Incentive Compatibility for Name Choice

Competent firms are indifferent between S_1 and S_3-names.

$$-V_1 + w_1 - c + \delta(1-\rho)V_1 = w_3 - c + \delta(1-\rho)V_2 + \delta\rho V_1 \implies V_1 = \frac{w_1 - w_3 - \delta(1-\rho)V_2}{1 + \delta\rho - \delta(1-\rho)}. \tag{2.4}$$

Competent firms are indifferent between S_2 and S_3-names.

$$-V_2 + w_2 - c + \delta(1-\rho)V_2 + \delta\rho V_1 = w_3 - c + \delta(1-\rho)V_2 + \delta\rho V_1 \implies V_2 = w_2 - w_3. \tag{2.5}$$

Incompetent firms must strictly prefer an S_3-name to an S_1-name.

$$w_3 + \delta V_1 > -V_1 + w_1 \implies V_1 > \frac{w_1 - w_3}{1+\delta}. \tag{2.6}$$

Incompetent firms must be indifferent between S_3 and S_2-names.

$$w_3 + \delta V_1 = -V_2 + w_2 + \delta V_1 \implies V_2 = w_2 - w_3. \tag{2.7}$$

Incentive for Competent Firms to Work Hard

Competent firms prefer working hard to being lazy irrespective of the name they enter with.

$$c < \delta(1-\rho)V_1 \qquad \text{and} \qquad c < \delta(1-\rho)(V_2 - V_1). \tag{2.8}$$

Market Clearing Conditions

S_1 names: $$\phi\sigma_1 = \phi\sigma_1(1-\rho) + \phi\sigma_2\rho + \phi(1-\sigma_1-\sigma_2)\rho + (1-\phi)\mu_2 + (1-\phi)(1-\mu_2)$$

$$\implies \qquad \sigma_1 = \frac{1-\phi+\phi\rho}{2\phi\rho}. \tag{2.9}$$

S_2 names: $$\phi\sigma_2 + (1-\phi)\mu_2 = \phi\sigma_2(1-\rho) + \phi(1-\sigma_1-\sigma_2)(1-\rho)$$

$$\implies \qquad \sigma_2 = \frac{1-\rho}{2} - \frac{(1-\rho)(1-\phi)}{2\phi\rho} - \frac{1-\phi}{\phi}\mu_2. \tag{2.10}$$

Wages w_1, w_2, w_3 just equal the expected utility to the consumer. It can be easily verified that the above equations and inequalities admit a non-empty solution set.

There are two interesting features of this example. First notice that when a consumer sees an S_1-name, she knows with certainty that the firm is a competent one. There is no uncertainty about the type, and yet competent firms still choose to work hard. This is contrary to standard reputation models in which it is the uncertainty about a player's type that forces effort exertion.

Second, notice that a richer market for names can extend the range of equilibria. From (2.9), we see that for σ_1 to lie in the interval $(0, 1)$, we need $1 - \phi + \phi\rho < 2\phi\rho$. In other words, the equilibrium conditions in this example imply that $\phi > \frac{1}{1+\rho}$. But these are precisely the conditions for which equilibria do not exist in the first example in a market with two names. So there exist SHE in the richer market with non-random transition rules which are not possible in the two-state market.

2.4. Markets with General Transition Rules

With two states and non-random transition rules our ability to sustain SHE seems limited. It is not possible to sustain "Trust S-Names" equilibria. Further, even if the cost of working hard is very small, it is not possible to sustain SHE if there are too many competent firms (high ϕ) or if the chance of failing (ρ) is high. Section 3.2 suggests one way of sustaining more SHE: by enriching the market with more states.

An alternate approach could be to consider more general transition rules. Why might this approach work? The following example illustrates why. With a non-random transition rule, the original name and outcome together determine with certainty the future name of the firm. Suppose now that N-names were 'disadvantaged' in the sense that even after a good outcome N-names found it harder to become S-names. (Formally, conditional on a good outcome the future of a firm with an N-name is determined by the realization on an independent idiosyncratic randomization device. With a strictly positive probability, N-names remain worthless even after a good outcome.) S-names do not suffer this disadvantage, i.e. conditional on a good outcome, an S-name remains an S-name with certainty. Conditional on a bad outcome, any name becomes worthless. If such a transition rule were prevalent, a competent firm would be willing to pay a strictly higher premium for S-names compared to an incompetent firm. An incompetent firm is willing to pay a price up to the increase in wages it gets from an S-name. A competent firm is willing to pay this and more. A C-firm is also willing to pay for the increase in expected continuation payoffs from an S-name (conditional on working hard). This would be enough to yield "Trust S-Names" equilibria, as sorting only requires that one type be willing to pay a strictly higher premium for an S-name than the other type.

This idea leads us to ask: If we use more general transition rules, what are the conditions under which SHE can be sustained without further enriching the market for names? This section addresses this question. I conjecture that a two-state market with random transitions is only an efficient way of representing a richer market with deterministic transitions. In other words, for any SHE with random transitions in the two-state market, it is possible to find a deterministic transition rule that can sustain an equivalent SHE in a market endowed with more names.

Definition 16. *A **general transition rule** is a function* $f : \{S, N\} \times \{G, B\} \to \Delta(\{S, N\})$, *where* $\Delta(\{S, N\})$ *represents the probability distributions over the states* $\{S, N\}$*. So, a general transition rule* f *can be described completely by a vector* $(\gamma_1, \gamma_2, \gamma_3, \gamma_4)$*, where*
$f(S, G) = (\gamma_1, 1 - \gamma_1)$*,* $f(S, B) = (\gamma_2, 1 - \gamma_2$*,* $f(N, G) = (\gamma_3, 1 - \gamma_3)$*,* $f(N, B) = (\gamma_4, 1 - \gamma_4)$*. Here,* $f(S, G) = (\gamma, 1 - \gamma)$ *denotes that conditional on outcome* G*, an* S*-name will transition to* S *with probability* γ *and to* N *with probability* $1 - \gamma$*.*

Notions of strategies and equilibrium are unaltered. We start with the following lemma.

Lemma 2. *In any SHE, some competent firms must buy* S*-names.*

Proof: Suppose there exists an equilibrium in which competent firms buy only N-names (the cheap name). Sorting implies that incompetent firms must use the valuable S-names. In other words, an I-type prefers (at least weakly) an S-name to an N-name.

$$-V_S + w_S + \delta\gamma_2 V_S \geq w_N + \delta\gamma_4 V_S. \tag{2.11}$$

Since in such an equilibrium $w_S < w_N$, the above expression implies that V_S is negative irrespective of the values of γ_2 and γ_4. This is a contradiction. □

To see the intuition, consider the incentives for I-firms. A name affects incentives via wages or via continuation payoffs. In an equilibrium where C-firms buy only N-names, an S-name would give lower wages than an N-name. So the only reason for an I-type to buy an S-name is that it gives higher continuation payoffs than an N-name, after accounting for the price paid for the S-name. At best, an S-name will give a continuation payoff δV_S, which means that net payoff from an S-name is really negative. At worst, an N-name will give a net payoff of 0. So, it cannot be that I-firms prefer buying S-names to N-names.

2.4.1. Characterizing the Region where SHE Exist

By the lemma above, C-firms must buy S-names and I-firms must buy N-names in SHE. This leaves us with only two possible kinds of sorting. These are in fact the types of sorting we observed in the examples in Section 3.

The first is "Trust S-names" sorting. Here, only competent firms buy successful names. C-firms use both N and S-names. I-firms use only N-names. So, competent firms signal their competence by buying S-names, and incompetent cannot impersonate competent firms by buying S-names because the price is too high to make it worthwhile. Consumers trust an S-name when they see it because they are certain it is owned by a competent firm.

The second type of sorting is "Mistrust N-names". Here, only incompetent firms enter with new names. S-names are used by both C and I-firms. Competent firms do not use N-names because if they work hard, the costly S-names still give them a higher expected payoff than

N-names. Consumers mistrust any N-name, because they are certain that it is owned by an incompetent firm.

Recall that our objective is to characterize the conditions under which it is possible to find transition rules which result in sorting high-effort equilibria. It turns out that there is a simple characterization of the conditions for existence. SHE exist if and only if the cost of effort is low enough.

Proposition 2. *Given ϕ, ρ, δ, c, a sorting high-effort equilibrium exists if and only if*

$$c < \bar{c} = \min \left\{ \frac{\delta(1-\phi)(1-\rho)^2}{1-\phi+\delta\phi\rho}, \frac{\delta(1-\phi)(1-\rho)^2}{\delta(1-\phi)+\phi\rho} \right\}.$$

The interested reader may refer to the appendix for the proof. Note that the upper bound $\bar{c}$ is decreasing in ϕ. In other words, as the proportion of competent firms increases, it becomes harder to maintain a high-effort sorting equilibrium. To see why, notice first that if almost all firms are competent, (as $\phi \to 1$) the consumer will pay very similar wages to firms with S and N-names (i.e. $(w_S - w_N) \to 0$). In all sorting equilibria, we have seen that the price of an S-name is always less than or equal to the difference in wages. As a result, the price of an S-name will approach zero as well. If there is no benefit from ending up with an S-name, the incentives for working hard will be lost.

The upper bound $\bar{c}$ is also decreasing in ρ. The intuition is straightforward. If the probability of a bad outcome conditional on working hard is reduced, competent firms have a better incentive to work hard. At the other extreme, if hard work resulted in a bad outcome for sure ($\rho = 1$), there would be no incentive to work hard, and a high-effort equilibrium would be impossible. Indeed, at $\rho = 1$, c needs to be negative.

Given any cost of effort $c < \bar{c}$, one may ask whether both types of SHE exist for this cost. It turns out that the conditions for existence of "Trust S-Names" equilibria alone are more restrictive. "Trust S-Names" equilibria exist if and only if the cost of effort is lower than a threshold $\frac{\delta(1-\phi)(1-\rho)^2}{1-\phi+\phi\rho}$. Notice that for impatient firms, this threshold is strictly smaller than $\bar{c}$. Hence, above this threshold cost, the only SHE that exist are of the "Mistrust N-Names" type. For perfectly patient firms, the thresholds are the same, and hence both types of SHE exist for any cost $c < \bar{c}$.

2.4.1.1. HE and SHE. As an aside, let us examine the relationship between observability and sorting. The sorting feature of equilibrium arises naturally in a model with observability of name trades. The rough intuition is as follows: For names to be valuable, consumers must be able to trust a good name even when they know that it has just been sold to a new owner. This is possible if one of two situations arise: Consumers may be sure that only competent firms can buy good names. Alternatively, consumers must believe that new names cannot be trusted, as only incompetent firms choose to enter with new names. Successful names are valuable (tradeable) if some form of sorting arises in equilibrium.

It turns out that under observability, sorting does not restrict the range of parameters where high-effort equilibria exist. Define high-effort equilibrium with no sorting as follows:

Definition 17. *A **high-effort equilibrium without sorting (HE)** is a steady-state equilibrium in which*

1. *There is no name that is chosen by one type of firm and not by the other, and*
2. *Competent firms strictly prefer to work hard on the equilibrium path.*

Proposition 3. *Given $\phi, \rho, \delta, c,$ a high-effort equilibrium (HE) exists if and only if*

$$c < \frac{\delta(1-\phi)(1-\rho)^2}{1-\phi+\phi\rho}.$$

The proof of this proposition is very similar to that of Proposition 2, and so is omitted. Note that the upper bound is lower in the case of HE. This implies that the requirement for separation in equilibrium is not a restrictive one. By relaxing the requirement of sorting, we cannot expand the range of parameter values where high-effort equilibria exist. In a sense, in high-effort equilibria under observability, we get the sorting feature for free.

Below we return to analyzing the SHE. The next two propositions characterize each type of sorting in terms of the transition rule, costs and primitives of the environment.

2.4.2. Trust S-Names Equilibria

What kind of transition rules can create "Trust S-names" equilibria? Intuitively, transition rules must display three properties:

(1) For C-firms to be indifferent between S and N-names, the price of an S-name must equal the sum of increase in wages and increase in expected continuation payoffs conditional on working hard. For I-firms to avoid S-names, the price must be higher than the increase in wages and the increase in continuation payoff conditional on a bad outcome.

(2) Transition rules must be such that the demand for S-names equal the supply.

(3) For C-firms to have a strict incentive to work hard, transitions must be such that the expected payoff from working hard is strictly higher than that from being lazy.

The following proposition describes these properties formally.

Proposition 4. ***(Trust S-Names Equilibria)*** *High-effort "Trust S-names" equilibria exist if and only if*

1. $\gamma_3 - \gamma_4 < \gamma_1 - \gamma_2$

2. $\gamma_4 < \frac{\phi}{1-\phi}(1 - (1-\rho)\gamma_1 - \rho\gamma_2)$

3. $c < \delta(1-\rho)(\gamma_3 - \gamma_4)\frac{(1-\phi)(1-\rho)}{[\phi(1-\sigma_S)+1-\phi][1+\delta(1-\rho)(\gamma_3-\gamma_1)+\delta\rho(\gamma_4-\gamma_2)]}$

where $\sigma_S = \frac{\phi(1-\rho)(\gamma_3-\gamma_4)+\gamma_4}{\phi(1-\rho)(\gamma_3-\gamma_1)+\phi\rho(\gamma_4-\gamma_2)+\phi}$.

So, for a given distribution of firms ϕ, success rate $(1-\rho)$ and discount factor δ, transition rules that satisfy (1) and (2) can sustain high-effort "Trust S-names" equilibria, if the cost of working hard is sufficiently low (i.e. (3) holds). The reader may refer to the appendix for the proof. An analogous result holds for "Mistrust N-names" equilibria.

2.4.3. Mistrust N-Names Equilibria

In "Mistrust N-Names" equilibria, what must transition rules look like?

(1) For I-firms to use S and N-names, the price of an S-name must be exactly equal to the increase in wages and increase in the continuation payoff conditional on a bad outcome. For C-firms to not use N-names, the price of an S-name must be less than the total increase in wages and continuation payoff conditional on hard work.

(2) The market for S-names must clear.

(3) C-firms must work hard on the equilibrium path i.e. when they buy S-names.

Proposition 5. ***("Mistrust N-Names" Equilibria)*** *High-effort "Mistrust N-names" equilibria exist if and only if*

1. $\gamma_3 - \gamma_4 < \gamma_1 - \gamma_2 < \frac{1-\gamma_2}{\phi(1-\rho)}$

2. $\gamma_4 > \frac{\phi}{1-\phi}(1-(1-\rho)\gamma_1 - \rho\gamma_2)$

3. $c < \delta(1-\rho)(\gamma_1-\gamma_2)\frac{\phi(1-\rho)}{[\phi+(1-\phi)\mu_S][1+\delta(\gamma_4-\gamma_2)]}$

where $\mu_S = \frac{\phi(1-\rho)\gamma_1+\phi\rho\gamma_2+(1-\phi)\gamma_4-\phi}{(1-\phi)(1+\gamma_4-\gamma_2)}$.

For a proof of the Proposition 5, the reader may refer to the appendix. The conditions for the existence of "Mistrust N-Names" equilibria turn out to be more permissive than those for "Trust S-Names". In "Mistrust N-Names" equilibria competent firms use only S-names. Since N-names are off the equilibrium path for C-firms, and it is no longer necessary to sustain high-effort conditional on an N-name. This makes it possible to sustain high-effort and sorting for some parameter ranges where "Trust S-Names" are unsustainable.

2.4.4. Non-Random Transitions as a Special Case

It can be easily verified that the example of SHE with non-random transition rules described in Section 3.1 satisfies the conditions in Proposition 5. From Propositions 4 & 5, it can be seen why these are the only non-random rules that work. Consider all possible non-random transition rules in the two-state market. To maintain incentives for competent firms to work hard, all rules which do not reward an S name for a good outcome are eliminated. (i.e. rules with $\gamma_1 \leq \gamma_2$). This leaves possible only four transition rules.

$$(1,0,1,0) \qquad (1,0,0,0) \qquad (1,0,1,1) \qquad (1,0,0,1)$$

$$(a) \qquad\qquad (b) \qquad\qquad (c) \qquad\qquad (d)$$

Rule (a) violates Condition 1 for both types of equilibria and cannot ensure sorting. Rule (b) destroys incentives for hard work in both types of equilibria. Rules (c) and (d) destroy effort incentives in the "Trust S-names" case. This leaves only rules (c) and (d) as options to sustain "Mistrust N-names" equilibria, which are precisely the ones described in Section 3.1.

2.4.4.1. Equilibrium when Clients Observe Full History of Outcomes. The SHE with non-random transitions satisfy another interesting property. These equilibria survive in the more standard infinitely repeated game setting where consumers can see not just the name but the full history of outcomes for any firm name that they encounter.

To elaborate, let us consider a slightly different environment, where a name is a complete history of outcomes (G or B) of the firms that owned it. So, when consumers see a name what they observe is the complete history of outcomes. Note that since play has occurred for ever into the past, consumers observe an infinite history of outcomes. The timing of the game is unaltered. As before, firms must choose a name before they enter. Retiring firms sell their name if possible. A new firm can choose to enter costlessly with a new name (with no history) N or with a name bought from a retiring firm. In this setting, we want to look for sorting high-effort equilibria. In other words, can we find equilibria where competent firms always exert high effort, and there is at least one kind of name which perfectly signals the type of the firm that uses it?

In this setting, since names are now just histories, firms can choose between potentially infinitely many different names. On the face of it, this would complicate the problem significantly. We can simplify the problem by considering certain equivalence classes of outcome histories. Denote the entire set of possible outcome histories by $\mathscr{H}$ and any arbitrary history by h. Consider two equivalence classes of $\mathscr{H}$, denoted by $\mathscr{N}$ and $\mathscr{S}$, defined as follows:

Let $N \in \mathscr{N}$.

$h \in \mathscr{S} \implies hB \in \mathscr{N}, hG \in \mathscr{S}$.

$h \in \mathscr{N} \implies hB, hG \in \mathscr{S}$.

To illustrate, the history NGB belongs to $\mathscr{N}$ while $NGBB$ belongs to $\mathscr{S}$.

Clients will treat any two histories in the same equivalence class in the same way. In effect, clients behave as if firms appear under two names $\mathscr{N}$ and $\mathscr{S}$. It can be easily verified that the following constitutes a steady-state equilibrium in this game.

- $\mathscr{N}$-names are worthless. $\mathscr{S}$-names sell at a price V_S as in Proposition 1.

- C-firms always enter with $\mathscr{S}$-names.

- I-firms use $\mathscr{S}$-names with probability μ_S as in Proposition 1 and $\mathscr{N}$-names with $1 - \mu_S$.

When a consumer sees a name from $\mathscr{N}$, she believes that the firm must be an incompetent one and so pays a wage of $w_N = 0$. When she sees a name from $\mathscr{S}$, she knows the firm could be a C or an I-firm, and pays her expected utility w_S as given in Proposition 1.

So, we have found an exact analog of the "Mistrust N-names" equilibrium derived in the economy with two kinds of names and non-random transition rules. Why is this a desirable property? In our model, a transition rule is a mechanism (conceivably managed by a mediator) that determines the future of a firm based on its original name, realized outcome and an independent idiosyncratic randomization. SHE that arise seem to depend critically on implementation of the transition rule. However, we see here that the SHE that are sustainable with non-random transitions are not really dependent on the existence of any mediator or transition mechanism. They can arise naturally in equilibrium in a standard repeated game environment where all players can observe the history of realized outcomes. This property does not hold for SHE under random transitions. However, I conjecture that for any SHE in the two-state market with random transitions, an equivalent one can be derived using non-random rules in a market with a richer set of names. In that case, we would be able to use a richer set of equivalence classes of histories and remove the dependence of SHE on specific transition rules.

2.5. Relaxing Observability

So far, we have considered an environment where name trades are fully observable, and derived sorting high-effort equilibria. An interesting extension would be a comparison of this environment with one with non-observability of name ownership changes, in terms of existence of SHE and welfare implications. A detailed examination of these issues is not included in the scope of this paper. In this section, I present some examples to illustrate how the welfare comparisons can go in either direction, based on the specific transitions rules being implemented and the primitives of the environment. In the model described so far, observability is automatic. In order to make a meaningful comparison between regimes with and without observability, we need to alter the environment.

2.5.1. Overlapping Generations Model

Consider an economy with overlapping generations of firms. As before there are two types of firms, competent and incompetent. A proportion ϕ is competent. Competent firms can choose to work hard and incompetent firms are incapable of working hard. There are two outcomes - good (G) and bad(B). Conditional on hard work, the probability of a good outcome is $(1-\rho)$, and conditional on laziness, a good outcome never occurs.

Each firm lives two periods. Firms can enter with N or S-names. Firms meet consumers once in each period. Retiring firms sell their name before retiring. Consider the general random transition rules described in Section 4. Conditional on a good outcome, an S-name remains an S-name with probability γ_1 and N-names become S with probability γ_2. Conditional on a bad outcome, S-names remain S-names with probability γ_3 and N-names become S-names with probability γ_4. The timing of the game is depicted in Figure 4.

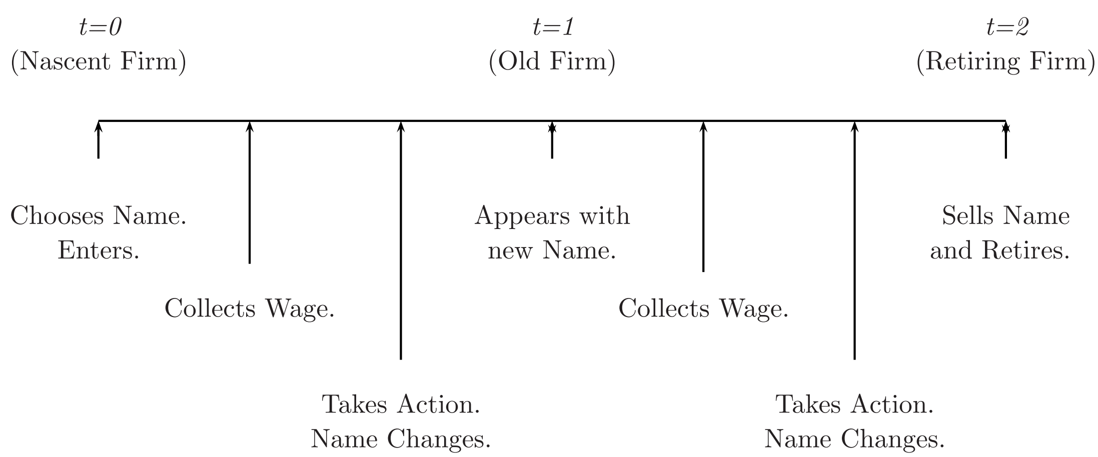

Figure 2.4. Timing of Game for Overlapping Generations of Firms

In any period there are three generations of firms - nascent firms who are just entering the market, old firms who have already lived for one period and retiring ones. Only retiring firms can sell names to nascent firms. Old firms cannot buy or sell names. (This is just a convenient assumption. The qualitative results are unaltered if old firms also bought names.)

At the time of entry, a firm decides what name to enter with. It can enter costlessly with a neutral (N) name, or can buy a costly name from a retiring firm. On entering, firms meet a consumer who pays the firm an upfront wage. Then the firm takes an action (work hard or be lazy), and conditional on the action choice and outcome, the firm's name changes. The firm meets a consumer again in the second period now with his altered name. Again, he collects his wage, takes an action, and his name changes. Before retiring the firm sells its name (if possible) to a nascent firm. For simplicity, we ignore discounting.

Consumers live only one period. They pay firms upfront, and so pay the expected utility based on what they observe. In this model, we can study the two regimes with and without observability of name ownership. In a regime of non-observability, consumers see only the name of the firm. Under observability, they see not just the name but also the age of the firm. The age of a nascent firm is 0 and the age of an old firm is 1. This is equivalent to full observability of name: for instance, a nascent S-name must be a firm which just bought the name. Denote the wages by $w_{N,0}, w_{N,1}, w_{S,0}, w_{S,1}$, where $w_{a,i}$ denotes the wage paid to a firm with name a and age i. Notions of strategies and equilibrium are extended in the natural way.

2.5.2. Observability and Welfare: Examples

When we compare regimes of observability and non-observability, it is clear that since we consider only high-effort equilibria the total surplus of consumers and firms is constant. The more interesting question is to ask separately whether the consumers are better off under any particular regime. Alternatively, which firms fare better under which regimes?

While a general analysis is postponed to later work, the examples presented here illustrate the different effects that may arise. To compare meaningfully, I choose examples where equilibria exist both under observability and non-observability. In the first example, observability is irrelevant. The wages paid are independent of the age of the firm. In later examples I show that it is possible for all firms to get better off under one regime or the other. I also present examples where the preferences of the two types of firms are opposed.

2.5.2.1. Observability makes no difference. Consider a market with the following transition rule. Conditional on a bad outcome, all names become N-names. Conditional on a good outcome, S-names remain S-names and N-names become S with probability $(1-\lambda) \in (0, 1)$. Consider a

"Trust S-Names" equilibrium. When consumers see an S-name, they know that it is a good firm, and will pay $w_{S,0} = w_{S,1} = 1 - \rho$. A nascent firm with an N-name can be competent or incompetent. Hence, the consumer pays her expected utility

$$w_{N,0} = \frac{\phi(1-\sigma_S)(1-\rho)}{\phi(1-\sigma_S) + 1 - \phi}.$$

For old firm with an N-name there are three possibilities. The firm may be an incompetent one, or a competent firm who had an N-name and remained N, or a competent firm who had an S-name but ended up with an N-name. Again, the consumer pays her expected utility

$$w_{N,1} = \frac{\phi(1-\sigma_S)(\rho + \lambda(1-\rho)) + \phi\sigma_S\rho}{1 - \phi + \phi(1-\sigma_S)(\rho + \lambda(1-\rho)) + \phi\sigma_S\rho}(1-\rho).$$

We can solve for the high-effort sorting equilibrium. It turns out that C-firms buy S-names with probability $\sigma_S = \frac{(1-\rho)(1-\lambda)}{1-\lambda+\lambda\rho}$. I-firms use only N-names. The equilibrium price V_S is

$$V_S = \frac{w_{S,0} - w_{N,0} + \lambda(1-\rho)(w_{S,1} - w_{N,1})}{1 - \lambda^2(1-\rho)^2}.$$

What happens if we impose the condition that ownership changes cannot be observed? It turns out that this makes no difference to the equilibrium in this market. To see why, recall that any firm with an S-name must be competent, and so earns a wage $w_{S,0} = w_{S,1} = 1 - \rho$. The wages earned by firms with N-names turns out to be

$$w_{N,0} = w_{N,1} = \frac{\phi\rho}{\phi\rho + (1-\phi)(1-\lambda+\lambda\rho)}.$$

Since the wages paid are always independent of the age of the firm, making ownership changes unobservable does not make any difference to firms or consumers.

2.5.2.2. All firms prefer one regime. Consider a market where the proportion of C-firms $\phi = 0.9$ and the cost of effort $c = 0.1$. The market has the following transition rule: Conditional on a good outcome, both S and N-names become S-names. Conditional on a bad outcome, S-names remain S with probability 0.1 and N-names become S with probability 0.2. (i.e. $\gamma_1 = \gamma_3 = 1, \gamma_2 = 0.1, \gamma_4 = 0.2$).

Example 1 (All firms prefer non-observability)**.** *Let $\rho = 0.1$. Under these conditions, under observability, there exists a "Trust S-names" equilibrium. If name changes were not observable, the same equilibrium survives, and it can be verified that all firms get better off.*

Example 2 (All firms prefer observability)**.** *Let $\rho = 0.01$. Under these conditions, under observability, there exists a "Mistrust N-names" equilibrium. If name changes were not observable, the same equilibrium survives, but now all firms can be shown to get worse off.*

The net payoff from buying or selling a name is weakly negative. The price of an S-name is higher in a regime with observability. This implies that under observability, name trading is more costly for any firm. A firm will prefer observability only if the sum of the wages it receives is high enough to cover the increased cost of name trade.

2.5.2.3. Different firms prefer different regimes. Now suppose the proportion the cost of working hard $c = 0.1$ and the probability of a bad outcome conditional on working hard is $\rho = 0.01$.

Example 3 (C-firms prefer non-observability, I-firms prefer observability)**.** *Suppose $\phi = 0.95$. Consider the following transition rule. Conditional on a good outcome, an S-name remains an S-name with probability* 1. *Conditional on a bad outcome, an S-name remains S with probability* 0.2. *The corresponding transition probabilities for an N-name are* 0.8 *and* 0.1. *Under observability there exists a "Trust S-names" equilibrium. If ownership changes were not observable, the same equilibrium survives. C-firms get better off while I-firms lose out. However, together the firms get better off and consequently, consumers get worse off.*

Example 4 (C-firms prefer observability, I-firms prefer non-observability)**.** *Let $\phi = 0.5$. Consider the transition rule as above with one difference. Now conditional on a good outcome, an*

S-name remains S with probability 0.95*. Under observability there exists a "Mistrust N-names" equilibrium. If ownership changes were not observable, the same equilibrium survives. Now, C-firms get worse off while I-firms get better off. Together, the firms get worse off and so consumers get better off.*

Observe that, in "Trust S-names equilibria", observability makes consumers weakly better off, while in "Mistrust N-names" equilibria with consumers get worse off under observability.

2.6. Related Literature

This paper contributes to the growing literature on firm reputations and the market for firm reputations. Earlier work closest in spirit to this paper would include Kreps (1990) and Tadelis (1999, 2002, 2003). In this and related work, the existence of a market for firm names plays two roles. A tradeable name provides incentives for short-lived agents to work hard. Also, a firm name acts as an assessment of firm's ability, and tradeable names allow firms to buy credibility with consumers.

Tadelis (1999, 2002, 2003) considers a general equilibrium framework where he establishes a link between the value of the name and price of the firm's services. Unlike this work, Tadelis (1999, 2003) considers a model of pure adverse selection. Tadelis (2002) includes moral hazard. One of the key insights from Tadelis's work is that non-observability of ownership changes is key to active name trading in all equilibria.

In the current paper, I depart from the literature by examining whether active trading of firm names can exist with observability of ownership changes. I present a model with both adverse selection and moral hazard and with full observability of name trading. I show that in this environment, all high-effort equilibria involve active trading of names. However, we cannot eliminate the 'bad' equilibrium in which all firms are lazy, consumers mistrust all firms, firms earn nothing and names are not traded. Since the 'commitment' type in my model is an incompetent lazy type, this equilibrium cannot be eliminated.

In a recent paper, Hakenes and Peitz (2007) also ask the question whether firm names can be traded when ownership changes are observable. They derive sorting equilibria with observability, but in a model of pure adverse selection. Further, they do not allow full observability and are able to obtain active trading only with partial observability. In their set-up, all consumers cannot observe names of firms. Firm reputations are observable locally to a subset of consumers.

The other feature that sets this paper apart is the sorting nature of equilibrium. Sorting of types does not arise in equilibrium in the Tadelis environment. This is a result of two opposing effects. On one hand, good firms value good names because they can work hard and maintain them. On the other hand, bad firms value existing good names because they cannot build a reputation for themselves. When good firms try to separate themselves by buying good names, the second effect overwhelms the first and bad firms value good names more than good firms.

Mailath and Samuelson (2001) also consider a market for firm reputations in which types cannot get sorted. They show that names of intermediate value are more likely to be bought by good firms and names with extreme reputations are more likely to be bought by bad firms. Very good names are more attractive to bad firms who will gain by depleting the high reputation. Intermediate names are bought by good firms, because they can build up the reputation. These names are not bought by bad firms, because there is less value to to be depleted. Good firms find it too expensive to buy a bad name and build up its reputation from scratch.

In this paper, I am able to derive equilibria with sorting of types. I show two types of sorting: some equilibria where competent firms can differentiate themselves by being the only ones buying valuable names, and other equilibria where incompetent firms give themselves away by being the only firms using worthless names.

2.7. Conclusion

In this paper, I raise two questions central to the literature on firm names and reputations. First, I ask whether firm names can be tradeable assets even when changes in name ownership are observable to the consumer. Second, I ask if the market for firm names can act as a sorting device and separate competent firms from incompetent ones.

I consider an infinite horizon economy with generations of firms and consumers interacting in each period. Firms can be competent or incompetent. Firms choose to enter the market under different names which they buy in a competitive market for names. Changes in name ownership are fully observable. There are two kinds of names available. Consumers buy a product from the firm for which they pay upfront. At the time of purchase, consumers only observe the name of the firm. After collecting their payment, firms provide their services. At the end of the period, each firm's name changes according to a fixed transition rule (potentially random) which determines the future of a name based on its original name, realized quality of services, and an idiosyncratic randomization. Before retiring, a firm can sell its name to a new entrant.

With full observability of name changes, I examine the existence of equilibria in which the market for names both makes competent firms work hard and sorts firms according to their type. I define a class of equilibria called sorting high-effort equilibria (SHE) where at least one type of firm has a strict incentive to not use one of the two names, and competent firms always work hard. In the main result of the paper, I characterize necessary and sufficient conditions for the existence of such sorting high-effort equilibria (SHE). I show that SHE exist provided the cost of effort is low enough for competent firms.

I also show that the market for firm reputations can act as an effective sorting device that separates competent firms from incompetent ones. Some names can perfectly signal the type of the firm that owns it. Two kinds of sorting may arise in equilibrium. It is possible for competent firms to separate themselves by being the only ones buying the valuable successful names. I call these situations "Trust S-Names equilibria". In these equilibria, when consumers see a successful name, they trust it to be a competent firm and pay the corresponding high price. The second type of sorting that arises is termed "Mistrust N-Names". Here, incompetent firms give themselves away by being the only firms using the cheap names. Consumers treat cheap names with mistrust and pay them corresponding low wages. Competent firms force this situation to arise by always buying S-names before entering the market.

CHAPTER 3

Large Games with Limited Individual Impact (Joint with Ehud Kalai)

3.1. Introduction

Consider a simple location choice game. N players decide simultaneously where to live in a linear city. Players are of different types (wealth levels) and want to live near other players who have similar wealth levels. Player types are drawn independently, and players do not observe other players' types before choosing their locations. So players choose their location based on their own wealth and the prior belief about others' wealth. In this context (and many similar situations), Bayesian Nash equilibrium may not be a good solution concept. If we view an equilibrium as a steady state of an ongoing rational interactive process, we should require some 'stability' in our solution concept. In this game, if relocation costs are low, players will want to revise their Bayesian Nash equilibrium actions after they observe the realized wealth levels of their opponents. A reasonable, intuitive notion of equilibrium in this game should imply some *'ex-post stability'* and preclude such relocation. Under what general conditions of the underlying game, will Bayesian Nash equilibria have this ex-post stability property? This is the central question of the paper.

Robustness of equilibria have been of interest both in theoretical and applied contexts. Ex-post stability relates to important issues in economic applications, and is known to imply several desirable properties of equilibria. For instance, ex-post stability implies that a purification result holds (Kalai 2004, Cartwright and Wooders, 2006). In certain large market games, ex-post stable Nash equilibria exhibit a strong rational expectations property, when types are drawn independently. In the implementation literature, it implies that the revelation principle holds. (Green and

Laffont, 1987). Recently, the robustness properties of equilibria have also been of much interest in the literature that lies in the interface of economics and computer science arising from work on distributed computing. This literature has been concerned with assessing the damage caused in systems by the presence of faulty behavior, and designing mechanisms (protocols) that are fault tolerant. It turns out that the notions of fault tolerance and asynchrony in distributed computing are closely related to the idea of robustness in equilibria in game theory. In fact, in this context, ex-post stability of equilibria can be viewed as a form of resilience to asynchrony in the world of distributed computing.

Large games or games with a large number of players have been studied because Bayesian equilibria in some classes of large games exhibit appealing stability properties like ex-post Nash robustness. Kalai (2004) establishes robustness of Bayesian equilibria in a particular class of large noncooperative games. He shows that if we restrict attention to games where payoff functions are equicontinuous, where action and type spaces are finite, types are drawn independently and players are anonymous (each player's utility is affected only by the aggregate actions of her rivals), Bayesian equilibria are *approximately ex-post stable*, provided the number of players is large enough. What is the notion of approximate ex-post stability? A strategy vector is (ϵ, ρ) ex-post stable if with very high probability (at least $1-\rho$), no player can gain, ex post, by more than some very small amount ϵ by deviating. In particular, if players play certain actions in equilibrium of a large game with incomplete information, they would have no incentive to revise these actions even with perfect hindsight knowledge of types and actions of their opponents. Further, equilibria in large games are robust also in the sense that the analyst's equilibrium predictions are less sensitive to the details of the game. Equilibrium predictions do not change across different extensive versions of the game. Immunity to alterations means that Nash equilibrium predictions are valid even in games whose structure is largely unknown to modelers or to players. Kalai (2004) establishes a strong form of convergence to ex-post stability. As the number of players gets larger, convergence to ex-post stability is at an exponential rate, simultaneously for all games and all equilibria. Gradwohl and Reingold (2007) provide some limited extensions of the above result. They show that the assumption of independence of types can be relaxed somewhat. Bayesian equilibria in the equicontinuous anonymous games can be ex-post Nash even when certain kinds of correlation between types are allowed.

However, the ex-post stability in the above-mentioned results relies on the anonymity of players, and on the restriction of games to have finite sets of types and actions. The results of Kalai (2004) cannot be applied to many important applications. The finiteness assumption eliminates important applications of economic interest such as market games, voting games, location choice games and auctions. Further, the anonymity restriction rules out many strategic situations[1], where player's payoffs can be affected asymmetrically on the types and actions of her different opponents. Our primary objective in this paper is to investigate whether the ex-post stability properties of Bayesian Nash equilibria hold more generally, when we consider infinite type and action spaces and non-anoymous games.

We consider a general family of large games where player's types and action spaces are infinite, and where players are not anonymous, but instead payoffs can depend on types and actions of specific rivals in potentially asymmetric ways. We show that if the strategy spaces are bounded and the payoff functions in the family of games satisfy a single regularity condition - a variant of uniform Lipschitz continuity - then every Bayesian Nash equilibrium of a large game is also an approximate ex-post Nash equilibrium. In fact, we show that if players use their equilibrium action, every Bayesian equilibrium is $(\epsilon, 0)$ ex-post Nash; i.e. with certainty, no player can gain by more than a very small amount by deviating ex-post.

[1]The anonymity assumption made in Kalai (2004) imposes anonymity on payoff functions but not on the symmetry or anonymity of the players. This implies any game could potentially be described by an alternative game where information about named players is incorporated into their types. However, the additional assumption of finite types restricts the generality of such alternate descriptions.

While this paper makes a substantive contribution by extending equilibrium robustness results to general settings with infinite types action spaces and non-anoymous games, we also make some technical contributions that we hope will be useful in the future. The regularity condition we impose on payoff functions is a notion of uniform Lipschitz continuity with respect to a different metric we call the scaled $L1$ metric. This metric that we introduce turns out to be very suitable for the analysis of large games. While the regularity condition offers some technical convenience, more significantly it implies a special property that we call *limited individual impact*. In these games, the unilateral impact that any player can have on the payoffs of any of her rivals is bounded and further this unilateral impact decreases with the total number of players in the game. This limited individual impact condition turns out to be important as it implies that an appropriate law of large numbers (McDiarmid's Bounded Differences Inequality) holds. This law of large numbers (deviation inequality) has not been used before in the economics literature, and we believe that it can be extremely useful in the analysis of large games.

The intuition of the main result is as follows. Consider a Bayesian equilibrium of a game where the Lipschitz continuity condition is satisfied. In particular this implies that this is a game with limited individual impact. By a law of large numbers, as the number of players increases, with a high probability every player gets a realized payoff arbitrarily close to her ex-ante expected payoff from her equilibrium strategies. Suppose the equilibrium were not ex-post stable. Then with positive probability, the realized outcome of the game (profile of type-action characters of one's rivals) will be one where player i has an incentive to deviate ex-post and make a 'significant' gain. Since players are playing equilibrium strategies, it must be the case that the deviation action gives i a lower ex-ante expected payoff, but at the realized profile, the payoff from deviation is 'significantly' higher than the ex-ante expected payoff from the deviant action. Since such an outcome can be realized in the game only with very small probability, the realized type-action character profile must be very 'near' some other profile, where such a profitable deviation is not possible. This implies that there are two 'nearby' type-action character profiles of one's rivals, where the same action gives a player 'significantly' different payoffs. But this is impossible if payoff functions satisfy Lipschitz continuity. So, the equilibrium must be approximately ex-post stable. To fix ideas, we present an example of such a game.

3.1.1. Example

Consider a city on the closed interval $[0,1]$. There are N people who must decide where to live in this interval i.e. they choose location $l_i \in [0,1]$. Players can observe who their opponents are and how old they are. Suppose each person i is characterized by her age $a(i)$ normalized to also lie in $[0,1]$. Players are of different types w_i, with types being drawn independently from a uniform distribution on $[0,1]$. Think of a player's type w_i to be a normalized measure of her wealth.

At the start of the game, each person learns her own type (wealth level). She also observes her opponents and their ages. Players must then simultaneously choose a location to live in, to maximize her own utility. Each player wants to live far away from other people. She would rather live closer to younger people than older people, and would rather live closer to people with similar wealth levels as her own. Formally, the utility function of players is given by

$$u_i(w_i, l_i, w_{-i}, l_{-i}, a_1, ..., a_N) = \frac{1}{N-1}\left(\sum_{j\neq i} a_j(|a_i - a_j|)(|w_i - w_j|)\right).$$

At this point, it is worthwhile to point out some specific features of this location choice game. Clearly, player types are drawn from a continuum, and actions spaces are infinite. Notice that in this game, players are not anonymous. Each player's impact on her rivals is different and depends on her age. As a result, each player's utility is not just determined by aggregate outcomes. However, the specific impact any player has on her rivals reduces with the number of players in the game. To see why, observe that the maximum difference any player can make on a rival's utility is $\frac{1}{N-1}$.

An equilibrium that is ex-post stable in this game would be one where once players choose where to live, they do not want to relocate. The results in this paper imply that if this location choice game were played by a large enough number of people, and if players chose their location according to a Bayesian Nash equilibrium of this game, they would not want to relocate even if after they have perfect knowledge of their rivals' location choices and wealth levels.

The rest of the paper is organized as follows. Section 2 describes the model. In Section 3, we discuss the main assumption of Lipschitz continuity and its implications for the class of games that we study. In Section 4, we state and prove the main result. Section 5 presents an example to demonstrate the main ideas of the paper. In Section 6, we discuss a related notion of information proofness of Bayesian equilibria. Section 7 concludes.

3.2. Model

Let $\mathscr{T}$ denote the space of feasible types of players, and let $\mathscr{A}$ denote the space of all feasible actions of players. We will consider a family of games $\Gamma = \Gamma(\mathscr{T}, \mathscr{A})$ of Bayesian games $G(N, T, A, \tau, \{u_i\})$ that can be described as follows.

- There are N players $\{i = 1, ..., N\}$.
- The type of each player i is drawn independently from a type-space T_i, where T_i is a compact subset of $\mathscr{T}$. Let T denote the type space of all players, i.e. $T := \prod_{i=1,...,N} T_i$. Let τ be a probability measure on the Borel subsets of T. Associated with τ is a marginal distribution on each T_i, which is denoted by τ_i. The distributions are common knowledge.
- Each player i chooses actions from her action space A_i which is a compact subset of $\mathscr{A}$. Let A denote the space of action profiles of all players,, i.e. $A := \prod_{i=1...,N} A_i$.

  We refer to a pair $(t_i, a_i) := c_i$ as a *type-action character* of player i. Denote the space of each player's *type-action characters* $T_i \times A_i$ as $\mathscr{C}_i$. Denote the space of *type-action character profiles* as $\mathscr{C} = \prod_i \mathscr{C}_i$. For any player i, we denote the space of type-action characters of her rivals as $\mathscr{C}_{-i} := \prod_{j \neq i} T_j \times A_j$. For the rest of the paper, we assume that each T_i and each A_i is a compact subset of $\mathbb{R}^k$ for some $k \geq 1$.

- Players' payoffs are given by bounded measurable functions $u_i : \mathscr{C} \to \mathbb{R}$.

  For any player i, for any (t_i, a_i), we can define induced payoff functions

$$u_i^{t_i,a_i} : \mathscr{C}_{-i} \to \mathbb{R} \text{ such that } u_i^{t_i,a_i}(t_{-i}, a_{-i}) = u_i(t_i, a_i, t_{-i}, a_{-i}).$$

Definition 18. *A family of games $\Gamma = \Gamma(\mathscr{T}, \mathscr{A})$ is said to have a **uniformly bounded strategy space** if $\mathscr{T}$ is a compact subset of $\mathbb{R}^{k_T}$ and $\mathscr{A}$ is a compact subset of $\mathbb{R}^{k_A}$ for some integers $k_T, k_A \geq 0$.*

In this paper, we consider families of games with uniformly bounded strategy spaces. Next, we define a metric which will be useful in our analysis.

Definition 19 (Scaled L1 Metric)**.** *Given an integer N, we define the **N-scaled L1 metric** on $\mathbb{R}^M$ as follows.*

$$\forall x, y \in \mathbb{R}^M, d(x, y) = \frac{1}{N} \sum_{m=1}^{M} |x_m - y_m|.$$

Henceforth, in any N-player game, we will use this metric on the space of any player's rival type-action characters; i.e. for any player i, we measure the distance between any two type-action character profiles of her rivals $c_{-i}, c'_{-i} \in \mathscr{C}_{-i}$ using this metric, where N is the total number of players in the game. Notice that this metric is very similar to the L1 metric except for the scaling by a factor of N. The scale factor N makes the distance between two type-action characters less sensitive to differences in magnitude of each coordinate.

We consider families of games with payoff functions that are Lipschitz continuous in the rivals' type-action character profile with respect to this new metric.

Definition 20. *Given $K \geq 0$, the payoff functions u_i in a family of games $\Gamma(\mathscr{T}, \mathscr{A})$ are said to be **uniformly K-Lipschitz continuous** in the rivals' type-action character with respect to* *the N-scaled L1 metric if for every N-player game in $\Gamma(\mathscr{T}, \mathscr{A})$, for every player i,*

$$|u_i(c_i, c_{-i}) - u_i(c_i, c'_{-i})| < K \; d(c_{-i}, c'_{-i})$$

where $d(.,.)$ is the 'scaled L1 metric'.

Notice that the above condition requires the Lipschitz bound is uniform for all N in the family of games. Let $\Gamma(\mathscr{T}, \mathscr{A}, K)$ denote a family of games consisting of games which also satisfy the above Lipschitz continuity condition.

As an aside, note that the equivalence of the L1 and Euclidean norms implies that this Lipschitz condition with respect to the scaled L1 metric implies the standard Lipschitz continuity with respect to the Euclidean metric.

Definition 21. *A **pure strategy** of a player i is a measurable function $s_i : T_i \to A_i$.*

Defining mixed strategies as maps from types to mixtures over pure strategies has the drawback that it is not well defined in games with a continuum of types (see Aumann (1964)). Here, we use the notion of distributional strategies as introduced by Milgrom and Weber (1985).

Definition 22. *A **distributional strategy** for player i is a probability measure σ_i on the Borel subsets of $T_i \times A_i$ for which the marginal distribution of T_i is τ_i. When players use distributional strategies, the expected payoff of player i is defined as follows:*

$$U_i(\sigma_1, \ldots, \sigma_N) = \int u_i(c) d\sigma(c).$$

A distributional strategy refers to an equivalence class of mixed strategies that gives rise to the same behavior. A mixed strategy induces a joint distribution across types and actions. Conversely, a joint distribution can be generated by many mixed strategies[2].

[2] Please see Milgrom Weber (1985) for more on the correspondence between behavioral strategies, mixed strategies and distributional strategies. A behavior strategy can be defined as function $\beta_i(B, \cdot) : T_i \to [0, 1]$ such that

Definition 23. *A profile of distributional strategies* σ^* *is an* ***equilibrium***[3] *if*

$$\forall i, \forall \sigma_i', U_i(\sigma_1^*, \ldots \sigma_i^*, \ldots \sigma_N^*) \geq U_i(\sigma_1^*, \ldots \sigma_i', \ldots, \sigma_N^*).$$

Since we are interested in ex-post Nash robustness and since we prove this property asymptotically as the number of players increases, we need to define a notion of approximate equilibrium.

Definition 24. *Let* $\epsilon > 0$ *(small). A profile of type-action characters* $c \equiv (t_1, a_1, ..., t_N, a_N)$ *is said to be an* ϵ***-best response*** *for player* i *if for all actions* a_i', $u_i(a_i', t_i, a_{-i}, t_{-i})) \leq u_i(a) + \epsilon$.

Definition 25. *A type-action character profile is said to be* ϵ***-Nash*** *if it is an* ϵ*-best response for all players.*

Definition 26. *A strategy profile* $(\sigma_1, ..., \sigma_N)$ *is said to be* ϵ***-ex-post Nash*** *if it yields* ϵ*-Nash type-action character profiles with probability* 1.

Note that an equilibrium is approximately ex-post Nash if it is stable in the sense that the realized actions, not the mixed strategies, constitute a Nash equilibrium of the complete information game.

3.3. Lipschitz Continuity and Limited Individual Impact

At this point, it is worthwhile to examine the Lipschitz continuity assumption more closely. In the next section, we will prove the ex-post stability of equilibria in the class of games $\Gamma_(\mathscr{T}, \mathscr{A}, K)$. To do this, we will use a special property of this family of games that we call the 'limited individual impact' property.

(i) for every $B \subset A_i$, the function $\beta_i(B, \cdot) : T_i \to [0, 1]$ is measurable and (ii) for every $t_i \in T_i$, the function $\beta_i(\cdot, t_i) : A_i \to [0, 1]$ is a probability measure. For any distributional strategy σ_i, the regular conditional distributions $\sigma_i(B|t_i)$ defined over the Borel subsets B of A_i can be understood as behavior strategies $\beta_i(B, t_i) = \sigma_i(B|t_i)$. Conversely, for any behavioral strategy β_i, the corresponding distributional strategy σ_i is defined for all Borel subsets of $T_i \times A_i$ by $\sigma_i(S \times B) = \int_S \beta_i(B, t_i \tau_i(dt_i)$

[3]In the context of this paper, we are not concerned about existence of equilibria, since our objective is to establish stability properties of equilibria where they exist. As as aside, note that independence of types and uniform continuity of payoff functions in a game delivers existence of an equilibrium in distributional strategies.

Definition 27. *Given* $\lambda \geq 0$*, players in games in* $\Gamma(\mathscr{T}, \mathscr{A})$ *are said to have* λ***-limited individual impact*** *if the set of payoff functions* $\{u_i\}$ *satisfies the following condition:*

$$\forall i, \forall \text{ type-action characters } c_i, \ \ |u_i^{c_i}(c_{-i}) - u_i^{c_i}(c'_{-i})| \leq \frac{\lambda}{N-1},$$

$$\text{whenever } c_{-i}, c'_{-i} \text{ differ only in one coordinate.}$$

It turns out that uniform Lipschitz continuity in the scaled $L1$ metric together with uniform boundedness of the strategy space implies the above property.

Lemma 3. *Let* $\Gamma(\mathscr{T}, \mathscr{A}, K)$ *be a family of games that satisfies the uniform* K*-Lipschitz continuity. Then, there exists a constant* $\lambda \geq 0$ *such that players in* $\Gamma(\mathscr{T}, \mathscr{A}, K)$ *satisfy the* λ*-limited individual impact condition.*

Consider any game in $\Gamma(\mathscr{T}, \mathscr{A}, K)$, and fix any player i. For any two type-action character profiles of player i's rivals c_i and c_{-i} that differ only in one coordinate, the distance (in the scaled L1 metric) between the two profiles is less than $\frac{B}{N}$ for some upper bound B. The existence of such an upper bound B follows directly from the uniform bounded strategy space of $\Gamma(\mathscr{T}, \mathscr{A}, K)$. Now Lipschitz continuity implies that the for any type-action character of player i, the difference in player i's payoffs at c_i and c_{-i} can be at most $\frac{KB}{N}$. This implies that the λ-limited individual impact condition holds with $\lambda = KB$. $\square$

The limited individual impact condition asserts that the influence that any player can unilaterally exert on the payoff of an opponent is uniformly bounded and decreases with the number of players in the game. Notice that this condition does not imply anonymity, as a player's utility is not just affected by the aggregate actions of her rivals. The impact of each player on another can be potentially asymmetric and dependent on the type and identity of the rival. Further, there are no conditions imposed on the difference that a player can make to his own payoffs by changing his action choice unilaterally.

3.4. Ex-post Robustness in Large Games

Theorem 3. ***(Ex-post Robustness in Large Games)*** *Let $G \in \Gamma(\mathscr{T}, \mathscr{A}, K)$. Let σ^* be an equilibrium of G. Then, the following condition holds.*
Given $\epsilon > 0$, $\exists \bar{N}$ such that $\forall N > \bar{N}$, the equilibrium σ^* is $\epsilon(1+K)$ ex-post Nash.

Remark 4. Approximate ex-post Nash with probability 1: *The above theorem delivers a strong robustness property. If the game is large enough, with probability 1, players have arbitrarily low ex-post regret. In this sense, this result is stronger than that in Kalai (2004), where approximate ex-post robustness is obtained with high probability, but not with certainty. Kalai (2004) shows that in large games, Bayesian equilibrium strategy profiles yield ϵ-Nash type-action character profiles with high probability. The stronger result in this paper is partly attributable to the stronger assumption of Lipschitz continuity[4].*

We prove the above theorem in the following steps:

Lemma 4. ***(McDiarmid's Independent Bounded Differences Inequality)***[5] *Let $X = (X_1, X_2, ..., X_M)$ be a family of independent random vectors with $X_k \in \mathscr{X}_k$ for each k. Suppose that the real-valued function g defined on $\prod \mathscr{X}_k$ satisfies*

$$|g(x) - g(x')| \leq l_k \text{ whenever } x \text{ and } x' \text{ differ only in the } k^{th} \text{ coordinate.}$$

Let μ be the expected value of the random variable $g(x)$. Then for any $t \geq 0$,

$$\Pr(g(x) - \mu \geq t) \leq e^{\frac{-2t^2}{\sum_{k=1}^{M} l_k^2}}.$$

Proof: This result and the proof can be found in McDiarmid (1989). □

A straightforward application of the above result in our framework yields the following lemma.

[4]Lipschitz continuity in rivals' type-action character profiles as defined in this paper implies the uniform equicontinuity condition in Kalai (2004).
[5]We thank Colin McDiarmid for pointing us to this result.

Lemma 5. *Let σ be a strategy profile of the game $G(N, T, A, \{u_i\})$.*

$$\forall i, \forall c_i,\ \Pr(|u_i^{c_i}(c_{-i}) - \mu_i^{c_i}| > \alpha) \leq 2e^{\frac{-2(N-1)\alpha^2}{\lambda^2}},$$

where $\mu_i^{c_i} = \mathbb{E}u_i^{c_i}(c_{-i})$.

Proof: For any player i, a realized type-action character profile of her opponents is a sequence of independent random vectors $(C_1, ..., C_{i-1}, C_{i+1}, ..., C_N)$. By the limited individual impact condition, we can find a constant λ such that the following condition holds.

$$\forall i,\ \forall c_{-i} \in \mathscr{C}_{-i}, \text{ whenever } c_{-i}, c'_{-i} \text{ differ only in one coordinate,}$$

$$|u_i^{c_i}(c_{-i}) - u_i^{c_i}(c'_{-i})| \leq \frac{\lambda}{N-1}.$$

Fix any $\alpha > 0$ and any c_i. Applying Lemma 2 to the function $u_i^{c_i}$, we can get the above result. □

The two lemmata above make transparent the role of the "limited individual impact". The Independent Bounded Differences inequality tells us that the probability that a function of independent random variables deviates from its mean by any quantity t is inversely proportional to the maximum impact that each random variable has on the value of the function. In our set-up, limited individual impact means that any player's impact on a rival's payoffs is bounded, and moreover decreases with N. This implies in turn that the probability of deviations from the mean vanish exponentially fast with N.

We need to establish one more intermediate result before we prove the main robustness theorem.

Lemma 6. *Let p be a probability measure defined over X, where X is a compact subset of $\mathbb{R}^k$.*

$$\text{Given } \epsilon > 0, \exists\ \delta > 0 : \text{for all events } S,\ p(S) > 1 - \delta \implies p(S^\epsilon) = 1,$$

where S^ϵ is defined as the union of open balls of radius ϵ around elements of S, i.e. $S^\epsilon = \bigcup_{x\in S} B_\epsilon(x)$.

Proof: Suppose not. Then $\exists\ \epsilon > 0$ such that $\forall\delta > 0$, we can find a set S_δ such that $p(S_\delta) > 1-\delta$ but $p(S_\delta^\epsilon) < 1$. This implies that $\exists\ T_\delta$ (let T_δ be the complement of S_δ^ϵ) such that $p(T_\delta) > 0$ and $p(T_\delta^\epsilon) < \delta$.

Now consider $\bigcup_{x\in X} B_{\frac{\epsilon}{3}}(x)$ which is the union of $\frac{\epsilon}{3}$-balls centred around elements of X. This is an open cover of X. By compactness there exists a minimal finite open sub-cover $\mathscr{O} = \bigcup_{j=1,\ldots,J} B_{\frac{\epsilon}{3}}(x_j)$. Consider the sub-collection of sets $\mathscr{O}^+ = \{O \in \mathscr{O} : p(O) \neq 0\}$. Define $\bar{\delta} = Min_{O\in\mathscr{O}^+}\, p(O)$. In particular for this $\bar{\delta}$, $\exists\ S_{\bar{\delta}}$ such that $p(S_{\bar{\delta}}) > 1 - \bar{\delta}$ but $p(S_{\bar{\delta}}^\epsilon) < 1$. Further for $T = (S_{\bar{\delta}}^\epsilon)^C$, $p(T) > 0$ and $p(T^\epsilon) < \bar{\delta}$. Since $p(T) > 0$, T must intersect one of the open balls in $\mathscr{O}^+$, and consequently T^ϵ must include one of the open balls in $\mathscr{O}^+$. This implies that $p(T^\epsilon) > \bar{\delta}$ which is a contradiction. □

3.4.1. Proof of Theorem

Now we use the above lemmata to prove Theorem 3. Consider an equilibrium σ^* of a game in $\Gamma(\mathscr{T}, \mathscr{A}, K)$. Let $\epsilon > 0$ be given. Fix any player i. Suppose type t_i is realized. Consider an action a_i of player i. The following must be true:

$$\Pr[u_i^{t_i,a_i}(c_{-i}) - \mu_i^{t_i,a_i} < \epsilon] \geq 1 - e^{-\frac{2(N-1)\epsilon^2}{\lambda^2}}. \tag{3.1}$$

By Lemma 6, for this $\epsilon > 0$, we can find $\delta > 0$ such that for any $S \in \mathscr{C}_{-i}$, $P(S) > 1 - \delta \implies P(S^\epsilon) = 1$. In particular this is true for S defined as $S = \{c_{-i} \in \mathscr{C}_{-i} : |u_i^{t_i,a_i}(c_{-i}) - \mu_i^{t_i,a_i}| < \epsilon\}$. If N is large enough, the RHS of the above equation is greater than $(1-\delta)$. By Lemma 6, $P(S^\epsilon) = 1$. Further, for any $x_{-i} \in S^\epsilon$, $\exists c_{-i} \in S$ such that

$$u_i^{t_i,a_i}(x_{-i}) - \mu_i^{t_i,a_i} < |u_i^{t_i,a_i}(x_{-i}) - u_i^{t_i,a_i}(c_{-i})| + |u_i^{t_i,a_i}(c_{-i}) - \mu_i^{t_i,a_i}|.$$

Since the payoff functions are K-Lipschitz continuous, we get that

$$u_i^{t_i,a_i}(x_{-i}) - \mu_i^{t_i,a_i} < \epsilon(1+K).$$

Since the type-action character (t_i, a_i) was chosen arbitrarily, we have shown that for any player i and for any realized type-action character (t_i, a_i), with probability 1, player i gets a payoff that is at most $\epsilon(1+K)$ different from the expected payoff conditional on her realized type-action character.

$$\Pr[u_i^{t_i,a_i}(x_{-i}) - \mu_i^{t_i,a_i} < \epsilon(1+K)] = 1. \tag{3.2}$$

Now, since σ^* is an equilibrium of the game, except on a set of type-action characters of measure zero, the expected utility to player i from playing σ_i^* is higher than that from playing any other action. This along with (3.2) above implies that with probability 1, player i does not have a profitable deviation that will give her a gain of more than $\epsilon(1+K)$. □

3.5. Example

In this section, we present an example to demonstrate the main ideas of the paper.

Example 5 (Location Choice Game)**.** *Consider a group of N people $\{1, ..., N\}$ choosing where to live in the closed interval $[0, 1]$. Each person's type is the place where she was born within the interval (denoted by β), and is drawn uniformly from $[0, 1]$. Each player must choose a point on the interval $[0, 1]$ where she will live. Each person $2, ..., N$ wants to live where she was born. Player 1 is different. She wants to live far away from the average choice of the other people, but wants to live close to one person, her friend player 2. Formally, the game is described as follows.*

- *Each player's type or birthplace β_i is drawn independently from the uniform distribution on $[0, 1]$.*
- *Each player must choose location $l_i \in [0, 1]$.*

- *Utility functions of the players are given by*

$$u_i(\tau_i, l_i, \tau_{-i}, l_{-i}) = \begin{cases} -|l_i - \tau_i| & \textit{if } i \in \{2, ..., N\} \\ \frac{4}{5}|l_i - \frac{1}{N-1}\sum_{j \neq i} l_j| - \frac{1}{5}\frac{1}{N-1}|l_i - l_2| & \textit{if } i = 1. \end{cases}$$

Observe that this game is not anonymous, as player 1 specifically cares about the action choice of player 2. First, consider this game being played by a small group of $N = 3$ players.

Claim: The following strategies constitute a Bayesian Nash equilibrium.

- Players 2 and 3 choose to live where they were born. (i.e. $l_i = \beta_i$ for $i = 2, 3$.)
- Player 1 chooses to live at $l_1 = 1$.

Proof: Player 2 and player 3 will clearly choose actions exactly equal to their realized types. Given player 1's realized type, she will choose location l to such that

$$l = argmax_{l \in [0,1]} \ \frac{4}{5}\mathbb{E}\left[|l - \frac{1}{2}(\beta_2 + \beta_3)|\right] - \frac{1}{5}\mathbb{E}\left[\frac{1}{2}|l - \beta_2|\right].$$

We know types are drawn independently from a uniform distribution over the unit interval. This implies that $\beta_1 + \beta_2$ has probability density function

$$f(x) = \begin{cases} x & 0 \leq x \leq 1 \\ 2 - x & 1 \leq x \leq 2 \\ 0 & \text{otherwise.} \end{cases}$$

Using this, it can be easily shown that player 1's expected utility gets maximized at $l = 1$.

However, choosing location 1 is not ex-post Nash. To see why, consider the following realization of types: players 1 and 2 are born at $\beta = 0$ and player 3 is born at $\beta = 1$. Ex-post, player 1 has a profitable deviation. She would rather change her choice from 1 to 0. This would increase her payoff from $\frac{3}{10}$ to $\frac{4}{10}$.

Now observe that the utility functions in this example satisfy the limited individual impact condition. The types and action choices of players 1 and 3 have no impact on their rivals utilities, and so the condition is satisfied trivially. Player 2 influences player 1's utility, but her influence reduces with the number of players in the game. It can be checked that each player's utility function is also Lipschitz continuous in her rivals' type-action characters. Theorem 3 implies that if a large enough number of players were playing this location choice game, player 1 would have no ex-post regret.

The intuition is clear in this example. Though player 1 cares about her distance from her friend player 2, as the number of players increases, she cares less about player 2. Being away from the average becomes relatively more important. By choosing location 1, she maximizes her expected distance from the average. As the number of players becomes very large, the realized distance from the average is very close to what she expects.

3.6. Information Proofness

If we view an equilibrium as a steady state of an ongoing rational interactive process, our notion of ex-post stability is compelling. In particular, it may be reasonable to assume that players have access to 'signals' or partial information about other players' actions or types, or that they alter the sequence of play or have revision possibilities. It turns out that our ex-post stability implies not only robustness to complete revelation of information about the final outcome of the game, but also robustness to partial revelation of information or revision possibilities. In this sense, we say that Bayesian Nash equilibria in large games with limited individual impact are in fact *"information-proof"*.

Information-proofness offers a different perspective on private information in such games. Being information-proof means that even if all or some information about the relevant choices in the game were publicly available, each player would be interested only in her own private information. Full decentralization in the use of information is compatible with individual incentives.

Formally, let $\psi_i : \mathscr{C}_{-i} \rightarrow M$ be a signaling function that maps the type-action character profiles of player i's rivals to an abstract message set M. Partial information revelation can be modeled by such a signaling function. For an equilibrium to be robust to partial revelation of information, we require that, conditioning on the observed signal and one's own type and action, players do not want to change their choice of action. It is straightforward to see why this would hold. Fix a player i. Given any outcome c, consider any event of positive probability that is compatible with the realized type-action character (t_i, a_i) and $\psi_i(c)$. Ex-post stability implies that for any such event, player i cannot significantly increase her payoff by changing her action.

3.7. Conclusion

This paper examines the robustness properties of Bayesian Nash equilibria in large games. Earlier literature (See (Kalai 2004)) has shown that in large finite games with semi-anonymous players and smooth payoff functions, all Bayesian Nash equilibria are structurally robust and in particular are ex-post stable. However, little was known on the robustness properties of equilibria in games with infinite types and actions, or in games which are not anonymous. This paper attempts to address this gap.

We study Bayesian games where the type and action spaces are infinite - they can be compact subsets of $\mathbb{R}^k$. Further, players are no longer anonymous and can affect their rivals' payoffs in asymmetric, player-specific ways. We impose a single regularity condition - a variant of Lipschitz continuity - on the payoff functions of players in a game, and show that this condition is enough to guarantee ex-post stability of Bayesian Nash equilibria in our general class of games if the number of players is large enough. Interestingly, it turns out that the regularity condition implies that the unilateral impact a player can have on any of her rival's payoffs is bounded and reduces with the size of the game.

We establish that all Bayesian Nash equilibria in this class of games are ex-post stable in a strong sense. If players play Bayesian equilibrium strategies, if the number of players is large enough, with probability 1, players have arbitrarily low ex-post regret. They would not want to change their action even with perfect knowledge of their rivals' realized types and actions. This is stronger than results in earlier literature, which obtain ex-post stability in large games with high probability but not with certainty. Further, we remove the restriction of finiteness, and prove that this robustness property holds even in games with infinite type and action spaces.

It is not clear if similar robustness properties of Bayesian Nash equilibrium would hold if types were not independent or if we had payoff functions that displayed discontinuities. This would be an interesting line of investigation, and particularly important for applications.

Bibliography

[1] R. Aumann. Mixed vs. behavior strategies in infinite extensive games. *Annals of Mathematics Studies*, 52:627–650, 1964.

[2] E. Cartwright and M. Wooders. On purification of equilibrium in bayesian games and ex-post nash equilibrium. October 2006.

[3] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*, volume 38 of *Applications of Mathematics*. Springer, 1998.

[4] Avinash Dixit. On modes of economic governance. *Econometrica*, 71(2):449–481, 2003.

[5] Avinash Dixit. Trade expansion and contract enforcement. *Journal of Political Economy*, 111(6):1293–1317, 2003.

[6] Glenn Ellison. Cooperation in the prisoner's dilemma with anonymous random matching. *Review of Economic Studies*, 61(3):567–88, July 1994.

[7] Jeffrey C. Ely, Johannes Hörner, and Wojciech Olszewski. Belief-free equilibria in repeated games. *Econometrica*, 73(2):377–415, 03 2005.

[8] 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.

[9] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.

[10] Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, 1986.

[11] Kenneth N. Gilpin. Brands still easier to buy than create. New York Times, September 14 1992.

[12] R. Gradwohl and O. Reingold. Partial exposure and correlated types in large games. 2007.

[13] J. R. Green and J. J. Laffont. Posterior implementability in a two-person decision problem. *Econometrica*, 55:69–94, 1987.

[14] A. Greif. Contract enforceability and economic institutions in early trade: the maghribi traders' coalition. *American Economic Review*, 83(3):525–48, June 1993.

[15] A. Greif. *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press, 2006.

[16] H. Hakenes and M. Peitz. Observable reputation trading. *International Economic Review*, 48(2):693–730, 2007.

[17] J.Y. Halpern. A computer scientist looks at game theory. *Games and Economic Behavior*, 45(1):114–131, 2003.

[18] J. E. Harrington. Cooperation in a one-shot prisoner's dilemma. *Games and Economic Behavior*, 8:364–377, 1995.

[19] Johannes Hörner and Wojciech Olszewski. The folk theorem for games with private almost-perfect monitoring. *Econometrica*, 74(6):1499–1544, 2006.

[20] E. Kalai. Private information in large games. December 2000.

[21] E. Kalai. Large robust games. *Econometrica*, 72(6):1631–1665, November 2004.

[22] E. Kalai. Partially-specified large games. In *WINE*, pages 3–13, 2005.

[23] Michihiro Kandori. Social norms and community enforcement. *Review of Economic Studies*, 59(1):63–80, January 1992.

[24] Daniel B Klein. Promise keeping in the great society: A model of credit information sharing. *Economics and Politics*, 4:117–136, 1992.

[25] D. Kreps. Corporate culture and economic theory. *J. Alt and K. Shepsle (eds.), Perspectives on Positive Political Economy*, pages 90–143, 1990.

[26] G.J. Mailath and L. Samuelson. Who wants a good reputation? *Review of Economic Studies*, 68(2):415–41, April 2001.

[27] G.J. Mailath and L. Samuelson. *Repeated Games and Reputations*. Oxford University Press, 2006.

[28] C. McDiarmid. On the method of bounded differences. In J. Siemons, editor, *Surveys in Combinatorics*, volume London Mathematical Society Lecture Note Series 141, pages 148–188, 1989.

[29] C. McDiarmid. *Probabilistic Methods for Algorithmic Discrete Math.*, chapter Concentration, pages 1–54. Springer Verlag, 1998.

[30] P.R. Milgrom and R.J. Weber. Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10(4):619–632, November 1985.

[31] M. Okuno-Fujiwara and A. Postlewaite. Social norms and random matching games. *Games and Economic Behavior*, 9(1):79–109, April 1995.

[32] 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.

[33] Glenn Rifkin and Jenna Smith. Quickly erasing "i" and "b" and "m". New York Times, April 12 2006.

[34] R.W. Rosenthal. Sequences of games with varying opponents. *Econometrica*, 47(6):1353–66, November 1979.

[35] W.A. Rudin. *Real and Complex Analysis*. McGraw Hill Publishing Company, 1987.

[36] L. Samuelson. The economics of relationships, 2005.

[37] D. Schmeidler. Equilibrium points on nonatomic games. *Journal of Statistical Physics*, 17(4):295–300, 1973.

[38] S. Tadelis. What's in a name? reputation as a tradeable asset. *American Economic Review*, 89(3):548–563, June 1999.

[39] S. Tadelis. The market for reputations as an incentive mechanism. *Journal of Political Economy*, 110(4):854–882, August 2002.

[40] S. Tadelis. Firm reputation with hidden information. *Economic Theory*, 21(2):635–651, 03 2003.

[41] Satoru Takahashi. Community enforcement when players observe partners' past play. 2007.

APPENDIX A

Appendix for Chapter 1

A.1. Sequential Equilibrium

Section 1.3.3 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 1.3.2. Off-equilibrium strategies are defined as follows.

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

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

- $\forall k_i^t = \{(\nu^1, a^1), \ldots, (\nu^{t-1}, a^{t-1}), \nu^t\} \in \mathscr{K}_i^t$ 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 misreported 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.

- $\forall k_i^t = \{(\nu^1, a^1), \ldots, (\nu^{t-1}, a^{t-1}), \nu^t\} \in \mathscr{K}_i^t$ with $\nu_i^t \neq i$,

  $\sigma^*[k_i^t] = \text{argmax}_{a_i \in A_i} U_i(a_i, \sigma_{-i}^* | \xi_i[k_i^t])$.

  In other words, at any t-period interim private history in which a player has misreported 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. $\forall k_i^t = \{(\nu^1, a^1), \ldots, (\nu^{t-1}, a^{t-1}), \nu^t\} \in \mathscr{K}_i^t$ 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 7. *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 a rival name j who has observed d deviations so far. By 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. □

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

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

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

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

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.[1] Since the infinite product converges, there exists a constant κ such that

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

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

[1]This 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)*

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

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

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,

(A.1)
$$P \leq \sum_{m=1}^{t} \eta^2 \frac{1}{e^m} \sum_{i}^{\sqrt{2m}} m^i$$

$$\leq \sum_{m=1}^{\infty} \eta^2 \frac{1}{e^m} \sum_{i}^{\sqrt{2m}} m^i$$

(A.2)
$$= \eta^2 \sum_{m=1}^{\infty} \frac{1}{e^m} \frac{m(-1 + m^{\sqrt{2m}})}{(-1 + \sqrt{m})(1 + \sqrt{m})}.$$

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 (A.2) is convergent. Denote the limit by Λ. Convergence follows from the observation that

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

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 $\eta \to 0$, 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. $\square$

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

Consider a payoff profile $(v_1, v_2) \in \operatorname{Int} \mathscr{F}^*$. We proceed as in the equilibrium construction of Theorem 1. Pick payoff profiles $w^{GG}, w^{GB}, w^{BG}, w^{BB}$ such that,

(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 $\operatorname{conv}(\{w^{GG}, w^{GB}, w^{BG}, w^{BB}\})$ 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 $\{a_1^{GG},\ldots,a_N^{GG}\}$, $\{a_1^{GB},\ldots,a_N^{GB}\}$, $\{a_1^{BG},\ldots,a_N^{BG}\}$, $\{a_1^{BB},\ldots,a_N^{BB}\}$ such that each vector w^{XY}, the average discounted payoff vector over the sequence $\{a_1^{XY},\ldots,a_N^{XY}\}$ 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 $\{a_1^{XY},\ldots,a_N^{XY}\}$ described above.

A.2.1. Defining Strategies at Complete Histories: Name Announcements

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

A.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 o $\Gamma_{i,j}$ 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 $\psi : A \to \{g, b, x, y\}$ 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.

$$\mathscr{S}_i := \Big\{ s_i \in S_i^T : \forall \hat{h}_i^t = \Big(a, \psi(a), (a_i^{P_2,P_1}, a_{-i}^{P_2,P_1}), \ldots, (a_i^{P_2,P_1}, a_{-i}^{P_2,P_1})\Big), a \in P_i \times G,$$

$$s_i[\hat{h}_i^1] = \psi_i([\hat{h}_i^1]) \text{ and } s_i[\hat{h}_i^t] = a_i^{P_2,P_1}, t \geq 2 \Big\}.$$

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

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

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

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

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

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

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

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

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

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

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

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 $s_i \in \mathscr{S}_i$ against s_{-i}^G strictly exceeds $\bar{v}_1$ and her average payoff from using any action plan $s_i \in S_i^T$ 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 $r_i^G \in \mathscr{S}_i$ 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 $\mathscr{S}_i$. Define $r_i^B \in S_i^T$ 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 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 as follows.
  - 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 action profile (a_1, b_1) and 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 $c(\tau)$, $\tau \in \{1, \ldots, T\}$ 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 $\mathscr{S}_{-i}$. 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.

A.2.2.1. Beliefs of Players. At any private history, each player believes that in every period, she met the true owners of the names she met, and that no player ever misreported her name.

A.2.3. Proof of Equilibrium

First we show that conditional on truthful reporting of names, the 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^G_{-i} with probability μ and s^B_{-i} with probability $(1-\mu)$ where μ solves $u = \mu\left[\epsilon\underline{v}_i + (1-\epsilon)\bar{v}_i\right] + (1-\mu)\left[(1-\epsilon)\underline{v}_i + \epsilon\bar{v}_i\right]$

Transition Function: For any history $\hat{h}^T_{-i}$ for player $-i$, if the realized action plan is s^G_{-i} then at the end of the block, the state transits to $\bar{v}_i + (1-\delta)\omega^G_i(\hat{h}^T_{-i})$. If the realized action plan is s^B_{-i} the new state is $\underline{v}_i + (1-\delta)\omega^B_i(\hat{h}^T_{-i})$.

It can be easily seen that given $-i$'s action plan, any action plan of player i whose restriction belongs to $\mathscr{S}_i$ is a best response. The average payoff within a block from playing r^G_i against s^G_{-i}

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 $k \in \{0, 1 \ldots, T\}$, 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 $(p_0, \ldots, p_T)$.

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

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

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

$$\lim_{q\to\infty}(1,0,\ldots,0)\cdot H^q = (\frac{4}{3T+4}, \frac{3}{3T+4}, \frac{3}{3T+4}, \ldots, \frac{3}{3T+4}).$$

Further, it can be easily shown that

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

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. 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. So, 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 holds and misreporting is not a profitable deviation.

Now consider incentives for truthtelling 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 $t = 2, 3, 4$. $\square$

A.3. Proof of Folk Theorem for Multilateral Matching

A.3.1. Structure of Equilibrium

In equilibrium, players report their names truthfully. Each player plays the equilibrium strategies separately against each possible playgroup that she can be matched to. On 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 $(v_1, \ldots, v_K) \in int(\mathscr{F}^*)$, 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 atleast 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 occured.) 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.

A.3.2. Preliminaries

Consider any payoff profile $(v_1, \ldots, v_K) \in Int(\mathscr{F}^*)$. 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 $[\underline{v}_1, \bar{v}_1] \times \ldots \times [\underline{v}_K, \bar{v}_K]$ is contained in the interior of conv $\left(\{w^P : P = (P_1, \ldots, P_K),\ P_i \in \{G, B\}\}\right)$ 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 $\{a_1^P, \ldots, a_N^P\}$, with $P = (P_1, \ldots, P_K)$, $P_i \in \{G, B\}$, 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 $\{a_1^P, \ldots, a_N^P\}$.

A.3.3. Name Announcements at Complete Histories

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

A.3.3.1. 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 $\Gamma_{i,\mathscr{G}_{-i}}$) with each possible playgroup $\mathscr{G}_{-i}$. If her current playgroup's name profile is $\mathscr{G}_{-i}$, she plays game $\Gamma_{i,\mathscr{G}_{-i}}$. Fix a player i and a playgroup $\mathscr{G}_{-i}$. Below, I describe how game $\Gamma_{i,\mathscr{G}_{-i}}$ is played. Let $\hat{h}_i^t$ denote a t-period history in the game $\Gamma_{\mathscr{G}_{-i}}$. It specifies the action profiles played in the last t interactions of i with the playgroup $\mathscr{G}_{-i}$.

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

The game $\Gamma_{i,\mathscr{G}_{-i}}$ between i and playgroup $\mathscr{G}_{-i}$ 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 $g_i \neq b_i \ \forall i \in \{1, \ldots, K\}$. Define the signature function $\psi : A \to A$ 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 $\psi(.)$ arbitrarily otherwise. Suppose the observed plans are $(P_1, \ldots, P_K)$.

Let $\tilde{P} = (P_K, P_1, \ldots, P_{K-1})$. Define a set of action plans of a T-period finitely repeated game as follows.

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

$\mathscr{S}_i$ 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

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

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

$$\forall \hat{h}_i^t = \left(a, \psi(a), a^{\tilde{P}}, \ldots, a^{\tilde{P}}\right), t \geq 1, \ s_i^G[\hat{h}_i^t] = a_i^{\tilde{P}}.$$

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

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

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

$$\forall \hat{h}_i^t = \left(a, \psi_i(a), a^{\tilde{P}}, \ldots, a^{\tilde{P}}\right), t \geq 1,\ s_i^B[\hat{h}_i^t] = a_i^{\tilde{P}},$$

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

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

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

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

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 $s_i \in \mathscr{S}_i$ against s_{-i}^G strictly exceeds $\bar{v}_1$ and her average payoff from using any strategy $s_i \in S_i^T$ against s_{-i}^B is strictly below $\underline{v}_1$.

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 $r_{i+1}^G \in \mathscr{S}_i$ 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 $\mathscr{S}_{i+1}$. Define $r^B_{i+1} \in S^T_i$ to be an action plan such that given any history $\hat{h}^t_{i+1}$, $r^B_{i+1}|\hat{h}^t_{i+1}$ gives the highest payoffs against s^B_i and s_j for $j \neq i, i+1$ among all action plans in S^T_{i+1}. Redefine $\bar{v}$ and $\underline{v}$ so that $U_{i+1}(r^G_{i+1}, s^G_i) = \bar{v}_{i+1}$ and $U_{i+1}(r^B_{i+1}, s^B_i) = \underline{v}_{i+1}$.

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

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

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

- Players always report their names truthfully.
- Each player plays the following strategies separately against each possible playgroup.
  - **Initial Period of $\Gamma_{i,\mathscr{G}_{-i}}$:** Player i plays s^G_i with probability μ_0 and s^B_i 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^B_i 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^G_i with probability μ and s^B_i with probability $(1-\mu)$ where the mixing probability μ is used to tailor $(i+1)$'s continuation payoff.

    For any history $\hat{h}^T_i$ 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 $c(t), t \in \{1, \ldots, T\}$ denote the calendar time period of the t^{th} period of the most recently elapsed block.

    If s^B_i was played in the last block, we specify the reward $\omega^B_{i+1}(\cdot)$ as

$$\omega^B_{i+1}(\hat{h}^T_i) := \sum_{\tau=1}^{T} \pi^B_\tau$$

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 $\theta_t^G \leq 0$ for all actions that are used by strategies in $\mathscr{S}_{i+1}$. 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 in 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.

A.3.3.2. Beliefs of Players. After every history, players believe that in every period so far, they met the true owners of the names they encountered.

A.3.3.3. 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 $\mathscr{G}_{-i}$. The partial strategy for player i in her game $\Gamma_{i,\mathscr{G}_{-i}}$ 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 $u = \mu\left[\epsilon\underline{v}_{i+1} + (1-\epsilon)\bar{v}_{i+1}\right] + (1-\mu)\left[(1-\epsilon)\underline{v}_{i+1} + \epsilon\bar{v}_{i+1}\right]$

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 $\mathscr{S}_{i+1}$ 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 $\left(\frac{1}{M^{K-1}}\right)^{T+1-\tau}$ 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 a player i and a rival playgroup $\mathscr{G}_{-i}$ to be $k \in \{1, \ldots, T\}$ 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 wants to impersonate i_2. Conditional on her private history, i_1 can form beliefs over the possible states that each of her playgroups is in with respect to i_2. Suppose i_1 has met the sequence of playgroups $\{\mathscr{G}_{-i}^1, \ldots, \mathscr{G}_{-i}^{t-1}\}$. She knows that the playgroup she meets in any period remains in the same state with i_2 in that period. Fix any playgroup $\mathscr{G}_{-i}$ whom i_1 can be matched to. Player i_1 has a belief over the possible states $\mathscr{G}_{-i}$ is in with respect to i_2. Represent i_1's beliefs over the states by a vector $(p_1, \ldots, p_n)$.

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

(A.3)

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

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

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

To derive the above expression, note that player i_1 knows that in periods when she met playgroup $\mathscr{G}_{-i}$ it is not possible that i_2 met the same playgroup. Hence in these periods, the state of play between i_2 and $\mathscr{G}_{-i}$ 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 $\mathscr{G}_{-i}$ conditioning on the event that they met for the first time ever in period τ. For any τ, the probability that i_2 and $\mathscr{G}_{-i}$ met for the first time at period τ is given by $(\frac{M-2}{M-1})^{\sum_{l=1}^{\tau-1}\left(1-\mathbb{I}_{\mathscr{G}_{-i}=\mathscr{G}^{l}_{-i}}\right)}\frac{1}{M-1}$. Finally player i_1 knows that i_2 and $\mathscr{G}_{-i}$ could not have met for the first time in a period when she herself met playgroup $\mathscr{G}_{-i}$, and so does not need to condition on such periods.

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

(A.4) $$\lim_{q\to\infty}(1,0,\ldots,0)\cdot H^q=(\frac{1}{T},\ldots,\frac{1}{T}).$$

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

$$\forall q \geq 1, \qquad [(1,0,\ldots,0)\cdot H^q]_2 > 0, \tag{A.5}$$

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

It follows from (A.4) and (A.5) that for any playgroup $\mathscr{G}_{-i}$ whom i_1 has not met at least in one period, there exists a lower bound $\phi > 0$ such that the probability of $\mathscr{G}_{-i}$ 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 $\mathscr{G}'_{-i}$ 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 $\mathscr{G}'_{-i}$ and with probability $\frac{1}{M-1}$ this $\mathscr{G}'_{-i}$ would have met i_2 at $t = 1$ and period t could be their signature period.)

Consequently, if player i_1 decides to 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.[2]

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

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

Player i's loss in continuation payoff $\geq$

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

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 $\mathscr{G}_{-i}$ 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 $\mathscr{G}_{-i}$. 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.

$$\left(1-\frac{\delta}{\delta+M^{K-1}(1-\delta)}\right)\gamma\leq$$

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

For δ close enough to 1, this inequality is satisfied, and so misreporting ones name is not a profitable deviation. Now consider incentives for truth-telling in the first period of the supergame. Suppose i_1 impersonates i_2 at $t=1$ and meets playgroup $\mathscr{G}_{-i}$. In the next period, with probability $\frac{\epsilon^K}{M^{K-1}}$, i_2 will meet the same playgroup $\mathscr{G}_{-i}$ and use the wrong signature, thus informing $\mathscr{G}_{-i}$ 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. $\square$

APPENDIX B

Appendix for Chapter 2

B.1. Proof of Proposition 1

By Lemma 1, we know that in SHE, competent firms playing only S-names. This implies that incompetent firms must use both S and N-names. The equilibrium conditions are as follows:

Incentive Compatibility for Name Choice

C-firms strictly prefer using S-names and working hard to using N-names.

$$-V_S + w_S - c + \delta(1-\rho)f_1 V_S + \delta\rho f_2 V_S > w_N - c + \delta(1-\rho)f_3 V_S + \delta\rho f_4 V_S. \tag{B.1}$$

$$-V_S + w_S - c + \delta(1-\rho)f_1 V_S + \delta\rho f_2 V_S > w_N + \delta f_4 V_S. \tag{B.2}$$

I-firms are indifferent between N-names and S-names.

$$-V_S + w_S + \delta f_2 V_S = w_N + \delta f_4 V_S. \tag{B.3}$$

Incentives for Competent Firms to Work Hard

$$c < \delta(1-\rho)(f_1 - f_2)V_S. \tag{B.4}$$

Equilibrium Determination

(B.5)

Market Clearing: $\phi+(1-\phi)\mu_S = \phi(1-\rho)f_1+\phi\rho f_2+(1-\phi)\mu_S f_2+(1-\phi)(1-\mu_S)f_4.$

(B.6)

$$\text{Wage Determination:} \qquad w_N = 0 \text{ and } w_S > w_N.$$

The effort constraint implies that $f_1 = 1$ and $f_2 = 0$. Further, I claim that $f_4 \neq f_2$. Suppose not, i.e. $f_4 = f_2 = 0$. Then (B.5) implies $\mu_S = \frac{-\phi\rho}{1-\phi}$ which is not possible. In other words, there would be a shortage of S-names and the market would not clear. Hence, we have $f_4 = 1 \neq f_2$. This leaves possible only two transition rules with $f_3 = 1$ or $f_3 = 0$. These are precisely the ones depicted in the figure. Consider the first transition rule in the figure (where $f_3 = 1$). Conditions (B.1)-(B.5) reduce to:

(B.7)
$$\mu_S = \frac{1 - \phi - \phi\rho}{2(1-\phi)}.$$

(B.8)
$$w_S = \frac{\phi(1-\rho)}{\phi + (1-\phi)\mu_S}.$$

(B.9)
$$V_S = \frac{w_S}{1+\delta}.$$

(B.10)
$$c < \delta(1-\rho)V_S.$$

It can easily be seen that if $\phi < \frac{1}{1+\rho}$, these conditions yield a non-empty set of sorting equilibria. Identical conditions arise for the second transition rule. □

B.2. Richer Market with Non-Random Transitions

Consider a richer market with three names. It turns out that there are only three non-random transition rules which allow sorting high-effort equilibria. Consider the automata below: SHE exist only under these transition rules. The nature of sorting differs based on the transition rule. With the first two rules, we get "Trust 1-Names" equilibria. In these equilibria, competent

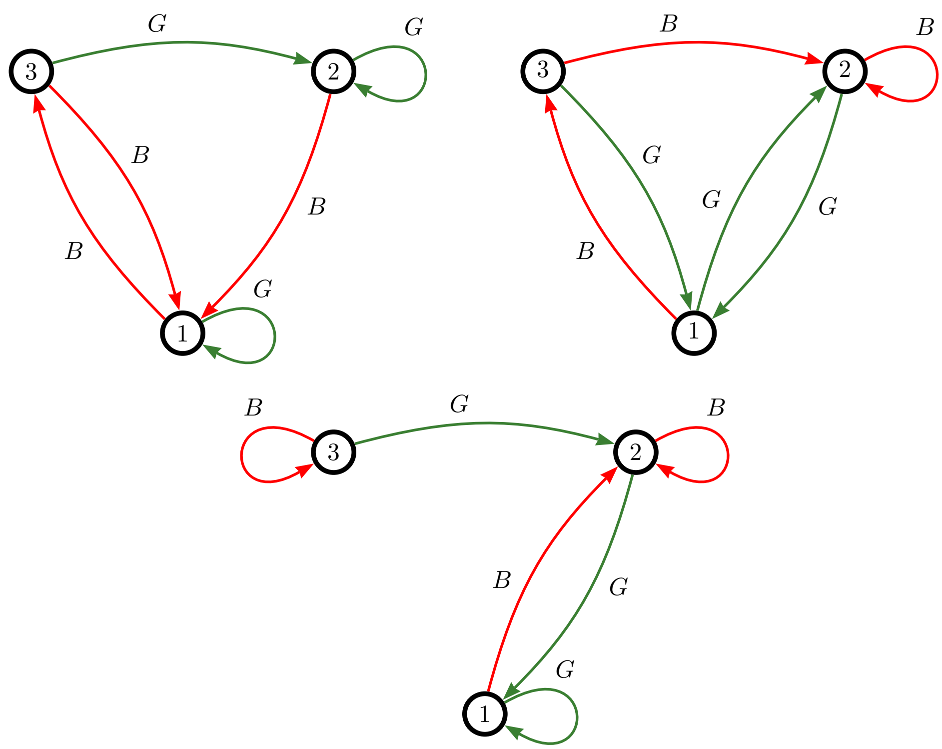

Figure B.1. Transitions for SHE in Market with Three States

firms buy all three names, but only competent firms buy S_1-names. Incompetent firms buy only S_2 and S_3-names. The last transition rule yields "Mistrust 3-Names". Here, competent firms never enter with costless S_3-names. I-firms are the only ones using S_3 names.

B.3. Proof of Proposition 2

Sufficiency: Fix ϕ, ρ, δ, c with $c < \min\left\{\frac{\delta(1-\rho)^2(1-\phi)}{1-\phi+\delta\phi\rho}, \frac{\delta(1-\phi)(1-\rho)^2}{\delta(1-\phi)+\phi\rho}\right\}$.

Case(i) $c < \frac{\delta(1-\rho)^2(1-\phi)}{1-\phi+\phi\rho}$.

Set the transition probabilities as follows:

$$\gamma_1 = \gamma_3 = 1, \gamma_2 = 0, \gamma_4 < \min\left\{\frac{\phi\rho}{1-\phi}, \frac{\delta(1-\rho)^2(1-\phi) - c(1-\phi+\phi\rho)}{\delta(1-\rho)^2(1-\phi) + c\delta\phi(1-\phi+\phi\rho)}\right\}.$$

I claim that there exists a "Trust S-Names" equilibrium where C-firms use S-names with probability $\sigma_S = \frac{\phi(1-\rho)(1-\gamma_4)+\gamma_4}{\phi(1+\rho\gamma_4)}$ and N-names with probability $(1-\sigma_S)$. Competent firms always choose to work hard. Equilibrium price $V_S = \frac{(1-\rho)(1-\phi)}{1+\delta\rho\gamma_4}\frac{1+\rho\gamma_4}{\rho+(1-\phi)(1-\rho)(1-\gamma_4)}$.

Notice first that $c < \frac{\delta(1-\rho)^2(1-\phi)}{1-\phi+\phi\rho}$ implies that γ_4 is well-defined. Further, since $\gamma_4 < \frac{\phi\rho}{1-\phi}$, the strategy σ_S is well-defined. Given the conjectured equilibrium strategies, the wages would be $w_S = 1-\rho$ and $w_N = \frac{(1-\rho)[\phi\rho-(1-\phi)\gamma_4]}{1-\phi+\phi\rho-\gamma_4(1-\rho)(1-\phi)}$. The incentive constraints for the C-type implies

$$-V_S + w_S - c + \delta(1-\rho)V_S = w_N - c + \delta(1-\rho)V_S + \delta\rho\gamma_4 V_S \implies V_S = \frac{w_S - w_N}{1+\delta\rho\gamma_4}. \tag{B.11}$$

It is easy to check that the proposed V_S satisfies the above. I-type's incentive constraints imply

$$V_S > \frac{w_S - w_N}{1+\delta\gamma_4}. \tag{B.12}$$

The conjectured V_S also satisfies this condition. Here, incentive compatibility of the competent type implies incentive compatibility for the incompetent type. The market clearing condition is also satisfied with these equilibrium strategies. Finally for the competent type to exert effort we need to check the following conditions.

$$c < \delta(1-\rho)V_S \qquad \text{and} \qquad c < \delta(1-\rho)(1-\gamma_4)V_S. \tag{B.13}$$

Clearly, it suffices to show that the second constraint holds. We need the following inequality:

$$c < \frac{\delta(1-\rho)^2(1-\phi)(1-\gamma_4)(1+\rho\gamma_4)}{[1+\delta\rho\gamma_4][\rho+(1-\phi)(1-\rho)(1-\gamma_4)]}.$$

$$\text{Now,} \qquad (1+\rho\gamma_4)(1+\delta\phi\gamma_4) = 1+\delta\rho\gamma_4 + \rho\gamma_4(1-\delta) + \delta\phi\gamma_4(1+\rho\gamma_4).$$

$$\text{So,} \qquad (1+\rho\gamma_4)(1+\delta\phi\gamma_4) \geq 1+\delta\rho\gamma_4.$$

$$\text{Further,} \qquad (1-\phi+\phi\rho) = \rho + (1-\phi)(1-\rho) \geq \rho + (1-\phi)(1-\rho)(1-\gamma_4)$$

$$\implies \frac{1}{(1+\delta\phi\gamma_4)(1-\phi+\phi\rho)} \leq \frac{1+\rho\gamma_4}{[\rho+(1-\phi)(1-\rho)(1-\gamma_4)][1+\delta\rho\gamma_4]}.$$

This implies that

$$\text{(B.14)} \quad \frac{\delta(1-\rho)^2(1-\phi)(1-\gamma_4)}{(1+\delta\phi\gamma_4)(1-\phi+\phi\rho)} \leq \frac{\delta(1-\rho)^2(1-\phi)(1-\gamma_4)[1+\rho\gamma_4]}{[\rho+(1-\phi)(1-\rho)(1-\gamma_4)][1+\delta\rho\gamma_4]} = \delta(1-\rho)(1-\gamma_4)V_S.$$

But recall that by definition of γ_4,

$$\gamma_4 < \frac{\delta(1-\rho)^2(1-\phi) - c(1-\phi+\phi\rho)}{\delta(1-\rho)^2(1-\phi) + c\delta\phi(1-\phi+\phi\rho)}$$

$$\implies c < \frac{\delta(1-\rho)^2(1-\phi)(1-\gamma_4)}{(1+\delta\phi\gamma_4)(1-\phi+\phi\rho)}.$$

So, by the above inequality, $c < \delta(1-\rho)(1-\gamma_4)V_S$.

Case(ii) $\frac{\delta(1-\rho)^2(1-\phi)}{1-\phi+\phi\rho} \leq c < \min\left\{\frac{\delta(1-\rho)^2(1-\phi}{1-\phi+\delta\phi\rho}, \frac{\delta(1-\phi)(1-\rho)^2}{\delta(1-\phi)+\phi\rho}\right\}$.
First consider cases where $\phi\rho < 1-\phi$. This implies that $\frac{\delta(1-\rho)^2(1-\phi)}{1-\phi+\phi\rho} \leq c < \frac{\delta(1-\rho)^2(1-\phi)}{1-\phi+\delta\phi\rho}$. Fix the following transition probabilities:

$$\gamma_1 = \gamma_3 = 1, \gamma_2 = 0, \gamma_4 \in \left(\frac{\phi\rho}{1-\phi}, \frac{\delta\phi(1-\rho)^2(1+\frac{\phi\rho}{1-\phi}) - c\phi(1-\rho)(1+\frac{\delta\phi\rho}{1-\phi})}{c(1+\frac{\delta\phi\rho}{1-\phi})}\right).$$

I claim that that there exists a "Mistrust N-Names" equilibrium where, incompetent firms use S-names with probability $\mu_S = \frac{(1-\phi)\gamma_4-\phi\rho}{(1-\phi)(1+\gamma_4)}$ and N-names with $1-\mu_S$.Competent firms choose to work hard on the equilibrium path. Equilibrium price $V_S = \frac{\phi(1-\rho)(1+\gamma_4)}{[1+\delta\gamma_4][\gamma_4+\phi(1-\rho)]}$.

First, we check that the transition probabilities chosen are well-defined.

$$c < \frac{\delta(1-\rho)^2(1-\phi)}{1-\phi+\delta\phi\rho}$$

$$\implies \phi\delta(1-\rho)^2(1-\phi) - \phi c(1-\phi+\delta\phi\rho) > 0$$

$$\implies \phi\delta(1-\rho)^2(1-\phi+\delta\phi\rho) - (1-\rho)\phi c(1-\phi+\delta\phi\rho) > 0$$

$$\implies (1+\frac{\phi\rho}{1-\phi})\phi\delta(1-\rho)^2-(1+\frac{\delta\phi\rho}{1-\phi})(1-\rho)\phi c>0.$$

So the upper bound for γ_4 is well-defined. Further,

$$c<\frac{\delta(1-\rho)^2(1-\phi)}{1-\phi+\delta\phi\rho}\implies \frac{\phi\rho}{1-\phi}<\frac{\delta\phi(1-\rho)^2(1+\frac{\phi\rho}{1-\phi})-c\phi(1-\rho)(1+\frac{\delta\phi\rho}{1-\phi})}{c(1+\frac{\delta\phi\rho}{1-\phi})}.$$

So, the interval from which γ_4 is chosen is well-defined. Given the conjectured strategies the wages are $w_N=0$ and $w_S=\frac{\phi(1-\rho)(1+\gamma_4)}{\gamma_4+\phi(1-\rho)}$. For I-firms to be indifferent between S and N-names, we need:

$$V_S=\frac{w_S-w_N}{1+\delta\gamma_4}.$$

We can check that V_S satisfies this condition. Hence, C-firms strictly prefer S-names to N-names. Since $\gamma_4>\frac{\phi\rho}{1-\phi}$, μ_S is well-defined, and satisfies the market clearing conditions. It only remains to check that competent firms have an incentive to work hard on the equilibrium path, i.e. $c<\delta(1-\rho)V_S$. To prove this, define a function

$$\psi(x)=\frac{\delta\phi(1-\rho)^2}{\gamma_4+\phi(1-\rho)}\frac{1+x}{1+\delta x}.$$

Notice that this function is strictly increasing in x. Recall that

$$\gamma_4<\frac{\delta\phi(1-\rho)^2(1+\frac{\phi\rho}{1-\phi})-c\phi(1-\rho)(1+\frac{\delta\phi\rho}{1-\phi})}{c(1+\frac{\delta\phi\rho}{1-\phi})}\implies c<\frac{\delta\phi(1-\rho)^2(1+\frac{\phi\rho}{1-\phi})}{(1+\frac{\delta\phi\rho}{1-\phi})(\gamma_4+\phi(1-\rho)}=\psi(\frac{\phi\rho}{1-\phi}).$$

By the monotonicity of $\psi(.)$, if $\gamma_4>\frac{\phi\rho}{1-\phi}$ then $c<\psi(\gamma_4)$. So, for γ in the specified range,

$$c<\psi(\gamma_4)=\frac{\delta\phi(1-\rho)^2(1+\gamma_4)}{(\gamma_4+\phi(1-\rho))(1+\delta\gamma_4)}=\delta(1-\rho)V_S.$$

This proves that we have found a "Mistrust N-Names" equilibrium.

Using a very similar argument as above, we can show that "Mistrust N-Names Equilibria" exist also in the case when $\phi\rho>1-\phi$.

Necessity: Propositions 4 & 5 characterize the types of SHE that exist. We can use these characterizations to show that a necessary condition for SHE to exist is that the cost of hard work is less than the upper bound.

Consider "Trust S-names" equilibria. Using the characterization in Proposition 4, we derive a maximal cost of effort for which it is possible to find transition rules to support SHE. Finding such an upper bound reduces to solving the following constrained maximization problem.

$$\max_{\gamma_1,\gamma_2,\gamma_3,\gamma_4} \frac{\delta(1-\rho)^2(1-\phi)(\gamma_3-\gamma_4)}{1+\delta(1-\rho)(\gamma_3-\gamma_1)+\delta\rho(\gamma_4-\gamma_2)} \frac{1+(1-\rho)(\gamma_3-\gamma_1)+\rho(\gamma_4-\gamma_2)}{1-(1-\rho)\gamma_1-\rho\gamma_2+(1-\phi)(1-\rho)(\gamma_3-\gamma_4)}$$

(B.15) subject to $\gamma_1-\gamma_2-\gamma_3+\gamma_4 > 0$

(B.16) $$-\gamma_4(1-\phi)+\phi(1-(1-\rho)\gamma_1-\rho\gamma_2) > 0$$

(B.17) $$\gamma_1,\gamma_2,\gamma_3,\gamma_4 \geq 0$$

(B.18) $$\gamma_1,\gamma_2,\gamma_3,\gamma_4 \leq 1.$$

The constraints are linear and so the constraint qualification condition holds. Further it can be verified that constraints (B.15) and (B.16) bind. So plugging back γ_4 and γ_3 and solving the reduced problem, we find that the objective function is maximized at

$$\gamma_1=\gamma_3=1 \qquad \gamma_2=\gamma_4=\frac{\phi\rho}{1-\phi+\phi\rho} \text{and the maximal value is } c^1_{max}=\frac{\delta(1-\phi)(1-\rho)^2}{1-\phi+\phi\rho}.$$

Next, consider "Mistrust N-names" equilibria. We will examine two cases:

Case A: Cost of hard work is low enough that it is sequentially rational for a C-firm to work hard on and off the equilibrium path. (i.e., $c < \min\{\delta(1-\rho)(\gamma_3-\gamma_4)V_S, \delta(1-\rho)(\gamma_1-\gamma_2)V_S\}$.)

Case B: Cost of work is low enough to sustain hard work only on the equilibrium path.

For equilibria under Case (A), the conditions of Proposition 5 imply
$c < \delta(1-\rho)(\gamma_3-\gamma_4)\frac{\phi(1-\rho)}{[\phi+(1-\phi)\mu_S][1+\delta(\gamma_4-\gamma_2)]}$. To find the maximal value of the RHS of this constraint, we solve the following problem:

$$\max_{\gamma_1,\gamma_2,\gamma_3,\gamma_4} \frac{\delta(1-\rho)^2\phi(\gamma_3-\gamma_4)}{1+\delta(\gamma_4-\gamma_2)} \frac{1+\gamma_4-\gamma_2}{\phi(1-\rho)(\gamma_1-\gamma_2)+\gamma_4}$$

(B.19) subject to $\gamma_1 - \gamma_2 - \gamma_3 + \gamma_4 > 0$

(B.20) $$\gamma_4(1-\phi) - \phi(1-(1-\rho)\gamma_1 - \rho\gamma_2) > 0$$

(B.21) $$1 - \gamma_2 - \phi(1-\rho)(\gamma_1 - \gamma_2) > 0$$

(B.22) $$\gamma_1, \gamma_2, \gamma_3, \gamma_4 \geq 0$$

(B.23) $$\gamma_1, \gamma_2, \gamma_3, \gamma_4 \leq 1.$$

Here constraints (B.19) and (B.20) bind. Solving the reduced problem we find the objective function is maximized at

$$\gamma_1 = \gamma_3 = 1 \qquad \gamma_2 = \gamma_4 = \frac{\phi\rho}{1-\phi+\phi\rho} \text{and the maximal value of } c \text{ is } c^2_{max} = \frac{\delta(1-\phi)(1-\rho)^2}{1-\phi+\phi\rho}.$$

Note that in the two maximization problems, the maximum and the maximizers are exactly the same. For the equilibria covered by Case (B), conditions of Proposition 5 imply
$\delta(1-\rho)(\gamma_3-\gamma_4)\frac{\phi(1-\rho)}{[\phi+(1-\phi)\mu_S][1+\delta(\gamma_4-\gamma_2)]} \leq c < \delta(1-\rho)(\gamma_1-\gamma_2)\frac{\phi(1-\rho)}{[\phi+(1-\phi)\mu_S][1+\delta(\gamma_4-\gamma_2)]}$.

Again, we solve the constrained maximization problem given by

$$\max_{\gamma_1,\gamma_2,\gamma_4} \frac{\delta(1-\rho)^2\phi(\gamma_1-\gamma_2)}{1+\delta(\gamma_4-\gamma_2)}\frac{1+(\gamma_4-\gamma_2)}{\gamma_4+\phi(1-\rho)(\gamma_1-\gamma_2)}$$

(B.24) subject to $$\gamma_4(1-\phi)-\phi(1-(1-\rho)\gamma_1-\rho\gamma_2)>0$$

(B.25) $$1-\gamma_2-\phi(1-\rho)(\gamma_1-\gamma_2)>0$$

(B.26) $$\gamma_1,\gamma_2,\gamma_4\geq 0$$

(B.27) $$\gamma_1,\gamma_2,\gamma_4\leq 1.$$

It can be verified that:

1. For ϕ,ρ such that $\phi<\frac{1}{1+\rho}$, the objective function is maximized at

$$\gamma_1=1 \qquad \gamma_2=0 \qquad \gamma_4=\frac{\phi\rho}{1-\phi} \qquad \text{and the maximal value of } c \text{ is } c^3_{max}=\frac{\delta(1-\phi)(1-\rho)^2}{1-\phi+\delta\phi\rho}.$$

2. For ϕ,ρ such that $\phi>\frac{1}{1+\rho}$, the objective function gets maximized at

$$\gamma_1=1 \qquad \gamma_2=1-\frac{1-\phi}{\phi\rho} \qquad \gamma_4=1 \quad \text{and the maximal value of } c \text{ is } c^4_{max}=\frac{\delta(1-\phi)(1-\rho)^2}{\delta(1-\phi)+\phi\rho}.$$

Inspecting the bounds yields that a necessary condition for sorting equilibria to exist is:

$$c<\min\left\{\frac{\delta(1-\phi)(1-\rho)^2}{1-\phi+\delta\phi\rho}\;,\;\frac{\delta(1-\phi)(1-\rho)^2}{\delta(1-\phi)+\phi\rho}\right\}. \qquad \square$$

B.4. Proof of Proposition 4

Consider "Trust S-names" equilibria. The equilbrium conditions are as follows:

Incentive Compatibility for Name Choice

C-firms must be indifferent between N and S-names.

$$-V_S + w_S - c + \delta(1-\rho)\gamma_1 V_S + \delta\rho\gamma_2 V_S = w_N - c + \delta(1-\rho)\gamma_3 V_S + \delta\rho\gamma_4 V_S.$$

(B.28) So, $$V_S = \frac{w_S - w_N}{1 + \delta(1-\rho)(\gamma_3 - \gamma_1) + \delta\rho(\gamma_4 - \gamma_2)}.$$

I firms must strictly prefer N-names to S-names.

(B.29) $$w_N + \delta\gamma_4 V_S > -V_S + w_S + \delta\gamma_2 V_S \implies V_S[1 + \delta(\gamma_4 - \gamma_2)] > w_S - w_N.$$

Incentives for Competent Firms to Work Hard

(B.30) $$c < \delta(1-\rho)V_S(\gamma_1 - \gamma_2).$$

(B.31) $$c < \delta(1-\rho)V_S(\gamma_3 - \gamma_4).$$

Equilibrium Determination

Market Clearing: $$\phi\sigma_S = \phi\sigma_S(1-\rho)\gamma_1 + \phi\sigma_S\rho\gamma_2 + \phi(1-\sigma_S)(1-\rho)\gamma_3 + \phi(1-\sigma_S)\rho\gamma_4 + (1-\phi)\gamma_4.$$

(B.32) So, $$\sigma_S = \frac{\phi(1-\rho)(\gamma_3 - \gamma_4) + \gamma_4}{\phi(1-\rho)(\gamma_3 - \gamma_1) + \phi\rho(\gamma_4 - \gamma_2) + \phi}.$$

Wage Determination: $$w_S = 1 - \rho > w_N = \frac{\phi(1-\sigma_S)(1-\rho)}{\phi(1-\sigma_S) + 1 - \phi}.$$

(B.33) $$V_S > 0.$$

(B.34) $$\sigma_S \in (0,1).$$

The incentive constraints of C-firms and I-firms together imply that (1) holds. The effort constraint for a C-firm with an S-name (B.30) is equivalent to condition (3). Consider the market clearing condition (B.32). Since the effort constraint (B.31) implies that $\gamma_3 > \gamma_4$, we know that the numerator of (B.32) is positive. So, $\sigma_S \in (0,1)$, implies that the denominator in (B.32) must be greater than the denominator. This implies (2) holds.

Conversely, assume conditions (1) - (3) in the proposition hold. I claim the following "Trust S-names" SHE exists. C-firms buy both S and N-names. They buy S-names with probability $\sigma_S = \frac{\phi(1-\rho)(\gamma_3-\gamma_4)+\gamma_4}{\phi(1-\rho)(\gamma_3-\gamma_1)+\phi\rho(\gamma_4-\gamma_2)+\phi}$. I-firms buy only N-names. So, wages are $w_S = 1-\rho$ and $w_N = \frac{\phi(1-\sigma_S)(1-\rho)}{\phi(1-\sigma_S)+1-\phi}$. S-names trade at a price $V_S = \frac{w_S - w_N}{1+\delta(\gamma_4-\gamma_2)}$.

For an equilibrium, we need to check (B.28) to (B.34). We know $V_S \geq 0$. By definition, (B.28) and (B.32) hold and (3) is equivalent to (B.31). Since $c > 0$, (1) and (3) imply (B.30) holds. (1) also implies (B.29) holds. Finally, (1) and (3) imply that the numerator in (B.32) is positive. Condition (2) implies that the numerator is strictly lesser than the denominator; so (B.34) holds.

□

B.5. Proof of Proposition 5

The proof is similar to Proposition 4. Suppose for given $\phi, \rho, \gamma_1, \gamma_2, \gamma_3, \gamma_4, \delta$ and c, there exists a "Mistrust N-names" equilibrium. The equilibrium conditions are as follows:

Incentive Compatibility for Name Choice

C-firms strictly prefer an S-name to an N-name.

$$-V_S + w_S - c + \delta(1-\rho)\gamma_1 V_S + \delta\rho\gamma_2 V_S > w_N - c + \delta(1-\rho)\gamma_3 V_S + \delta\rho\gamma_4 V_S.$$

(B.35) $$V_S < \frac{w_S - w_N}{1 + \delta(1-\rho)(\gamma_3 - \gamma_1) + \delta\rho(\gamma_4 - \gamma_2)}.$$

$$\text{Further,} \; -V_S + w_S - c + \delta(1-\rho)\gamma_1 V_S + \delta\rho\gamma_2 V_S > w_N + \delta\gamma_4 V_S.$$

(B.36) $$\text{So,} \qquad V_S(1 + \delta\gamma_4 - \delta(1-\rho)\gamma_1 - \delta\rho\gamma_2) < w_S - w_N - c.$$

I-firms are indifferent between N-names and S-names.

(B.37) $$V_S[1 + \delta(\gamma_4 - \gamma_2)] = w_S - w_N.$$

Incentives for Competent Firms to Work Hard

(B.38) $$c < \delta(1-\rho)V_S(\gamma_1 - \gamma_2).$$

Equilibrium Determination

(B.39) $$\text{Market Clearing} \implies \mu_S = \frac{\phi(1-\rho)\gamma_1 + \phi\rho\gamma_2 + (1-\phi)\gamma_4 - \phi}{(1-\phi)(1+\gamma_4-\gamma_2)}.$$

$$\text{Wage Determination} \qquad w_N = 0 < w_S = \frac{\phi(1-\rho)}{\phi + (1-\phi)\mu_S}.$$

(B.40) $$V_S > 0. \text{ and}$$

(B.41) $$\mu_S \in (0,1).$$

Clearly, we need $\mu_S \in (0,1)$. This implies that the second inequality in Conditions 1 and 2 hold. (B.38) is equivalent to Condition 3. Finally, the incentive constraints (B.35) and (B.37) together imply that the first inequality in Condition 1 holds. This proves the necessary conditions.

Conversely, suppose that there for given $\phi, \rho, \gamma_1, \gamma_2, \gamma_3, \gamma_4, \delta$ and c, the conditions of the proposition are satisfied. I claim that there exists a "Mistrust N-names" equilibrium in which incompetent firms buy S-names with probability μ_S given by (B.39) and where the price of an S-name, V_S is given by (B.37). To verify that this is an equilibrium, we need to check for conditions (B.35) through (B.41). Conditions (B.37) through (B.40) are trivially satisfied.

$$\gamma_3 - \gamma_1 < \gamma_4 - \gamma_2 \implies 1 + \delta(1-\rho)(\gamma_3 - \gamma_1) + \delta\rho(\gamma_4 - \gamma_2) < 1 + \delta(\gamma_4 - \gamma_2)$$

$$\implies V_S[1 + \delta(1-\rho)(\gamma_3 - \gamma_1) + \delta\rho(\gamma_4 - \gamma_2)] < V_S[1 + \delta(\gamma_4 - \gamma_2)]$$

$$\implies (B.35) \text{ is satisfied.}$$

(A4) and (B.35) imply that the other IC of the C-type (B.36) is also satisfied. Condition (3) and (B.38) are equivalent. Conditions (1) and (2) ensure that (B.41) holds. It remains to be shown that the second incentive constraint for the C-type (B.36) also holds. Consider two cases:

Case(i) $c < \delta(1-\rho)(\gamma_3 - \gamma_4)\frac{\phi(1-\rho)}{[\phi+(1-\phi)\mu_S][1+\delta(\gamma_4-\gamma_2)]}$: c is low enough for C-firms to work hard in every state. (B.36) is implied by the first IC constraint of the C-type and we are done.

Case (ii) $c \geq \delta(1-\rho)(\gamma_3 - \gamma_4)\frac{\phi(1-\rho)}{[\phi+(1-\phi)\mu_S][1+\delta(\gamma_4-\gamma_2)]}$: Here, a C-firm works hard only if it buys an S-name. Then, the IC of I-type and (3) imply (B.36) holds. This proves sufficiency. □