Efficient Contracting in Network Financial Markets*

Darrell Duffie and Chaojun Wang
Stanford University

October 22, 2016

Abstract

We model bargaining in over-the-counter network markets over the terms and prices of contracts. Of concern is whether bilateral non-cooperative bargaining is sufficient to achieve efficiency in this multilateral setting. For example, will market participants assign insolvency-based seniority in a socially efficient manner, or should bankruptcy laws override contractual terms with an automatic stay? We provide conditions under which bilateral bargaining over contingent contracts is efficient for a network of market participants. Examples include seniority assignment, close-out netting and collateral rights, secured debt liens, and leverage-based covenants. Given the ability to use covenants and other contingent contract terms, central market participants efficiently internalize the costs and benefits of their counterparties through the pricing of contracts. We provide counterexamples to efficiency for less contingent forms of bargaining coordination.

JEL Classifications: D47, D60, D70, G12, K22

Keywords: Over-the-counter, network, contracting, bargaining, efficiency, game theory

*This version is preliminary and incomplete. We are grateful for early conversations with Martin Oehmke, David Skeel, and Jeff Zwiebel. Comments are welcome.

1 Introduction

For a network market setting, we propose a theory of bilateral bargaining over the terms and pricing of contracts that may be contingent on the contracts signed by other pairs of agents in the network. We state conditions under which contingent bilateral contracting is socially efficient, subject to the available sets of contracts. We provide counterexamples to efficiency in settings with less effective forms of bargaining coordination. We develop two solution concepts: an extensive form alternating-offers bargaining game equilibrium refinement, and an axiomatic-bargaining solution. We show that the solutions arising from these approaches coincide in the case of three-player networks. (Our extensive-form alternating-offers game treats only the case of three-player networks.)

An example application depicted in Figure 1 is the contracting between a debtor firm and a creditor, and between the same debtor firm and a derivatives counterparty. The unique trembling-hand perfect equilibrium in our basic alternating-offers contingent contract bargaining game specifies socially efficient actions by the three firms. Efficiency arises through the ability of the debtor to internalize the costs and benefits of its two counterparties through the pricing of contracts with each of them. For example, if a particular change in contract terms would have large benefits for the swap counterparty, and could be accommodated at a small total cost to the debtor and creditor firms, then this change in contract terms will be chosen in the course of pairwise contingent contract bargaining, given that the debtor firm can extract a compensating payment from the swap counterparty that provides a sufficient incentive to the debtor and creditor.

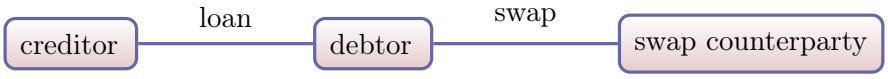

Figure 1 – An illustrative three-firm financial contracting network.

An illustrative issue of concern in this setting is whether the equilibrium contract terms would efficiently assign recovery priority to the creditor and the swap counterparty in the event of the debtor's insolvency. Assignment of priority in either direction is currently permitted under U.S. bankruptcy law through exemptions for qualified financial contracts such as swaps, repos, clearing agreements, and securities lending contracts. This exemption allows for enforceable *ipso facto* clauses that assign the right to terminate a contract in the event of insolvency and to liquidate collateral. This sort of clause is standard in current swap and repo contracts.

There has been a debate over allowing qualified financial contracts to include ipso facto clauses, unrestricted by bankruptcy law. Roe (2013) suggests that contractual assignment to swaps of the right to terminate and keep collateral in the event of the debtor's insolvency should be unenforceable or subject to rejection under an automatic stay. This issue is modeled in a different setting by Bolton and Oehmke (2015), who instead assume price-taking competitive markets and rule out the negotiation of covenants regarding the assignment of priority. Like Bolton and Oehmke (2015), our model does not cover externalities such as firesales, a tradeoff discussed by Duffie and Skeel (2012).

In our model, each pair of directly connected firms bargains over contractual terms. In the example setting illustrated in Figure 1, the equilibrium contract prices reflect the relative distress costs of the two counterparties, allowing the debtor to efficiently internalize its counterparties' distress costs and assign contractual priority efficiently. For example, if the creditor suffers greater distress from loss of default priority than does the swap counterparty, then in the naturally selected equilibrium, the pricing of the swap contract will include a price concession that is sufficient to "convince" the swap counterparty to give up priority. The creditor would in this case be willing to accept a lower interest rate in order to receive effective seniority. Conversely, if the creditor is better equipped to suffer losses at the debtor's default than the swap counterparty, then in equilibrium the debtor will offer a high enough interest rate to the creditor to encourage the creditor to agree to loss of priority, and the debtor will receive a correspondingly higher upfront payment from the swap counterparty. The debtor's shareholders have no direct concern with seniority at the debtor's own default,

and are therefore able to act as a conduit by which the creditor and the swap counterparty can indirectly compensate each other for priority assignment.

Our results are based on an extension to network settings of the alternating-offers bargaining game of Rubinstein (1982). Bargaining is conducted by each pair of directly connected nodes, which we call "firms." Our model allows for incomplete information. While a given pair of firms is bargaining, they are unaware of the bargaining offers and responses being made elsewhere in the network. In order to isolate natural equilibria, we therefore extend the notion of trembling-hand perfect equilibrium of Selten (1975) to this network bargaining setting. The trembling-hand perfect equilibrium choices are socially efficient by virtue of the assumed ability to sign contracts whose terms are contingent on the terms of other contracts. For instance, in the setting of our illustrative example, the creditor and the debtor can choose to make the terms of their loan agreement contingent on the terms of the swap contract chosen by the debtor at its swap counterparty.

The efficiency of the trembling-hand perfect equilibrium contract terms does not depend on some aspects of the bargaining protocol, such as which pair of counterparties writes contingent contracts and which pair of counterparties uses only simple contracts. In practice, covenants in a given contract normally restrict the terms of future contracts, but our setting is static.

We show that equilibrium contract prices converge, as exogenous breakdown probabilities go to zero, to those associated with the unique axiomatic solution that satisfies, on top of some obvious bargaining axioms, two newly proposed axioms, "multilateral stability" and "bilateral optimality." In particular, the non-cooperatively bargained prices do not depend on which pair of firms writes contingent contracts. Our axioms apply to more general types of networks, under conditions that rule out general cross-network externalities.

An extensive literature on network bargaining games includes some prior work that, like ours, focuses on non-cooperative bilateral bargaining. Until now this literature has studied settings in which there are two key impediments to socially efficient outcomes: *(i)* general cross-network externalities, and *(ii)* coordination failures that arise from a restriction to contracts that are not contingent on other bilateral contracts. We assume an absence

of general network externalities and we allow contracts to have unlimited cross-contract contingencies, such as covenants. These "ideal" conditions are not to be expected in practical settings. We nevertheless believe that it is valuable to characterize a theoretical benchmark network market setting in which bilateral contracting is socially efficient, subject to the restrictions imposed by the feasible sets of contractible actions. Even in our "ideal" setting, our analysis suggests that apparently reasonable changes to our proposed bargaining protocol can lead to additional equilibria that are not efficient.

Several papers provide non-cooperative bilateral bargaining foundations for the Myerson-Shapley outcomes and values, as defined by Myerson (1977a). In the first of these papers, Stole and Zwiebel (1996) provide a non-cooperatove foundation for the Myerson-Shapley values as those arising in the unique sub-game perfect equilibrium of a network game in which a firm negotiates bilateral labor contracts with each of its potential workers. Their bargaining protocol, like ours, is based on the Rubinstein alternating offers game. In their case, however, breakdown in a given bilateral bargaining encounter results in a re-start of the bargaining of the firm with other workers, in which any previously "agreed" labor contract is discarded. In this sense, the labor contracts are non-binding. The work of Stole and Zwiebel (1996) is extended to more general settings by de Fontenay and Gans (2013). In a different setting, Navarro and Perea (2013) provide a bilateral bargaining foundation for Myseron values, with a sequential bilateral bargaining protocol in which pairs of linked nodes bargain over their share of the total surplus created by the connected component of the graph in which they participate.

2 A Simple Network with Three-Node in a Line

We begin with a simple three-firm network. Firms 1 and 2 bargain over the terms of one contract. Firms 2 and 3 bargain over the terms of another contract. The contracts specify the actions to be taken by each firm. Firm i takes an action in a given finite set S_i. We denote by $S = \prod_{i=1}^{3} S_i$ the set of all possible action vectors. For each action $s_2 \in S_2$ of the central firm, there is a limited subset $C_1(s_2) \subseteq S_1$ of feasible actions that can be undertaken

by firm 1, and a limited set $C_3(s_2) \subseteq S_3$ of feasible actions for firm 3. That is, C_i is a correspondence on S_2 into the non-empty subsets of S_i. These correspondences describe compatibility conditions on actions taken by bilaterally contracting counterparties. We write $(s_1, s_2) \in C_{12}$ and $(s_2, s_3) \in C_{23}$ to indicate that a pair of actions is compatible. For a given action vector s, s is said to be *feasible* if both pairwise compatibility conditions are satisfied. We denote by C the set of feasible action vectors.

In applications, the actions of two linked firms could represent the terms of their bilateral contract, other than the initial compensating payment. These terms could include, for example, maturity date, stipulated future actions or payments (contingent perhaps on future states of the world), seniority assignment, and so on. Each pair of contracting firms is also able to exchange an initial monetary payment. Firm 2 pays firms 1 and 3 the amounts y_1 and y_3 respectively. These amounts are any real numbers, positive or negative. Later we will specify intervals by which the payments are bounded. Equivalently, firm i pays firm 2 the amount $-y_i$. In summary, firms choose actions $s \in S = S_1 \times S_2 \times S_3$ and compensation amounts $y \in \mathbb{R}^2$, with respective quasi-linear utilities

$$\begin{aligned} u_1(y, s) &= f_1(s_1, s_2) + y_1 \\ u_2(y, s) &= f_2(s_1, s_2, s_3) - y_1 - y_3 \\ u_3(y, s) &= f_3(s_2, s_3) + y_3, \end{aligned} \tag{1}$$

for some $f_1 : S_1 \times S_2 \to \mathbb{R}$, $f_2 : S \to \mathbb{R}$, and $f_3 : S_2 \times S_3 \to \mathbb{R}$, as illustrated in Figure 2. It is important for our efficiency results that a firm's utility depends only on its direct compensation, and on the actions of itself and its direct bilateral counterparty. Nevertheless, the compatibility condition $s_3 \in C_3(s_2)$ and the dependence of $f_3(s_2, s_3)$ on s_2 imply that firms 1 and 2 contracting over (s_1, s_2) has an influence over firm 3 through the choice of s_2. A symmetric situation applies to contracting between firms 2 and 3.

In the event of a failure to reach contractual agreement, there are some pre-arranged "outside options," which can be viewed as the "status quo." We let $(s^0, y^0) \in S \times \mathbb{R}^2$ be

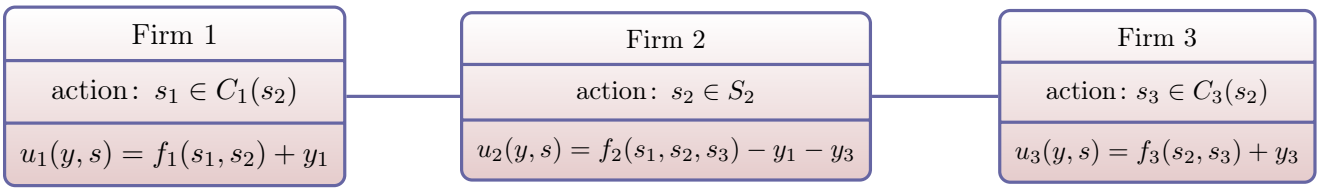

Figure 2 – Actions and utilities in the three-firm financial contracting network.

the status-quo actions and payments. Without loss of generality via a normalization, we let $y^0 = (0,0)$, $f_1\left(s_1^0, s_2^0\right) = 0$, $f_3\left(s_2^0, s_3^0\right) = 0$, and $f_2\left(s_1^0, s_2^0, s_3^0\right) = 0$.

If the bargaining between Nodes 1 and 2 break down, an event that can arise in the extensive-form bargaining game to be defined, then Node 2 has a limited set of actions that can be taken with Node 3, and likewise with a breakdown between Nodes 2 and 3. Specifically, in the event of a bargaining breakdown between Nodes 1 and 2, the action of Node 1 is its status quo action s_1^0, whereas the pair (s_2, s_3) of actions of Nodes 2 and 3 must be chosen from

$$\mathcal{S}_{2,3}^B = \left\{(s_2, s_3) : s_2 \in S_2^1, s_3 \in C_3(s_2)\right\},$$

where S_2^1 is a given non-empty subset of S_2 with the property that any action in S_2^1 is compatible with s_1^0. That is, $s_1^0 \in C_1(s_2)$ for any $s_2 \in S_2^1$. Likewise, in the event of a breakdown in the bargaining between Nodes 2 and 3, the action of Node 3 is its status quo action s_3^0, whereas the actions of Nodes 1 and 2 must be chosen from

$$\mathcal{S}_{1,2}^B = \left\{(s_1, s_2) : s_2 \in S_2^3, s_1 \in C_1(s_2)\right\},$$

where S_2^3 is a given non-empty subset of S_2 with the property that any action in S_2^3 are compatible with s_3^0. By assumption, $\mathcal{S}_{2,3}^B$ and $\mathcal{S}_{1,2}^B$ are not empty.

We assume for simplicity that each stated argmax is a singleton (that, is the associated maximization problem has a unique solution). This is generically true in the space of

utilities.[1] The socially optimal result is

$$(s_1^{**}, s_2^{**}, s_3^{**}) = \underset{s \in C}{\text{argmax}} \quad U(s).$$

where $U : S \to \mathbb{R}$ be the "social welfare function" defined by

$$U(s) = f_1(s_1, s_2) + f_2(s_1, s_2, s_3) + f_3(s_2, s_3).$$

We also adopt the following two standing assumptions regarding outside option values. First, for $i = 1$ and $i = 3$, we have

$$f_i\left(s_i^0, s_2^i\right) \quad = \quad f_i\left(s_i^0, s_2^0\right), \qquad s_2^i \in S_2^i,$$

This condition means that that if there is a breakdown between Node 2 and another Node i, then the utility of Node i does not vary across the restricted choices remaining to Node 2. Basically, if they don't sign a contract, Node 2 cannot help or hurt Node i, no matter what Node 2 chooses or whether Node 2 also takes its outside option.

The second assumption is

$$\max_{s_1 \in C_1(s_2)} U(s_1, s_2, s_3) \geq U\left(s_1^0, s_2^{1B}, s_3^B\right), \qquad (s_2, s_3) \in C_{23}, \tag{2}$$

where

$$\left(s_2^{1B}, s_3^B\right) = \underset{\left(s_2^{1b}, s_3^b\right) \in \mathcal{S}_{2,3}^B}{\text{argmax}} \; f_2\left(s_1^0, s_2^{1b}, s_3^b\right) + f_3\left(s_2^{1b}, s_3^b\right).$$

This condition means that Node 1 always has a feasible choice that is a strict welfare improvement over any breakdown option.

Our main result is a protocol for contingent pairwise bargaining under which the unique extensive form trembling hand perfect equilibrium, for any sufficiently small trembles, achieves the socially optimal actions s^{**}. We also provide alternative plausible bargaining approaches

[1] That is, fixing all other primitives of the model, we can view the vectors of utilities of the firms (or of a subset of firms) defined by the utilities f_1, f_2, and f_3 as elements of a Euclidean space. A condition is said to hold "generically" in a Euclidean space if it holds except for a closed subset of zero Lebesgue measure.

that do not lead to this efficient result.

3 Simple Illustrative Example

For an extremely simple illustrative example, we may imagine a situation in which firm 2 is negotiating credit agreements with firms 1 and 3. The creditor firms 1 and 3 each begin with 1 unit of cash at time zero. Firm 2 initially has $c < 1$ in cash, and has the opportunity to undertake a project that requires 2 units of cash. At time 1, the project will pay some amount $A > 2$ with success probability p, and otherwise pays B, where $1 < B < 2$. For some negotiated note discount y_i, creditor firm i will provide firm 2 with $1 - y_i$ in cash at time zero in return for a note promising to pay 1 at time 1. Without loss of generality for this example, we can take $S_1 = S_3 = \{0\}$ and $C_1(s_2) = C_3(s_2) = \{0\}$. Firm 2 chooses from $S_2 = \{0, 1\} \times \{0, 1\} \times \{1, 3\}$, each element of which specifies, respectively, whether credit is taken from 1, whether credit is taken from 3, and whether firm 1 or firm 3 receives seniority. If there is no agreement, all firms consume their initial cash. If firm 2 can negotiate funding from each of the creditors, then at time zero it will invest 2 in the risky project and consume all of its excess cash, which is $c - y_1 - y_3$. At time 1, firm 2 will consume $A - 2$ if the risky project is successful and nothing otherwise. If the project is funded, then at time 0 firms 1 and 3 will consume y_1 and y_3 respectively. At time 1, these creditors will each consume 1 if the project is successful. Otherwise, the senior creditor will consume 1 and the junior creditor will consume $B - 1$.

Firm i has utility $c_0 + \gamma_i E(c_1)$ for consumption c_0 in period 0 and c_1 in period 1. We suppose that $\gamma_2 < \gamma_1 < \gamma_3$. The status-quo (breakdown) actions is taken to be $s_2^0 = (0, 0, 3)$. (Which of the creditors is senior in the event of no funding is irrelevant, and is taken to be firm 3 without loss of generality.) The model's primitive set of parameters is thus $(A, B, p, c, \gamma_1, \gamma_2, \gamma_3, s_2^0)$.

We suppose that the project is worth funding, in terms of total utility, no matter which creditor is senior. The unique efficient outcome, subject to the limited available forms of credit agreements, is therefore to fund the project and for firm 3 to be the senior creditor.

That is, $s_2^{**} = (1, 1, 3)$. We assume that $y_1 + y_3 < c$ for any discounts y_1 and y_3 that are individually rational for firms 2. This is the condition $c > \gamma_2 p(A-2)$. With this, c is irrelevant and can be ignored when calculating an equilibrium. We also assume that the set of discounts (y_1, y_3) that are individually rational for all firms is not empty, even if the "wrong" creditor, firm 1, is senior. This is the condition $\gamma_2 p(A-2) > 2 - \gamma_1 - \gamma_3(p + (1-p)(B-1))$.

After normalizing by subtracting the initial cash utility of 1, firms 1 and 3 receive nothing in the event that the project is not funded fully, and otherwise receive utility $y_i + f_i(1, 1, j)$ where

$$\begin{aligned} f_i(1, 1, i) &= \gamma_i \\ f_i(1, 1, i') &= \gamma_i(p + (1-p)(B-1)), \end{aligned} \tag{3}$$

depending on whether firm $i' \neq i$ obtains seniority. (We have suppressed from the notation the dependence of f_i on s_1 and s_3, since this is trivial.)

After normalizing by subtracting the initial cash utility of c, firm 2 receives nothing in the event that the project is not funded fully, and otherwise receives utility

$$-y_1 - y_3 + f_2(1, 1, j) = -y_1 - y_3 + \gamma_2 p(A-2),$$

regardless of which firm j obtains seniority.

4 The Axiomatic Solution

Appendix C provides foundations for an axiomatic solution of network bilateral bargaining problems. As we will show, the axiomatic solution coincides with the proposed equilibrium for the associated non-cooperative extensive-form bargaining game.

In addition to *(i)* axioms for two-node networks that support the bargaining solution of Nash (1950), our axioms are *(ii)* multilateral stability, *(iii)* independence of irrelevant actions, and *(iv)* bilateral optimality. Under these axioms, we show that there is a uniquely defined solution, which we call the *axiomatic solution*, and moreover the axiomatic solution outcome is socially efficient. We briefly motivate the axiomatic solution here, and provide details in Appendix C. A key objective of the paper is to show conditions under

which the unique extensive-form trembling-hand-perfect equilibrium for our non-cooperative extensive-form alternating-offers bargaining game associated with contingent network contracting reaches the same payments and the same (efficient) actions as those uniquely specified by the axiomatic solution. This can be viewed as an extension to network games of the non-cooperative-game foundation established by Binmore, Rubinstein, and Wolinsky (1986) for Nash bargaining.

Formally, our axioms concern the properties of a "solution" $F : \Sigma \to \Omega$, a function that maps the space Σ of network bilateral bargaining problems to the space Ω of associated actions and payments. These spaces Σ and Ω are formally defined in Appendix C.

Our first axiom is that when applied to a network bilateral bargaining problem σ whose graph merely consists of two connected nodes, the solution $F(\sigma)$ is effectively the Nash Bargaining Solution, specifying actions $s^* = (s_1^*, s_2^*)$ and a payment y_{12}^* that solve

$$\begin{aligned} &\max_{s \in C} \{f_1(s) + f_2(s)\} \\ y_{1,2}^* &= \frac{1}{2}\left([f_1(s^*) - \underline{u}_1] - [f_2(s^*) - \underline{u}_2]\right), \end{aligned}$$

where $\underline{u}_1$ and $\underline{u}_2$ are the respective outside option values of nodes 1 and 2. In Appendix C, we discuss underlying axioms for two-player games that support this Nash Bargaining Solution.

Roughly speaking, a solution F satisfies our second axiom, *multilateral stability*, if, for any given network bilateral bargaining problem σ, whenever one "freezes" the actions and payments among a subset of pairs of directly connected nodes, and then applies F to the bargaining problem σ_{sub} induced for the remaining sub-network G_{sub}, the solution $F(\sigma_{\text{sub}})$ of the sub-network bilateral bargaining problem σ_{sub} coincides on the sub-network with that prescribed by the solution $F(\sigma)$ to the entire network problem σ. Multilateral stability is illustrated in Figure 3.

Our third axiom, independence of irrelevant actions, states roughly that if the solution specifies some outcome (s^*, y^*) for a network bilateral bargaining problem, and if we alter this network bilateral bargaining problem merely by reducing the set of feasible actions while still admitting s^* as feasible, then (s^*, y^*) remains the solution.

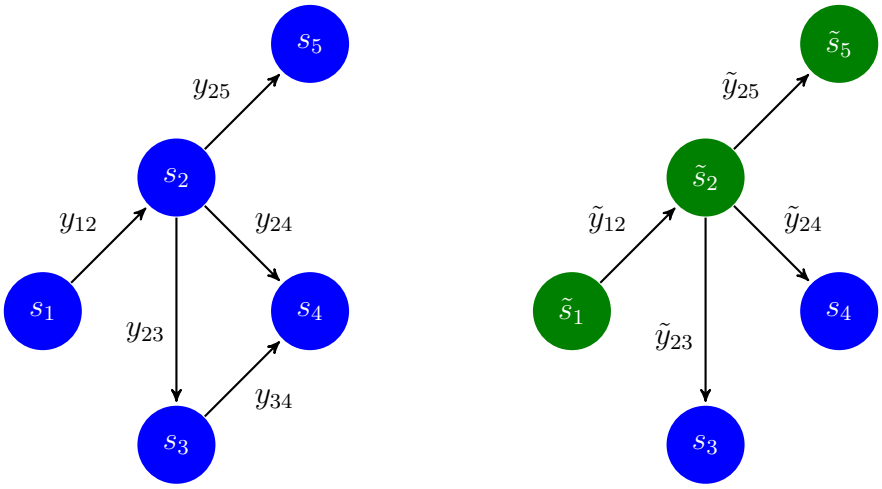

Figure 3 – *Multilateral stability.* By freezing the bargaining outcome of nodes 3 and 4 (blue) off the sub-network formed by edges other than $(3, 4)$, the solution on the induced sub-network game coincides with the blue solution. That is, $\tilde{s}_i = s_i$ and $\tilde{y}_{ij} = y_{ij}$.

A solution F satisfies the final axiom, bilateral optimality, if for any given network bilateral bargaining problem σ, when any two directly connected nodes maximize the sum of their total payoff under the assumption that the remaining nodes will react according the solution F applied to their sub-network, then the maximizing actions they would choose are consistent with the solution $F(\sigma)$ of the entire network bilateral bargaining problem. Bilateral optimality is illustrated in Figure 4.

A result stated in Appendix C implies that the axiomatic solution (y^a, s^a) associated with our our present three-firm network setting is given by the efficient actions $s^a = s^{**}$ and the upfront payments $y^a = (y^a_1, y^a_3)$ that uniquely solve the equations

$$\begin{aligned} u_1(y^a, s^{**}) - \underline{u}_{12} &= u_2(y^a, s^{**}) - \underline{u}_{21} \\ u_3(y^a, s^{**}) - \underline{u}_{32} &= u_2(y^a, s^{**}) - \underline{u}_{23}, \end{aligned} \tag{4}$$

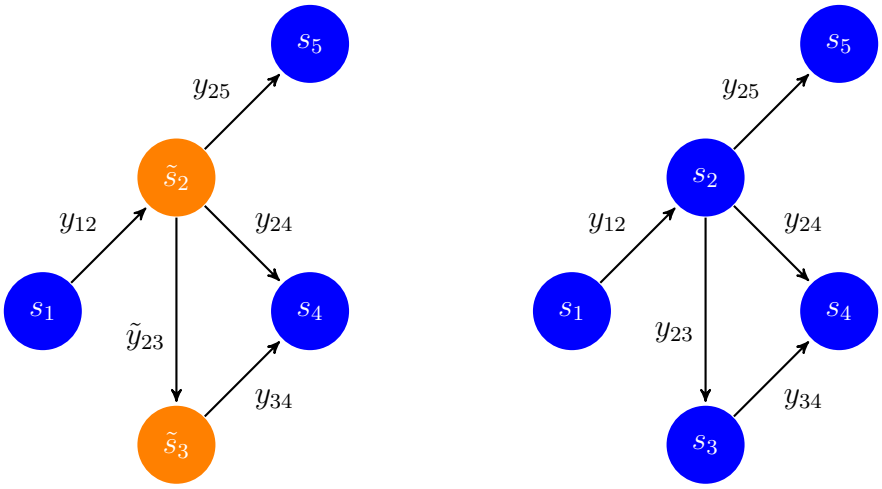

Figure 4 – Suppose the orange nodes, 2 and 3, maximize their total payoff assuming that the remaining network will react according to the (blue) solution. *Bilateral optimality* implies that the actions $(\tilde{s}_2, \tilde{s}_3)$ coincides with applying the blue solution to the whole network. That is, $(\tilde{s}_2, \tilde{s}_3) = (s_2, s_3)$.

where $\underline{u}_{ij}$ is the outside-option value of node i in its bargaining with node j. Here

$$\begin{aligned} \underline{u}_{12} &= \underline{u}_{32} = 0 \\ \underline{u}_{21} &= \frac{1}{2} \max_{(s_2,s_3)\in\mathcal{S}^B_{2,3}} U\left(s_1^0, s_2, s_3\right) \\ \underline{u}_{23} &= \frac{1}{2} \max_{(s_1,s_2)\in\mathcal{S}^B_{1,2}} U\left(s_1, s_2, s_3^0\right). \end{aligned} \tag{5}$$

The main result of the paper, stated in Section 8, provides simple conditions under which this axiomatic solution is also the unique extensive-form trembling-hand equilibrium of the non-cooperative contingent-contract bilateral bargaining game to be described. In this sense, bilateral bargaining over complete contingent contracts is socially efficient in our network market setting. It also follows that our non-cooperative equilibrium solution concept for extensive-form bilateral bargaining over contingent contracts satisfies multilateral stability, irrelevance of independent actions, and bilateral optimality.

5 Counterexamples

This section explores variants of the model definition or solution concept that do not lead to efficient equilibrium outcomes. Our objective here is to promote an understanding of the dependence of our main efficiency results on our assumptions. For this purpose, we will restrict attention to the simple special case in which $f_1 = 0$, $f_3 = 0$, and $S_2 = \{s_2^0\}$. That is, utility is obtained only by the central node, and the treatment of the central node is fixed. We will simply write f for f_2.

5.1 Bargaining Without Communication

We will first consider bilateral bargaining without the ability of either pair of connected nodes to contract on the bargaining outcome of the other pair. That is to say, when Nodes 1 and 2 bargain, Nodes 1 and 2 are unable write a contract that depends on the action s_3 of Node 3 or the payment y_3. Likewise, when Nodes 2 and 3 bargain, Node 3 cannot contract on the action s_1 of Node 1 or the payment y_1. The behavior of Nodes 1 and 3 in such a game depends on their beliefs about the contracted action and payment of other nodes.

In this setting, it is common to restrict attention to passive beliefs in which, after observing a deviation, each node continues to believe that other nodes receive their equilibrium offers. This is typical in Hart, Tirole, Carlton, and Williamson (1990) and Segal (1999). Let $(s_1^*, s_3^*; y_1^*, y_3^*)$ denote the equilibrium outcome. With passive beliefs, if Node i is offered $(s_i, y_i) \neq (s_i^*, y_i^*)$, he still believes that other nodes make their equilibrium choices of treatments and prices. We suppose that the outcomes of bilateral negotiations are given by the Nash Bargaining Solution (NBS). The passive-beliefs equilibrium must therefore be a pairwise stable Nash bargaining solution. That is, (s_i^*, y_i^*) is the Nash solution to the bargaining problem between Node i and Node 2, under the belief that $\left(s_j^*, y_j^*\right)$ is the agreed choice by

Node j and Node 2. Hence (s^*, y^*) solves[2]

$$\begin{cases} \max\limits_{s_1 \in S_1, y_1 \in \mathbb{R}} \; y_1 \cdot [f(s_1, s_3^*) - y_3^* - y_1] \\ \max\limits_{s_3 \in S_3, y_3 \in \mathbb{R}} \; y_3 \cdot [f(s_1^*, s_3) - y_1^* - y_3] \,. \end{cases} \tag{6}$$

From (6), it is straightforward to characterize pairwise-stable NBS by

$$\begin{aligned} s_1^* &= \underset{s_1 \in S_1}{\operatorname{argmax}} \; f(s_1, s_3^*) \\ s_3^* &= \underset{s_3 \in S_3}{\operatorname{argmax}} \; f(s_1^*, s_3) \\ y_1^* = y_3^* &= \frac{1}{3} \, f(s_1^*, s_3^*). \end{aligned}$$

In a pairwise stable NBS, each node has the same utility

$$u_1 = u_2 = u_3 = \frac{1}{3} \, f(s_1^*, s_3^*).$$

We can see that the efficient vector of treatments s^{**} is indeed consistent with a pairwise stable NBS. Not every pairwise stable NBS, however, is necessarily efficient. This is so because (s_1^*, s_3^*) merely solves

$$\begin{aligned} s_1^* &\in \underset{s_1 \in S_1}{\operatorname{argmax}} \; f(s_1, s_3^*) \\ s_3^* &\in \underset{s_3 \in S_3}{\operatorname{argmax}} \; f(s_1^*, s_3), \end{aligned} \tag{7}$$

whereas (s_1^{**}, s_3^{**}) jointly maximizes $f(s_1, s_3)$. Depending on the utility function f, there may be other pairwise stable NBS, which are Pareto ranked. In a pairwise stable NBS, Nodes 1 and 3 cannot be sure of making efficient choices because of the inability to contract based on communication between the two spoke-end nodes in the solution concept. Node i cannot be certain that node j will choose the efficient treatment s_j^{**}. Suppose a pair of treatments (s_1^*, s_3^*) satisfies (7). If Node i believes that Node j chooses s_j^*, then the outcome of her

[2]One may argue that $B_i = \{(f(s_i, s_j^*) - y_j^* - y_i, y_i) : s_i \in S_i, y_i \in \mathbb{R}\}$ is not convex, whereas the Nash solution requires convexity. Indeed, the payoff pairs form a finite number of parallel lines in the Euclidean plane. One can convexify this set by filling in the gaps between the lines. Then the axiom of Independence of Irrelevant Alternatives implies that the unique solution is given by maximizing the utility products in equation (6).

bargaining with the central node is s_i^*.

One can construct a non-cooperative alternating-offers game whose Perfect Bayesian Nash equilibria with passive beliefs coincide with the pairwise stable NBS, in the limit as players become infinitely patient. This is shown by de Fontenay and Gans (2013). In such a game, the central node is assumed to be able to bargain only over prices, but not the actions to be chosen by the spoke-end nodes. Thus the central node cannot efficiently coordinate the actions of the spoke ends. In practice, a borrowing firm is typically able to credibly assign higher seniority to one lender over another, in return for a low interest rate. A junior lender, through accounting disclosure or covenants, is typically able to receive information on its relative and absolute loss of priority and demand a correspondingly high interest rate. In other cases, however, firms in a network may fail to coordinate their contracts and use price negotiations to promote efficient outcomes due to a lack verifiable and contractible information.

5.2 Distortion of Outside Option Values

We now allow Nodes 1 and 2 to sign enforceable contingent contracts. For each action s_3 chosen by Node 3, Nodes 1 and 2 can choose a different action-payment pair (s_1, y_1). Nodes 2 and 3 sign a simple, non-contingent binding contract (s_3, y_3), to which both nodes commit. This setting is equivalent to the following 2-stage game: Nodes 2 and 3 first bargain over (s_3, y_3) in Stage 1. Then, in Stage 2, Node 2 bargains with Node 1 over $(s_1(s_3), y_1(s_3))$, with common knowledge of the action s_3 chosen by Node 3.

We show that the socially efficient outcome (s_1^{**}, s_3^{**}) may not be an equilibrium.

Suppose in Stage 1, Nodes 2 and 3 choose (s_3, y_3). Then, in Stage 2, a breakdown between Nodes 1 and 2 leads to payoffs of $u_{12}^0(s_3) = 0$ for Node 1 and $u_{21}^0(s_3) = f(s^0, s_3) - y_3$ for Node 2. Therefore $u_{12}^0(s_3)$ and $u_{21}^0(s_3)$ are the respective outside option values for Nodes 1 and 2 in their bilateral bargaining in Stage 2. Likewise, the outside option values for Nodes 2 and 3 in Stage 1 are $u_{23}^0 = 0.5f(s_1^*(s_3^0))$ and $u_{32}^0 = 0$, where $s_1^*(s_3) = \text{argmax}_{s_1 \in S_1} f(s_1, s_3)$.

By the same argument used to determine the pairwise stable NBS, the Nash bargaining

outcome $(s_1(s_3), y_1(s_3))$ between Nodes 1 and 2 in Stage 2 is

$$s_1^*(s_3) = \underset{s_1 \in S_1}{\operatorname{argmax}} f(s_1, s_3),$$

whereas the payment $y_1(s_3)$ is determined by

$$\begin{cases} y_1(s_3) = f(s_1^*(s_3), s_3) - y_1(s_3) - y_3 - u_{21}^0(s_3) \\ y_3 = f(s_1^*(s_3), s_3) - y_1(s_3) - y_3 - u_{23}^0. \end{cases}$$

Hence

$$y_1(s_3) = \frac{1}{2}\left[f(s_1^*(s_3), s_3) - f\left(s_1^0, s_3\right)\right],$$

$$2y_3 = \frac{1}{2}\left[f(s_1^*(s_3), s_3) + f\left(s_1^0, s_3\right)\right] - \frac{1}{2}f\left(s_1\left(s_3^0\right), s_3^0\right).$$

Therefore, in Stage 2, when choosing s_3, Nodes 2 and 3 s_3 both receive the payment

$$g(s_3) = \frac{1}{2}\left(\frac{1}{2}\left(f(s_1^*(s_3), s_3) + f\left(s_1^0, s_3\right)\right] - \frac{1}{2}f\left(s_1\left(s_3^0\right), s_3^0\right)\right).$$

One can easily choose f so that the maximum of g is not attained at $s_3 = s_3^{**}$. In this case, the socially efficient outcome (s_1^{**}, s_3^{**}) cannot be an equilibrium.

Indeed, by committing to s_3, when bargaining with Node 1, Node 2 has an outside option value $u_{21}^0(s_3) = f(s^0, s_3) - y_3$ that depends on s_3. In this sense, the choice of s_3 in Stage 1 may distort the outside option value of Node 2 in Stage 2. A low value of u_{21}^0 forces Node 2 to make a high payment to Node 1, which is detrimental to both Nodes 2 and 3. Thus the distortion caused by this outside option value u_{21}^0 can create an incentive for inefficient equilibrium outcomes.

5.3 Incentives to Lie Distort the Distribution of Surplus

We again allow nodes to sign contingent contracts. For each action s_3 chosen by Node 3, Nodes 1 and 2 choose some action-payment pair (s_1, y_1), and vice versa. We now assume, however, that the contingent contracts are not enforceable. We will see, not surprisingly, that the central node may have an incentive to misreport to one end node the outcome of

its bargain with the other end node.

We assume that the outside option values for both bilateral bargaining problems (that between Nodes 1 and 2, and that between Nodes 2 and 3) are all 0. This assumption allows us to isolate the effect of dishonesty by the central node. As this rules out the distortion of outside option values through commitment to a contract, we consider only binding contracts.

Consider the following 2-stage game. In Stage 1, Nodes 2 and 3 first bargain over (s_3, y_3). In Stage 2, Node 2 communicates to Node 1, not necessarily truthfully, that the action chosen by Node 3 is $\tilde{s}_3$. Then Nodes 1 and 2 bargain over $(s_1(\tilde{s}_3), y_1(\tilde{s}_3))$. Truthful communication by Node 2 and the socially efficient outcome (s_1^{**}, s_3^{**}) need not be an equilibrium of the game. In order to see this, suppose this is in fact an equilibrium. In Stage 1, Nodes 2 and 3 choose the treatment s_3^{**} along with some payment $y_3 \in \mathbb{R}$. If Node 1 believes the report from Node 2 that Node 3 agreed to take the action $\tilde{s}_3$, then following the earlier determination of pairwise stable NBS, the Nash bargaining outcome $(s_1(\tilde{s}_3), y_1(\tilde{s}_3))$ between Nodes 1 and 2 in Stage 2 would be:

$$s_1^*(\tilde{s}_3) = \underset{s_1 \in S_1}{\operatorname{argmax}} f(s_1, \tilde{s}_3).$$

The associated payment $y_1(\tilde{s}_3)$ is determined by

$$\begin{cases} y_1(\tilde{s}_3) = f(s_1^*(\tilde{s}_3), \tilde{s}_3) - y_1(\tilde{s}_3) - y_3 \\ \quad y_3 = f(s_1^*(\tilde{s}_3), \tilde{s}_3) - y_1(\tilde{s}_3) - y_3. \end{cases}$$

Hence $y_1(\tilde{s}_3) = \frac{1}{3} f(s_1^*(\tilde{s}_3), \tilde{s}_3)$. The value to Node 2 associated with reporting $\tilde{s}_3$ is

$$f\left(s_1^*(\tilde{s}_3), s_3^{**}\right) - y_1(\tilde{s}_3) - y_3 = f\left(s_1^*(\tilde{s}_3), s_3^{**}\right) - \frac{1}{3} f(s_1^*(\tilde{s}_3), \tilde{s}_3) - y_3.$$

We define $g : S_3 \to \mathbb{R}$ by

$$g(\tilde{s}_3) = f\left(s_1^*(\tilde{s}_3), s_3^{**}\right) - \frac{1}{3} f(s_1^*(\tilde{s}_3), \tilde{s}_3).$$

One can choose f so that the maximum of g is not attained at $\tilde{s}_3 = s_3^{**}$. It is therefore not credible that Node 2 correctly reports. This could destroy the socially efficient outcome

(s_1^{**}, s_3^{**}) for being an equilibrium.

We have shown that if Node 1 believes the report of Node 2, then Node 2 will not report truthfully, so there is no truth-telling equilibrium. There are other possibilities, based on a definition of equilibrium in which Node 1 does not necessarily believe Node 2, but rather makes an inference about s_3 based on the report $\tilde{s}_3$. In this setting, it is conceivable that an equilibrium may not exist, or that there may be either efficient or inefficient equilibria. We intend to go further into this in the next version of the paper.

6 Contingent-Contract Network Bargaining

Our main objective now is to extend and apply the bargaining protocol of Rubinstein (1982) and Binmore, Rubinstein, and Wolinsky (1986). The associated unique extensive-form trembling-hand perfect equilibrium outcome of the corresponding negotiation game converges, as the breakdown probability goes to zero, to the axiomatic solution (s^a, y^a), and in particular achieves the socially efficient choice $s^a = s^{**}$. As explained in the previous section, not all plausible extensions of the Rubinstein model to our network setting have this efficiency property.

We allow the actions negotiated by Nodes 1 and 2 to be contractually contingent on the actions chosen by Nodes 2 and 3 from $\mathcal{S}_{2,3} = \{(s_2, s_3) : s_2 \in S_2, s_3 \in C_3(s_2)\}$. As we shall see, Nodes 2 and 3 may experience a breakdown in their negotiation, a contingency that we label $B_{2,3}$. Thus, the set of conceivable contingencies is $\mathcal{S}_{2,3} \cup \{B_{2,3}\}$. The contingent action to be negotiated between Nodes 1 and 2 is $\left(s_1(\cdot), s_1^b, s_2^{1b}\right)$, where s_1 is chosen from

$$\mathcal{C} = \{s_1 : \mathcal{S}_{2,3} \to S_1 : \ s_1(s_2, s_3) \in C_1(s_2)\},$$

and $\left(s_1^b, s_2^{1b}\right)$ is chosen from $\mathcal{S}_{1,2}^B$. That is, $s_1(\cdot)$ is a "menu" of actions such that, for each pair of conceivable actions $(s_2, s_3) \in \mathcal{S}_{2,3}$ by Nodes 2 and 3, $s_1(s_2, s_3)$ specifies an action, compatible with s_2, to be taken by Node 1. In the event $B_{2,3}$ that bargaining between Nodes 2 and 3 breaks down, the pair of actions to be taken by Nodes 1 and 2 is $\left(s_1^b, s_2^{1b}\right)$. Nodes 1 and 2 also bargain, separately, over a contingent payment $\left(y_1, y_1^b\right)$, where y_1 is a mapping

from $\mathcal{S}_{2,3}$ to $\mathbb{R}$, and y_1^b is a real number that specifies the amount of payment from Node 2 to 1 in the breakdown event $B_{2,3}$. Thus, the contract between Nodes 1 and 2 takes the form of some contingent action $\left(s_1, s_1^b, s_2^{1b}\right)$ and some contingent payment $\left(y_1, y_1^b\right)$.

On the other hand, the contract to be chosen by Nodes 2 and 3 specifies $(s_2, s_3) \in \mathcal{S}_{2,3}$ and an associated payment $y_3 \in \mathbb{R}$, as well as also some $\left(s_2^{3b}, s_3^b\right) \in \mathcal{S}_{2,3}^B$ and associated payment $y_3^b \in \mathbb{R}$ to be taken in the event of a breakdown between Nodes 1 and 2.

The proposed four-stage extensive-form network bargaining game is defined as follows.

Stage a: In Stage a, Nodes 1 and 2 bargain over their contingent action. Node 1 is the first proposer and offers a contingent action s_1 in $\mathcal{C}$ and some $\left(s_1^b, s_2^{1b}\right) \in \mathcal{S}_{1,2}^B$ in period 0. (The identity of the first proposer is irrelevant for the ultimate solution concept.) For each contingency $(s_2, s_3) \in \mathcal{S}_{2,3}$, Node 2 either accepts or rejects the offered action $s_1(s_2, s_3)$. Likewise, for the contingency $B_{2,3}$, Node 2 either accepts or rejects $\left(s_1^b, s_2^{1b}\right)$. Acceptance closes the bargaining between Nodes 1 and 2 over the action s_1 (respectively, $\left(s_1^b, s_2^{1b}\right)$) contingent on (s_2, s_3) (respectively, on $B_{2,3}$). Agreement or rejection at one contingency does not bind behavior at any other contingency. Rejection at a particular contingency (including $B_{2,3}$) leads, with a given probability $\eta \in (0, 1)$, to a breakdown of the negotiation over that contingency. If Nodes 1 and 2 break down when they are bargaining over $\left(s_1^b, s_2^{1b}\right)$, the resulting actions for both Nodes 1 and 2 are the exogenous status-quo choices, $\left(s_1^b, s_2^{1b}\right) = (s_1^0, s_2^0)$. These breakdown events are independent across contingencies. The process continues to the next period, when Node 2 is the proposer and Node 1 responds, as illustrated in Figure 5. This alternating-offers procedure is iterated until agreement or breakdown. Let $\mathcal{S}_{2,3}^t$ be the set of contingencies (including $B_{2,3}$) that are still open for negotiation. That is, $\mathcal{S}_{2,3}^t$ is the set of contingencies for which Nodes 1 and 2 have reached neither agreement nor breakdown by the beginning of period t. In period t, Nodes 1 and 2 bargain over $s_1(s_2, s_3)$ for these remaining contingencies (s_2, s_3) in $\mathcal{S}_{2,3}^t$, and over $\left(s_1^b, s_2^{1b}\right)$ for the contingency $B_{2,3}$ if $B_{2,3} \in \mathcal{S}_{2,3}^t$. The bargaining between Nodes 1 and 2 concludes at the first time by which they have reached an agreement or have broken down for all of contingencies in $\mathcal{S}_{2,3} \cup \{B_{2,3}\}$. This is a finite time, almost surely. The breakdown probability η is an exogenous parameter of the model.

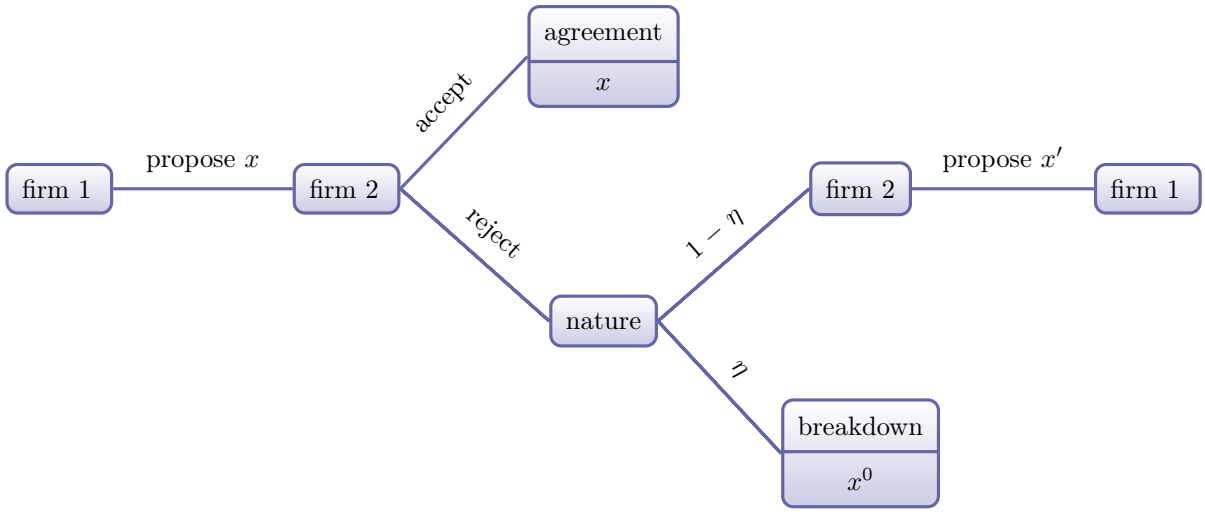

Figure 5 – The first four stages of a generic Rubinstein alternating-offers game.

We will later be interested in the limit behavior as $\eta \to 0$. The result of Stage a of the bargaining game is some random set $\Xi_{2,3} \subseteq \mathcal{S}_{2,3} \cup \{B_{2,3}\}$ on which Nodes 1 and 2 reach agreement, and for each contingency in $\Xi_{2,3}$, the agreed action $s_1(s_2, s_3)$ for Node 1, as well as the agreed pair of actions $\left(s_1^b, s_2^{1b}\right)$ in the contingency $B_{2,3}$, if it is in $\Xi_{2,3}$.

Stage b: In Stage b, Nodes 2 and 3 bargain without contingencies over some $\left(s_2, s_3, s_2^{3b}, s_3^b\right)$ such that $(s_2, s_3) \in \mathcal{S}_{2,3}$ and $\left(s_2^{3b}, s_3^b\right) \in \mathcal{S}_{2,3}^B$. One should understand $\left(s_2^{3b}, s_3^b\right)$ to be the actions of Nodes 2 and 3 that apply in the event that Nodes 1 and 2 break down at the contingency (s_2, s_3), actions that Nodes 2 and 3 would have liked to choose. Simultaneous with the bargaining in Stage a, Nodes 2 and 3 play a similar alternating-offers bargaining game. We suppose that Node 3 proposes first. (Again, the identity of the first proposer does not matter in the limit as $\eta \to 0$.) The two negotiations, between Nodes 1 and 2 in Stage a, and between Nodes 2 and 3 in Stage b, are not coordinated in any way. Specifically, Stage b strategies cannot depend on information from ongoing play or breakdowns in Stage a, and vice versa. Let $A_{2,3} \in \{Y, N\}$ be the binary variable indicating whether Node 2 and 3 reach an agreement (Y) or not (N) over (s_2, s_3). If Nodes 2 and 3 break down when bargaining over $\left(s_2^{3b}, s_3^b\right)$, the resulting actions for both Node 2 and 3 are their exogenous status-quo choices, $\left(s_2^{3b}, s_3^b\right) = (s_2^0, s_3^0)$.

Stage a': Once Nodes 1 and 2 have finished bargaining over their contingent actions in Stage a, they bargain in Stage a' over the corresponding payments $y_1(s_2, s_3)$ and y_1^b. Stage a precedes Stage a' so that the play of Stage a is available information to both players at the beginning of Stage a', information including the bargaining outcome of Stage a. If $B_{2,3} \notin \Xi_{2,3}$, that is, if Nodes 1 and 2 did not reach an agreement over $\left(s_1^b, s_2^{1b}\right)$ in Stage a, then y_1^b is not subject to negotiation and is fixed to be the null payment $y_1^b = 0$. Otherwise when bargaining over the payments $y_1(s_2, s_3)$ and y_1^b, Nodes 1 and 2 use the same form of alternating-offer game. For any contingency in $\Xi_{2,3}$, breakdown of the associated payment bargaining leads to the status-quo action s_1^0 (respectively, actions (s_1^0, s_2^0)) and the null payment, $y_1^0 = 0$. That is, unless they can agree on the payment $y_1(s_2, s_3)$ (respectively, y_1^b), the contingent action $s_1(s_2, s_3)$ (respectively, actions $\left(s_1^b, s_2^{1b}\right)$) agreed in Stage a is discarded. We let $\Xi_{2,3}^t$ be the set of contingencies that are still open for negotiation at the beginning of period t in Stage a'. The result of Stage a' is a random set $\mathcal{A}_{2,3} \subseteq \Xi_{2,3}$ at which there is ultimately agreement on both an action $s_1(s_2, s_3)$ and a payment $y_1(s_2, s_3)$, as well as actions $\left(s_1^b, s_2^{1b}\right)$ and a payment y_1^b for the contingency $B_{2,3}$.

Stage b': Similarly, following Stage b, Nodes 2 and 3 bargain over the payment y_3 associated with (s_2, s_3) and the payment y_3^b associated with $\left(s_2^{3b}, s_3^b\right)$ to be made in the event that Nodes 1 and 2 do not reach agreement at the choice (s_2, s_3) of Nodes 2 and 3. As with the paired Stages a and b, the negotiations and breakdowns in Stages a' and b' are carried out independently.

In summary, given the results $\left(s_1(\cdot), s_1^b, s_2^{1b}\ ;\ y_1, y_1^b; \mathcal{A}_{2,3}\right)$ and $\left(s_2, s_3, s_2^{3b}, s_3^b\ ;\ y_3, y_3^b, A_{2,3}\right)$ of the first four stages, the actions and payments of the game are determined as follows. If $A_{2,3} = Y$ and $(s_2, s_3) \in \mathcal{A}_{2,3}$, then the ultimate actions and payments are

$$[s_1(s_2, s_3), s_2, s_3\ ;\ y_1(s_2, s_3), y_3(s_2, s_3)].$$

If, instead, $A_{2,3} = Y$ and $(s_2, s_3) \notin \mathcal{A}_{2,3}$, the outcome of the bargaining game is then

$$\left[s_1^0, s_2^{3b}, s_3^b\ ;\ y_1^0 = 0, y_3^b\right].$$

Finally, in the event $A_{2,3} = N$, the final outcome is $\left[s_1^b, s_2^{1b}, s_3^0\ ;\ y_1^b, y_3^0 = 0\right]$. This combination of the above four stages into final results is illustrated in Figure 6.

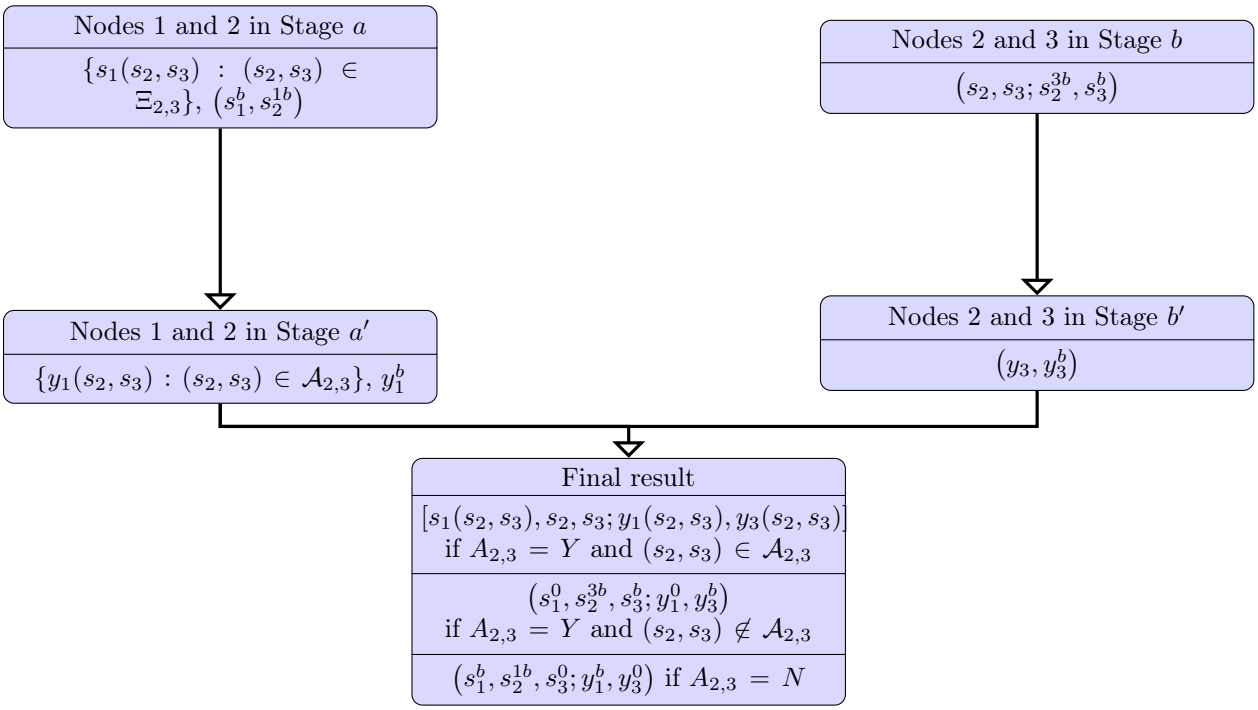

Figure 6 – The stages of the bilateral bargaining encounters

To complete the description of the extensive form game, we can use individual rationality to bound, without loss of generality, payments into specified compact intervals. For any contingency $(s_2, s_3) \in \Xi_{2,3}$, a monetary payment $y_1(s_2, s_3)$ from Node 2 to 1 that is being negotiated in Stage a' must be individually rational for both nodes. This is so for Node 1 if and only if

$$f_1(s_1(s_2, s_3), s_2) + y_1(s_2, s_3) \geq f_1\left(s_1^0, s_2^0\right) = 0.$$

Likewise, a payment y_3 that is being negotiated in Stage a' is individually rational for Node 3 if and only if

$$f_3(s_2, s_3) + y_3 \geq f_3\left(s_2^0, s_3^0\right) = 0.$$

Individual rationality for Node 2 is the condition

$$f_2(s_1(s_2,s_3),s_2,s_3) - y_1(s_2,s_3) - y_3 \geq f_2\left(s_1^0, s_2^0, s_3^0\right) = 0.$$

So, it must be the case that

$$\begin{array}{rcccl} y_1(s_2,s_3) & \leq & f_2(s_1(s_2,s_3),s_2,s_3) - y_3 & \leq & f_2(s_1(s_2,s_3),s_2,s_3) + f_3(s_2,s_3), \\ y_3 & \leq & f_2(s_1(s_2,s_3),s_2,s_3) - y_1(s_2,s_3) & \leq & f_2(s_1(s_2,s_3),s_2,s_3) + f_1(s_1(s_2,s_3),s_2) \\ & & & \leq & f_2(s_1^*(s_2,s_3),s_2,s_3) + f_1(s_1^*(s_2,s_3),s_2), \end{array}$$

where

$$s_1^*(s_2,s_3) = \underset{s_1 \in C_1(s_2)}{\operatorname{argmax}}\ f_1(s_1,s_2) + f_2(s_1,s_2,s_3).$$

Let $\underline{y}_1$ and $\bar{y}_1$ be the real-valued functions, on $C_{1,2}$ and S, respectively, determined by

$$\underline{y}_1(s_1,s_2) = -f_1(s_1,s_2), \tag{8a}$$

$$\bar{y}_1(s_1,s_2,s_3) = f_2(s_1,s_2,s_3) + f_3(s_2,s_3). \tag{8b}$$

Let $\underline{y}_3$ and $\bar{y}_3$ be the real-valued functions on $C_{2,3}$ determined by

$$\underline{y}_3(s_2,s_3) = -f_3(s_2,s_3), \tag{9a}$$

$$\bar{y}_3(s_2,s_3) = f_2(s_1^*(s_2,s_3),s_2,s_3) + f_1(s_1^*(s_2,s_3),s_2,s_3). \tag{9b}$$

Given individual rationality, we can therefore restrict the set of potential payment between Nodes 1 and 2 to the interval $[\underline{y}_1(s_1(s_2,s_3),s_2), \bar{y}_1(s_1(s_2,s_3),s_2,s_3)]$, on the contingency (s_2,s_3), given an agreed contingent treatment $s_1(\,\cdot\,)$ from Stage a. Likewise, the potential payment y_3 between Nodes 2 and 3 is restricted to $[\underline{y}_3(s_2,s_3), \bar{y}_3(s_2,s_3)]$ when the agreed treatments from Stage b is (s_2,s_3).

Individual rationality also bounds the payments y_1^b and y_3^b by $\left[\underline{y}_1^B\left(s_1^b, s_2^{3b}\right), \bar{y}_1^B\left(s_1^b, s_2^{3b}\right)\right]$ and $\left[\underline{y}_3^B\left(s_2^{1b}, s_3^b\right), \bar{y}_3^B\left(s_2^{1b}, s_3^b\right)\right]$, where $\underline{y}_1^B$ and $\bar{y}_1^B$ are the real-valued functions on $\mathcal{S}_{1,2}^B$ deter-

mined by

$$f_1\left(s_1^b, s_2^{3b}\right) + \underline{y}_1^B\left(s_1^b, s_2^{3b}\right) = 0, \tag{10a}$$

$$f_2\left(s_1^b, s_2^{1b}, s_3^0\right) - \bar{y}_1^B\left(s_1^b, s_2^{3b}\right) = 0, \tag{10b}$$

and $\underline{y}_3^B$ and $\bar{y}_3^B$ are the real-valued functions on $\mathcal{S}_{2,3}^B$ determined by

$$f_3\left(s_2^{1b}, s_3^b\right) + \underline{y}_3^B\left(s_2^{1b}, s_3^b\right) = 0, \tag{11a}$$

$$f_2\left(s_1^0, s_2^{1b}, s_3^b\right) - \bar{y}_3^B\left(s_2^{1b}, s_3^b\right) = 0. \tag{11b}$$

Any payment that is not in these intervals is not individually rational for at least one of the nodes, thus is ruled out from the bargaining game.

An extensive-form bargaining game is defined in this manner for each list (η, S, C, f, s^0) of model parameters, where $S = (S_1, S_2, S_3, S_2^1, S_2^3)$, $C = (C_1, C_3)$, $f = (f_1, f_2, f_3)$, and $s^0 = (s_1^0, s_2^0, s_3^0)$.

We could have merged Stages a and a' (and likewise have merged Stages b and b') without strategic difference. We split the game into these stages, however, in order to take advantage of the refinement associated with extensive form trembling hand perfection, which we turn to next. As we explain in Section 8.3, a failure to split the game into stages would admit additional "weird" and inefficient equilibria that survive the equilibrium refinement.

7 Solution Concept

Our solution concept is a variant of extensive form trembling hand perfection, due to Selten (1975), to which we add two requirements. The first of these is that minimum tremble probabilities do not depend on strategically irrelevant information. Despite the simplicity of its motivation, the ultimate mathematical definition of this property is long and complicated, so relegated to Appendix F. The second requirement is adapted from the refinement concept of Milgrom and Mollner (2016), *extended proper equilibrium*, under which a costless deviation

by one player must be more likely than a costly deviation by the same or another player. This section provides a complete definition of the final resulting solution concept, "restricted equilibrium."

7.1 Multistage Games, Strategies, and Nash Equilibria

In this subsection we define n-person extensive-form games in which the action spaces can be discrete or continuum, and in which the time horizon can be finite or countably infinite. We combine the treatments of finite and infinite time horizon games. Thus when $N = \infty$, the notation of the form $S_1, S_1, \ldots, S_N$ or $t \leq N$ mean $S_1, S_1, \ldots$ and $t < \infty$ respectively.

A *noncooperative n-person game in extensive form*, or simply a *multistage game*, consists of the following objects:

(i) A finite set I of players.

(ii) A number N, which is either a positive integer or ∞, called *horizon*. We let $W = \{(i,t) : i \in I, t \in \{1, \ldots, N\}\}$. We write it for (i,t).

(ii) A (finite or infinite) sequence $S_1, S_2, \ldots, S_N$ of nonempty Borel spaces called *state spaces*.

(iii) For every $it \in W$, a nonempty Borel space A_{it} called *action space*. Here, A_{it} is the space of potential actions of player i in period t.

(iv) For every $it \in W$, a nonempty compact metric space[3] Z_{it} called *information space*. An element $z_{it} \in Z_{it}$ is called an *information set*.

We apply the subscript "$< t$" to denote the projection that maps a set of the form $A = \prod_{k=1}^{N} A_k$ to $A_{<t} = \prod_{k=1}^{t-1} A_k$, and likewise maps an element $(a_k)_{k=1}^{N}$ to $a_{<t} = (a_k)_{k=0}^{t-1}$. We similarly apply the subscripts "$\leq t$" and "$> t$" for projection to periods weakly before and strictly after t respectively. We define the space $H_t = S_{\leq t} \times A_{\leq t}$ of partial histories to time t, the space $\Theta_t = S_{\leq t} \times A_{<t}$ of preplays to time t ($\Theta_1 = S_1$), as well as the space $H = S \times A$

[3] A compact metric space is necessarily separable and complete, thus a Borel space.

of (complete) histories, where $S = \prod_{t \geq 1} S_t$. The remaining elements of a multistage game are as follows.

(v) A sequence $(p_t)_{t \geq 1}$, where $p_1 \in \Delta(S_1)$ and, for $t > 0$, p_t is a Borel-measurable probability kernel from H_{t-1} to S_t. These kernels are called *state transition kernels*. For simplicity, we assume that the probability measure $p_t(\cdot \,|\, h_{t-1})$ has full support for all h_{t-1} in H_{t-1}. All of the following results remain valid without this assumption, by ignoring states that are not in the support.

(vi) For every player i, a sequence $(\zeta_{it})_{t \geq 1}$ of Borel-measurable *information functions*, where $\zeta_{i1} : S_1 \to Z_{i1}$ and, for $t > 0$, $\zeta_{it} : \Theta_t \to Z_{it}$. This means that, given a preplay $\theta_t \in \Theta_t$ of the game to period t, at time t player i is "at the information set" (or "has the information") $\zeta_{it}(\theta_t)$.

(vii) For every player i, a Borel measurable *payoff function* $u_i : H \to \mathbb{R}$.

These primitives determine a multistage game $\Gamma = (I, N, S, A, Z, p, \zeta, u)$. Roughly speaking, this game is played as follows. In period 1, nature chooses a state $s_1 \in S_1$ whose probability distribution is p_1. For each i, player i is informed of $z_{i1} = \zeta_{i1}(s_1)$, representing the information given to player i before his move in period 1. Player i then chooses an action a_{i1} in A_{i1}, possibly using a behavioral (mixed) strategy, which we will formally define. The initial partial history $h_1 = (a_1, s_1) \in H_1$ is thus determined. Nature then chooses a state s_1 with the probability distribution $p_1(\cdot \,|\, h_1)$. For each i, player i is then informed of $\zeta_{i1}(s_1, a_1)$ and chooses an action $a_{i1} \in A_{i1}$ in period 1. The partial history $h_1 = ((s_1, s_1), (a_1, a_1)) \in H_1$ is thus determined, and the game continues recursively in this manner, period by period. A play of the game results in a complete history $h = (s, a) \in H$, under which the payoff of player i is $u_i(h)$. This differs from the classical "tree" form[4] of Kuhn (1953).

[4]When one represents a finite extensive-form game in the classical tree model of Kuhn (1953), a nonterminal node of the game tree corresponds to some partial history of the form h_t or $(s_{\leq t}, a_{<t})$, whereas a terminal node corresponds to a complete history $h \in H$. Furthermore, an information set (a subset of nodes) corresponds to the inverse image $\zeta_{it}^{-1}(z_{it})$ of an information set $z_{it} \in Z_{it}$, via the information function ζ_{it}. As Aumann pointed out, on page 511 of Aumann (1964), "not all finite extensive games in the sense of Kuhn ... are included in the above definition; however all games of perfect recall are included,... The condition for

A multistage game is said to have *perfect recall* if, for all $t \in W$ and all $\tau < t$, there is a Borel measurable functon $r_{it\tau} : Z_{it} \to Z_{i\tau} \times A_{i\tau}$ such that

$$r_{it\tau}\left(\zeta_{it}(s_{\leq t}, a_{<t})\right) = (\zeta_{i\tau}(s_{\leq \tau}, a_{<\tau}), a_{i\tau}), \quad s \in S, \quad a \in A.$$

That is, any player can infer from the information that he currently possesses his information set and his action chosen at a previous date. The functions $r_{it\tau}$ are called the *recall functions*. From now on we assume that a multistage game always has perfect recall.

Given a multistage game $\Gamma = (I, N, S, A, Z, p, \zeta, u)$, a *behavioral strategy* of player i is a sequence $(\sigma_{it})_{it \in W}$ of universally measurable probability kernels σ_{it} from Z_{it} to A_{it}. (The notion of universal measurability is reviewed in Appendix A.) We let Σ_{it} denote the set of behavioral strategies of player i at time t, and then let $\Sigma_i = \prod_{t \geq 1} \Sigma_{it}$ denote the set of behavioral strategies of player i. The set of *behavioral strategy profiles*[5] is $\Sigma = \prod_{i \in I} \Sigma_i$.

We say that $\sigma = (\sigma)_{it \in W}$ is a *Borel measurable strategy profile* if for every $it \in W$, σ_{it} is a Borel measurable probability kernel from the information space Z_{it} to the action space A_{it}. So when we refer to a strategy profile σ without the qualification "Borel measurable", we mean that σ is universally measurable, but not necessarily Borel measurable.

A strategy profile $\sigma = (\sigma_i)_{i \in I}$ determines,[6] via the Ionescu-Tulcea Extension Theorem, a unique probability measure P^σ on the space $(H, \mathscr{B}(H))$ of histories, such that the finite dimensional distribution of P^σ on H_t is given by $\otimes_{k=0}^{t} \left[p_k \otimes \prod_{i \in I}(\sigma_{ik} \circ \zeta_{ik})\right]$. The expected utility of player i determined by the strategy profile σ is

$$U_i(\sigma) = \mathrm{E}^\sigma(u_i) \equiv \int_H u_i(h)\, \mathrm{P}^\sigma(dh),$$

where E^σ denotes expectation with respect to P^σ.

a Kuhn game to be included is that the game can be 'serialized' timewise,... the possibility of serialization is not at all equivalent with perfect recall (but the latter implies the former)."

[5] We do not provide here a formal definition of "mixed strategies" because (1) it is not needed for the subsequent development of the model; (2) Aumann (1964) showed that Kuhn's theorem remains valid for infinite extensive games as defined above, so that it is sufficient to consider behavioral strategies when a game has perfect recall. We sometimes omit the qualification "behavioral" when the notation makes this obvious.

[6] See Proposition 7.45 in Bertsekas and Shreve (1978).

A *Nash equilibrium* is a strategy profile $\sigma = (\sigma_i)_{i \in I} \in \Sigma$ such that

$$U_i(\sigma) \geq U_i\left(\sigma_i', \sigma_{-i}\right) \tag{12}$$

for every $i \in I$ and for every behavioral strategy $\sigma_i' \in \Sigma_i$ of player i, where σ_{-i} as usual denotes $(\sigma_j)_{j \in I \setminus \{i\}}$.

7.2 Perturbed Game and Trembling Hand Perfect Equilibrium

We now extend the notions of "perturbed game" and "trembling hand perfect equilibrium," due to Selten (1975), to multistage games that may have continuum action spaces. For a topological space E, let $\widehat{\Delta}(E)$ be the set of strictly positive probability measures on E. That is, $\chi \in \widehat{\Delta}(E)$ if and only if for, any non-empty open subset $O \subset E$, we have $\chi(O) > 0$.

For a multistage game $\Gamma = (I, N, S, A, Z, p, \zeta, u)$, a *perturbed game* $\widehat{\Gamma}$ is a triple (Γ, ϵ, χ) defined by:

- *Minimum tremble probabilities,* defined by some $\epsilon = (\epsilon_{it})_{it \in W}$, where ϵ_{it} is a Borel measurable mapping from Z_{it} to $(0, 1]$.
- *Reference strategy profile* $\chi = (\chi_{it})_{it \in W}$, where χ_{it} is a Borel measurable probability kernel from Z_{it} to A_{it}, and $\chi_{it}(z_{it}) \in \widehat{\Delta}(A_{it})$ for every $z_{it} \in Z_{it}$.

The perturbed game $\widehat{\Gamma}$ has the same game structure and utility functions as those of the original game Γ, but the behavioral strategies of $\widehat{\Gamma}$ are restricted as follows. A behavioral strategy profile σ of $\widehat{\Gamma}$ is a behavioral strategy profile of Γ such that, for each information set z_{it} and each Borel measurable subset $B \subset A_{it}$ of actions, the probability $\sigma_{it}(B \,|\, z_{it})$ assigned to actions in B is bounded below by $\epsilon_{it}(z_{it})\chi_{it}(B \,|\, z_{it})$. As ϵ_{it} is always strictly positive, this means that a strategy for $\widehat{\Gamma}$ is "bounded away from" being a pure strategy.

Letting $\widehat{\Sigma}_i$ denote the set of behavioral strategies of player i in the perturbed game $\widehat{\Gamma}$, the associated set of behavioral strategy profiles is $\widehat{\Sigma} = \prod_{i \in I} \widehat{\Sigma}_i$. A Nash equilibrium of the perturbed game $\widehat{\Gamma}$ is then a strategy profile $\sigma = (\sigma_i)_{i \in I} \in \widehat{\Sigma}$ such that, for every player i,

$$U_i(\sigma) \geq U_i\left(\sigma_i', \sigma_{-i}\right) \tag{13}$$

for every $\sigma_i' \in \widehat{\Sigma}_i$.

A sequence $\widehat{\Gamma}^1, \widehat{\Gamma}^2 \ldots$ where $\widehat{\Gamma}^n = (\Gamma, \epsilon^n, \chi^n)$ is a perturbed game of Γ, is called a *test sequence* for Γ if (1) $\sup \epsilon^n \to 0$, where, for any minimum tremble probabilities ϵ,

$$\sup \epsilon \equiv \sup_{it \in W} \sup_{z_{it} \in Z_{it}} \epsilon_{it}(z_{it}),$$

and (2) there exists a reference strategy profile $\chi = (\chi_{it})_{it \in W}$ such that for every $n \geq 0$, $it \in W$ and z_{it}, $\chi_{it}(z_{it})$ is absolute continuous with respect to $\chi^n_{it}(z_{it})$.

A behavioral strategy profile $\sigma \in \Sigma$ of Γ is a *limit equilibrium point* for $\left(\widehat{\Gamma}^n\right)_{n \geq 0}$ if, for each n, a Borel measurable Nash equilibrium $\sigma^n \in \widehat{\Sigma}^n$ of $\widehat{\Gamma}^n$ can be found such that, as n goes to infinity, $\sigma^n_{it}(z_{it})$ converges weak* to $\sigma_{it}(z_{it})$ for every $it \in W$ and every information set $z_{it} \in Z_{it}$. That is, σ^n converges weak*, pointwise, to σ. If, in addition, σ is an Nash equilibrium[7] for Γ, then we say that σ is an *extensive form trembling hand perfect equilibrium* of Γ.

7.3 Extended Properness in Extensive Form Games

We now adapt to our setting the refinement concept of Milgrom and Mollner (2016), *extended proper equilibrium*, under which a costless deviation by one player must be more likely than a costly deviation by the same or another player. We extend this notion from finite games in normal form to extensive-form games in which the action spaces can be finite, countably infinite, or a continuum, and in which the time horizon can be finite or countably infinite.

We fix a multistage game $\Gamma = (I, N, S, A, Z, p, \zeta, u)$. Given a strategy $\sigma_i \in \Sigma_i$ of player i, a time t, and a potentially different time-t strategy $\sigma'_{it} \in \Sigma_{it}$, let σ_i / σ'_{it} denote the strategy which is $\sigma_{i\tau}$ in any period $\tau \neq t$ and which is σ'_{it} in period t. For some $a_{it} \in A_{it}$, we let σ_i / a_{it} be the strategy $\sigma_i / x_{it}(a_{it})$, where $x_{it}(a_{it})$ is the time-t pure strategy that maps any

[7] In Selten's original treatment of finite games, it is not required, as a matter of definition, that the limit strategy profile σ is a Nash equilibrium for it to be extensive form perfect. This is because in finite games, being the weak* limit of perturbed Nash equilibria implies being a Nash equilibrium itself. However this implication no longer holds in infinite games of our setting without some sort of continuity conditions on the payoff function. For example, see Carbonell-Nicolau (2011b,a) for counter examples and such continuity conditions.

information sets $z_{it} \in Z_{it}$ to the fixed action a_{it}. For a strategy profile $\sigma \in \Sigma$, we let σ/σ'_{it} denote the strategy profile $(\sigma_i/\sigma'_{it}, \sigma_{-i})$, and likewise let σ/a_{it} denote the strategy profile $(\sigma_i/a_{it}, \sigma_{-i})$.

Given a strategy profile $\sigma \in \Sigma$ and some $it \in W$, consider a Borel measurable $M \subset Z_{it}$ with the property that $\mathrm{P}^\sigma(H_M) > 0$, where $H_M = \{(s,a) : \zeta_{it}(s_{\leq t}, a_{<t}) \in M\}$. Given an action $a_{it} \in A_{it}$, we define $L_i^\sigma(a_{it} \,|\, M)$ to be the expected loss for player i from playing a_{it} instead of a best response in period t against σ, conditional on the event that player i is in some information set in M. That is,

$$L_i^\sigma(a_{it} \,|\, M) = \sup_{\tilde{a}_{it} \in A_{it}} \mathrm{E}^{\sigma/\tilde{a}_{it}}(u_i \,|\, H_M) - \mathrm{E}^{\sigma/a_{it}}(u_i \,|\, H_M).$$

We also define $p_{it}^\sigma(a_{it} \,|\, M)$ as the probability that player i chooses a_{it} in period t, conditional on the event that the information set of player i at time t is in M. That is, letting

$$H(a_{it}) = \{(s', a') \in H : a'_{it} = a_{it}\},$$

we will write $p_{it}^\sigma(a_{it} \,|\, M) = \mathrm{P}^\sigma[H(a_{it}) \,|\, H_M]$.

We now introduce a notion of approximate equilibrium that will subsequently play a role in the final solution concept. For any metric space such as a typical information-set space Z_{it}, we let $B(x,r)$ denote as usual the open ball of radius r centered at a point x.

Definition 1. For strictly positive scalars λ and δ, a (λ, δ)-*extended proper equilibrium of* Γ *in extensive form* is a strategy profile $\sigma \in \Sigma$ with the following property. For every $it \in W$, there exists a function $\alpha_{it} : Z_{it} \to \mathbb{R}^+$ satisfying two properties:

1. For every $z_{it} \in Z_{it}$, the ball $M = B(z_{it}, \alpha_{it}(z_{it})) \subset Z_{it}$ centered at z_{it} is reached with positive probability, in the sense that $P^\sigma(H_M) > 0$.

2. For any player j and time τ, and any information sets $z_{it} \in Z_{it}$ and $z_{j\tau} \in Z_{j\tau}$, if there are Borel measurable $M_i \subset B(z_{it}, \alpha_{it}(z_{it}))$ and $M_j \subset B(z_{j\tau}, \alpha_{j\tau}(z_{j\tau}))$ that are reached with positive probability and satisfy $L_i^\sigma(a_{it} \,|\, M_i) > \lambda$ and $L_j^\sigma(a_{j\tau} \,|\, M_j) < \lambda$, then we must have $p_{it}^\sigma(a_{it} \,|\, M_i) \leq \delta\, p_{j\tau}^\sigma(a_{j\tau} \,|\, M_j)$.

Property 2 of the above definition requires that whenever a deviation is costly (in that its expected loss is larger than the threshold value λ), then the probability that this deviation is being played is at most a multiple δ of the probability of a costless deviation.

The notion of a (λ, δ)-extended proper equilibrium is similar in spirit to the concept due to Milgrom and Mollner (2016) of approximate equilibrium solution. Our version is prompted by the need to treat a continuum action space.

To conclude this section, we provide the solution concept for our main result concerning network bargaining problems. This definition refers to the property "independence of strategically irrelevant information," which was motivated earlier in this section and is completely defined in Appendix F.

Definition 2 (Restricted Equilibrium)**.** Given a multistage game Γ, a behavioral strategy profile σ is a *restricted equilibrium* if σ is a Nash equilibrium, and if, for all strictly positive scalar λ, sufficiently small, there exists a sequence σ^n of behavioral strategy profiles and a sequence $\{\delta^n\}$ of strictly positive reals converging to zero such that: satisfying the following properties.

1. There exists a test sequence $\widehat{\Gamma}^n = (\Gamma, \epsilon^n, \chi^n)$ for Γ, such that each $\widehat{\Gamma}^n$ respects independence of strategically irrelevant information.

2. For each n, σ^n is a Borel measurable Nash equilibrium of $\widehat{\Gamma}^n$ and is also a (λ, δ^n)-extended proper equilibrium of Γ in extensive form.

3. σ^n converges weak*, pointwise, to σ, as $n \to \infty$.

If σ is a restricted equilibrium of Γ, then a test sequence $(\widehat{\Gamma}^n)_{n\geq 0}$ associated with σ as in the definition above is called a *restricted test sequence* for σ, and a converging sequence $(\sigma^n)_{n\geq 0}$ of equilibria associated with σ is called a *restricted trembling sequence* for σ.

8 Equilibrium Network Bargaining Solution

The game described in Section 6 can now be treated as a three-agent multistage game with perfect recall in which players have a continuum of potential actions. The specific

action or state spaces, depending on the stage of the game, are finite (consisting of proposed treatments, or {accept, reject}) or are real intervals (consisting of proposed payments). Any finite space is given its discrete topology, real intervals are given the topology defined by the usual "distance" metric $m(x, y) = |x - y|$, and a product space is given its product topology.

We show existence and uniqueness of restricted equilibrium of this extensive-form network market bilateral bargaining game, and explicitly calculate the unique associated equilibrium treatments and payment outcomes, showing that they coincide in the limit (as Nature's breakdown probability η goes to zero) with the treatments and payments of the axiomatic solution. In particular, the equilibrium treatments maximize total social surplus.

8.1 Equilibrium Strategies

We first define candidate equilibrium strategies, beginning with some notation. The candidate equilibrium contingent action $s_1^*(\,\cdot\,)$ in $\mathcal{C}$ is defined by

$$s_1^*(s_2, s_3) = \underset{s_1 \in C_1(s_2)}{\operatorname{argmax}} \; f_1(s_1, s_2) + f_2(s_1, s_2, s_3). \tag{14}$$

We will later use the fact that

$$s_1^*(s_2, s_3) = \underset{s_1 \in C_1(s_2)}{\operatorname{argmax}} \, U(s_1, s_2, s_3), \tag{15}$$

recalling that U is the social welfare function.

Recall that

$$\left(s_1^B, s_2^{1B}\right) = \underset{(s_1, s_2) \in \mathcal{S}_{1,2}^B}{\operatorname{argmax}} \; f_1\left(s_1, s_2\right) + f_2\left(s_1, s_2, s_3^0\right)$$

and

$$\left(s_2^{3B}, s_3^B\right) = \underset{(s_2, s_3) \in \mathcal{S}_{2,3}^B}{\operatorname{argmax}} \; f_2\left(s_1^0, s_2, s_3\right) + f_3\left(s_2, s_3\right).$$

We let $\mathcal{S} = \{(s_1, s_2, s_3) : s_1 \in C_1(s_2),\ s_2 \in S_2, s_3 \in C_3(s_2)\}$ be the set of feasible actions. For

each $\eta \in [0,1]$, we define $y_1^\eta : \mathcal{S} \to \mathbb{R}$ and $y_3^\eta : \mathcal{S}_{2,3} \to \mathbb{R}$ so that

$$\begin{aligned} &f_2(s_1(s_2,s_3),s_2,s_3) - y_1^\eta(s_1,s_2,s_3) - y_3^\eta(s_2,s_3) - \underline{u}_{21}^\eta \\ &\quad = (1-\eta)\left[f_1(s_1(s_2,s_3),s_2) + y_1^\eta(s_1,s_2,s_3) - \underline{u}_{12}^\eta\right] \end{aligned} \tag{16}$$

and

$$\begin{aligned} &f_2(s_1^*(s_2,s_3),s_2,s_3) - y_1^\eta(s_1^*(s_2,s_3),s_2,s_3) - y_3^\eta(s_2,s_3) - \underline{u}_{23}^\eta \\ &\quad = (1-\eta)\left[f_3(s_2,s_3) + y_3^\eta(s_2,s_3) - \underline{u}_{32}^\eta\right], \end{aligned} \tag{17}$$

where

$$\begin{aligned} \underline{u}_{12}^\eta &= \underline{u}_{32}^\eta = 0 \\ \underline{u}_{21}^\eta &= \frac{1-\eta}{2-\eta}\, U\left(s_1^0, s_2^{3B}, s_3^B\right) \\ \underline{u}_{23}^\eta &= \frac{1-\eta}{2-\eta}\, U\left(s_1^B, s_2^{1B}, s_3^0\right). \end{aligned} \tag{18}$$

One can view $\underline{u}_{ij}^\eta$ as the outside option value of Node i in the bilateral bargaining between Nodes i and j, given the exogenous breakdown probability η.

In order to see that the payments y_1^η and y_3^η are uniquely well defined, first take $s_1 = s_1^*(s_2,s_3)$ in (16). In this case, (16) and (17) together uniquely determine $y_1^\eta(s_1^*(s_2,s_3),s_2,s_3)$ and $y_3^\eta(s_2,s_3)$ for any (s_2,s_3) in $\mathcal{S}_{2,3}$. Then, for any $s_1 \in C_1(s_2)$, one can solve for the unique payment $y_1^\eta(s_1,s_2,s_3)$ from (16).

Continuing to build needed notation, let $\tilde{y}_1^\eta : \mathcal{S} \to \mathbb{R}$ and $\tilde{y}_3^\eta : \mathcal{S}_{2,3} \to \mathbb{R}$ be such that

$$\begin{aligned} &(1-\eta)\left[f_2(s_1(s_2,s_3),s_2,s_3) - \tilde{y}_1^\eta(s_1,s_2,s_3) - y_3^\eta(s_2,s_3) - \underline{u}_{21}^\eta\right] \\ &\quad = f_1(s_1(s_2,s_3),s_2) + \tilde{y}_1^\eta(s_1,s_2,s_3) - \underline{u}_{12}^\eta \end{aligned} \tag{19}$$

and

$$\begin{aligned} &(1-\eta)\left[f_2(s_1^*(s_2,s_3),s_2,s_3) - y_1^\eta(s_1^*(s_2,s_3),s_2,s_3) - \tilde{y}_3^\eta(s_2,s_3) - \underline{u}_{23}^\eta\right] \\ &\quad = f_3(s_2,s_3) + \tilde{y}_3^\eta(s_2,s_3) - \underline{u}_{32}^\eta. \end{aligned} \tag{20}$$

Next, let $y_1^{B\eta}$ and $\tilde{y}_1^{B\eta}$ be the real-valued functions on $\mathcal{S}_{1,2}^B$ satisfying

$$f_2\left(s_1^b, s_2^{1b}, s_3^0\right) - y_1^{B\eta}\left(s_1^b, s_2^{1b}\right) = (1-\eta)\left[f_1\left(s_1^b, s_2^{1b}\right) + y_1^{B\eta}\left(s_1^b, s_2^{1b}\right)\right] \tag{21}$$

and

$$(1-\eta)\left[f_2\left(s_1^b, s_2^{1b}, s_3^0\right) - \tilde{y}_1^{B\eta}\left(s_1^b, s_2^{1b}\right)\right] = f_1\left(s_1^b, s_2^{1b}\right) + \tilde{y}_1^{B\eta}\left(s_1^b, s_2^{1b}\right). \tag{22}$$

Let $y_3^{B\eta}$ and $\tilde{y}_3^{B\eta}$ be the real-valued functions on $\mathcal{S}_{2,3}^B$ satisfying

$$f_2\left(s_1^0, s_2^{3b}, s_3^b\right) - y_3^{B\eta}\left(s_2^{3b}, s_3^b\right) = (1-\eta)\left[f_3\left(s_2^{3b}, s_3^b\right) + y_3^{B\eta}\left(s_2^{3b}, s_3^b\right)\right] \tag{23}$$

and

$$(1-\eta)\left[f_2\left(s_1^0, s_2^{3b}, s_3^b\right) - \tilde{y}_3^{B\eta}\left(s_2^{3b}, s_3^b\right)\right] = f_3\left(s_2^{3b}, s_3^b\right) + \tilde{y}_3^{B\eta}\left(s_2^{3b}, s_3^b\right). \tag{24}$$

Our candidate equilibrium strategy profile $\sigma^{*\eta}$ is defined as follows:

- In Stage a, Nodes 1 and 2 offer each other, at any period t, the contingent treatment contract $s_1^*(s_2, s_3)$ for all contingencies $(s_2, s_3) \in \mathcal{S}_{2,3}^t$, and $\left(s_1^B, s_2^{1B}\right)$ for the breakdown contingency $B_{2,3}$ if $B_{2,3}$ is still open for negotiation in period t, that is if $B_{2,3} \in \mathcal{S}_{2,3}^t$. These offers are immediately accepted. Given any other feasible contingent offer $s_1 : \mathcal{S}_{2,3}^t \to S_1$ and $\left(s_1^b, s_2^{1b}\right) \in \mathcal{S}_{1,2}^B$, Nodes 1 and 2 both accept $s_1(s_2, s_3)$ (respectively, $\left(s_1^b, s_2^{1b}\right)$) at any (s_2, s_3) for which $s_1(s_2, s_3) = s_1^*(s_2, s_3)$ (respectively, if $\left(s_1^b, s_2^{1b}\right) = \left(s_1^B, s_2^{1B}\right)$), and otherwise reject $s_1(s_2, s_3)$ (respectively, $\left(s_1^b, s_2^{1b}\right)$).

- Suppose $(s_1(\,\cdot\,), \Xi_{2,3})$ is the outcome of Stage a. (Recall that $\Xi_{2,3}$ is the set of all contingencies on which Nodes 1 and 2 reached an agreement after Stage a). At each period t in Stage aa, contingent on any (s_2, s_3) in $\Xi_{2,3}^t$, Node 1 offers the payment

  $$\mathbb{1}\left\{y_1^\eta(s_1, s_2, s_3) > \underline{y}_1(s_1, s_2)\right\} y_1^\eta(s_1, s_2, s_3) + \mathbb{1}\left\{y_1^\eta(s_1, s_2, s_3) \leq \underline{y}_1(s_1, s_2)\right\} \bar{y}_1(s_1, s_2, s_3)$$

  if it is the turn of Node 1 to make an offer, while Node 2 offers

  $$\max\left\{\tilde{y}_1^\eta(s_1, s_2, s_3), \underline{y}_1(s_1, s_2)\right\}.$$

  If $y_1^\eta > \underline{y}_1$, then Node 1 accepts offers that are at least $\tilde{y}_1^\eta$ and rejects the rest; If $y_1^\eta \leq \underline{y}_1$, then Node 1 accepts all offers that are strictly larger than $\underline{y}_1$ but rejects $\underline{y}_1$.

If $y_1^{\eta} \geq \underline{y}_1$, Node 2 accepts offers that are at most y_1^{η} and rejects the rest. If $y_1^{\eta} < \underline{y}_1$, Node 2 rejects all offers.

If $B_{2,3} \in \Xi_{2,3}$, then contingent on $B_{2,3}$, Node 1 offers the payment

$$\mathbb{1}\left\{y_1^{B\eta}\left(s_1^b, s_2^{1b}\right) > \underline{y}_1^B\left(s_1^b, s_2^{1b}\right)\right\} y_1^{B\eta}\left(s_1^b, s_2^{1b}\right) + \mathbb{1}\left\{y_1^{B\eta}\left(s_1^b, s_2^{1b}\right) \leq \underline{y}_1^B\left(s_1^b, s_2^{1b}\right)\right\} \bar{y}_1^B\left(s_1^b, s_2^{1b}\right)$$

if it is the turn of Node 1 to make an offer, while Node 2 offers

$$\min\left\{\tilde{y}_1^{B\eta}\left(s_1^b, s_2^{1b}\right), \underline{y}_1^B\left(s_1^b, s_2^{1b}\right)\right\}.$$

If $y_1^{B\eta} > \underline{y}_1^B$, then Node 1 accepts offers that are at least $\tilde{y}_1^{B\eta}$ and rejects the rest. If $y_1^{B\eta} \leq \underline{y}_1^B$, then Node 1 accepts all offers that are strictly larger than $\underline{y}_1^B$ but rejects $\underline{y}_1^B$. If $y_1^{B\eta} \geq \underline{y}_1^B$, Node 2 accepts offers that are at most y_1^{η} and rejects the rest. If $y_1^{B\eta} < \underline{y}_1^B$, Node 2 rejects all offers.

- In Stage b, Nodes 2 and 3 offer the simple contingent treatment contract $\left(s_2^{**}, s_3^{**}\,;\, s_2^{3B}, s_3^B\right)$ at each period. Both of these nodes accept this simple contingent treatment contract, and only this one.

- In Stage bb, given any simple contingent treatment contract $\left(s_2, s_3\,;\, s_2^{3b}, s_3^b\right)$ agreed in Stage b, Node 3 offers the simple contingent payment

$$\Big[\mathbb{1}\left\{y_3^{\eta}(s_2, s_3) > \underline{y}_3(s_2, s_3)\right\} y_2^{\eta}(s_2, s_3) + \mathbb{1}\left\{y_3^{\eta}(s_2, s_3) \leq \underline{y}_1(s_2, s_3)\right\} \bar{y}_3(s_2, s_3),$$
$$\mathbb{1}\left\{y_3^{B\eta}\left(s_2^{3b}, s_3^b\right) < \underline{y}_3^B\left(s_2^{3b}, s_3^b\right)\right\} y_3^{B\eta}\left(s_2^{3b}, s_3^b\right) + \mathbb{1}\left\{y_3^{B\eta}\left(s_2^{3b}, s_3^b\right) \leq \underline{y}_3^B\left(s_2^{3b}, s_3^b\right)\right\} \bar{y}_3^B\left(s_2^{3b}, s_3^b\right)\Big]$$

while Node 2 offers the simple contingent payment

$$\left[\min\left\{\tilde{y}_3^{\eta}(s_2, s_3), \underline{y}_3(s_2, s_3)\right\},\ \min\left\{\tilde{y}_3^{B\eta}\left(s_2^{3b}, s_3^b\right), \underline{y}_3^B\left(s_2^{3b}, s_3^b\right)\right\}\right].$$

The first element of each payment pair applies when Nodes 1 and 2 reach agreement, while the second applies when Nodes 1 and 2 break down.

When Nodes 1 and 2 reach agreement:

If $y_3^\eta > \underline{y}_3$, then Node 3 accepts offers that are at least $\tilde{y}_3^\eta$ and rejects the rest. If $y_3^\eta \leq \underline{y}_3$, then Node 3 accepts all offers that are strictly larger than $\underline{y}_3$ but rejects $\underline{y}_3$. If $y_3^\eta \geq \underline{y}_3$, Node 2 accepts offers that are at most y_3^η and rejects the rest. If $y_3^\eta < \underline{y}_3$, Node 2 rejects all offers.

When Nodes 1 and 2 break down:
If $y_3^{B\eta} > \underline{y}_3^B$, then Node 3 accepts offers that are at least $\tilde{y}_3^{B\eta}$ and rejects the rest. If $y_3^{B\eta} \leq \underline{y}_3^B$, then Node 3 accepts all offers that are strictly larger than $\underline{y}_3^B$ but rejects $\underline{y}_3^B$. If $y_3^{B\eta} \geq \underline{y}_3^B$, Node 2 accepts offers that are at most y_3^η and rejects the rest. If $y_3^{B\eta} < \underline{y}_3^B$, Node 2 rejects all offers.

8.2 The Main Result

Our main result provides the following sense in which contingent bilateral contracting is efficient.

Theorem 1. *Fix a contracting network (S, C, f, s^0). There is an $\eta^* > 0$ such that the following is true. For any breakdown probability $\eta \in (0, \eta^*)$, the pure strategy profile $\sigma^{*\eta}$ is a restricted equilibrium of the network market bargaining game $\Gamma(\eta, S, C, f, s^0)$. For any restricted equilibrium of this game, with probability 1, the (same) deterministic outcome $\left[s^{**}, \left(y_1^\eta(s^{**}), y_3^\eta(s_2^{**}, s_3^{**})\right)\right]$ is immediately implemented as an accepted offer in period 1 of the respective bargaining stages.*

This theorem is proved in Appendix B.

Proposition 1. *As $\eta \to 0$, the unique restricted equilibrium payments of the game $\Gamma(\eta, S, C, f, s^0)$ converge to the axiomatic solution payments y^a.*

8.3 Why Separate Stages a and aa, or Stages b and bb?

In Section 6, we mentioned that although splitting Stage a from Stage aa, or Stage b from Stage bb, makes no strategic difference, this separation allows us to take advantage of the refinement associated with extensive form trembling-hand perfection. If we were to merge

Stages a and aa, or Stages b and bb, then in general there would be extensive form trembling-hand perfect equilibria that are not efficient. This is explained as follows.

If we were to merge Stages b and bb while Stages a and aa separate, then Nodes 1 and 2 would have the same equilibrium behavior in Stage a and aa. The outcome of those two stages is that Nodes 1 and 2 agree on the contingent contract (s^*, s_1^B, s_2^{1B}). Now, suppose we wish to sustain an equilibrium in which Nodes 2 and 3 sign a contract $(\bar{s}_2, \bar{s}_3) \in \mathcal{S}_{2,3}$ that is different from the socially efficient choice (s_2^{**}, s_3^{**}). When Stages b and bb are merged into a single stage, nodes 2 and 3 offer a pair of actions $(s_2, s_3) \in \mathcal{S}_{2,3}$ and a payment y_3 in each round. Conditional on making a "trembling mistake," Nodes 2 and 3 could then make "unreasonable" mistakes in their payment. For example, for any (s_2, s_3) that is not $(\bar{s}_2, \bar{s}_3)$, the bargaining outcome of Nodes 2 and 3 could be $(s_2, s_3; y_3 = -M)$ with some probability ϵ. That is, conditional on signing the contract (s_2, s_3), Node 3 would need to pay the maximum payment M to Node 2. This payment would place Node 2 in an extremely bad bargaining position when he bargains with Node 1 on the contingency (s_2, s_3). This is so because if Node 2 wishes to secure this large payment M from Node 3, then Node 2 would be forced to share his surplus with Node 1. Thus the total payoff of Nodes 2 and 3 as a whole would be less for the contract choice (s_2, s_3) than for $(\bar{s}_2, \bar{s}_3)$. This is so for *any* $(s_2, s_3) \neq (\bar{s}_2, \bar{s}_3)$. This shows that an equilibrium in which Nodes 2 and 3 sign an arbitrary contract $(\bar{s}_2, \bar{s}_3)$ can be sustained by allowing uneasonable payment trembles.

If we were to merge Stages a and aa, the same argument shows that Nodes 1 and 2 need not agree, in equilibrium, on the efficient contingent contract s_1^*. It may be that some other natural refinement concept would not require a separation of these bargaining stages in order to achieve efficiency.

Appendices

A Auxiliary Measure Theoretic Facts

A.1 Borel and Analytic Sets, Semi-analytic Functions

For an extensive-form game that allows players a continuum of potential actions, some care must be taken regarding the measurability of mixed and behavioral strategies. Aumann (1964) was the first to address this concern.

A topological space is said to be a *Borel space*, if it is topologically homeomorphic to a Borel subset of a Polish space. By Kuratowski's Theorem,[8] Borel spaces can be easily classified by isomorphism via cardinality. That is, a Borel space is either finite or denumerable with the discrete structure, or it is isomorphic with the unit interval $[0, 1]$. For any topological Y, we denote its Borel σ-algebra $\mathscr{B}(Y)$, and denote by $\Delta(Y)$ the set of probability measures on $\mathscr{B}(Y)$ endowed with the weak* topology. When Y is a Borel space, $\Delta(Y)$ is a Borel space[9]. The property of being a Borel space is preserved when taking countable Cartesian products.[10]

However the Borel σ-algebra has a deficiency in the context of optimization that it does not preserve measurability under projection. It is well known that if Y and Z are Borel spaces, and if B a Borel subset of $Y \times Z$, then the projection of B on Y need not be Borel. To circumvent this difficulty, the literature of stochastic optimal control works with an enriched σ-algebra, the *universal σ-algebra*[11]. We introduce this concept into the context of noncooperative games in extensive form.

The *universal σ-algebra* $\mathscr{U}(Y)$ of a Borel space is the intersection of all completions of

[8] A formal statement and proof of Kuratowski Theorem is provided by Bertsekas and Shreve (1978), Corollary 7.16.1.

[9] See Section 7.4 in Bertsekas and Shreve (1978).

[10] See Proposition 7.13 in Bertsekas and Shreve (1978).

[11] Interested readers are referred to Appendix A of Bertsekas (2012) or Section 7 of Bertsekas and Shreve (1978) for a detailed treatment.

$\mathscr{B}(Y)$ with respect to all probability measures. That is,

$$\mathscr{U}(Y) = \bigcap_{\mathbb{P}\in\Delta(Y)} \mathscr{B}^{\mathbb{P}}(Y)$$

where $\mathscr{B}^{\mathbb{P}}(Y)$ is the complete σ-algebra with respect to the probability measure $\mathbb{P}$. Clearly, we have $\mathscr{B}(Y) \subset \mathscr{U}(Y)$. A probability measure $\mathbb{P}$ on $(Y, \mathscr{B}(Y))$ has a unique extension to a probability measure $\overline{\mathbb{P}}$ on $(Y, \mathscr{U}(Y))$. We write simply $\mathbb{P}$ instead of $\overline{\mathbb{P}}$.

Let Y and Z be Borel spaces, and consider a function $g : Y \to Z$. We say that g is *universally measurable* if $g^{-1}(B) \in \mathscr{U}(Y)$ for every $B \in \mathscr{B}(Z)$. A *probability kernel* q from Y to Z is function from Y to $\Delta(Z)$. We sometimes denote by $q(dz \,|\, y)$ the probability measure $q(y)$ on Z. A probability kernel q is Borel measurable (universally measurable) if and only if for each Borel set $B \in \mathscr{B}(Z)$, the function $q(B \,|\, y)$ is Borel measurable (universally measurable, respectively) in y. (See Proposition 7.26 in Bertsekas and Shreve (1978).)

A proof of Kuratowski's theorem can be found in Chapter I, Section 3 of Parthasarathy (1972).

Proposition 2 (Kuratowski's theorem)**.** *Let X be a Borel space, Y a separable metrizable space, and $\varphi : X \to Y$ be one-to-one and Borel measurable. Then $\varphi(X)$ is a Borel subset of Y and φ^{-1} is Borel measurable. In particular, if Y is a Borel space, then X and $\varpi(X)$ are isomorphic Borel spaces.*

Suppose $(\Omega, \mathscr{F}, \mathrm{P})$ is a probability space, $(\Omega_1, \mathscr{F}_1)$ is a measurable space, and X a measurable function from Ω to Ω_1. We will simply write "for P almost every $x \in X$" to mean "$\forall x \in N$, where $N \in \mathscr{F}_1$ is such that $\mathrm{P}(X \in N) = 0$". If $\widetilde{\mathrm{P}}$ is another probability measure on $(\Omega, \mathscr{F})$ such that X has the same probability distribution under P and $\widetilde{\mathrm{P}}$, that is, if for any $A \in \mathscr{F}_1$,

$$\mathrm{P}(X \in A) = \widetilde{\mathrm{P}}(X \in A),$$

then it is clear that "for P almost every $x \in X$" is means the same as "for $\widetilde{\mathrm{P}}$ almost every $x \in X$".

Now suppose Y is an extended real-valued random variable on Ω for which either $\mathrm{E}^{\mathrm{P}}(Y^+)$

or $\mathrm{E}^{\mathrm{P}}(Y^-)$ is finite, so that the conditional expectation $\mathrm{E}^{\mathrm{P}}(Y \,|\, \mathscr{G})$ of Y given a sub-sigma-algebra $\mathscr{G}$ is well defined. Suppose $g : \Omega_1 \to \mathbb{R}$ is measurable and $\mathrm{E}^{\mathrm{P}}(Y \,|\, X) = g(X)$, P almost surely. We define, for every $x \in \Omega_1$,

$$\mathrm{E}^{\mathrm{P}}(Y \,|\, X = x) = g(x).$$

For any measurable $\tilde{g} : \Omega_1 \to \mathbb{R}$ such that $\mathrm{E}^{\mathrm{P}}(Y \,|\, X) = \tilde{g}(X)$, P almost surely, we have

$$\mathrm{E}^{\mathrm{P}}(Y \,|\, X = x) = \tilde{g}(x)$$

for P almost every $x \in \Omega_1$.

Proposition 3. *Let $(\Omega, \mathcal{F})$ be a measurable space, and Y is an extended real-valued random variable on Ω for which either $\mathrm{E}(Y^+)$ or $\mathrm{E}(Y^-)$ is finite. Let $(\Omega_1, \mathscr{F}_1)$ be a measurable spaces, and X is a measurable function from Ω to Ω_1. Suppose P and $\widetilde{\mathrm{P}}$ are two probability measures on $(\Omega, \mathcal{F})$. If (X, Y) has the same probability distribution under P and $\widetilde{\mathrm{P}}$, that is, if for every $A \in \mathscr{F}_1$, $B \in \mathscr{F}$,*

$$\mathrm{P}(X \in A, Y \in B) = \widetilde{\mathrm{P}}(X \in A, Y \in B),$$

then

$$\mathrm{E}^{\mathrm{P}}(\, Y \,|\, X = x) = \mathrm{E}^{\widetilde{\mathrm{P}}}(\, Y \,|\, X = x), \tag{25}$$

for P and $\widetilde{\mathrm{P}}$ every $x \in \Omega_1$.

Proof. For any bounded measurable $f : \Omega \to \mathbb{R}$, we have

$$\begin{aligned} & \int_{\Omega_1} \mathrm{E}^{\widetilde{\mathrm{P}}}(Y \,|\, X = x)\, f(x)\; \mathrm{P} \circ X^{-1}(dx) \\ = & \int_{\Omega_1} \mathrm{E}^{\widetilde{\mathrm{P}}}(Y \,|\, X = x)\, f(x)\; \widetilde{P} \circ X^{-1}(dx) \\ = & \; \mathrm{E}^{\widetilde{\mathrm{P}}}\left[\mathrm{E}^{\widetilde{\mathrm{P}}}(Y \,|\, X) f(X)\right] = \mathrm{E}^{\widetilde{\mathrm{P}}}\left[Y f(X)\right] = \mathrm{E}^{\mathrm{P}}\left[Y f(X)\right]. \end{aligned}$$

Therefore the function $\tilde{g} : x \to \mathrm{E}^{\widetilde{\mathrm{P}}}(\, Y \,|\, X = x)$ satisfies $\tilde{g}(X) = \mathrm{E}^{\mathrm{P}}(Y \,|\, X)$, P almost surely. Thus Equation (25) holds for P every $x \in \Omega_1$. Likewise, Equation (25) holds for $\widetilde{\mathrm{P}}$ every

$x \in \Omega_1$ □

Proposition 4 (Propostion 7.44 in Bertsekas and Shreve (1978))**.**
Let X, Y and Z be Borel spaces. Suppose $f : X \to Y$ and $g : Y \to Z$ are Borel measurable (universally measurable). Then the composition $g \circ f$ is Borel measurable (universally measurable, respectively).

Proposition 5 (Proposition 7.26 and Lemma 7.28 in Bertsekas and Shreve (1978))**.**
Let X and Y be Borel spaces, $\mathscr{E}$ is a collection of subsets of Y which generates $\mathscr{B}(Y)$ and is closed under finite intersections, and q a probability kernel from X to Y. Then q is Borel measurable (universally measurable) if and only if the mapping $X \ni x \mapsto q(E \,|\, x) \in [0,1]$ is Borel measurable (universally measurable, respectively) for every $E \in \mathscr{E}$.

Proposition 6. *Let X, Y and Z be Borel spaces, and let q be a Borel-measurable (universally measurable) probability kernel from X to $Y \times Z$. Then there exists Borel measurable (universally measurable) probability kernels $\mathfrak{b}$ from $X \times Y$ to Z, and m from X to Y such that*

$$q(\underline{Y} \times \underline{Z} \,|\, x) = \int_{\underline{Y}} \mathfrak{b}(\underline{Z} \,|\, x, y) m(dy \,|\, x), \qquad \forall\, \underline{Y} \in \mathscr{B}(Y), \quad \underline{Z} \in \mathscr{B}(Z)$$

Proof. The Borel measurability part is precisely Corollary 7.27.1 in Bertsekas and Shreve (1978). The universal measurability part is an immediate application of Proposition 7.27 in Bertsekas and Shreve (1978) along with the fact that $\mathscr{U}(X) \otimes \mathscr{U}(Y) \subset \mathscr{U}(X \times Y)$. □

Proposition 7 (Proposition 7.29 and 7.46 in Bertsekas and Shreve (1978))**.**
Let X and Y be Borel spaces and q a Borel-measurable (universally measurable) probability kernel from X to Y. If $f : X \times Y \to \mathbb{R}$ is Borel measurable (universally measurable) and bounded either above or below, then the function $X \ni x \mapsto \int f(x,y) q(dy \,|\, x) \in \mathbb{R}$ is Borel measurable (universally measurable, respectively).

Corollary 1 (Extension of Corollary 7.29.1 and 7.46.1 in Bertsekas and Shreve (1978))**.**
Let X and Z be Borel spaces, and let $f : X \times Y \to [-\infty, \infty]$ be Borel measurable (universally measurable) and bounded either above or below. The function $\theta_f : X \times \Delta(Y) \to [-\infty, \infty]$

given by

$$\theta_f(x,p) = \int_Y f(x,y)p(dy)$$

is Borel measurable (universally measurable, respectively).

Proof. Define a Borel measurable probability kernel on $X \times \Delta(Y)$ to Y by $q(dy \,|\, x,p) = p(dy)$ and apply Proposition 7. □

A subset A of a Borel space X is said to be *analytic* if there exits a Borel space Y and a Borel subset B of $X \times Y$ such that $A = \text{proj}_X(B)$, where proj_X is the projection mapping from $X \times Y$ to X. It is clear that every Borel subset of a Borel space is analytic. It is also true that every analytic set is universally measurable[12].

Let X be a Borel space and let $f : X \to [-\infty, \infty]$ be a function. We say that f is *upper semianalytic* if the level set

$$\{x \in X \,|\, f(y) > c\}$$

is analytic for every $c \in \mathbb{R}$. Likewise, f is *lower semianalytic* if $\{x \in X \,|\, f(x) < c\}$ is analytic for every $c \in \mathbb{R}$. Every upper or lower semianalytic function is universally measurable. Moreover, upper (lower) semianalycity is preserved under partial maximization (minimization, respectively), and under integration with respect to a Borel measurable probability kernel:

Proposition 8 (Proposition 7.47 in Bertsekas and Shreve (1978))**.**
Let X and Y be Borel spaces, and consider a function f from $X \times Y$ to $[-\infty, \infty]$. Let $\bar{f} : X \to [-\infty, \infty]$ and $\underline{f} : X \to [-\infty, \infty]$ be defined by

$$\bar{f}(x) = \sup_{y \in Y} f(x,y),$$
$$\underline{f}(x) = \inf_{y \in Y} f(x,y).$$

If f is upper semianalytic, then $\bar{f}$ is upper semianalytic; If f is lower semianalytic, then $\underline{f}$ is lower semianalytic.

[12] for a proof, see Bertsekas and Shreve (1978), Corollary 7.42.1.

Proposition 9 (Proposition 7.48 in Bertsekas and Shreve (1978))**.**
Let X and Y be Borel spaces, and let $q(dy \,|\, x)$ be a probability kernel from X to Y. Consider a function f from $X \times Y$ to $[-\infty, \infty]$ that is bounded either above or below. If q is Borel measurable and f is upper (lower) semianalytic, then the function $\ell : X \to [-\infty, \infty]$ given by

$$\ell(x) = \int_Y g(x, y) q(dy \,|\, x)$$

is upper (lower, respectively) semianalytic.

Corollary 2 (Extension of Corollary 7.48.1 in Bertsekas and Shreve (1978))**.**
Let X and Z be Borel spaces, and let $f : X \times Y \to [-\infty, \infty]$ be upper (lower) semianalytic and bounded either above or below. The function $\theta_f : X \times \Delta(Y) \to [-\infty, \infty]$ given by

$$\theta_f(x, p) = \int_Y f(x, y) p(dy)$$

is upper (lower, respectively) semianalytic.

Proof. Define a Borel measurable probability kernel on $X \times \Delta(Y)$ to Y by $q(dy \,|\, x, p) = p(dy)$ and apply Proposition 9. □

We conclude this subsection by stating the universally measurable selection theorem. A proof can be found in Bertsekas and Shreve (1978) Proposition 7.50:

Theorem 2 (Universally Measurable Selection Theorem)**.** *Let X and Y be Borel spaces, $B \subset X \times Y$ an analytic set, and let $f : (X, Y) \to [-\infty, \infty]$ be upper semianalytic. Define $f^* : proj_X(D) \to [-\infty, \infty]$ by*

$$f^*(x) = \sup_{y \in B_x} f(x, y).$$

For any $\epsilon > 0$, there exists a universally measurable function $\varphi : proj_X(D) \to Y$ such that for every $x \in X$, $\varphi(x) \in B_x$ and

$$f(x, \varphi(x)) \geq f^\epsilon(x) := (f^*(x) - \epsilon) \, \mathbb{1}_{f^*(x) < \infty} + \frac{1}{\epsilon} \, \mathbb{1}_{f^*(x) = \infty}.$$

B Appendix: Proof of Theorem 1

We now provide a proof of Theorem 1.

B.1 Filtering in an Multistage Game

We first provide the following basic results about filtering, that is adapted from Section 10.3.1 in Bertsekas and Shreve (1978).

Given a multistage game $\Gamma = (I, N, S, A, Z, p, \zeta, u)$. Suppose a strategy profile σ is Borel measurable. Then given σ_{-i}, player i faces a single person imperfect state information stochastic control problem. Recall that $\zeta_{it} : \Theta_t \to Z_{it}$ is the (Borel measurable) information function of player i in period t. ζ_{it} can be lifted to be a Borel measurable function $\hat{\zeta}$ on H by $\bar{\zeta}(s, a) = \zeta(s_{\leq t}, a_{<t})$. We write simply ζ_{it} instead of $\bar{\zeta}_{it}$. Likewise, if g is a function defined on a domain which is a projected space of H, we sometimes lift g to be a function on H and write g for the lifted function. We first fix a Borel-measurable strategy profile σ.

Lemma 1. *For every $it \in W$, there exist Borel-measurable probability kernels $\mathfrak{b}_{it}(d\theta_t \,|\, p, z_{it})$ from $\Delta(\Theta_t) \times Z_{it}$ to Θ_t which satisfy*

$$\int_{\underline{\Theta}_t} \mathbb{1}\{\zeta(\theta_t) \in \underline{Z}_{it}\}\, p(d\theta_t) = \int_{\underline{Z}_{it}} \mathfrak{b}_{it}(\underline{\Theta}_t \,|\, p, z_{it}) \left(p \circ \zeta_{it}^{-1}\right)(dz_{it}) \tag{26}$$

Proof. For fixed $p \in \Delta(\Theta_t)$, define a probability measure q on $\Theta_t \times Z_{it}$ by specifying its values measurable rectangles to be

$$q(\underline{\Theta}_t \times \underline{Z}_{it} \,|\, p) = \int_{\underline{\Theta}_t} \mathbb{1}\{\zeta(\theta_t) \in \underline{Z}_{it}\}\, p(d\theta_t).$$

By Proposition 5 and corollary 1, q is Borel-measurable probability kernel from $\Delta(\Theta_t)$ to $\Theta_t \times Z_{it}$. By Proposition 6, this probability kernel can be decomposed into its marginal from Z_{it} given $\Delta(\Theta_t)$ and a Borel-measurable probability kernel $\mathfrak{b}(d\theta_t \,|\, p, z_{it})$ on Θ_t given $\Delta(\Theta_t) \times Z_{it}$ such that Equation (26) holds. □

It is customary to view $\mathfrak{b}_{it}$ as a belief updating operator of player i in period t: given a

prior distribution of the preplay θ_t of the game, player i observes z_{it} and updates his belief about θ_t to be the posterior distribution $\mathfrak{b}_{it}(d\theta_t \,|\, p, z_{it})$.

For every $it \in W$, consider the function $\bar{f}_{it} : \Delta(\Theta_t) \times A_{it} \to \Delta(\Theta_{t+1})$ defined by

$$\bar{f}_{it}(q, a_{it})(\underline{\Theta}_{t+1}) = \int_{\left(\theta_t, a_{it}, (a_{jt})_{j\neq i}, s_{t+1}\right) \in \underline{\Theta}_t} q(d\theta_t) \prod_{j \in I\setminus\{i\}} \sigma_{jt}(da_{jt} \,|\, \zeta_{jt}(\theta_t)) \tag{27}$$
$$p_{t+1}\left[ds_{t+1} \,\middle|\, \left(\theta_t, a_{it}, (a_{jt})_{j\in I\setminus\{i\}}\right)\right], \quad \forall\, \underline{\Theta}_{t+1} \in \mathscr{B}(\Theta_{t+1}).$$

Equation (27) is called the *one-stage prediction equation*. If player i has a posterior distribution q_{it} about Θ_t and takes an action a_{it} in period t, then his a priori belief of Θ_{t+1} is $\bar{f}_{it}(q_{it}, a_{it})$. The mapping $\bar{f}$ is Borel measurable (Propositions 5 and 7).

Define the probability kernels q_{it} from $Z_{it} \times A_{it}$ to Θ_{t+1} $(t \geq 1)$ recursively by

$$q_{i1}(z_{i1}, a_{i1}) = \bar{f}_{i1}(\mathfrak{b}_{i1}(p_1, z_{i1}), a_{i1}), \tag{28}$$

$$q_{it}\left(z_{it}, a_{it}\right) = \bar{f}_{it}\left(\mathfrak{b}_{it}\left[q_{i(t-1)}\left(r_{it(t-1)}\left(z_{it}\right)\right), z_{it}\right], a_{it}\right) \quad \text{where } r_{it(t-1)} \text{ is the recall function.} \tag{29}$$

Note that for each $t \geq 1$, q_t is Borel measurable by composition of Borel measurable functions.

Equations (26) to (28) are called *filtering equations* corresponding to the multistage game Γ and the Borel measurable strategy profile σ. The probability kernels q_{it} provide a version of the conditional distribution about the preplay of the game given player i current information, as the following lemma shows.

Let $\mathfrak{a}_{it}$ be the projection mapping from H to A_{it}. That is, $\mathfrak{a}_{it}(s, a) = a_{it}$.

Lemma 2. *Given a multistage game Γ. For every Borel measurable strategy profile σ, $it \in W$ and $\underline{\Theta}_{t+1} \in \mathscr{B}(\Theta_{t+1})$, we have*

$$\mathrm{P}^\sigma[\underline{\Theta}_{t+1} \,|\, \zeta_{it}, \mathfrak{a}_{it}] = q_{it}(\underline{\Theta}_{t+1} \,|\, \zeta_{it}, \mathfrak{a}_{it}) \qquad \mathrm{P}^\sigma \text{ almost surely.}$$

Proof. The proof works by induction in t, and uses the filtering equations and Fubini's theorem. We omit the details and refer to Lemma 10.4 in Bertsekas and Shreve (1978) for a complete proof of a very similar result. □

Fix some Borel measurable strategies $\sigma_{t+1}, \dots$ and the state transition kernels $p_{t+2}, \dots$ determine, a probability measure q_t on Θ_{t+1} determines, via the Ionescu-Tulcea Extension Theorem, a unique probability measure $\kappa^\sigma(dh \,|\, q_t)$ on the space $(H, \mathscr{B}(H))$ of histories, such that the finite dimensional distribution of $\kappa^\sigma(dh \,|\, q_t)$ on $\Theta_{t+1+\tau}$ is given by $q_{it} \otimes_{k=1}^{\tau} \left[\prod_{i \in N} \left(\sigma_{i(t+k)} \circ \zeta_{i(t+k)}\right) \otimes p_{t+k+1}\right]$ for every $\tau \geq 0$. The function that maps every $q_t \in \Delta(\Theta_{t+1})$ to $\kappa^\sigma(dh \,|\, q_t)$ is Borel measurable, as the following lemma shows.

Lemma 3. *The function $\kappa^\sigma : \Delta(\Theta_{t+1}) \to \Delta(H)$ defined above is Borel measurable.*

Proof. Fix any Borel rectangle $\prod_{k \leq t+1+\tau} \underline{S}_k \times \prod_{i \in N, k \leq t+\tau} \underline{A}_{ik}$ in $\Theta_{t+1+\tau}$, the function

$$\Delta(\Theta_{t+1}) \ni q_t \mapsto \kappa^\sigma \left(\prod_{k \leq t+1+\tau} \underline{S}_k \times \prod_{i \in N, k \leq t+\tau} \underline{A}_{ik} \,\middle|\, q_t \right) \in [0,1]$$

is Borel measurable by Corollary 1. Thus κ^σ is Borel measurable by Proposition 5. □

Since both $q_{it} : Z_{it} \times A_{it} \to \Delta(\Theta_{t+1})$ and $\kappa : \Delta(\Theta_{t+1}) \to \Delta(H)$ are Borel measurable, their composition $\kappa \circ q_{it}$ defines a Borel measurable probability kernel from $Z_{it} \times A_{it}$ to H. We immediately obtain the following corollary.

Corollary 3. *Suppose σ is a Borel measurable strategy profile and $g : H \to \mathbb{R}$ is upper semianalytic and bounded. Then for every $it \in W$, there exists a bounded upper semianalytic function $\hat{g} : Z_{it} \times A_{it} \to \mathbb{R}$ such that*

$$\hat{g}(z_{it}, a_{it}) = \mathrm{E}^\sigma[g \,|\, \zeta_{it} = z_{it}, \mathfrak{a}_{it} = a_{it}] \tag{30}$$

for P^σ almost every (z_{it}, a_{it}).

Proof. Define $\hat{g} : Z_{it} \times A_{it} \to \mathbb{R}$ by

$$\hat{g}(z_{it}, a_{it}) = \int g(h) \; (\kappa^\sigma \circ q_{it})\,(dh \,|\, z_{it}, a_{it}).$$

By Proposition 9, $\hat{g}$ is upper semianalytic. Equation (30) is an immediate consequence of Lemma 2. □

One can see from the construction of the function $\hat{g}$ in Corollary 3 and the filtering equations Equations (26) to (28) that $\hat{g}$ does not depend on $(\sigma_{i1}, \ldots, \sigma_{it})$. This result is intuitive in the following sense: player i can infer perfectly his previous actions before date t from his current information set z_{it}, due to perfect recall; The action of player i in the current stage is also conditioned upon; Thus he does not need to rely on his previous strategies $\left(\sigma_{i1}, \ldots, \sigma_{i(t-1)}\right)$ nor his current strategy σ_{it} to compute the conditional expectation. We combine Corollary 3 and this remark to obtain the following proposition.

Proposition 10. *Suppose σ is a Borel measurable strategy profile and $g : H \to \mathbb{R}$ is upper semianalytic and bounded. Then for every $it \in W$, there exists a bounded upper semianalytic function $\hat{g} : Z_{it} \times A_{it} \to \mathbb{R}$ such that for every $(\sigma'_{i\tau})_{\tau \leq t}$,*

$$\hat{g}(z_{it}, a_{it}) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau \leq t}}[g \,|\, \zeta_{it} = z_{it}, \mathfrak{a}_{it} = a_{it}] \tag{31}$$

for $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau \leq t}}$ almost every (z_{it}, a_{it}).

Similar to Proposition 10, we have the following lemma:

Proposition 11. *Suppose σ is a Borel measurable strategy profile and $g : H \to \mathbb{R}$ is upper semianalytic and bounded. Then for every $it \in W$, there exists a bounded upper semianalytic function $\hat{g} : Z_{it} \to \mathbb{R}$ such that for every $(\sigma'_{i\tau})_{\tau < t}$,*

$$\hat{g}(z_{it}) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau < t}}[g \,|\, \zeta_{it} = z_{it}] \tag{32}$$

for $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau < t}}$ almost every z_{it}.

Note the difference between the two propositions above, Equation (32) holds only for every $(\sigma'_{i\tau})_{\tau < t}$, as opposed to $(\sigma'_{i\tau})_{\tau \leq t}$ in Equation (31). This is because in Equation (32) the action of player i in period t not conditioned upon, thus σ_{it} needs to be specified for the computation of the conditional expectation.

If the strategy profile σ and the function g are only a universally measurable, then same construction above lead to a universally measurable function $\hat{g}$ with the same property, via

the universally measurable part of Propositions 4 and 7 and corollary 1. Formally, we have the following propositions.

Proposition 12. *Suppose σ is a universally measurable strategy profile and $g : H \to \mathbb{R}$ is universally measurable and bounded. Then for every $it \in W$, there exists a bounded universally measurable function $\hat{g} : Z_{it} \times A_{it} \to \mathbb{R}$ such that for every $(\sigma'_{i\tau})_{\tau \le t}$,*

$$\hat{g}(z_{it}, a_{it}) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau \le t}}[g \mid \zeta_{it} = z_{it}, \mathfrak{a}_{it} = a_{it}]$$

for $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau \le t}}$ almost every (z_{it}, a_{it}).

Proposition 13. *Suppose σ is a universally measurable strategy profile and $g : H \to \mathbb{R}$ is universally measurable and bounded. Then for every $it \in W$, there exists a bounded universally measurable function $\hat{g} : Z_{it} \to \mathbb{R}$ such that for every $(\sigma'_{i\tau})_{\tau < t}$,*

$$\hat{g}(z_{it}) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau < t}}[g \mid \zeta_{it} = z_{it}]$$

for $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau < t}}$ almost every z_{it}.

B.2 A Sufficient Condition and A Necessary Condition for Nash Equilibrium

Now that we finish establishing the basic results about filtering, we provide a sufficient and then a necessary condition for a strategy profile to be a Nash equilibrium.

Given a multistage game $\Gamma = (I, N, S, A, Z, p, \zeta, u)$. For every t, let ϑ_t be the projection from H to Θ_t. We define, for every time t and every preplay $\theta_t \in \Theta_t$ of the game,

$$\begin{aligned} \bar{u}_{i|t}(\theta_t) &= \sup_{\{h \in H : \vartheta_t(h) = \theta_t\}} u_i(h), \\ \underline{u}_{i|t}(\theta_t) &= \inf_{\{h \in H : \vartheta_t(h) = \theta_t\}} u_i(h). \end{aligned}$$

Since u_i is Borel measurable, thus upper and lower semianalytic, therefore by Proposition 8, $\bar{u}_{i|t}$ and $\underline{u}_{i|t}$ are upper and lower semianalytic functions on Θ_t respectively, thus universally

measurable. Both functions can be lifted to be an upper and lower semianalytic function respectively, on H.

We say that u_i is *lower convergent* if for each history $h \in H$,

$$\lim_{t\to\infty} \underline{u}_{i|t}(\vartheta_t(h)) = u_i(h),$$

and there is a uniform lower bound on u_i. The notion of upper convergence of a u_i is symmetrically defined.

For every $it \in W$ and strategy profile σ, there exists (Proposition 13) a universally measurable function $U^{\sigma}_{i|t} : Z_{it} \to \mathbb{R}$ such that for every $(\sigma'_{i\tau})_{\tau<t}$,

$$U^{\sigma}_{i|t}(z_{it}) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}\left[u_i \,|\, \zeta_{it} = z_{it}\right],$$

for $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}$ almost every z_{it}. $U^{\sigma}_{i|t}(z_{it})$ is the expected utility of player i conditional on observing z_{it}.

Fix a perturbed game $\widehat{\Gamma} = (\Gamma, \epsilon, \chi)$. Consider a strategy profile $\sigma = (\sigma_i)_{i\in I} \in \widehat{\Sigma}$. We say that the strategy σ_i of player i is *unimprovable* with respect to the strategies σ_{-i} of all other players in $\widehat{\Gamma}$ if and only if for every $t \geq 1$ and for every strategy $\sigma'_i \in \widehat{\Sigma}_i$ of player i,

$$U^{\sigma}_{i|t}\left(\zeta_{it}\right) \geq \mathrm{E}^{\sigma/\sigma'_i}\left[U^{\sigma}_{i|t+1}\left(\zeta_{i(t+1)}\right) \,\middle|\, \zeta_{it}\right] \qquad \mathrm{P}^{\sigma/\sigma'_i} \text{ almost surely.} \tag{33}$$

Remark 1. Unimprovability, also called the one-shot deviation principle, was originally formulated by Blackwell (1965) in the context of dynamic programming. In the context of extensive games, after fixing the strategy of the other players to be σ_{-i}, player i faces a potentially infinite-horizon dynamic programming problem. To say that a strategy σ_i is unimprovable is to say that it cannot be improved upon in one step by any other strategy σ'_i. If at any time t, when presuming that he will revert to the strategy σ_i beginning in next period, he does as well as possible by using σ_i in this period. for every σ'_i, $t, k \geq 0$, and every partial history $h_{t-1} \in H_{t-1}$.

By induction and the law of iterated expectations, this easily extends to k-period unim-

provability, as follows.

Lemma 4. *If σ_i is unimprovable with respect to σ_{-i} in $\widehat{\Gamma}$, then for every $t, k \geq 0$ and for every strategy $\sigma_i' \in \widehat{\Sigma}_i$ of player i,*

$$U^{\sigma}_{i|t}\left(\zeta_{it}\right) \geq \mathrm{E}^{\sigma/\sigma_i'}\left[U^{\sigma}_{i|t+k}\left(\zeta_{i(t+k)}\right) \,\middle|\, \zeta_{it}\right] \qquad \mathrm{P}^{\sigma/\sigma_i'} \text{ almost surely.} \tag{34}$$

Proof. Suppose inequality 34 holds for some $k \geq 1$, since

$$U^{\sigma}_{i|(t+k)}\left(\zeta_{i(t+k)}\right) \geq \mathrm{E}^{\sigma/\sigma_i'}\left[U^{\sigma}_{i|t+k+1}\left(\zeta_{i(t+k+1)}\right) \,\middle|\, \zeta_{i(t+k)}\right] \qquad \mathrm{P}^{\sigma/\sigma_i'} \text{ almost surely,}$$

we then have

$$\begin{aligned} U^{\sigma}_{i|t}\left(\zeta_{it}\right) \geq & \ \mathrm{E}^{\sigma/\sigma_i'}\left[\mathrm{E}^{\sigma/\sigma_i'}\left[U^{\sigma}_{i|t+k+1}\left(\zeta_{i(t+k+1)}\right) \,\middle|\, \zeta_{i(t+k)}\right] \,\middle|\, \zeta_{it}\right] \\ \geq & \ \mathrm{E}^{\sigma/\sigma_i'}\left[U^{\sigma}_{i|t+k+1}\left(\zeta_{i(t+k+1)}\right) \,\middle|\, \zeta_{it}\right] \qquad \mathrm{P}^{\sigma/\sigma_i'} \text{ almost surely.} \end{aligned}$$

This establishes k-period unimprovability by induction. □

Now we are ready state a sufficient condition for being a Nash equilibrium.

Proposition 14 (A sufficient condtion)**.** *Suppose $\Gamma = (I, N, S, A, Z, p, \zeta, u)$ is a multistage game in which every utility function u_i is lower convergent, and $\widehat{\Gamma} = (\Gamma, \epsilon, \chi)$ is a perturbed game of Γ (if $\epsilon \equiv 0$, then $\widehat{\Gamma} = \Gamma$ by convention). If σ is a strategy profile such that σ_i is unimprovable with respect to σ_{-i} in $\widehat{\Gamma}$ for all $i \in N$, then σ is a Nash equilibrium of $\widehat{\Gamma}$.*

Proof. Suppose $\sigma \in \widehat{\Sigma}$ has the stated unimprovability property. For every $t, k \geq 0$ and for every $\sigma_i' \in \widehat{\Sigma}_i$, from the definition of the function $U^{\sigma}_{i|t+k}$ we have,

$$\begin{aligned} & U^{\sigma}_{i|t+k}\left(z_{i(t+k)}\right) \\ = & \ \mathrm{E}^{\sigma/(\sigma_{i\tau}')_{\tau<t+k}}\left[u_i \,\middle|\, \zeta_{i(t+k)} = z_{i(t+k)}\right] && \text{for } \mathrm{P}^{\sigma/(\sigma_{i\tau}')_{\tau<t+k}} \text{ and } \mathrm{P}^{\sigma/\sigma_i'} \text{ almost every } z_{i(t+k)}, \\ \geq & \ \mathrm{E}^{\sigma/(\sigma_{i\tau}')_{\tau<t+k}}\left[\underline{u}_{i|t+k} \,\middle|\, \zeta_{i(t+k)} = z_{i(t+k)}\right] && \text{for } \mathrm{P}^{\sigma/(\sigma_{i\tau}')_{\tau<t+k}} \text{ and } \mathrm{P}^{\sigma/\sigma_i'} \text{ almost every } z_{i(t+k)}, \\ \geq & \ \mathrm{E}^{\sigma/\sigma_i'}\left[\underline{u}_{i|t+k} \,\middle|\, \zeta_{i(t+k)} = z_{i(t+k)}\right] && \text{for } \mathrm{P}^{\sigma/(\sigma_{i\tau}')_{\tau<t+k}} \text{ and } \mathrm{P}^{\sigma/\sigma_i'} \text{ almost every } z_{i(t+k)}. \end{aligned}$$

where the last inequality follows from Proposition 3, as the probability distributions of both

$\underline{u}_{i|t+k}$ and $\zeta_{i(t+k)}$ are the same under $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau<t+k}}$ and $\mathrm{P}^{\sigma/\sigma'_i}$. This is also why all equalities and inequalities above hold for $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau<t+k}}$ and $\mathrm{P}^{\sigma/\sigma'_i}$ almost every $z_{i(t+k)}$.

Thus inequality (34) implies that,

$$\begin{aligned} U^{\sigma}_{i|t}(\zeta_{it}) \geq & \ \mathrm{E}^{\sigma/\sigma'_i}\left[\mathrm{E}^{\sigma/\sigma'_i}\left[\underline{u}_{i|t+k}\,\middle|\,\zeta_{i(t+k)}\right]\,\middle|\,\zeta_{it}\right] \\ = & \ \mathrm{E}^{\sigma/\sigma'_i}\left[\underline{u}_{i|t+k}\,\middle|\,\zeta_{it}\right] \qquad\qquad \mathrm{P}^{\sigma/\sigma'_i} \text{ almost surely.} \end{aligned}$$

As $k \to \infty$, $\underline{u}_{i|t+k}$ converges monotonically (upwards) to u_i (since u_i is lower convergent), so an application of monotone convergence implies that

$$U^{\sigma}_{i|t}(\zeta_{it}) \geq \mathrm{E}^{\sigma/\sigma'_i}[u_i \,|\, \zeta_{it}] = U^{\sigma/\sigma'_i}_{i|t}(\zeta_{it}) \qquad \mathrm{P}^{\sigma/\sigma'_i} \text{ almost surely.}$$

In particular, $U^{\sigma}_{i|1}(z_{i1}) \geq U^{\sigma/\sigma'_i}_{i|1}(z_{i1})$ for $\mathrm{P}^{\sigma/\sigma'_i}$ almost every z_{i1}. Since the probability distribution of ζ_{i1} does not change with the strategy profile (it is always $p_1 \circ \zeta_{i1}^{-1}$), thus $U^{\sigma}_{i|1}(z_{i1}) \geq U^{\sigma/\sigma'_i}_{i|1}(z_{i1})$ for p_1 almost every z_{i1}. Thus,

$$U_i(\sigma) = \int_{Z_{i1}} U^{\sigma}_{i|1}(z_{i1})\left(p_1 \circ \zeta_{i1}^{-1}\right)(dz_{i1}) \geq \int_{Z_{i1}} U^{\sigma/\sigma'_i}_{i|1}(z_{i1})\left(p_1 \circ \zeta_{i1}^{-1}\right)(dz_{i1}) = U_i(\sigma'_i, \sigma_{-i}).$$

Therefore σ is a Nash equilibrium of $\widehat{\Gamma}$. □

Now we provide a necessary condition for a strategy profile to be a Nash equilibrium.

Definition 3. We say that a utility function u_i is *convergent uniformly in probability* if $\forall\, \epsilon > 0$,

$$\sup_{\sigma \in \Sigma} \mathrm{P}^{\sigma}\left[\bar{u}_{i|t} - \underline{u}_{i|t} > \epsilon\right] \to 0.$$

Proposition 15 (A nessary condition)**.** *Suppose $\Gamma = (I, N, S, A, Z, p, \zeta, u)$ is a multistage game in which every utility function u_i is convergent uniformly in probability. Let $\widehat{\Gamma} = (\Gamma, \epsilon, \chi)$ be a perturbed game of Γ. If σ is a Borel measurable Nash equilibrium of $\widehat{\Gamma}$, then for every $it \in W$,*

$$U^{\sigma}_{i|t}(z_{it}) \geq \sup_{\sigma'_i \in\, \widehat{\Sigma}_i} U^{\sigma/\sigma'_i}_{i|t}(z_{it}) \qquad \textit{for } \mathrm{P}^{\sigma} \textit{ almost every } z_{it}. \tag{35}$$

We need some intermediate results to establish the necessary condition above.

Lemma 5. *If σ is a Nash equilibrium of $\widehat{\Gamma}$, then for every $\sigma_i' \in \widehat{\Sigma}_i$,*

$$U_{i|1}^{\sigma}(z_{i1}) \geq U_{i|1}^{\sigma/\sigma_i'}(z_{i1}), \qquad \textit{for almost every } z_{i1}.$$

Proof. Let $\underline{Z}_{i1} = \left\{z_{i1} \in Z_{i1} : U_{i|1}^{\sigma}(z_{i1}) < U_{i|1}^{\sigma/\sigma_i'}(z_{i1})\right\}$. Intuitively, $\underline{Z}_{i1}$ is the set of places where player i can improve by adopting the strategy σ_i' in place of σ. $\underline{Z}_{i1}$ is a universally measurable subset of Z_{i1}. Define a strategy $\tilde{\sigma}_i \in \hat{\Sigma}_i$ of player i as follows: for every $t \geq 1$,

$$\begin{aligned} \tilde{\sigma}_{it}(z_{it}) &= \sigma_{it}(z_{it}) \quad \text{if } r_{it0}(z_{it}) \notin \underline{Z}_{i1}, \\ \tilde{\sigma}_{it}(z_{it}) &= \sigma_{it}'(z_{it}) \quad \text{if } r_{it0}(z_{it}) \in \underline{Z}_{i1} \end{aligned}$$

where r_{it0} is one of the recall functions of player i. Then

$$\begin{aligned} U_i(\tilde{\sigma}_i, \sigma_{-i}) = & \int_{z_{i1} \notin \underline{Z}_{i1}} U_{i|1}^{\sigma}(z_{i1}) \left(p_1 \circ \zeta_{i1}^{-1}\right)(dz_{i1}) + \int_{z_{i1} \in \underline{Z}_{i1}} U_{i|1}^{\sigma/\sigma_i'}(z_{i1}) \left(p_1 \circ \zeta_{i1}^{-1}\right)(dz_{i1}) \\ = & \int_{Z_{i1}} U_{i|1}^{\sigma}(z_{i1}) \left(p_1 \circ \zeta_{i1}^{-1}\right)(dz_{i1}) + \int_{z_{i1} \in \underline{Z}_{i1}} \left[U_{i|1}^{\sigma/\sigma_i'}(z_{i1}) - U_{i|1}^{\sigma}(z_{i1})\right] \left(p_1 \circ \zeta_{i1}^{-1}\right)(dz_{i1}) \\ = & U(\sigma) + \int_{z_{i1} \in \underline{Z}_{i1}} \left[U_{i|1}^{\sigma/\sigma_i'}(z_{i1}) - U_{i|1}^{\sigma}(z_{i1})\right] \left(p_1 \circ \zeta_{i1}^{-1}\right)(dz_{i1}). \end{aligned}$$

Since $U_i(\tilde{\sigma}_i, \sigma_{-i}) \leq U_i(\sigma)$, then the integral on the right hand side of the equation above must be non positive. But the integrand $\left[U_{i|1}^{\sigma/\sigma_i'}(z_{i1}) - U_{i|1}^{\sigma}(z_{i1})\right]$ is strictly positive when $z_{i1} \in \underline{Z}_{i1}$, so it must be that $\left(p_1 \circ \zeta_{i1}^{-1}\right)(\underline{Z}_{i1}) = 0$. This completes the proof. □

For every $it \in W$, let

$$B_{it} = \{(z_{it}, p) \in Z_{it} \times \Delta(A_{it}) : \ \ p(B) \geq \epsilon_{it}(z_{it})\chi_{it}(B \,|\, z_{it}),\ \forall B \subset A_{it}\}.$$

Lemma 6. *For every $it \in W$, B_{it} is a Borel subset of $Z_{it} \times \Delta(A_{it})$.*

Proof. Define $\varphi : Z_{it} \times \Delta(A_{it}) \to \Delta(A_{it})$ by

$$\varphi(z_{it}, p) = (1 - \epsilon_{it}(z_{it}))p + \epsilon_{it}(z_{it})\chi_{it}(z_{it}).$$

The function φ is Borel measurable, as $\epsilon_{it}(\cdot)$ and $\chi_{it}(\cdot)$ are. Let $\widehat{\varphi} : Z_{it} \times \Delta(A_{it}) \to Z_{it} \times \Delta(A_{it})$ be defined as $\widehat{\varphi}(z_{it}, p) = (z_{it}, \varphi(z_{it}, p))$. Then $\widehat{\varphi}$ is Borel measurable and bijective, and

$\widehat{\varphi}(Z_{it} \times \Delta(A_{it})) = B_{it}$. By Kuratowski's theorem (Proposition 2), B_{it} is a Borel subset of $Z_{it} \times \Delta(A_{it})$. □

Suppose σ is a Borel measurable Nash equilibrium of $\widehat{\Gamma}$. Without lose of generality, we will show Equation (35) for $t = 0$. Let $g : H \to \mathbb{R}$ be universally measurable and bounded. For every $it \in W$ and $\theta_t \in \Theta_t$, there exists, by Proposition 12, a bounded universally measurable function $\hat{g} : Z_{it} \times A_{it} \to \mathbb{R}$ such that for every $(\sigma'_{i\tau})_{\tau \leq t}$,

$$\hat{g}(z_{it}, a_{it}) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau \leq t}}[g \,|\, \zeta_{it} = z_{it}, \mathfrak{a}_{it} = a_{it}], \qquad \text{for } \mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau \leq t}} \text{ almost every } (z_{it}, a_{it}). \quad (36)$$

Define $f : B_{it} \to \mathbb{R}$ by

$$f(z_{it}, p) = \int_{A_{it}} \hat{g}(z_{it}, a_{it}) p(da_{it}). \quad (37)$$

The function f is universally measurable and bounded, because $\hat{g}$ is (Corollary 1).

For every $\sigma'_{it} \in \widehat{\Sigma}_{it}$, the operator $J^{\sigma'_{it}}$ mapping g into $J^{\sigma'_{it}}(g) : \Theta_t \to \mathbb{R}$ is defined by

$$J^{\sigma'_{it}}(g)\,(\theta_t) = f\left(\zeta_{it}(\theta_t), \sigma'_{it}\left[da_{it} \,|\, \zeta_{it}(\theta_t)\right]\right).$$

for every $\theta_t \in \Theta_t$. The function $J^{\sigma'_{it}}(g)$ is bounded and universally measurable by Proposition 4. Thus the operator $J^{\sigma'_{it}}$ maps a bounded universally measurable function $g : H \to \mathbb{R}$ to a bounded universally measurable function $J^{\sigma'_{it}}(g) : H \to \mathbb{R}$. From the definition of $J^{\sigma'_{it}}(g)$ and the law of iterated expectation, it follows that for every $(\sigma'_{i\tau})_{\tau < t}$,

$$J^{\sigma'_{it}}(g)\,(\vartheta_t(h)) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau \leq t}}\left[g \,|\, \zeta_{it}\right](h), \qquad \mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau \leq t}} \text{ almost surely.} \quad (38)$$

Lemma 7. *For every $t \geq 1$, $(\sigma'_{i\tau})_{\tau \leq t} \in \prod_{\tau \leq t} \widehat{\Sigma}_{i\tau}$ and bounded universally measurable $g : H \to \mathbb{R}$, we have*

$$\left(J^{\sigma'_{i1}} \ldots J^{\sigma'_{it}}\right)(g) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau \leq t}}[g \,|\, \zeta_{i1}], \qquad \textit{almost surely.}$$

where $J^{\sigma'_{i1}} \ldots J^{\sigma'_{it}}$ denotes the composition of $J^{\sigma_{i1}}, \ldots, J^{\sigma_{it}}$.

Proof. We proceed by induction on t. The case where $t = 0$ reduces to the definition of $J^{\sigma'_{i1}}$.

Suppose Lemma 7 holds for some $t-1 \geq 0$. Since $J^{\sigma'_{it}}(g)$ is defined on Θ_t, thus it has the same probability distribution under $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}$ and $\mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau\leq t}}$. Thus we have

$$\begin{aligned}
& \left(J^{\sigma_{i1}} \dots J^{\sigma_{it}}\right)(g) & \\
= \; & \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}\left[J^{\sigma'_{it}}(g) \,|\, \zeta_{i1}\right] & \text{(induction hypothesis at } t-1) \\
= \; & \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau\leq t}}\left[J^{\sigma'_{it}}(g) \,|\, \zeta_{i1}\right] & \text{(Proposition 3)} \\
= \; & \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau\leq t}}\left[\mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau\leq t}}\left[g \,|\, \zeta_{it}\right] \,|\, \zeta_{i1}\right] & \text{(Equation (38))} \\
= \; & \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau\leq t}}[g \,|\, \zeta_{i1}] & \mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau\leq t}} \text{ almost surely.}
\end{aligned}$$

Since both sides of the equation are functions defined on Θ_t, their probability distributions are unaffected by the underlying probability measure. Therefore the equations above hold almost surely. This establishes Lemma 7 by induction. □

Given a $z_{it} \in Z_{it}$, let $B_{z_{it}}$ be the section of B_{it} at z_{it}. That is,

$$B_{z_{it}} = \{p \in \Delta(A_{it}) : (z_{it}, p) \in B_{it}\}.$$

The operator J_{it} mapping g into $J_{it}(g) : \Theta_t \to \mathbb{R}$ is defined by

$$J_{it}(g)(\theta_t) = \sup_{p \in B_{\zeta_{it}(\theta_t)}} f(\zeta_{it}(\theta_t), p).$$

If the function g is upper semianalytic, then the function $\hat{g}$ in Equation (36) can be taken to be upper semianalytic (Proposition 10), and thus the function f in Equation (37) is also upper semianalytic (Corollary 2). Therefore the function $J_{it}(g)$ is upper semianalytic (Propositions 4 and 8). Thus the operator J maps a bounded upper semianalytic function $g : H \to \mathbb{R}$ to a bounded upper semianalytic function $J_{it}(g) : H \to \mathbb{R}$.

It is clear that

$$\left(J^{\sigma'_{i1}} \dots J^{\sigma'_{it}}\right)(g) \leq \left(J_{i1} \dots J_{it}\right)(g)$$

for every $t \geq 1$, $(\sigma'_{i1}, \dots, \sigma'_{it}) \in \prod_{\tau\leq t} \widehat{\Sigma}_{i\tau}$ and bounded upper semianalytic function $g : H \to \mathbb{R}$. But the inequality can be arbitrarily tight, as the following result shows.

Lemma 8. *For every $t \geq 1$, bounded upper semianalytic function $g : H \to \mathbb{R}$ and $\epsilon > 0$,*

there exists $\sigma'_{it} \in \widehat{\Sigma}_{it}$ *such that for every* $h \in H$,

$$J^{\sigma'_{it}}(g)(h) \geq J_{it}(g)(h) - \epsilon$$

Proof. An application of Theorem 2 to the upper semianalytic function f defined in Equation (37) implies the existence of a universally measurable probability function: $\sigma'_{it} : Z_{it} \to \Delta(A_{it})$ such that for every $z_{it} \in Z_{it}$, $\sigma'_{it} \in B_{z_{it}}$ and for every $z_{it} \in Z_{it}$,

$$f(z_{it}, \sigma'_{it}(z_{it})) \geq \sup_{p \in B_{z_{it}}} f(z_{it}, p) - \epsilon.$$

Thus $\sigma'_{it} \in \widehat{\Sigma}_{it}$ and for every $h \in H$,

$$J^{\sigma'_{it}}(g)(h) \geq J_{it}(g)(h) - \epsilon$$

□

Lemma 9. *For every* $t \geq 1$, *bounded upper semianalytic function* $g : H \to \mathbb{R}$ *and* $\epsilon > 0$, *there exists* $(\sigma'_{i1}, \ldots, \sigma'_{it}) \in \prod_{\tau \leq t} \widehat{\Sigma}_{i\tau}$ *such that for every* $h \in H$,

$$\left(J^{\sigma'_{i1}} \ldots J^{\sigma'_{it}}\right)(g)(h) \geq (J_{i1} \ldots J_{it})(g)(h) - \epsilon$$

Proof. We proceed by induction on t. The case where $t = 0$ reduces to Lemma 8. Suppose it is true for some $t - 1 \geq 0$. Then there exists $\left(\sigma'_{i1}, \ldots, \sigma'_{i(t-1)}\right) \in \prod_{\tau < t} \widehat{\Sigma}_{i\tau}$ such that

$$\left(J^{\sigma'_{i1}} \ldots J^{\sigma'_{i(t-1)}}\right)(J_{it}(g)) \geq (J_{i1} \ldots J_{it})(J_{it}(g)) - \epsilon.$$

By Lemma 8, there exists $\sigma'_{it} \in \widehat{\Sigma}_{it}$ such that

$$J^{\sigma'_{it}}(g) \geq J_{it}(g) - \epsilon.$$

Combining the two inequalities above, we obtain

$$\left(J^{\sigma'_{i1}} \ldots J^{\sigma'_{it}}\right)(g) \geq (J_{i1} \ldots J_{it})(g) - 2\epsilon.$$

This establishes Lemma 9 by induction. □

Let $J^* : Z_{i1} \to \mathbb{R}$ be defined by

$$J^*(z_{i1}) = \sup_{\sigma_i' \in \widehat{\Sigma}_i} U_{i|1}^{\sigma/\sigma_i'}(z_{i1}).$$

Lemma 10. *For every $t \geq 1$, $J^*(\zeta_{i1}) \leq (J_{i1} \dots J_{i(t-1)})\left(\bar{u}_{i|t}\right)$ almost surely.*

Proof. For every $\sigma_i' \in \widehat{\Sigma}_i$,

$$U_{i|1}^{\sigma/\sigma_i'}(\zeta_{i1}) = \mathrm{E}^{\sigma/\sigma_i'}\left[U_{i|t}^{\sigma/\sigma_i'}(\zeta_{it}) \,|\, \zeta_{i1}\right] \leq \mathrm{E}^{\sigma/\sigma_i'}\left[\bar{u}_{i|t}(\zeta_{it}) \,|\, \zeta_{i1}\right], \qquad \text{almost surely.}$$

As the probability distribution of $\bar{u}_{i|t}(\zeta_{it})$ and ζ_{i1} does not depend on the choice of $(\sigma_{it}', \dots,)$, thus

$$\mathrm{E}^{\sigma/\sigma_i'}\left[\bar{u}_{i|t}(\zeta_{it}) \,|\, \zeta_{i1}\right] = \mathrm{E}^{\sigma/(\sigma_{i\tau}')_{\tau<t}}\left[\bar{u}_{i|t}(\zeta_{it}) \,|\, \zeta_{i1}\right] = \left(J^{\sigma_{i1}'} \dots J^{\sigma_{i(t-1)}'}\right)(\bar{u}_{i|t}) \leq \left(J_{i1} \dots J_{i(t-1)}\right)(\bar{u}_{i|t}).$$

almost surely. Therefore $J^*(\zeta_{i1}) \leq \left(J_{i1} \dots J_{i(t-1)}\right)(\bar{u}_{i|t})$ almost surely. □

Now we are ready to show Proposition 15.

Proof. Let σ be a Borel measurable Nash equilibrium of $\hat{\Gamma}$ as in the statement of Proposition 15. For every $i \in I$, let $M = \sup u_i - \inf u_i$. Since u_i is convergent uniformly in probability, thus for every $\epsilon > 0$, there exists $t \geq 1$ such that

$$\mathrm{P}^{\sigma'}\left(\bar{u}_{i|t} - \underline{u}_{i|t} > \epsilon\right) < \epsilon \tag{39}$$

for every $\sigma' \in \widehat{\Sigma}$. Fix such a t, then there exists $(\sigma_{i1}', \dots, \sigma_{i(t-1)}') \in \prod_{\tau<t} \widehat{\Sigma}_{i\tau}$ such that

$$\left(J^{\sigma_{i1}'} \dots J^{\sigma_{i(t-1)}'}\right)(\bar{u}_{i|t}) \geq \left(J_{i1} \dots J_{i(t-1)}\right)(\bar{u}_{i|t}) - \epsilon \geq J^*(\zeta_{i1}), \qquad \text{almost surely}$$

by Lemmas 9 and 10. On the other hand,

$$\begin{aligned}
& \left(J^{\sigma'_{i1}} \dots J^{\sigma'_{i(t-1)}}\right)(\bar{u}_{i|t}) & \\
= \;& \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}[\bar{u}_{i|t} \,|\, \zeta_{i1}] & \text{(Lemma 7)} \\
\leq \;& \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}[\underline{u}_{i|t} \,|\, \zeta_{i1}] + \epsilon + M\epsilon & \text{(Inequality 39)} \\
\leq \;& \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}\left[U_{i|t}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}(\zeta_{it}) \,|\, \zeta_{i1}\right] + (M+1)\epsilon & \\
= \;& U_{i|1}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}(\zeta_{i1}) + (M+1)\epsilon & \text{almost surely.}
\end{aligned}$$

Therefore

$$U_{i|1}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}(\zeta_{i1}) + (M+1)\epsilon \geq J^*(\zeta_{i1}) \qquad \text{almost surely.}$$

Let

$$A_\epsilon = \left\{U_{i|1}^{\sigma}(\zeta_{i1}) + (M+1)\epsilon < J^*(\zeta_{i1})\right\}.$$

Since $U_{i|1}^{\sigma}(\zeta_{i1}) \geq U_{i|1}^{\sigma/(\sigma'_{i\tau})_{\tau<t}}(\zeta_{i1})$ almost surely (Lemma 5), we have $\mathrm{P}^\sigma(A_\epsilon) = 0$. As ϵ is arbitrary, thus $\mathrm{P}^\sigma(A) = 0$, where

$$A = \bigcup_{k\to\infty} A_{1/k} = \left\{U_{i|1}^{\sigma}(\zeta_{i1}) < J^*(\zeta_{i1})\right\}.$$

Therefore $U_{i|1}^{\sigma}(\zeta_{i1}) \geq J^*(\zeta_{i1}) = \left(\sup_{\sigma'_i \in \widehat{\Sigma}_i} U_{i|1}^{\sigma/\sigma'_i}\right)(\zeta_{i1})$ almost surely, completing the proof. □

Remark 2. One difference between the sufficient and the necessary conditions, is that between the qualifications for "almost surely". In the sufficient condition, inequality 33 holds $\mathrm{P}^{\sigma/(\sigma_{i\tau})_{\tau\leq t}}$ almost surely, where $(\sigma'_{i\tau})_{\tau\leq t}$ vary across $\prod_{\tau\leq t}\Sigma_{i\tau}$. This effectively requires that the inequality holds "everywhere", whether or not it is ruled out by the underlying strategy profile σ. Formally, unimprovability is equivalent to the following: for every $it \in W$, for every $\sigma'_{it} \in \Sigma_{it}$, there exists a universally measurable function $g : Z_{it} \to \mathbb{R}$ such that for every $(\sigma'_{i\tau})_{\tau<t}$,

$$\begin{aligned}
& g(z_{it}) = \mathrm{E}^{\sigma/(\sigma'_{i\tau})_{\tau\leq t}}\left[U_{i|t+1}^{\sigma}\left(\zeta_{i(t+1)}\right) \,\middle|\, \zeta_{it} = z_{it}\right] && \text{for } \mathrm{P}^{\sigma/(\sigma'_{i\tau})_{\tau\leq t}} \text{ almost every } z_{it}, \\
& U_{i|t}^{\sigma}(z_{it}) \geq g(z_{it}) && \text{for every } z_{it}.
\end{aligned}$$

In this sense, it is inspirit similar to subgame perfection. It is stronger than the necessary condition, which only requires inequality Equation (35) to hold P^σ almost surely. That is, every player's continuation strategy has to be optimal only at points that the strategy profile σ allows. Optimality is not required at points that σ rules out. For this reason, the necessary condition is not sufficient for a strategy profile to be a Nash equilibrium, because a player can have intertemporal coordination failure.

To differentiate the two concepts of optimality, we will refer to the type in the sufficient condition *"everywhere optimality"*.

Suppose a multistage game $\Gamma = (I, N, S, A, Z, p, \zeta, u)$ is such that players can perfectly observe the complete preplay of the game before he acts in any period, that is, $Z_{it} = \Theta_{it}$ and $\zeta_{it} : \Theta_{it} \to Z_{it}$ is the identity mapping for every $it \in W$, then every preplay $\theta_t \in \Theta_t$ induces a subgame $\Gamma_{|\theta_t}$ of Γ. We apply the subscript "$|\theta_t$" to denote the components of the subgame $\Gamma_{|\theta_t}$. For example, $I_{|\theta_t}$ is the set of players in Γ_{θ_t}, and $u_{i|\theta_t}$ is the utility function of player i in Γ_{θ_t}. Let $\widehat{\Gamma} = (\Gamma, \epsilon, \chi)$ be a perturbed game of Γ, a subgame $\widehat{\Gamma}_{|\theta_t} = (\Gamma_{|\theta_t}, \epsilon_{|\theta_t}, \chi_{|\theta_t})$ of $\widehat{\Gamma}$ consists of a subgame Γ_{θ_t} of Γ, and the restrictions $\epsilon_{|\theta_t}$ of ϵ and $\chi_{|\theta_t}$ of χ to the information sets in $\Gamma_{|\theta_t}$. Let $\widehat{\Sigma}_{i|\theta_t}$ be the set of strategies of player i in $\widehat{\Gamma}_{|\theta_t}$. If $\sigma = (\sigma_i)_{i \in N}$ be a strategy profile of Γ. We denote the strategy profile induced by σ on a subgame $\Gamma_{|\theta_t}$ by $\sigma_{|\theta_t}$. If σ is a strategy profile of $\widehat{\Gamma}$, then the induced strategy profile $\sigma_{|\theta_t}$ is a strategy profile of the subgame $\widehat{\Gamma}_{\theta_t}$. Whenever we consider several multistage games with possibly different utility functions simultaneously, we add the game itself as an argument to any utility related functions to avoid confusion. Thus given any multistage game Γ, and a strategy profile σ of Γ, $U_i(\sigma; \Gamma)$ denotes the expected utility of player i under the strategy profile σ. The same applies to functions like $U_{i|t}$.

A multistage game with the property above is called a game with *quasi-perfect information*. The setting where a multistage game has perfect information is a special case. An immediate corollary of the necessary condition Proposition 15 follows:

Corollary 4. *Suppose $\Gamma = (I, N, S, A, Z, p, \zeta, u)$ is a multistage game with quasi-perfect information in which every utility function u_i is convergent uniformly in probability. Let*

$\widehat{\Gamma} = (\Gamma, \epsilon, \chi)$ *be a perturbed game of* Γ. *If* σ *is a Borel measurable Nash equilibrium of* $\widehat{\Gamma}$, *then for every* $t \geq 0$ *and* P^σ *almost every* $\theta_t \in \Theta_t$, $\sigma_{|\theta_t}$ *is a Borel measurable Nash equilibrium of* $\widehat{\Gamma}_{|\theta_t}$.

Proof. For every $it \in W$, one can take the function $U^\sigma_{i|t} : \Theta_t \to \mathbb{R}$ to be

$$U^\sigma_{i|t}(\theta_t; \Gamma) = U_i\left(\sigma_{|\theta_t}; \Gamma_{|\theta_t}\right).$$

Then by Proposition 15, there exists a universally measurable subset $B_i \subset \Theta_t$ for every $i \in I$ such that $\mathrm{P}^\sigma[\vartheta_t \in B_i] = 1$ and for each $\theta_t \in B_i$,

$$U_i\left(\sigma_{|\theta_t}; \Gamma_{|\theta_t}\right) \geq \sup_{\sigma'_{i|\theta_t} \in \widehat{\Sigma}_{i|\theta_t}} U_i\left(\sigma'_{i|\theta_t}, \sigma_{-i|\theta_t}; \Gamma_{|\theta_t}\right).$$

Let $B = \cap_{i \in I} B_i$. Then we have $\mathrm{P}^\sigma(\vartheta_t \in B) = 1$ and for each $\theta_t \in B$, the inequality above holds. That is, $\sigma_{|\theta_t}$ is a Borel measurable Nash equilibrium of $\widehat{\Gamma}_{|\theta_t}$ for P^σ almost every $\theta_t \in \Theta_t$. □

B.3 Generic Rubinstein alternating-offers game

We begin with a generic two-player bargaining game, defined as follows. Two players, 1 and 2, bargain over an outcome in some compact metric set K, the set of possible agreements. Failure to agree, a "breakdown," results an outcome denoted by D. As an example, we may have $K = \left\{(x_1, x_2) \in \mathbb{R}^2_+ : x_1 + x_2 = 1\right\}$ and $D = (0, 0)$. Endowing $K \cup D$ with its Borel σ-algebra, each player i has a measurable Borel utility function $\hat{u}_i : K \cup \{D\} \to \mathbb{R}$.

Rubinstein (1982) suggested the following dynamic bargaining game. In each period, one of the players proposes an outcome x in K. The other player accepts the offer, an action denoted Y, or rejects (denoted N). Acceptance ends the bargaining and the agreement is implemented. Rejection leads, with some given probability η, to a breakdown of the negotiation, denoted B, in which case the outcome of the game is D. Otherwise, the game advances to the next period, a result called A, when the offer is made by the other player. We suppose that Player 1 is the first to propose.

The resulting multistage game is denoted Γ. We use the same notation as in the description of an abstract multistage game, unless otherwise specified.

B.3.1 Histories, Lower Convergence of Utility Functions and Behavioral Strategies

A complete history $h = (x_t, r_t, s_t)_{t\geq 1}$ of everything that happens over the course of the game consists of a sequence of offers $x = (x_t)_{t\geq 1} \in K^{\mathbb{Z}_+}$, a sequence of responses $r = (r_t)_{t\geq 1} \in \{\text{Y,N}\}^{\mathbb{Z}_+}$, and a sequence of choices by nature $s = (s_t)_{t\geq 1} \in \{\text{A,B}\}^{\mathbb{Z}_+}$. The complete set of histories is $H = (K \times \{\text{Y,N}\} \times \{\text{A,B}\})^{\mathbb{Z}_+}$. Each utility function u_i induces a Borel measurable utility $u_i : H \to \mathbb{R}$. Specifically, given a history $h = (x, r, s)$, let $\tau = \inf\{t \geq 1, r_t = Y \text{ or } s_t = B\}$. That is, τ is the period in which the bargaining process ends. If $\tau < \infty$, then each player's utility is well defined on such a history with $u_i(h) = \hat{u}_i(x_\tau)\mathbb{1}\{r_\tau = Y\} + u_i(D)\mathbb{1}\{r_\tau = N\}$. If $\tau = \infty$, we let $u_i(h) = \hat{u}_i(D)$ by convention.

Proposition 16. *Each utility function u_i is lower convergent and convergent uniformly in probability.*

Proof. Because $0 \leq u_i \leq 1$, we may take 0 and 1 as a uniform lower and upper bounds on $\hat{u}_i$. Given a history $h = (x, r, s)$, let $\tau = \inf\{t \geq 1, r_t = \text{Y or } s_t = \text{B}\}$ be the period in which the bargaining ends. If τ if finite, then for all $t > \tau$, we have $\underline{u}_i(h_t) = u_i(h)$. If τ is infinite, then by definition $\lim_{t\to\infty} \underline{u}_{i|t}(h_t) = u_i(h) = \hat{u}_i(D)$ for every $t \geq 1$. Therefore $\lim_{t\to\infty} \underline{u}_{i|t}(h_t) = u_i(h)$, that is, the utility function u_i is lower convergent.

On the other hand, $\bar{u}_{i|t}$ and $\underline{u}_{i|t}$ are different only on the event $\{\tau > t\}$, an event with probability less than $(1-\eta)^t$ under any strategy profile. Therefore the utility function u_i is convergent uniformly in probability. □

To simplify notation, we will simply denote a partial history $(x_{\leq t}, r_{\leq t}, s_{\leq t})$ by its offer history $x^t = (x_1, \ldots, x_t)$. We can do so because $r_\tau = Y$ or $s_\tau = B$ for some $\tau \leq t$, the game would have ended by period t and no more description is needed for what happens thereafter. Therefore an offer history $x^t = (x_1, \ldots, x_t)$ denotes the partial history $(x^t, r_{\leq t}, s_{\leq t})$ where $r_\tau = N$ and $s_\tau = A$ for all $\tau \leq t$.

We denote by μ_t an offer strategy in period t, and by ρ_t a response strategy. Thus μ_t is a universally measurable probability kernel from K^{t-1} to K, mapping a history of prior rejected offers $x^{t-1} = (x_1, \ldots, x_{t-1}) \in K^{t-1}$ to a probability measure on K, and ρ_t is a universally measurable probability kernel from K^t to $\{Y, N\}$, mapping a history of prior and current offers x^t to a probability measure on {Y,N}. A behavioral strategy σ_1 of Player 1 is a sequence $(\mu_1, \rho_2, \mu_3, \rho_4, \ldots)$ of such mappings. A strategy σ_2 of Player 2 is likewise a sequence of such functions of the form $(\rho_1, \mu_2, \rho_3, \mu_4, \ldots)$.

B.3.2 Perturbed Game, Independence of Strategically Irrelevant Information

The Rubinstein's game Γ is a multistage game with perfect information. All subgames that starts with an offer of Player 1 (Player 2) are strategically isomorphic, and all subgames that starts with an response of Player 1 (Player 2) and have the same current offer are strategically isomorphic. Then we have the following proposition:

Proposition 17. *If a perturbed game $\widehat{\Gamma} = (\Gamma, \epsilon, \chi)$ of Γ respects independence of strategically irrelevant information, then*

$$\begin{array}{lll} \epsilon_{it}\left(x^{t-1}\right) = \epsilon_{i\tau}\left(x^{\tau-1}\right), & \chi_{it}\left(x^{t-1}\right) = \chi_{i\tau}\left(x^{\tau-1}\right), & \text{if } t, \tau \text{ and } i \text{ are of the same parity.} \\ \epsilon_{it}\left(x^{t}\right) = \epsilon_{i\tau}(x^{\tau}), & \chi_{it}\left(x^{t}\right) = \chi_{i\tau}(x^{\tau}), & \text{if } t, \tau \text{ are of the same parity, different from } i. \end{array}$$

Recall the notation $\widehat{\Gamma}_{|x^{t-1}}$, which denotes the subgame of $\widehat{\Gamma}$ which follows the history of prior rejected offers x^{t-1} and starts with a player making an offer. The following corollary is an immediate consequence:

Corollary 5. *Suppose $\widehat{\Gamma}$ is a perturbed game of Γ. If Γ respects independence of irrelevant information, then for all $t \geq 1$, $x^{t-1} \in K^{k-1}$ and $x_1' \in K$,*

$$\begin{array}{ll} \widehat{\Gamma}_{|x^{t-1}} = \widehat{\Gamma} & \text{if } t \text{ is odd.} \\ \widehat{\Gamma}_{|x^{t-1}} = \widehat{\Gamma}_{|x_1'} & \text{if } t \text{ is even.} \end{array}$$

We now consider the original Rubinstein's game in which two players bargain to split a "pie" of size 1. The goal of the next subsection is to show that this game has a unique outcome that is induced by a restricted equilibrium, and that this outcome is the same as the

subgame perfect equilibrium outcome. We will also provide some extension of these results in preparation of our proof of Theorem 1.

B.4 The Original Rubinstein's alternating-offers game

The original Rubinstein's alternating-offers game corresponds to the case where

$$K = \left\{(x_1, x_2) \in \mathbb{R}^2_+ : x_1 + x_2 = 1\right\}$$

and $D = (0, 0)$. We assume quasi-linear utilities for both players, that is $u_i(x) = x_i$, and $u_i(D) = 0$ for $i = 1, 2$. This form of Rubinstein's alternating-offers game can be viewed as a special case of our three-node bargaining game.[13] Let Γ_R denote the original Rubinstein's alternating-offers game.

The game Γ_R exhibits certain symmetry between Player 1 and 2, in that every subgames that starts with an offer of Player 1 is strategically isomorphic to every subgames that starts with an offer of Player 2, and every subgames that starts with an response of Player 1 is strategically isomorphic to every subgames that starts with an response of Player 2 with the same current offer. For each t, let $i(t)$ be the player designated to offer in period t, and denote by $j(t)$ the responding player. We have the following characterization of a perturbed game of Γ_R that respects independence of strategically irrelevant information.

Proposition 18. *A perturbed game $\widehat{\Gamma} = (\Gamma_R, \epsilon, \chi)$ of Γ_R respects independence of strategically irrelevant information if and only if*

$$\begin{array}{llll} \epsilon_{it}\left(x^{t-1}\right) = \epsilon_{11}, & \chi_{it}\left(x^{t-1}\right) = \chi_{11}, & \text{if } i = i(t). \\ \epsilon_{it}\left(x^{t}\right) = \epsilon_{21}(x_1), & \chi_{it}\left(x^{t}\right) = \chi_{21}(x_1), & \text{if } i = j(t). \end{array}$$

The following corollary is an immediate consequence:

Corollary 6. *Suppose $\widehat{\Gamma}$ is a perturbed game of Γ_R. If Γ respects independence of irrelevant*

[13] Indeed, if we let the treatment sets $S_2 = \left\{s_2^0\right\}$ and $S_3 = \left\{s_3^0\right\}$ to be singleton sets (so that we freeze the treatments of Node 2 and 3), and let $C_1\left(s_2^0\right) = \left\{s_1^0, s_1\right\}$ be a binary set, then the three node bargaining problem is reduced to the original Rubinstein's alternating offer game between Nodes 1 and 2.

information, then for all $t \geq 1$ *and* $x^{t-1} \in K^{k-1}$,

$$\widehat{\Gamma}_{|x^{t-1}} = \widehat{\Gamma}.$$

To simplify notation, we reparametrize the offer set K by the unit interval $[0,1]$ in the following obvious way: in each period, one player offers a share $y \in [0,1]$ to his opponent, which is equivalent to offering $(1-y, y) \in K$ in the old parametrization.

Let

$$v_1 = \frac{1}{2-\eta}, \quad v_2 = \frac{1-\eta}{2-\eta}.$$

It is known that there is a unique subgame perfect equilibrium to this game. In equilibrium each player proposes v_2 whenever it is his turn to make an offer, and accepts an offer y of the other player if and only if $y \geq v_2$. The equilibrium outcome is that Player 1 proposes v_2 in period 1, and Player 2 immediately accepts this offer. Thus the equilibrium payoff is v_1 for Player 1, and v_2 for Player 2.

We write this equilibrium in the form of a behavioral strategy profile. Let $\sigma^* = (\mu^*, \rho^*)$ be the behavioral strategy profile defined by the following: for every $t \geq 1$ and for every offer history $y^t \in [0,1]^t$ up to period t,

- $\mu_t^*(y^{t-1}) = \delta_{v_2}$ (δ_{v_2} is the Borel probability measure on $[0,1]$ corresponding to the point mass at v_2.);
- $\rho_t^*(\{Y\} \,|\, y^t) = \begin{cases} 1 & \text{if } y_t \geq v_2 \quad \text{(recall that } y_t \text{ is the current offer in period } t\text{)} \\ 0 & \text{if } y_t < v_2 \end{cases}$

Given a behavioral strategy profile σ, let $p_a(\sigma)$ denote the probability that the two players conclude the game by reaching an agreement when playing σ, and $U_i(\sigma|\text{agreement}) = E^\sigma[\hat{u}_i(h)|\cup_{t\geq 1} A_t]$ be the expected payoff of Player i under σ, conditional on the two players

eventually reaching an agreement.[14] Then we have

$$U_i(\sigma) = p_a(\sigma)\, U_i(\sigma|\text{agreement}),$$

and on a separate account,

$$U_1(\sigma|\text{agreement}) + U_2(\sigma|\text{agreement}) = 1.$$

We can now state the following proposition:

Theorem 3. *Rubinstein's alternating offer game has a restricted equilibrium. In particular, the pure strategy profile σ^* is a restricted equilibrium. Every restricted equilibrium σ induces the same outcome as σ^*. Moreover,*

(a) There exists a restricted trembling sequence $(\sigma^n)_{n\geq 0}$ for σ^* such that the followings are true:

(i) for every n, $U_i(\sigma^n|agreement) = v_i$ $(i = 1, 2)$.

(ii) As $n \to \infty$, the probability p_a^n of an eventual agreement converges to 1.

(iii) As a consequence, the unconditional payoff $U_i(\sigma^n)$ of Player i under σ^n converges to v_i.

(b) For each restricted trembling sequence $(\sigma^n)_{n\geq 0}$ for some restricted equilibrium, the followings are true:

(i') As $n \to \infty$, the unconditional payoff $U_i(\sigma^n)$ of Player i under σ^n converges to v_i;

(ii') The probabilities p_a^n converge to 1;

(iii') As a consequence, $U_i(\sigma^n|agreement) \to v_i$.

Remark 3. (a) and (b) will be useful for the proof of Theorem 1.

[14] Given a behavioral strategy profile, the event that the two players reach an agreement is $\cup_{t\geq 1} A_t$ where A_t is the event that the players reach an agreement in period t and can be written as in **??**. Likewise the event that the bargaining process eventually breaks down is $\cup_{t\geq 1} B_t(\sigma)$. Both are measurable events. Moreover, if σ is a totally mixed behavioral strategy profile, then $\mathrm{P}^\sigma[\cup_{t\geq 1} A_t] > 0$, $\mathrm{P}^\sigma[\cup_{t\geq 1} B_t] > 0$. Hence conditional probabilities given either agreement or breakdown events are uniquely defined.

Proof. Existence and (a): We will show that σ^* is a restricted equilibrium of Rubinstein's game. For this, we fix an arbitrary positive scalar $\lambda \leq (1-\eta)/2$, a positive sequence $\epsilon^n \to 0$, and a strictly positive measure ξ on $[0,(1-\eta-2\lambda)v_1]$. For each n, let $\chi^n \in \Delta([0,1])$ be defined by

$$\chi^n = (\epsilon^n)^2\xi + \left(1-\alpha^n-(\epsilon^n)^2\right)\mathcal{U}(0,\gamma^n) + \alpha^n\,\mathcal{U}(\gamma^n,1),$$

where $\mathcal{U}(a,b)$ denotes the uniform distribution on $[a,b]$, α^n and γ^n are two positive numbers in $(0,1)$ that we will determine. By definition, χ^n is a strictly positive distribution on the unit interval $[0,1]$, and will serve as the reference measure for the n^{th} perturbed game along the sequence.

For each n and each t, let μ_t^n and ρ_t^n be defined by

$$\mu_t^n\left(y^{t-1}\right) = (1-\epsilon^n)\,\delta_{\gamma^n} + \epsilon^n\chi^n,$$

for every history of offers $y^t \in [0,1]^t$, and

$$\rho_t^n\left(\{Y\}\,|\,y^t\right) = \begin{cases} (\epsilon^n)^2 & \text{if } y_t \in [0,\gamma^n-2\lambda v_1), \\ \epsilon^n & \text{if } y_t \in [\gamma^n-2\lambda v_1,\gamma^n), \\ 1-\epsilon^n & \text{if } y_t \in [\gamma^n,\gamma^n+2\lambda v_1), \\ 1-(\epsilon^n)^2 & \text{if } y_t \in [\gamma^n+2\lambda v_1,1]. \end{cases}$$

Recall that $i(t)$ is the player designated to offer in period t, and $j(t)$ is the the responding player. The combination $\sigma^n = (\mu^n,\rho^n)$ defines a behavioral strategy profile for the perturbed game $\widehat{\Gamma}^n$ whose minimum probabilities ϵ_{it}^n and reference measures χ_{it}^n are as follows: for every offer history $y^t \in [0,1]^t$,

$$\epsilon_{it}^n\left(y^t\right) = \begin{cases} \epsilon^n, \\ 2\epsilon^n, \\ 2(\epsilon^n)^2, \end{cases} \qquad \chi_{it}^n\left(y^t\right) = \begin{cases} \chi^n & \text{if } i = i(t), \\ \frac{1}{2}\delta_{\{Y\}}+\frac{1}{2}\delta_{\{N\}} & \text{if } i = j(t) \text{ and } y_t \in [\gamma^n-2\lambda v_1,\gamma^n+2\lambda v_1), \\ \frac{1}{2}\delta_{\{Y\}}+\frac{1}{2}\delta_{\{N\}} & \text{if } i = j(t)t \text{ and } y_t \notin [\gamma^n-2\lambda v_1,\gamma^n+2\lambda v_1), \end{cases}$$

where $\frac{1}{2}\delta_{\{Y\}}+\frac{1}{2}\delta_{\{N\}}$ is the probability measure on the set of responses $\{Y,N\}$ which assigns $1/2$ probability to each response. The sequence $(\widehat{\Gamma}^n)_{n\geq 0}$ is a test sequence for Γ. For each n, since the minimum probabilities ϵ_{it}^n and the reference measures χ_{it}^n are constant and identical

across all pairs (i,t) such that $i = i(t)$, and only depend on the current offer y_t in a identical manner across all pairs (i,t) such that $i = j(t)$, the perturbed game $\widehat{\Gamma}^n$ respects independence of strategically irrelevant information.

We look for values of α^n and γ^n such that all strategies μ_t^n and ρ_t^n at each period $t (t \geq 1)$ are everywhere optimal against σ^n in $\widehat{\Gamma}^n$. As the unimprovability property is sufficient for being a Nash equilibrium (Proposition 14), this would imply that σ^n is a Nash equilibrium of $\widehat{\Gamma}^n$. Then we verify that σ^n is a $(2\lambda, \epsilon^n)$-extended proper equilibrium of Γ_R in extensive form. Upon verifying that σ^n converges weak*, pointwise, to σ^*, we would prove that σ^* is a restricted equilibrium of Γ_R.

Notice that the Rubinstein's game Γ_R and the strategy profiles σ^n are both stationary and symmetric between Player 1 and 2. Therefore the value is the same for both the offering and the responding player in all subsequent continuation games. Let $v_1^n = U_1(\sigma^n)$ be the value of the offering player, and $v_2^n = U_2(\sigma^n)$ be that of the responding player. Then

$$\begin{aligned} v_1^n =& (1-\epsilon^n)\left[(1-\epsilon^n)(1-\gamma^n) + \epsilon^n(1-\eta)v_2^n\right] \\ &+ \epsilon^n\Big[(1-\alpha^n)(1-\epsilon^n)(1-\eta)v_2^n \\ &\qquad + \epsilon^n(1-\gamma^n+\lambda v_1) + \alpha^n(1-\gamma^n)/2\Big] + \alpha^n O\left((\epsilon^n)^2\right) + O\left((\epsilon^n)^3\right), \end{aligned} \tag{40}$$

where $O(\epsilon^n)$ denotes as usual a term that is asymptotically bounded by ϵ^n. To derive this value for v_1^n, consider period 1 when

- with probability $1-\epsilon^n$, Player 1 offers γ^n to Player 2. This offer is
  - accepted with probability $1-\epsilon^n$, in which case Player 1 receives $1-\gamma^n$, and
  - rejected with probability ϵ^n, in which case, with probability $1-\eta$, Player 1 receives its value in the subsequent continuation game (That value is v_2^n by symmetry and stationarity);
- With probability ϵ^n, Players 1 offers an amount y with distribution χ^n. Such an offer is
  - rejected with probability $(1-\alpha^n)(1-\epsilon^n) + \alpha^n O(\epsilon^n) + O\left((\epsilon^n)^2\right)$, and

– conditional on the offer y being accepted and being in the interval $[\gamma^n - 2\lambda v_1, \gamma^n)$, an event with probability $\epsilon^n + \alpha^n O(\epsilon^n) + O\left((\epsilon^n)^3\right)$, Player 1 receives an expected value of $1 - \gamma^n + \lambda v_1$; whereas conditional on y being accepted and being in the interval $[\gamma^n, 1]$, an event with probability $\alpha^n + \alpha^n O(\epsilon^n)$, Player 1 receives an expected value of $(1 - \gamma^n)/2$. The probability that y is in the interval $[0, \gamma^n)$ is $O\left((\epsilon^n)^2\right)$.

This leads to the value v_1^n of Player 1 as showed in equation (40) above. Likewise, we can obtain that

$$\begin{aligned} v_2^n = &(1-\epsilon^n)\left[(1-\epsilon^n)\gamma^n + \epsilon^n(1-\eta)v_1^n\right] \\ &+ \epsilon^n\Big[(1-\alpha^n)(1-\epsilon^n)(1-\eta)v_1^n \\ &\quad + \epsilon^n(\gamma^n - \lambda v_1) + \alpha^n\left(1+\gamma^n\right)/2\Big] + \alpha^n O\left((\epsilon^n)^2\right) + O\left((\epsilon^n)^3\right), \end{aligned} \tag{41}$$

Consider the everywhere optimality of ρ_1^n as a response to σ^n: In period 1, knowing that his value in the subsequent continuation game is $(1-\eta)v_1^n$ if he rejects an offer from Player 1, Player 2's best response is such that it rejects, with maximum probability, all offers that are strictly less than $(1-\eta)v_1^n$, and accepts all offers that are strictly larger than $(1-\eta)v_1^n$ with maximum probability. Thus, ρ_1^n is everywhere optimal if and only if $\gamma^n = (1-\eta)v_1^n$. That is,

$$\begin{aligned} v_1^n = &(1-\epsilon^n)^2\left(1-(1-\eta)v_1^n\right) + (2-\alpha^n)(1-\epsilon^n)\epsilon^n\,(1-\eta)v_2^n \\ &+ \epsilon^n\left[\epsilon^n(1-(1-\eta)v_1^n + \lambda v_1) + \alpha^n\left(1-(1-\eta)v_1^n\right)/2\right] + \alpha^n O\left((\epsilon^n)^2\right) + O\left((\epsilon^n)^3\right), \end{aligned} \tag{42}$$

$$\begin{aligned} v_2^n = &(1-\epsilon^n)^2\,(1-\eta)v_1^n + (2-\alpha^n)(1-\epsilon^n)\epsilon^n\,(1-\eta)v_1^n \\ &+ \epsilon^n\left[\epsilon^n((1-\eta)v_1^n - \lambda v_1) + \alpha^n\left(1+(1-\eta)v_1^n\right)/2\right] + \alpha^n O\left((\epsilon^n)^2\right) + O\left((\epsilon^n)^3\right), \end{aligned} \tag{43}$$

This condition is also necessary and sufficient for the everywhere optimality of ρ_t^n for all $t \geq 1$, by stationarity and symmetry of both $\widehat{\Gamma}^n$ and σ^n.

Equations (42) and (43) form a linear system of equations with unknowns (v_1^n, v_2^n). Adding equations (42) and (43), we obtain

$$\begin{aligned} v_1^n + v_2^n &= (1-\epsilon^n)^2 + (\epsilon^n)^2 + \alpha^n\epsilon^n + \alpha^n O\left((\epsilon^n)^2\right) + O\left((\epsilon^n)^3\right) \\ &\quad + (2 - 2\epsilon^n - \alpha^n)\epsilon^n(1-\eta)(v_1^n + v_2^n), \\ \iff \quad v_1^n + v_2^n &= 1 - (2-\alpha^n)\eta\epsilon^n - 2(1-2\eta)\eta\epsilon^n + \alpha^n O\left((\epsilon^n)^2\right) + O\left((\epsilon^n)^3\right). \end{aligned} \tag{44}$$

Recall that p_a^n denotes the probability that the two players conclude the game by reaching an agreement. Conditional on an eventual agreement, the total payoff of the two players is 1. Conditional on a breakdown, the total payoff is 0. Thus, we have $v_1^n + v_2^n = p_a^n$. Therefore

$$p_a^n = 1 - (2-\alpha^n)\eta\epsilon^n - 2(1-2\eta)\eta(\epsilon^n)^2 + \alpha^n O\left((\epsilon^n)^2\right) + O\left((\epsilon^n)^3\right). \tag{45}$$

Substituting equation (44) into equation (42), we obtain

$$\begin{aligned} v_1^n = &(1-\epsilon^n)^2\left(1-(1-\eta)v_1^n\right) + (2-2\epsilon^n-\alpha^n)\epsilon^n(1-\eta)(p_a^n - v_1^n) \\ &+ \epsilon^n\left[\epsilon^n(1-(1-\eta)v_1^n + \lambda v_1) + \alpha^n\left(1-(1-\eta)v_1^n\right)/2\right] + \alpha^n O\left((\epsilon^n)^2\right) + O\left((\epsilon^n)^3\right). \end{aligned} \tag{46}$$

Equation (46) is a linear equation in v_1^n with a unique solution in $[0,1]$. We shall choose α^n so that

$$v_1^n = p_a^n v_1.$$

This particular choice for α^n would imply that, conditional on an eventual agreement, the expected payoff $U_1(\sigma^n|\text{agreement})$ of Player 1 is v_1 (thus, the corresponding conditional expected payoff $U_2(\sigma^n|\text{agreement})$ of Player 2 is $1 - U_1(\sigma^n|\text{agreement}) = v_2$). We now show that such a choice for the value of α^n exists.

We denote the righthand side of equation (46) by $g(v_1^n; \alpha^n)$, and let $f(\alpha^n) = p_a^n v_1$. Both $f(\alpha^n)$ and $g(p_a^n v_1; \alpha^n)$ are continuous in α^n. For n sufficiently large,

$$f(1) \simeq v_1 - \eta v_1\,\epsilon^n > v_1 - (1/2+\eta)v_1\,\epsilon^n \simeq g(p_a^n v_1; 1).$$

$$\begin{aligned} f(0) \simeq & \quad v_1 - 2\eta v_1\,\epsilon^n - 2\eta(1-2\eta)v_1\,(\epsilon^n)^2 \\ < & \quad v_1 - 2\eta v_1\,\epsilon^n + [\lambda - 2\eta(1-2\eta)]v_1\,(\epsilon^n)^2 \simeq g(p_a^n v_1; 0). \end{aligned}$$

By the Intermediate Value Theorem, there exists a $\alpha^n \in (0,1)$ such that $p_a^n v_1 = f(\alpha^n) = g(p_a^n v_1; \alpha^n)$. We fix such a α^n, implying that $v_1^n = p_a^n v_1$ is the solution to equation (46). Since $p_a^n \to 1$, v_1^n converges from below to v_1, and equation (44) then implies $v_2^n \to v_2$.

Given these values of α^n and γ^n, now we verify the optimality of μ_1^n against σ^n: In period 1, if Player 1 offers some $y \in [0, (1-\eta)v_1^n - 2\lambda v_1)$, then such an offer is rejected with probability $1-(\epsilon^n)^2$, in which case Player 1 gets the same continuation value $(1-\eta)v_2^n$ for all offers y in this interval. Conditional on y being accepted, the expected payoff of Player 1 is decreasing in y. Thus the unconditional payoff of Player 1 is decreasing in y in this interval. Likewise, the unconditional payoff of Player 1 is decreasing in y in all other three intervals $[(1-\eta)v_1^n - 2\lambda v_1, (1-\eta)v_1^n)$, $[(1-\eta)v_1^n, (1-\eta)v_1^n + 2\lambda v_1)$ and $[(1-\eta)v_1^n + 2\lambda v_1, 1]$. Thus we need only compare the payoffs of Player 1 at $y = 0$, $(1-\eta)v_1^n - 2\lambda v_1$, $(1-\eta)v_1^n$ and $(1-\eta)v_1^n + 2\lambda v_1$. The unconditional payoff of Player 1 at these values of y are respectively

$(\epsilon^n)^2 + (1-(\epsilon^n)^2)(1-\eta)v_2^n$	$\epsilon^n(1-(1-\eta)v_1^n + \lambda v_1) + (1-\epsilon^n)(1-\eta)v_2^n$
$(1-\epsilon^n)(1-(1-\eta)v_1^n) + \epsilon^n(1-\eta)v_2^n$	$(1-(\epsilon^n)^2)(1-(1-\eta)v_1^n - \lambda v_1) + (\epsilon^n)^2(1-\eta)v_2^n$

When $n \to \infty$, these payoffs converge to

$(1-\eta)v_2$	$(1-\eta)v_2$	v_1	$(1-\lambda)v_1$

As $(1-\eta)v_2 < v_1$, we know that for n sufficiently large, offering $y = (1-\eta)v_1^n$ yields a strictly higher payoff than offering any other three values of y. On the other hand, an offer μ_1 must be bounded below by $\epsilon^n\chi^n$ in the perturbed game $\widehat{\Gamma}^n$, thus an offer is optimal if and only if its probability distribution is $(1-\epsilon^n)\delta_{(1-\eta)v_1^n} + \epsilon^n\chi^n$, which is precisely μ_1^n. This proves the optimality of μ_1^n. By stationarity and symmetry of $\widehat{\Gamma}^n$ and σ^n, we also obtain the everywhere optimality of μ_t^n for all $t \geq 1$.

Given that all strategies μ_t^n and ρ_t^n at each period t $(t \geq 1)$ are everywhere optimal against σ^n in $\hat{\Gamma}^n$, the strategy profile $\sigma^n = (\sigma_1^n, \sigma_2^n)$ is such that σ_i^n is unimprovable with respect to σ_{-i}^n in the perturbed game $\widehat{\Gamma}^n$ $(i = 1, 2)$. Thus σ^n is a Nash equilibrium of $\widehat{\Gamma}^n$,

by Proposition 14. The strategy profile σ^n is Borel measurable, therefore σ^n is a Borel measurable Nash equilibrium.

Now we verify that σ^n is a $(2\lambda v_1, \epsilon^n)$ extended proper equilibrium of Γ_R in extensive form. For every t that is odd, and every history of offers $y^{t-1} \in [0,1]^{t-1}$, the optimal offer in period t is $y_t^* = (1-\eta)v_1^n$. On the other hand, the offer strategy μ_t^n assigns 0 probability to any offer that is not optimal:

$$\mu_t^n\left(\{y_t\} \,|\, y^{t-1}\right) = 0, \qquad \forall\, y_{t-1} \in [0,1]^{t-1},\, y_t \neq y_t^*.$$

Thus the offer strategies μ_t^n in each period t satisfy the requirement that the probability of a costly deviation being played is at most ϵ^n times the probability of a costless deviation. For every t that is even, and every history of offers $y^t \in [0,1]^t$ (including the current offer y_t), the optimal response in period t is to accept if $y_t > y_t^*$, and reject if $y_t < y_t^*$. The expected loss of a deviation is $|y_t^* - y_t|$. On the other hand, when $|y_t^* - y_t| < 2\lambda v_1$, a deviation is being played with probability ϵ^n under the response $\rho_t^n(\cdot \,|\, y^t)$; when $|y_t^* - y_t| > 2\lambda v_1$, a deviation is being played with probability $(\epsilon^n)^2$. This shows that the offer strategies ρ_t^n satisfy the same requirement that the probability of a costly deviation being played is at most ϵ^n times the probability of a costless deviation. Therefore the strategy profile σ^n is a $(2\lambda v_1, \epsilon^n)$ extended proper equilibrium of Γ_R in extensive form.

Finally, when letting n go to infinity, $\mu_t^n\left(y^{t-1}\right) = (1-\epsilon^n)\,\delta_{(1-\eta)v_1^n} + \epsilon^n \chi^n$ weakly* converges to $\mu_t^*\left(y^{t-1}\right) = \delta_{v_2}$ for every $t \geq 1$ and for every offer history $y^{t-1} \in [0,1]^{t-1}$ up to period t; Likewise, $\rho_t^n\left(y^t\right)$ weakly* converges to $\rho_t^*\left(y^t\right)$ because v_1^n converges from below to v_1. Therefore, we conclude that $\sigma^* = (\mu^*, \rho^*)$ is a restricted equilibrium, and that (a) holds.

Unique equilibrium outcome and (b):

Step 1: Consider a strategy profile $\sigma = (\mu, \rho)$, for a $y_1 \in [0,1]$, replace the offer μ_1 of Player 1 in period 1 by δ_{y_1}. Denote the new strategy profile by σ/y_1. If σ is an Nash equilibrium of

a perturbed game $\widehat{\Gamma} = (\Gamma, \epsilon, \chi)$, then we have

$$U_1(\sigma^n) \geq (1 - \epsilon_{11}^n) \sup_{y_1 \in [0,1]} U_1(\sigma^n / y_1) \geq (1 - \sup \epsilon^n) \sup_{y_1 \in [0,1]} U_1(\sigma^n / y_1). \tag{47}$$

This is because one can let the time 1 offer to be $(1 - \epsilon_{11})\delta_{\tilde{y}_1} + \epsilon_{11}\chi_{11}$ where $\tilde{y}_1 \in [0,1]$ is such that $U_1(\sigma^n/\tilde{y}_1) \geq \sup_{y_1 \in [0,1]} U_1(\sigma^n/y_1) - \delta$ for some $\delta > 0$. This implies that

$$U_1(\sigma^n) \geq (1 - \epsilon_{11}) \left(\sup_{y_1 \in [0,1]} U_1(\sigma^n / y_1) - \delta \right).$$

As δ can be arbitrarily small, one obtains equation (47).

Suppose σ is a restricted equilibrium of Γ_R, and $\widehat{\Gamma}^n = (\Gamma_R, \epsilon^n, \chi^n)$ is a restricted test sequence for σ. In particular, $\sup \epsilon^n \to 0$ and each $\widehat{\Gamma}^n$ respects independence of strategically irrelevant information. For any multistage game Γ, let $\mathcal{E}(\Gamma)$ be the set of Borel measurable Nash equilibria of Γ. Then it is clear that $\mathcal{E}(\widehat{\Gamma}^n) \neq \emptyset$. Let

$$\begin{aligned} M &= \limsup_{n \to \infty} \sup_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_1(\sigma^n), \\ m &= \liminf_{n \to \infty} \inf_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_1(\sigma^n). \end{aligned}$$

Because $\mathcal{E}(\widehat{\Gamma}^n) \neq \emptyset$, we have $M > -\infty$ and $m < \infty$. We first show that

$$M = m = \frac{1}{2 - \eta}. \tag{48}$$

In order to do this, we will show that

$$m \geq 1 - (1 - \eta)M, \tag{49}$$

$$M \leq 1 - (1 - \eta)m. \tag{50}$$

Then, equations (49) and (50) would together imply that $M = m = \dfrac{1}{2 - \eta}$.

Fix an arbitrary $\delta > 0$. There exists n_δ, such that for all $n \geq n_\delta$, we have $U_1(\sigma^n) < M + \delta$ for any $\sigma^n \in \mathcal{E}\left(\widehat{\Gamma}^n\right)$, and $\sup \epsilon^n < \delta$. Fix some $\sigma^n = (\mu^n, \rho^n) \in \mathcal{E}\left(\widehat{\Gamma}^n\right)$. Let y_1 be an offer from Player 1 to Player 2 in period 1, and let $\bar{y} = (1-\eta)(M+\delta)$. It follows from Corollaries 4

and 6 that for μ_1^n almost every[15] $y_1 \in [\bar{y}, 1]$, the value of Player 2 in the continuation game is strictly less than $\bar{y}$ if he rejects the offer y_1. On the other hand, the necessary condition Proposition 15 implies that for μ_1^n almost every[16] $y_1 \in [\bar{y}, 1]$, the response $\rho_1^n(\cdot \mid y_1)$ of Player 2 to the offer y_1 is optimal. Thus for μ_1^n almost every $y_1 \in [\bar{y}, 1]$, Player 2 accepts y_1 with maximum probability:

$$\rho_1^n(\{Y\} \mid y_1) > 1 - \sup \epsilon^n > 1 - \delta.$$

Since μ_1^n is a strictly positive measure on unit interval $[0, 1]$, thus there is a dense subset $D \subset [\bar{y}, 1]$ such that for every $y_1 \in D$, Player 2 accepts y_1 with probability at least $1 - \delta$. Thus $\forall\, y_1 \in D$,

$$U_1(\sigma^n / y_1) \geq (1 - \delta)(1 - y_1)$$

It follows from equation (47) that

$$U_1(\sigma^n) \geq (1 - \delta) \sup_{y_1 \in [0,1]} U_1(\sigma^n / y_1) \geq (1 - \delta)^2 \sup_{y_1 \in D} (1 - y_1) \geq (1 - \delta)^2 [1 - (1 - \eta)(M + \delta)].$$

We therefore have for any $\delta > 0$, there exists n_δ^* such that for every $n \geq n_\delta$,

$$\inf_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_1(\sigma^n) \geq (1 - \delta)^2 [1 - (1 - \eta)(M + \delta)].$$

Hence

$$m = \liminf_{n \to \infty} \inf_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_1(\sigma^n) \geq 1 - (1 - \eta)M.$$

establishing equation (49).

We show equation (49) with a similar argument. Fix a $\delta > 0$, let $\underline{y} = (1 - \eta)(m - \delta)$. There exists n_δ such that for all $n \geq n_\delta$ and $\sigma^n = (\mu^n, \rho^n) \in \mathcal{E}(\widehat{\Gamma}^n)$, the followings are true: (1) for μ_1^n almost every $y_1 \in [0, 1]$, the value of Player 2 in the continuation game is strictly greater than $\underline{y}$ if he rejects the offer y_1. Let $\pi_2(y_1)$ denote this continuation value of Player 2 as a function of $y_1 \in [0, 1]$, and let $\pi_1(y_1)$ denote the value of Player 1 in the same

[15] This means, "there exists a universally measurable $A \subset [\bar{y}, 1]$ such that $\mu_1^n(A) = \mu_1^n([\bar{y}, 1])$ and for every $y_1 \in A$".

[16] This means, "there exists a universally measurable $B \subset [\bar{y}, 1]$ such that $\mu_1^n(B) = \mu_1^n([\bar{y}, 1])$ and for every $y_1 \in B$".

continuation game. Then for μ_1^n almost every $y_1 \in [0,1]$,

$$\begin{aligned}&\pi_2(y_1) > \underline{y},\\ &\pi_1(y_1) \leq (1-\eta) - \pi_2(y_1) < 1 - \eta - \underline{y},\end{aligned}$$

(2) For μ_1^n almost every $y_1 \in \left[0, \underline{y}\right]$, Player 2 rejects y_1 with maximum probability:

$$\rho^n(\{N\} \,|\, y_1) = 1 - \sup \epsilon^n > 1 - \delta.$$

Hence for every $n \geq n_\delta$,

$$\begin{aligned}U_1(\sigma^n) = & \int_{[0,1]} \mu_1^n(dy_1)\, \rho_1^n(\{N\} \,|\, y_1)\, \pi_1(y_1)\\ & + \int_{[0,\underline{y}]} \mu_1^n(dy_1)\, \rho_1^n(\{Y\} \,|\, y_1)\, (1-y_1) + \int_{(\underline{y},1]} \mu_1^n(dy_1)\, \rho_1^n(\{Y\} \,|\, y_1)\, (1-y_1)\\ \leq & \left(1-\eta-\underline{y}\right) \int_{[0,1]} \mu_1^n(dy_1)\, \rho_1^n(\{N\} \,|\, y_1)\\ & + \delta + \left(1-\underline{y}\right) \int_{(\underline{y},1]} \mu_1^n(dy_1)\, \rho_1^n(\{Y\} \,|\, y_1)\\ \leq & \left(1-\eta-\underline{y}\right) \mu_1^n\left(\left[0,\underline{y}\right]\right) + \left(1-\underline{y}\right) \mu_1^n\left(\left(\underline{y},1\right]\right) + \delta\\ \leq & \; 1 - \underline{y} + \delta.\end{aligned}$$

Therefore for every $n \geq n_\delta$,

$$\inf_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_1(\sigma^n) \leq 1 - (1-\eta)(m-\delta) + \delta.$$

Hence $M \leq 1 - (1-\eta)m$, establishing equation (50). Therefore Equation (48) holds.

Step 2: Now we show (b). Equation (48) implies that for any restricted trembling sequence $(\sigma^n)_{n\geq 0}$ of some restricted equilibrium σ,

$$U_1(\sigma^n) \to v_1. \tag{51}$$

Then we have

$$U_2(\sigma^n) \leq 1 - U_1(\sigma^n) \to v_2. \tag{52}$$

On the other hand, by rejecting all offers $y_1 \in [0,1]$ from Player 1 with maximum probability, Player 2 can ensure a payoff of at least

$$\int_{[0,1]} \mu_1^n(dy_1)\,(1-\sup\epsilon^n)\,\pi_2(y_1).$$

where as in step 1, $\pi_2(y_1)$ is the value of Player 2 in the continuation game after he rejects the offer y_1. Then as step 1 show, for any δ, there exists n_δ such that for all $n \geq n_\delta$, $\pi_2(y_1) > v_2 - \delta$ for μ_1^n almost every $y_1 \in [0,1]$ and $\sup\epsilon^n < \delta$. Hence the integral above is bounded below by $(1-\delta)(v_2-\delta)$, for all $n \geq n_\delta$. Thus for all $n \geq n_\delta$,

$$U_2(\sigma^n) \geq (1-\delta)(v_2-\delta). \tag{53}$$

Equations (52) and (53) together imply that

$$U_2(\sigma^n) \to v_2.$$

Thus

$$p_a(\sigma^n) = p_a^n\,[U_1(\sigma^n|\text{agreement}) + U_2(\sigma^n|\text{agreement})] = U_1(\sigma^n) + U_2(\sigma^n) \to 1$$

Therefore $U_i(\sigma^n|\text{agreement}) \to v_i$ for $i = 1,2$, establishing (b).

Step 3: Suppose $\sigma = (\mu,\rho)$ is a restricted equilibrium of Γ_R, now we show that σ induces the same outcome as σ^*. Let $\Gamma^n = (\Gamma_R, \epsilon^n, \chi^n)$ be a restricted test sequence for σ, and $(\sigma^n)_{n\geq 0}$ a corresponding restricted trembling sequence for σ associated with $(\widehat{\Gamma}^n)_{n\geq 0}$. As in step 1, we fix a $\delta > 0$, there exists n_δ such that for all $n \geq n_\delta$,

$$\rho_1^n(\{Y\}\,|\,y_1) > 1 - \sup\epsilon^n > 1-\delta^2 \qquad \text{for } \mu_1^n \text{ almost every } y_1 \in [v_2+\delta, 1], \tag{54}$$

$$\rho_1^n(\{N\}\,|\,y_1) > 1 - \sup\epsilon^n > 1-\delta^2 \qquad \text{for } \mu_1^n \text{ almost every } y_1 \in [0, v_2-\delta]. \tag{55}$$

Since μ_1^n is a strictly positive measure, equation (54) implies that there exists a $y^* \in [v_2 +$

$\delta, v_2 + 2\delta]$ such that

$$\rho_1^n(\{Y\} \,|\, y^*) > 1 - \delta^2.$$

We show that μ_1^n assigns minimal probability to offers $y_1 \in [v_2 + 3\delta, 1]$ and $y_1 \in [0, v_2 - \delta]$. That is,

$$\mu_1^n([v_2 + 3\delta, 1]) = \epsilon_{11}^n \, \chi_{11}^n([v_2 + 3\delta, 1]), \tag{56}$$

$$\mu_1^n([0, v_2 - 3\delta]) = \epsilon_{11}^n \, \chi_{11}^n([0, v_2 - \delta]). \tag{57}$$

Since the probability measure μ_1^n is bounded below by $\epsilon_{11}^n \, \chi_{11}^n$, it can be decomposed as

$$\mu_1^n = \epsilon_{11}^n \, \chi_{11}^n + (1 - \epsilon_{11}^n)\mu$$

for some probability measure μ on the unit interval $[0, 1]$. Let $\beta = \mu([v_2 + 3\delta, 1])$. If $\beta > 0$, we define a probability measure ν on $[0, 1]$ as

$$\nu = \epsilon_{11}^n \, \chi_{11}^n + (1 - \epsilon_{11}^n)\mu_{[0, v_2 + 3\delta)} + (1 - \epsilon_{11}^n)\beta\delta_{y^*},$$

where $\mu_{[0, v_2 + 3\delta)}$ is the measure on $[0, 1]$ defined by $\mu_{[0, v_2 + 3\delta)}(B) = \mu([0, v_2 + 3\delta) \cap B)$ for every Borel subset B of $[0, 1]$. That is, ν "moves" all the extra mass of μ_1^n in the interval $[v_2 + 3\delta, 1]$ relative to the minimally required measure $\epsilon_{11}^n \, \chi_{11}^n$ to the point y^*. Thus ν is a valid Borel measurable offer strategy of Player 1 in period 1 in the perturbed game $\widehat{\Gamma}^n$. Let σ^n/ν be the strategy profile $(\nu, \rho_2, \mu_3, \dots)$. Then

$$\begin{aligned}
& U_1(\sigma^n/\nu) - U_1(\sigma^n) \\
\geq \; & (1 - \epsilon_{11}^n)\beta(1 - y^*) \, \rho_1^n(\{Y\} \,|\, y^*) \\
& \quad - \int_{[v_2 + 3\delta, 1]} \mu(dy_1) \, \rho_1^n(\{Y\} \,|\, y_1) \, (1 - y_1) - \int_{[v_2 + 3\delta, 1]} \mu(dy_1) \, \rho_1^n(\{N\} \,|\, y_1) \\
\geq \; & (1 - \delta^2)\beta(1 - y^*)(1 - \delta^2) - (1 - v_2 - 3\delta)\beta - \delta^2\beta \\
\geq \; & \beta\delta(1 - 3\delta - 2\delta^2).
\end{aligned}$$

Thus if $\beta > 0$, $U_1(\sigma^n/\nu) - U_1(\sigma^n) > 0$ when δ is sufficiently small. This contradicts the fact that σ^n is a Nash equilibrium of $\widehat{\Gamma}^n$. Therefore $\beta = 0$, establishing equation (56). One can

show equation (57) in a similar way.

Now we show that $\mu_1 = \delta_{v_2}$. Since μ_1^n converges weak* to μ, and $(v_2 + 3\delta, 1]$ is a relative open subset of $[0, 1]$, we have by Portemanteau Theorem,

$$\mu_1((v_2 + 3\delta, 1]) \leq \liminf \mu_1^n((v_2 + 3\delta, 1]) = 0.$$

Likewise, we have

$$\mu_1([0, v_2 - \delta)) \leq \liminf \mu_1^n([0, v_2 - \delta)) = 0.$$

As δ can be arbitrarily small, we have $\mu_1 = \delta_{v_2}$.

Now we consider ρ_2. Since $(\widehat{\Gamma}^n)_{n\geq 0}$ is a test sequence for Γ, there exists a strictly positive probability measure χ_{11} on $[0, 1]$ that is absolute continuous with respect to χ_{11}^n for every n. Then inequality 54 implies that for all $n \geq n_\delta$, there exists a Borel subset $A_n \subset [v_2 + \delta, 1]$ such that $\chi(A_n) = \chi([v_2 + \delta, 1])$ and

$$\rho_1^n(\{Y\} \,|\, y_1) > 1 - \sup \epsilon^n \qquad \forall\, y_1 \in A_n.$$

Letting $A_\delta = \cap_{n\geq n_\delta} A_n$, we have $\chi(A_\delta) = \chi([v_2 + \delta, 1])$ for $n \geq n_\delta$,

$$\rho_1^n(\{Y\} \,|\, y_1) > 1 - \sup \epsilon^n \qquad \forall\, y_1 \in A_\delta.$$

Therefore

$$\rho_2(\{Y\} \,|\, y_1) = \lim_{n\to\infty} \rho_1^n(\{Y\} \,|\, y_1) = 1 \qquad \forall\, y_1 \in A_\delta.$$

Since χ is a strictly positive probability measure and $\chi(A_\delta) = \chi([v_2 + \delta, 1])$, A_δ is a dense subset in $[v_2 + \delta, 1]$. Then $A = \cup_{\delta\to 0} A_\delta$ is dense in $[v_2, 1]$, and we have $\rho_2(\{Y\} \,|\, y_1) = 1$ for all $y_1 \in A$. That is, ρ_2 accepts all offers in A with certainty.

We already know that μ_1 offers v_2 with certainty. If ρ_2 does not accept v_2 with certainty (that is, ρ_2 either randomizes between accept and reject, or ρ_2 rejects v_2 with certainty), then it must be that the value $\pi_2(v_2)$ of Player 2 in the continuation game after he rejects v_2 is at least v_2 (recall that since $\sigma = (\mu, \rho)$ is a restricted equilibrium of Γ_R, σ is a Nash equilibrium of Γ_R to begin with). Thus the value $\pi_1(v_2)$ of Player 1 in the same continuation

game is at most $\pi_1(v_2) \leq 1 - \eta - \pi_2(v_2) \leq v_1 - \eta$. However, by offering some $y_1 \in A$, Player 1 is guaranteed to secure a payoff of $1 - y_1$. Thus Player 1 can achieve a payoff that is arbitrarily close to v_1. This contradicts the fact that σ is a Nash equilibrium of Γ_R. Thus it must be that ρ_2 accepts v_2 with certainty. Therefore σ induces the same outcome as σ^*. □

Remark 4. In the proof of uniqueness and (b) above, we did not use the requirement that a costly deviation is less likely than a costless one. That is, Theorem 3 still holds if the last condition in the definition of restricted equilibrium Definition 2 was removed. However this requirement will be useful later, when the total size of the pie to be shared is non-positive.

We now prepare some careful corollaries and extensions to Theorem 3, which will be useful for the proof of Theorem 1. In the original Rubinstein's alternating offer game, the set of possible payments is constrained to be the interval $[0, 1]$, because any payment that is not in this interval would not be individually rational for one of the players. However the proof of Theorem 3 did not use this fact other than $[0, 1]$ is a compact interval. Thus Theorem 3 still holds if the set of possible payments is some larger compact interval that contains $[0, 1]$. Therefore we have the following corollary:

Corollary 7. *Theorem 3 still holds if the set of possible payments is some compact interval that contains* $[0, 1]$.

Let $\Gamma(\pi, \underline{u}_1, \underline{u}_2)$ be the Rubinstein's game in which the size of the "pie" to be shared by the two players is π, and the outside option values of the two players are $\underline{u}_1$ and $\underline{u}_2$ respectively. We always assume that Player 1 is the one that starts offering at time 1. We first consider the case where $\pi > \underline{u}_1 + \underline{u}_2$. The associated set of possible payments in the game is assumed to be some compact interval that contains the payment interval $[0, \pi - \underline{u}_1 - \underline{u}_2]$ prescribed by individual rationality. Let $\sigma^*(\pi, \underline{u}_1, \underline{u}_2)$ be the unique subgame perfect equilibrium of $\Gamma(\pi, \underline{u}_1, \underline{u}_2)$.

An immediate corollary to Theorem 3 is the following:

Corollary 8 (Constant Rubinstein's Game)**.** *Consider the Rubinstein's game* $\Gamma(\pi, \underline{u}_1, \underline{u}_2)$ *with* $\pi > \underline{u}_1 + \underline{u}_2$. *The game has a restricted equilibrium. In particular, the pure strategy*

profile $\sigma^*(\pi, \underline{u}_1, \underline{u}_2)$ *is a restricted equilibrium. Every restricted equilibrium* σ *induces the same outcome as* σ^**. Moreover,*

(a) There exists a restricted trembling sequence $(\sigma^n)_{n\geq 0}$ *for* $\sigma^*(\pi, \underline{u}_1, \underline{u}_2)$ *such that the followings are true:*

(i) for every n*, the expected payoff of Player* i *under* σ^n*, conditional on the game reaching an eventual agreement, is constant and equal to* $v_i(\pi - \underline{u}_1 - \underline{u}_2) + \underline{u}_i$ *(*$i = 1, 2$*).*

(ii) As $n \to \infty$*, the probability of an eventual agreement converges to 1.*

(iii) As a consequence of (ii), the unconditional payoff of Player i *under* σ^n *converges to* $v_i(\pi - \underline{u}_1 - \underline{u}_2) + \underline{u}_i$*.*

(b) For each restricted trembling sequence $(\sigma^n)_{n\geq 0}$ *for some restricted equilibrium, the followings are true:*

(i') As $n \to \infty$*, the unconditional payoff of Player* i *under* σ^n *converges to* $v_i(\pi - \underline{u}_1 - \underline{u}_2) + \underline{u}_i$*.*

(ii') The probability of an eventual agreement converge to 1.

(iii') As a consequence of (i) and (ii), the expected payoff of Player i *under* σ^n*, conditional on the game reaching an eventual agreement, converges to* $v_i(\pi - \underline{u}_1 - \underline{u}_2) + \underline{u}_i$ *(*$i = 1, 2$*).*

In Corollary 8 (Constant Rubinstein's Game), there is a fixed game $\Gamma(\pi, \underline{u}_1, \underline{u}_2)$ that does not depend on n. This is why we name Corollary 8 the "Constant Rubinstein's Game". We next consider the situation where the the total payoff π^n and the outside option values $\underline{u}_i^n$ $(i = 1, 2)$ depends on n and converges to some $(\pi, \underline{u}_1, \underline{u}_2)$. Recall the assumption that $\pi > \underline{u}_1 + \underline{u}_2$. So for n sufficiently large, we have $\pi^n > \underline{u}_1^n + \underline{u}_2^n$. Consider the sequence of multistage games $\Gamma^n = \Gamma(\pi^n, \underline{u}_1^n, \underline{u}_2^n)$. The associated set of possible payments in all those games is assumed to be some compact interval $[0, \beta]$ (independent of n) that contains the intervals $[0, \pi^n - \underline{u}_1^n - \underline{u}_2^n]$ for *all* n. Thus the set of strategy profiles of Γ^n are the same for

all n. The utility functions, however, are different across n. We look for a strategy profile σ that is "almost" a restricted equilibrium of $\Gamma(\pi, \underline{u}_1, \underline{u}_2)$, in that satisfies the conditions in Definition 2 upon replacing the environment game $\Gamma(\pi, \underline{u}_1, \underline{u}_2)$ by Γ^n along the converging sequence of strategy profiles. We will establish existence and uniqueness results that are similar to the ones in Corollary 8. The only difference is that, along the converging sequence of strategy profiles, the environment game is the sequence of games $\Gamma^n = \Gamma(\pi^n, \underline{u}_1^n, \underline{u}_2^n)$ in the place of the constant game $\Gamma(\pi, \underline{u}_1, \underline{u}_2)$.

We first introduce the following definition. Consider a sequence of generic Rubinstein's game $(\Gamma^n)_{n\geq 0}$. For each n, the set of all conceivable agreements is some set K, and the breakdown outcome is D. Each player i's utility function is some $u_i^n : K \cup \{D\} \to \mathbb{R}$. When $n \to \infty$, u_i^n converges to some u_i pointwise. Let Γ be the Rubinstein's game with parameters $(K, D, (u_i)_{i=1,2})$. The sequence of games $\Gamma^n = \Gamma(\pi^n, \underline{u}_1^n, \underline{u}_2^n)$ is an example of such "converging games" , with the "limiting game" $\Gamma = \Gamma(\pi, \underline{u}_1, \underline{u}_2)$.

Definition 4. Let $(\Gamma^n)_{n\geq 0}$ and Γ be given as above. Define $\mathcal{R}((\Gamma^n)_{n\geq 0})$ to be the set of strategy profiles of Γ such that $\sigma \in \mathcal{R}((\Gamma^n)_{n\geq 0})$ if and only if σ is a Nash equilibrium of the limiting Rubinstein's game Γ, and there exists $\bar{\lambda} > 0$ such that for all strictly positive scalar $\lambda < \bar{\lambda}$, there exists a sequence σ^n of strategy profiles satisfying the following properties.

1. There exists a sequence $\widehat{\Gamma}^n = (\Gamma^n, \epsilon^n, \chi^n)$ such that, $(\Gamma, \epsilon^n, \chi^n)_{n\geq 0}$ constitutes a test sequence for Γ, and each $\widehat{\Gamma}^n$ respects independence of strategically irrelevant information.

2. For each n, σ^n is a Borel measurable Nash equilibrium of $\widehat{\Gamma}^n$ and also a (λ, δ^n)-extended proper equilibrium of Γ^n in extensive form, for some strictly positive scalar sequence δ^n converging to zero.

3. σ^n converges weak*, pointwise, to σ, as $n \to \infty$.

For every $\sigma \in \mathcal{R}((\Gamma^n)_{n\geq 0})$ and every λ sufficiently, let $\mathcal{T}[\lambda, \sigma; (\Gamma^n)_{n\geq 0}]$ be the set of such sequences σ^n of converging equilibria.

Given a strategy profile σ, for each n, recall that $U_i(\sigma;\Gamma^n)$ is the expected utility of Player i under σ in the game Γ^n. Let $p_a(\sigma;\Gamma^n)$ be the probability that the two players conclude the game by reaching an agreement in Γ^n when playing σ, and $U_i(\sigma;\Gamma^n|\text{agreement})$ be the expected utility of Player i under σ in Γ^n, conditional on the two players eventually reaching an agreement.

Lemma 11. *Suppose $(\pi^n)_{n\geq 0}$ is a sequence of total payoffs converging to π, and $(\underline{u}_i^n)_{n\geq 0}$ are two sequences of outside option values converging to $\underline{u}_i$ $(i=1,2)$, $\pi > \underline{u}_1+\underline{u}_2$ and $\pi > \underline{u}_1^n+\underline{u}_2^n$ for all n. Let $\Gamma^n = \Gamma(\pi^n,\underline{u}_1^n,\underline{u}_2^n)$. The unique sub-game perfect equilibrium $\sigma^*(\pi,\underline{u}_1,\underline{u}_2)$ of the game $\Gamma(\pi,\underline{u}_1,\underline{u}_2)$ is in the set $\mathcal{R}((\Gamma^n)_{n\geq 0})$. Every strategy profile in $\mathcal{R}((\Gamma^n)_{n\geq 0})$ induces the same outcome as $\sigma^*(\pi,\underline{u}_1,\underline{u}_2)$. Moreover,*

(a) For λ sufficiently small, there exists a sequence $(\sigma^n)_{n\geq 0}$ in $\mathcal{T}\left[\lambda,\sigma^*(\pi,\underline{u}_1,\underline{u}_2);(\Gamma^n)_{n\geq 0}\right]$ such that the followings are true:

(i) for every n, $U_1(\sigma^n;\Gamma^n|\text{agreement}) = v_1(\pi-\underline{u}_1-\underline{u}_2)+\underline{u}_1$ for Player 1.

(ii) As $n\to\infty$, the probability $p_a(\sigma^n;\Gamma^n)$ of an eventual agreement converges to 1.

(iii) $U_i(\sigma^n;\Gamma^n)$ converges to $v_i(\pi-\underline{u}_1-\underline{u}_2)+\underline{u}_i$ $(i=1,2)$.

(b) For each sequence $(\sigma^n)_{n\geq 0}\in\mathcal{T}\left[\sigma,(\Gamma^n)_{n\geq 0}\right]$ for some $\sigma\in\mathcal{R}((\Gamma^n)_{n\geq 0})$, the followings are true:

(i') As $n\to\infty$, the unconditional payoff $U_i(\sigma^n;\Gamma^n)$ converges to $v_i(\pi-\underline{u}_1-\underline{u}_2)+\underline{u}_i$ $(i=1,2)$.

(ii') The probabilities $p_a(\sigma^n;\Gamma^n)$ converge to 1;

(iii') As a consequence, $U_i(\sigma^n;\Gamma^n|\text{agreement})\to v_i(\pi-\underline{u}_1-\underline{u}_2)+\underline{u}_i$ $(i=1,2)$.

The proof is very similar to that of Theorem 3. There is only one small adjustment that need to be made when showing existence and (a) which is the following: In Theorem 3, we simply took an arbitrary positive sequence $\epsilon^n\to 0$. In the lemma above, we should take a

positive sequence $\epsilon^n \to 0$ such that

$$|\pi^n - \pi| + |\underline{u}_1^n - \underline{u}_1| + |\underline{u}_2^n - \underline{u}_2| = o\left((\epsilon^n)^2\right),$$

That is, the deviation of the total payoffs π^n and the outside option values $\underline{u}_i$ from their respective limits $(\pi, \underline{u}_1, \underline{u}_2)$ are very small relative to the minimum tremble probabilites ϵ^n asymptotically. Then all the calculation in the proof of Theorem 3 remains valid, upon replacing the total payoff by π and the outside option values by $\underline{u}_i$. The proof of uniqueness and (b) are simply a reiteration of that of Theorem 3, which we do not repeat here.

In the setting of Lemma 11, we next consider the situation where $\underline{u}_1^n = \underline{u}_1$ and $\pi \leq \underline{u}_1 + \underline{u}_2$. This means the outside option value of Player 1 is independent of n and the total payoff is less than the combined outside option values in the limit. Thus agreement on any payment would make at least one of the players worse off than taking his outside value. Therefore agreement should not happen in equilibrium since it violates individual rationality. To ease exposition, we reparametrize the set of possible payments by the payoff of Player 1. That is, instead of saying that "Player 1 offers a payment y to Player 2", we say "Player 1 proposes that he gets $a = \pi - y$ and Player 2 gets $u_2(a) = y$". Under the new parametrization with a, $u_1(a) = a$ and $u_2(a) = \pi - a$. We suppose the set of possible payoffs to Player 1 is some compact interval $[\underline{u}_1, \beta]$ $(\beta > \underline{u}_1)$ for all n. That is, Player 1 always gets at least his outcome option value.

Lemma 12. *In the setting of Lemma 11, suppose $\underline{u}_1^n = \underline{u}_1$ and If $\pi \leq \underline{u}_1 + \underline{u}_2$.*

(a) There exists a strategy profile in the set $\mathcal{R}((\Gamma^n)_{n\geq 0})$, which is the following: Player 1 proposes that he gets β in payoff and Player 2 gets $\pi - \beta$, while Player 2 proposes that Player 1 gets $\underline{u}_1$ in each round. Player 1 accepts offers in which his payoff is strictly larger than $\underline{u}_1$ and rejects $\underline{u}_1$, while Player 2 rejects all offers from Player 1. We denote this strategy profile also by $\sigma^*(\pi, \underline{u}_1, \underline{u}_2)$, the same notation as for the case where $\pi > \underline{u}_1 + \underline{u}_2$.

(b) Every strategy profile in $\mathcal{R}((\Gamma^n)_{n\geq 0})$ induces the same payoff as $\sigma^*(\pi, \underline{u}_1, \underline{u}_2)$, in which

both players get their respective outside option values $\underline{u}_i$ ($i = 1, 2$). For each λ sufficiently small, and for each sequence $(\sigma^n)_{n\geq 0} \in \mathcal{T}[\lambda, \sigma, (\Gamma^n)_{n\geq 0}]$ for some $\sigma \in \mathcal{R}((\Gamma^n)_{n\geq 0})$, the followings are true:

(i) As $n \to \infty$, the unconditional payoff $U_i(\sigma^n; \Gamma^n)$ converges to $\underline{u}_i$ ($i = 1, 2$).

(ii) $\limsup_{n\to\infty} U_1(\sigma^n; \Gamma^n \mid \text{agreement}) \leq \underline{u}_1 + 2\lambda$ for Player 1.

Proof. (a). When $\pi < \underline{u}_1 + \underline{u}_2$, we show that $\sigma^*(\pi, \underline{u}_1, \underline{u}_2) \in \mathcal{R}((\Gamma^n)_{n\geq 0})$. For this, we fix an arbitrary positive scalar $\lambda < (\beta - \underline{u}_1)/2$, a positive sequence $\epsilon^n \to 0$ such that

$$|\pi^n - \pi| + |\underline{u}_2^n - \underline{u}_2| = O(\epsilon^n),$$

and a strictly positive measure χ on $[\underline{u}_1, \beta]$. For each $n \geq 0$ and $t \geq 1$, let μ_t^n and ρ_t^n be defined by

$$\mu_{2t-1}^n\left(a^{2t-2}\right) = (1-\epsilon^n)\delta_\beta + \epsilon^n\chi, \qquad \mu_{2t}^n\left(a^{2t-1}\right) = \left[1-(\epsilon^n)^2\right]\delta_{\underline{u}_1} + (\epsilon^n)^2\chi,$$

and

$$\rho_{2t-1}^n\left(\{Y\} \mid a^{2t-1}\right) = (\epsilon^n)^2, \qquad \rho_{2t}^n\left(\{Y\} \mid a^{2t}\right) = \begin{cases} \epsilon^n & \text{if } a_{2t} \in [\underline{u}_1, \gamma^n), \\ 1-\epsilon^n & \text{if } a_{2t} \in [\gamma^n, \gamma^n + \lambda), \\ 1-(\epsilon^n)^2 & \text{if } a_{2t} \in [\gamma^n + \lambda, \beta). \end{cases}$$

for every history of offers $a^{2t-2} \in [\underline{u}_1, \beta]^{2t-2}$, $a^{2t-1} \in [\underline{u}_1, \beta]^{2t-1}$.

The combination $\sigma^n = (\mu^n, \rho^n)$ defines a behavioral strategy profile for the perturbed game $\widehat{\Gamma}^n$ of Γ^n whose minimum probabilities ϵ_{it}^n and reference measures χ_{it}^n are as follows: for every offer history $a^t \in [\underline{u}_1, \beta]^t$,

$$\epsilon_{1t}^n\left(a^t\right) = \begin{cases} \epsilon^n, \\ 2\epsilon^n, \\ 2(\epsilon^n)^2, \end{cases} \qquad \chi_{1t}^n\left(a^t\right) = \begin{cases} \chi & \text{if } t \text{ is odd}, \\ \frac{1}{2}\delta_{\{Y\}} + \frac{1}{2}\delta_{\{N\}} & \text{if } t \text{ is even and } a_t \in [\underline{u}_1, \gamma^n + \lambda), \\ \frac{1}{2}\delta_{\{Y\}} + \frac{1}{2}\delta_{\{N\}} & \text{if } t \text{ is even and } a_t \in [\gamma^n + \lambda, \beta], \end{cases}$$

and

$$\epsilon_{2t}^n\left(a^t\right) = \begin{cases} 2(\epsilon^n)^2, \\ (\epsilon^n)^2, \end{cases} \qquad \chi_{2t}^n\left(a^t\right) = \begin{cases} \frac{1}{2}\delta_{\{Y\}} + \frac{1}{2}\delta_{\{N\}} & \text{if } t \text{ is odd}, \\ \chi & \text{if } t \text{ is even}, \end{cases}$$

One can verify that the sequence $(\Gamma, \epsilon^n, \chi^n)_{n\geq 0}$ is a test sequence for Γ, and for each n, the perturbed game $\widehat{\Gamma}^n = (\Gamma^n, \epsilon^n, \chi^n)$ of Γ^n respects independence of strategically irrelevant information.

Let $v_1^n = U_1(\sigma^n; \Gamma^n)$. As in the proof of Theorem 3, σ^n is a Borel measurable Nash equilibrium of $\widehat{\Gamma}^n$ if $\gamma^n = (1-\eta)v_1^n + \eta\underline{u}_1$. In this case, v_1^n is determined by the following linear system of equations with unknowns $(v_1^n, \tilde{v}_1^n)$:

$$\begin{aligned} v_1^n &= \eta\underline{u}_1 + (1-\eta)\tilde{v}_1^n + O(\epsilon^n), \\ \tilde{v}_1^n &= \eta\underline{u}_1 + (1-\eta)v_1^n + O(\epsilon^n). \end{aligned}$$

Thus v_1^n converges to $\underline{u}_1$, so does $\gamma^n = (1-\eta)v_1^n + \eta\underline{u}_1$. One can also verify that σ^n is a (λ, ϵ^n) extended proper eqiulibrium of Γ^n in extensive form. Finally, when letting n go to infinity, σ^n converges weak*, pointwise, to $\sigma^*(\pi, \underline{u}_1, \underline{u}_2)$. Therefore, we conclude that $\sigma^*(\pi, \underline{u}_1, \underline{u}_2) \in \mathcal{R}((\Gamma^n)_{n\geq 0})$.

Proof of (b). The first claim is easy to show: If $\pi < \underline{u}_1 + \underline{u}_2$, every strategy profile $\sigma \in \mathcal{R}((\Gamma^n)_{n\geq 0})$ must be a Borel Nash equilibrium of the limiting Rubinstein's game $\Gamma = \Gamma(\pi, \underline{u}_1, \underline{u}_2)$. In the game Γ, both player can secure their outside option values $\underline{u}_i$ $(i = 1, 2)$ by rejecting all offers from the other player in every round. This implies that $U_i(\sigma; \Gamma) \geq \underline{u}_i$. On the other hand,

$$U_1(\sigma; \Gamma) + U_2(\sigma; \Gamma) = p_a(\sigma; \Gamma)\pi + [1 - p_a(\sigma; \Gamma)](\underline{u}_1 + \underline{u}_2) \leq \underline{u}_1 + \underline{u}_2.$$

Thus we have $U_i(\sigma; \Gamma) = \underline{u}_i$ $(i = 1, 2)$. That is, σ induces the same payoff as $\sigma^*(\pi, \underline{u}_1, \underline{u}_2)$, in which both players get their respective outside option values $\underline{u}_i$ $(i = 1, 2)$.

For the rest, we will proceed in three steps as in the proof of the uniqueness part of Theorem 3. There is one major difference, which is that we make use of the requirement that a costly deviation is less likely than a costless one in step 3.

Step 1: Suppose $\sigma \in \mathcal{R}((\Gamma^n)_{n\geq 0})$, and $\widehat{\Gamma}^n = (\Gamma^n, \epsilon^n, \chi^n)$ is a sequence of perturbed games of Γ^n associated with σ as in Definition 4. Since $\widehat{\Gamma}^n$ respects irrelevance of strategically information, we have, from Corollary 5, the continuation game $\widehat{\Gamma}^n_{|a_1}$ after the offer a_1 in

period 1 is rejected are identical for all $a_1 \in [\underline{u}_1, \beta]$. To simplify, we denote this continuation game by $\widehat{\Gamma}_2^n$, and let Γ_2^n denote the subgame of Γ which starts in period 2 with an offer of Player 2. Let

$$m_2 = \liminf_{n\to\infty} \inf_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_2(\sigma^n; \Gamma^n),$$
$$\widetilde{m}_2 = \liminf_{n\to\infty} \inf_{\sigma'^n \in \mathcal{E}(\widehat{\Gamma}_2^n)} U_2^n\left(\sigma'^n; \Gamma_2^n\right).$$

Fix an arbitrary $\delta > 0$. There exists n_δ such that for all $n \geq n_\delta$, we have $U_2^n\left(\sigma^n; \Gamma^n\right) > m_2 - \delta$ for all $\sigma^n \in \mathcal{E}\left(\widehat{\Gamma}_2^n\right)$, and $\sup \epsilon^n < \delta$. Fix some $\sigma'^n = (\mu_2^n, \rho_2^n, \mu_3^n, \rho_3^n, \dots) \in \mathcal{E}(\widehat{\Gamma}_2^n)$. In the game $\widehat{\Gamma}_2^n$, Player 2 first makes an offer a_2 to Player 1, for μ_2^n almost every a_2, the value of Player 1 in the continuation game after he rejects a_2 is at least

$$\underline{v}_1\left(\widehat{\Gamma}^n\right) \equiv \eta \underline{u}_1 + (1-\eta) \inf_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_1^n(\sigma^n; \Gamma^n). \tag{58}$$

by Corollaries 4 and 5. In the game $\widehat{\Gamma}^n$, Player 1 gets at least his outside option value $\underline{u}_1$, and strictly better whenever there is an agreement and the agreement is not $\underline{u}_1$. For any strategy profile σ of $\widehat{\Gamma}^n$ (σ does not need to be a Nash equilibirium), σ is bounded below by some reference strategy profile, hence $U_1^n(\sigma)$ is bounded below by some value $m_1^n > \underline{u}_1$ independent of the strategy profile σ. Therefore,

$$\underline{v}_1\left(\widehat{\Gamma}^n\right) > \eta \underline{u}_1 + (1-\eta) m_1^n. \tag{59}$$

Since the response $\rho_2(\cdot \,|\, a_2)$ of Player 1 to the offer a_2 is optimal for μ_2^n almost every a_2 by Proposition 15. Thus for μ_2^n almost every $a_2 \in [\underline{u}_1, \eta\underline{u}_1 + (1-\eta)m_1^n]$, Player 1 rejects a_2 with maximum probability :

$$\rho_2(\{N\} \,|\, a_2) > 1 - \sup \epsilon^n > 1 - \delta.$$

Hence by Corollaries 4 and 5, for μ_2^n almost every $a_2 \in [\underline{u}_1, \eta\underline{u}_1 + (1-\eta)m_1^n]$,

$$U_2(\sigma'^n / a_2; \Gamma_2^n) \geq \delta(\pi - \beta) + (1-\delta)\eta\underline{u}_2^n + (1-\delta)(1-\eta)(m_2 - \delta).$$

Since μ_2^n is a strictly positive measure on $[\underline{u}_1, \beta]$, thus there is a dense subset $D \subset [\underline{u}_1, \eta\underline{u}_1 +$

$(1-\eta)m_1^n]$ such that the inequality above holds for every $a_2 \in D$. It follows that

$$\begin{aligned} U_2(\sigma'^n;\Gamma_2^n) \geq & \ (1-\delta)\sup_{a_2\in[\underline{u}_1,\beta]} U_2(\hat{\sigma}^n/a_2;\Gamma_2^n)+\delta(\pi-\beta) \\ \geq & \ (2-\delta)\delta(\pi-\beta)+(1-\delta)^2\,\eta\underline{u}_2^n+(1-\delta)^2(1-\eta)(m_2-\delta). \end{aligned}$$

We therefore have for any $\delta>0$, there exists n_δ^* such that for every $n \geq n_\delta$,

$$\inf_{\sigma'^n\in\mathcal{E}(\widehat{\Gamma}_2^n)} U_2(\sigma'^n;\Gamma_2^n) \geq (2-\delta)\delta(\pi-\beta)+(1-\delta)^2\eta\underline{u}_2^n+(1-\delta)^2(1-\eta)(m_2-\delta).$$

Hence

$$\widetilde{m}_2 = \liminf_{n\to\infty} \inf_{\sigma'^n\in\mathcal{E}(\widehat{\Gamma}_2^n)} U_2(\sigma'^n;\Gamma_2^n) \geq \eta\underline{u}_2+(1-\eta)m_2. \tag{60}$$

Likewise, in the game $\widehat{\Gamma}^n$, Player 2 can do at least as well as rejecting all offers a_1 from Player 1 in period 1 with maximum probability. Thus by Corollaries 4 and 5,

$$U_2(\sigma^n;\Gamma^n) \geq \delta(\pi-\beta)+(1-\delta)\eta\underline{u}_2^n+(1-\delta)(1-\eta)\widetilde{m}_2.$$

Therefore

$$m_2 \geq \eta\underline{u}_2+(1-\eta)\widetilde{m}_2. \tag{61}$$

Inequalities (60) and (61) together imply that

$$m_2 = \widetilde{m}_2 = \underline{u}_2. \tag{62}$$

Step 2: Since for any n and any strategy profiles σ of Γ^n, σ' of Γ_2^n, we have

$$U_2(\sigma;\Gamma^n) \leq \underline{u}_2^n, \qquad U_2(\sigma';\Gamma_2^n) \leq \underline{u}_2^n.$$

(62) then implies that,

$$\lim_{n\to\infty}\inf_{\sigma^n\in\mathcal{E}(\widehat{\Gamma}^n)} U_2(\sigma^n;\Gamma^n) = \lim_{n\to\infty}\inf_{\sigma'^n\in\mathcal{E}(\widehat{\Gamma}_2^n)} U_2(\sigma'^n;\Gamma_2^n) = \underline{u}_2.$$

Let

$$\underline{v}_2(\widehat{\Gamma}^n) = \eta \underline{u}_2^n + (1-\eta) \inf_{\sigma'^n \in \mathcal{E}(\widehat{\Gamma}_2^n)} U_2(\sigma'^n; \Gamma_2^n), \tag{63}$$

it immediately follows that

$$\lim_{n\to\infty} \underline{v}_2(\widehat{\Gamma}^n) = \underline{u}_2 \tag{64}$$

Letting $M_1^n = \sup_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_1(\sigma^n; \Gamma^n)$, it follows that

$$\limsup_{n\to\infty} M_1^n \leq \limsup_{n\to\infty} \left[\underline{u}_1 + \underline{u}_2^n - \inf_{\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)} U_2(\sigma^n; \Gamma^n) \right] = \underline{u}_1.$$

On the other hand, we already see that $M_1^n > \underline{u}_1$ for all n, therefore

$$\lim_{n\to\infty} M_1^n = \underline{u}_1. \tag{65}$$

In particular, we have for every $\sigma^n \in \mathcal{E}(\Gamma^n)$,

$$\lim_{n\to\infty} U_1(\sigma^n; \Gamma^n) = \underline{u}_1, \qquad \lim_{n\to\infty} U_2(\sigma^n; \Gamma^n) = \underline{u}_2.$$

Let $\bar{v}_1(\widehat{\Gamma}^n) = \eta \underline{u}_1 + (1-\eta) M_1^n$. From (65) we have

$$\underline{u}_1 < \eta \underline{u}_1 + (1-\eta) m_1^n < \underline{v}_1(\widehat{\Gamma}^n) \leq \bar{v}_1(\widehat{\Gamma}^n) \to \underline{u}_1.$$

Thus $\bar{v}_1(\widehat{\Gamma}^n) \to \underline{u}_1$.

Step 3: Now fix a λ sufficiently small, and fix a sequence $\sigma^n = (\mu^n, \rho^n) \in \mathcal{T}[\sigma, \lambda; \Gamma^n]$ of converging equilibria associated with the sequence $(\widehat{\Gamma}^n)_{n\geq 0}$ of perturbed games. That is, for each n, $\sigma^n \in \mathcal{E}(\widehat{\Gamma}^n)$ and σ^n is also a (λ, δ^n) extended proper equilibrium of Γ^n in extensive form, for some strictly positive scalar sequence δ^n converging to zero. For every $it \in W$ such that $i = j(t)$ (that is, player i is the one responding in period t), suppose $\alpha_{it} : [\underline{u}_1, \beta]^t \to \mathbb{R}^+$ is the function that has the properties stated in the definition of an extended proper equilibrium in extensive form (Definition 1).

For any $(a_1, a_2) \in [\underline{u}_1, \beta]^2$ of the first-two-period offers, recall that $\sigma^n_{|(a_1,a_2)}$ is the strategy profile induced on the subgame $\Gamma^n_{|(a_1,a_2)}$ by σ^n. By Corollaries 4 and 5, the strategy profile

$\sigma^n_{|(a_1,a_2)}$ is a Nash equilibrium of $\widehat{\Gamma}^n$ for P^{σ^n} almost every $(a_1,a_2) \in [\underline{u}_1,\beta]^2$. The value of Player 1 in the continuation game after he rejects some offer (a_1,a_2) in period 2 is

$$\eta \underline{u}_1 + (1-\eta) U_1^{\sigma^n}\left(\sigma_{|(a_1,a_2)}, \Gamma^n\right).$$

Thus this continuation value is upper bounded by $\bar{v}(\widehat{\Gamma}^n)$ for P^{σ^n} almost every (a_1,a_2).

Now we evaluate the "cost" to Player 1 from accepting offers in period 2, relative to his best response. We shall see that this cost is less than λ, for n sufficiently large. As $\bar{v}_1(\widehat{\Gamma}^n) \to \underline{u}_1$, there exists n_λ such that for all $n \geq n_\lambda$,

$$\bar{v}_1(\widehat{\Gamma}^n) < \underline{u}_1 + \lambda.$$

Let $r_2(Y)$ denote the time-2 pure response strategy that accepts all offers $(a_1,a_2) \in [\underline{u}_1,\beta]^2$ with certainty, and $r_2(N)$ the one that rejects all offers. Fix some $(a_1,a_2) \in [\underline{u}_1,\beta]^2$, recall that $B(a_{\leq 2}, \alpha_{12}(a_{\leq 2}))$ is the open ball centered at $a_{\leq 2}$ with radius $\alpha_{12}(a_{\leq 2}))$. Let

$$M(a_{\leq 2}) = B(a_{\leq 2}, \alpha_{12}(a_{\leq 2})) \cap \left([\underline{u}_1,\beta] \times \left[\underline{u}_1, \bar{v}_1(\widehat{\Gamma}^n)\right]\right).$$

For any subset $\tilde{M}(a_{\leq 2}) \subset M(a_{\leq 2})$, if the probability $\mathrm{P}^{\sigma^n}\left(H_{\tilde{M}(a_{\leq 2})}\right)$ that Player 1 receives an offer in the set $\tilde{M}(a_{\leq 2})$ in period 2 is strictly positive, then conditional on receiving such an offer, his expected payoff from accepting the offer is lower bounded by

$$\mathrm{E}^{\sigma^n / r_2(Y)}\left[u_1 \,|\, H_{\tilde{M}(a_{\leq 2})}\right] \geq \underline{u}_1,$$

while his expected payoff from rejecting the offer is upper bounded by

$$\mathrm{E}^{\sigma^n / r_2(N)}\left[u_1 \,|\, H_{\tilde{M}(a_{\leq 2})}\right] \leq \bar{v}_1(\widehat{\Gamma}^n).$$

Therefore conditional on an offer in $\tilde{M}(a_{\leq 2})$, the expected loss $L_i^{\sigma^n}(r_2(Y) \,|\, \tilde{M}(a_{\leq 2}))$ for Player 1 from accepting the offer is upper bounded by

$$L_i^{\sigma^n}(r_2(Y) \,|\, \tilde{M}(a_{\leq 2})) = \left\{ \mathrm{E}^{\sigma^n / r_2(N)}\left[u_1 \,|\, H_{\tilde{M}(a_{\leq 2})}\right] - \mathrm{E}^{\sigma^n / r_2(Y)}\left[u_1 \,|\, H_{\tilde{M}(a_{\leq 2})}\right] \right\}^+ \leq \bar{v}_1(\widehat{\Gamma}^n) - \underline{u}_1 < \lambda.$$

Now we evaluate the "cost" to Player 2 from accepting offers in period 1, when facing an offer $a_1 \in [\underline{u}_1 + 2\lambda, \beta]$. We shall see that accepting offers in this region costs more than λ to Player 2, for n sufficiently large. Recall that $\underline{v}_2(\widehat{\Gamma}_2^n)$ is given by (63), and we have $\underline{v}_2(\widehat{\Gamma}_2^n) \to \underline{u}_2$ from (64). There exists n_λ such that for all $n \geq n_\lambda$,

$$\underline{v}_2(\widehat{\Gamma}_2^n) + \underline{u}_1 - \pi^n > -\lambda.$$

Fix some $\tilde{a}_1 \in [\underline{u}_1 + 2\lambda, \beta]$, let

$$M(\tilde{a}_1) = B(\tilde{a}_1, \alpha_{21}(\tilde{a}_1)) \cap [\underline{u}_1 + 2\lambda, \beta].$$

For any subset $\tilde{M}(\tilde{a}_1) \subset M(\tilde{a}_1)$, if the probability $\mathrm{P}^{\sigma^n}(H_{\tilde{M}(\tilde{a}_1)})$ that Player 2 receives an offer in the set $\tilde{M}(\tilde{a}_1)$ in period 1 is strictly positive, then conditional on receiving such an offer, his expected payoff from accepting the offer is upper bounded by

$$\mathrm{E}^{\sigma^n / r_1(Y)} \left[u_2 \,|\, H_{\tilde{M}(\tilde{a}_1)} \right] \leq \pi^n - \underline{u}_1 - 2\lambda,$$

while his expected payoff from rejecting the offer is upper bounded by

$$\mathrm{E}^{\sigma^n / r_1(N)} \left[u_2 \,|\, H_{\tilde{M}(\tilde{a}_1)} \right] \geq \underline{v}_2(\widehat{\Gamma}_2^n).$$

Therefore conditional on an offer in $\tilde{M}(\tilde{a}_1)$, the expected loss $L_i^{\sigma^n}(r_1(Y) \,|\, \tilde{M}(\tilde{a}_1))$ for Player 2 from accepting the offer is lower bounded by

$$L_i^{\sigma^n}(r_1(Y) \,|\, \tilde{M}(\tilde{a}_1)) = \left\{ \mathrm{E}^{\sigma^n / r_1(N)} \left[u_2 \,|\, H_{\tilde{M}(\tilde{a}_1)} \right] - \mathrm{E}^{\sigma^n / r_1(Y)} \left[u_2 \,|\, H_{\tilde{M}(\tilde{a}_1)} \right] \right\}^+ \geq \underline{v}_2(\widehat{\Gamma}_2^n) - \pi^n + \underline{u}_1 + 2\lambda > \lambda.$$

As σ^n is a (λ, δ^n) extended proper equilibrium of Γ^n in extensive form, then for any $\tilde{M}(\tilde{a}_1) \subset M(a_1)$, $\tilde{M}(a_{\leq 2}) \subset M(a_{\leq 2})$ such that $\mathrm{P}^{\sigma^n}(H_{\tilde{M}(\tilde{a}_1)}) > 0$ and $\mathrm{P}^{\sigma^n}(H_{\tilde{M}(a_{\leq 2})}) > 0$, we have

$$\mathrm{P}^{\sigma^n} \left[H_1(Y) \,|\, H_{\tilde{M}(\tilde{a}_1)} \right] \leq \delta^n \mathrm{P}^{\sigma^n} \left[H_2(Y) \,|\, H_{\tilde{M}(a_{\leq 2})} \right]$$

where $H_t(Y) = \{(a, r, s) : r_t = Y\}$. That is, the probability that Player 2 accepts an offer in

$\tilde{M}(\tilde{a}_1)$, conditional on him receiving such an offer, is at most a multiple δ^n of the probability that Player 1 accepts an offer in $\tilde{M}(a_{\leq 2})$, conditional on him receiving such an offer. It then follows that

$$\mathrm{P}^{\sigma^n}\left[H_1(Y)\cap H_{\tilde{M}(\tilde{a}_1)}\right]\leq \delta^n \mathrm{P}^{\sigma^n}\left[H_2(Y)\,|\,H_{\tilde{M}(a_{\leq 2})}\right]\mathrm{P}^{\sigma^n}(H_{\tilde{M}(\tilde{a}_1)}).$$

This inequality holds even when $\mathrm{P}^{\sigma^n}(H_{\tilde{M}(\tilde{a}_1)})=0$, as both sides of the inequality are zero.

The collection $\{B(\tilde{a}_1,\alpha_{21}(\tilde{a}_1))\}_{a_1\in[\underline{u}_1+2\lambda,\beta]}$ of open balls forms an open cover of the compact set $[\underline{u}_{1+2\lambda},\beta]$, thus there exists an finite subcover

$$\cup_{\tilde{a}_1\in A_1}B(\tilde{a}_1,\alpha_{21}(\tilde{a}_1))=[\underline{u}_1+2\lambda,\beta]$$

for some finite set $A_1\subset[\underline{u}_1+2\lambda,\beta]$. This implies that $\cup_{\tilde{a}_1\in A_1}M(\tilde{a}_1)=[\underline{u}_1+2\lambda,\beta]$. Therefore by letting one can find a finite partition of the interval

$$\cup_{\tilde{a}_1\in A_1}\tilde{M}(\tilde{a}_1)[\underline{u}_1+2\lambda,\beta]$$

such that for each $a_1\in A_1$, $\tilde{M}(\tilde{a}_1)\subset M(\tilde{a}_1)$. Then

$$\begin{aligned}\mathrm{P}^{\sigma^n}\left[H_1(Y)\cap H_{[\underline{u}_1+2\lambda,\beta]}\right] &= \textstyle\sum_{a_1\in A_1}\mathrm{P}^{\sigma^n}\left[H_1(Y)\cap H_{\tilde{M}(\tilde{a}_1)}\right]\\ &\leq \delta^n\mathrm{P}^{\sigma^n}\left[H_2(Y)\,|\,H_{\tilde{M}(a_{\leq 2})}\right]\textstyle\sum_{a_1\in A_1}\mathrm{P}^{\sigma^n}\left(H_{\tilde{M}(\tilde{a}_1)}\right)\\ &\leq \delta^n\mathrm{P}^{\sigma^n}\left[H_2(Y)\,|\,H_{\tilde{M}(a_{\leq 2})}\right]\mathrm{P}^{\sigma^n}\left(H_{[\underline{u}_1+2\lambda,\beta]}\right).\end{aligned}$$

Thus for every Borel subset $\tilde{M}(a_{\leq 2})\subset M(a_{\leq 2})$, we have

$$\frac{1}{\delta^n}\mathrm{P}^{\sigma^n}\left[H_1(Y)\,|\,H_{[\underline{u}_1+2\lambda,\beta]}\right]\mathrm{P}^{\sigma^n}\left[H_{\tilde{M}(a_{\leq 2})}\right]\leq \mathrm{P}^{\sigma^n}\left[H_2(Y)\cap H_{\tilde{M}(a_{\leq 2})}\right].$$

This inequality holds even when $\mathrm{P}^{\sigma^n}\left[H_{\tilde{M}(a_{\leq 2})}\right]=0$. By considering a finite partition of the compact set $[\underline{u}_1,\beta]\times[\underline{u}_1,\underline{u}_1+2\lambda]$, we can obtain

$$\mathrm{P}^{\sigma^n}\left[H_1(Y)\,|\,H_{[\underline{u}_1+2\lambda,\beta]}\right]\leq \delta^n\mathrm{P}^{\sigma^n}\left[H_2(Y)\,|\,H_{[\underline{u}_1,\beta]\times[\underline{u}_1,\underline{u}_1+2\lambda]}\right]. \tag{66}$$

That is, the probability that Player 2 accepts an offer in $[\underline{u}_1 + 2\lambda, \beta]$, conditional on him receiving such an offer, is at most a multiple δ^n of the probability that Player 1 accepts an offer in $[\underline{u}_1 + 2\lambda, \beta]$ in period 2, conditional on him receiving such an offer. To simplify notation, we let

$$p^n = \mathrm{P}^{\sigma^n}\left[H_2(Y) \,|\, H_{[\underline{u}_1,\beta]\times[\underline{u}_1,\underline{u}_1+2\lambda]}\right]$$

Let H_2^c be the event that bargaining continues in period 2,

$$H_2^c = \{(a, r, s) : r_1 = N, s_1 = A\},$$

then we have $H_{[\underline{u}_1,\beta]\times[\underline{u}_1,\underline{u}_1+2\lambda]} \subset H_2^c$ by definition. By a similar argument, one can upper bound, as follows, the probability that Player 2 offers some $a_2 \in [\underline{u}_1 + 2\lambda, \beta]$ to Player 1 in period 2 conditional on the event H_2^c that bargaining continues in period 2, because offering such an a_2 is a costly deviation for Player 2:

$$\mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1,\beta]\times[\underline{u}_1,\underline{u}_1+2\lambda]} \,|\, H_2^c\right] \leq \delta^n p^n. \tag{67}$$

We upper bound the probability of H_2^c as follows:

$$\begin{aligned}\mathrm{P}^{\sigma^n}(H_2^c) = \mathrm{P}^{\sigma^n}(H_1(N))(1-\eta) \geq& (1-\eta)\mathrm{P}^{\sigma^n}\left(H_1(N) \,|\, H_{[\underline{u}_1+2\lambda,\beta]}\right)\mathrm{P}^{\sigma^n}\left(H_{[\underline{u}_1+2\lambda,\beta]}\right)\\ \geq& (1-\eta)(1-\delta^n p^n)\,\mathrm{P}^{\sigma^n}\left(H_{[\underline{u}_1+2\lambda,\beta]}\right).\end{aligned} \tag{68}$$

The last inequality follows from (66).

Now we upper bound the the expected utility of Player 1, conditional on the two players reaching an agreement in the first two periods of the bargaining game. The event that the two players reach an agreement in period 1 is simply $H_1(Y)$. The event that the two players reach an agreement in period 2 is $H_2^c \cap H_2(Y)$. Let $\mathfrak{a}_t$ be the projection mapping a complete

history (a, r, s) to the period-t offer a_t. Thus

$$\begin{aligned}
& \mathrm{E}^{\sigma^n}(u_1 \,|\, \{\text{The two players reach an agreement in the first two periods}\}) \\
= \; & \frac{\mathrm{E}^{\sigma^n}\left[\mathfrak{a}_1 \mathbb{1}_{H_1(Y)}\right] + \mathrm{E}^{\sigma^n}\left[\mathfrak{a}_2 \mathbb{1}_{H_2^c \cap H_2(Y)}\right]}{\mathrm{P}^{\sigma^n}[H_1(Y)] + \mathrm{P}^{\sigma^n}[H_2^c \cap H_2(Y)]} \\
& \mathrm{E}^{\sigma^n}\left[\mathfrak{a}_1 \mathbb{1}_{H_1(Y)}\right] \\
\leq \; & (\underline{u}_1 + 2\lambda)\, \mathrm{P}^{\sigma^n}\left[H_1(Y) \,|\, H_{[\underline{u}_1, \underline{u}_1 + 2\lambda)}\right]\, \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1, \underline{u}_1 + 2\lambda)}\right] \\
& + \beta\, \mathrm{P}^{\sigma^n}\left[H_1(Y) \,|\, H_{[\underline{u}_1 + 2\lambda, \beta]}\right]\, \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1 + 2\lambda, \beta]}\right] \\
& \mathrm{E}^{\sigma^n}\left[\mathfrak{a}_2 \mathbb{1}_{H_2^c \cap H_2(Y)}\right] \\
\leq \; & (\underline{u}_1 + 2\lambda)\, \underbrace{\mathrm{P}^{\sigma^n}\left[H_2(Y) \,|\, H_{[\underline{u}_1, \beta] \times [\underline{u}_1, \underline{u}_1 + 2\lambda]}\right]}_{p^n}\, \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1, \beta] \times [\underline{u}_1, \underline{u}_1 + 2\lambda]}\right] \\
& + \beta\, \mathrm{P}^{\sigma^n}\left[H_2(Y) \,|\, H_{[\underline{u}_1, \beta] \times [\underline{u}_1 + 2\lambda, \beta]}\right]\, \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1, \beta] \times [\underline{u}_1 + 2\lambda, \beta]}\right]
\end{aligned}$$

By letting

$$\begin{aligned}
w_1^n &= \mathrm{P}^{\sigma^n}\left[H_1(Y) \,|\, H_{[\underline{u}_1, \underline{u}_1 + 2\lambda)}\right]\, \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1, \underline{u}_1 + 2\lambda)}\right] + p^n \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1, \beta] \times [\underline{u}_1, \underline{u}_1 + 2\lambda]}\right], \\
w_2^n &= \mathrm{P}^{\sigma^n}\left[H_1(Y) \,|\, H_{[\underline{u}_1 + 2\lambda, \beta]}\right]\, \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1 + 2\lambda, \beta]}\right] + \mathrm{P}^{\sigma^n}\left[H_2(Y) \,|\, H_{[\underline{u}_1, \beta] \times [\underline{u}_1 + 2\lambda, \beta]}\right] \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1, \beta] \times [\underline{u}_1 + 2\lambda, \beta]}\right],
\end{aligned}$$

we have

$$\mathrm{E}^{\sigma^n}(u_1 \,|\, \{\text{The two players reach an agreement in the first two periods}\}) \leq \frac{(\underline{u}_1 + 2\lambda) w_1^n + \beta w_2^n}{w_1 + w_2}.$$

It follows from (66), (67) and (68) that

$$\begin{aligned}
& w_1^n \geq p^n \mathrm{P}^{\sigma^n}\left[H_{[\underline{u}_1, \beta] \times [\underline{u}_1, \underline{u}_1 + 2\lambda]} \,|\, H_2^c\right]\, \mathrm{P}^{\sigma^n}(H_2^c) \geq p^n (1 - \delta^n p^n) \mathrm{P}^{\sigma^n}(H_2^c) \\
& w_2^n \leq \delta^n p^n\, \mathrm{P}^{\sigma^n}(H_2^c)\, [1 + 1/(1 - \delta^n p^n)(1 - \eta)]
\end{aligned}$$

Since a weighted average of x and y $(x \leq y)$ increases with the weight on y and decreases with the weight on x: $\frac{x w_1 + y w_2}{w_1 + w_2} \leq \frac{x \tilde{w}_1 + y \tilde{w}_2}{\tilde{w}_1 + \tilde{w}_2}$ if $\tilde{w}_1 \leq w_1$, $\tilde{w}_2 \geq w_2$, we obtain

$$\begin{aligned}
& \mathrm{E}^{\sigma^n}(u_1 \,|\, \{\text{The two players reach an agreement in the first two periods}\}) \\
\leq \; & \frac{(\underline{u}_1 + 2\lambda)(1 - \delta^n p^n) + \beta \delta^n [1 + 1/(1 - \delta^n p^n)(1 - \eta)]}{(1 - \delta^n p^n) + \delta^n [1 + 1/(1 - \delta^n p^n)(1 - \eta)]} \xrightarrow{n \to \infty} \underline{u}_1 + 2\lambda.
\end{aligned}$$

Thus for every $\epsilon > 0$, there exists n_ϵ such that for all $n \geq n_\epsilon$,

$$\mathrm{E}^{\sigma^n}(u_1 \,|\, \{\text{The two players reach an agreement in the first two periods}\}) \leq \underline{u} + 2\lambda + \epsilon$$

Likewise, one can show that for every $t \geq 1$,

$$\mathrm{E}^{\sigma^n}(u_1 \,|\, \{\text{The two players reach an agreement in period } 2t-1 \text{ or } 2t\}) \leq \underline{u} + 2\lambda + \epsilon$$

for all $n \geq n_\epsilon$. Thus $U_1(\sigma^n; \Gamma^n \,|\, agreement) \leq \underline{u} + 2\lambda + \epsilon$ for all $n \geq n_\epsilon$, implying

$$\limsup_{n\to\infty} U_1(\sigma^n; \Gamma^n \,|\, agreement) \leq \underline{u} + 2\lambda$$

This completes the proof. □

We further introduce the following lemma, which will be useful in determining equilibrium choice of contingent treatment in Stage a and Stage b in our three node game. Formally, let $\Pi = \{\pi_1, \cdots, \pi_l\}$ be a set of total payoffs. Without loss of generality, suppose π_1 is the unique maximum value of Π. Suppose for each $j = 1, \cdots, l$, there are two sequences $\left(v_{1j}^n\right)_{n\geq 0}$ and $\left(v_{2j}^n\right)_{n\geq 0}$ such that

$$v_{ij}^n \to v_{ij} := v_i(\pi_j - \underline{u}_1 - \underline{u}_2)^+ + \underline{u}_i$$

for $i = 1, 2$. Suppose there are two sequences of outside option values $\underline{u}_i^n$ converging to $\underline{u}_i$ $(i = 1, 2)$.

For each $n \geq 0$, consider the following alternating game $\Gamma\left[\left(v_{ij}^n\right)_{\substack{i=1,2\\ j=1,\cdots,l}}, (\underline{u}_i^n)_{i=1,2}\right]$: In each period, one of the players proposes a choice of total payoff $\pi_j \in \Pi$. The other player accepts or rejects this proposal. Acceptance ends the bargaining and the agreement is implemented. Rejection leads, with some given probability η, to a breakdown of the negotiation. Absent breakdown, the game proceeds to the next period, when offers are made in alternating order. If some agreement $\pi_j \in \Pi$ is implemented, then Player i's payoff is v_{ij}^n $(i = 1, 2)$; Otherwise breakdown leads to a payoff of $\underline{u}_i^n$ for Player i $(i = 1, 2)$. We suppose Player 1 is the first to propose. This is a Rubinstein's alternating offer game, with finite a set of possible

agreements. For simplicity of notation, we write Γ^n for the game $\Gamma\left[\left(v_{ij}^n\right)_{\substack{i=1,2\\ j=1,\cdots,l}}, (\underline{u}_i^n)_{i=1,2}\right]$ when stating the next lemma, and Γ for $\Gamma\left[(v_{ij})_{\substack{i=1,2\\ j=1,\cdots,l}}, (\underline{u}_i)_{i=1,2}\right]$.

Lemma 13. *Under the above setting and notation for Γ^n and Γ,*

- *If $\pi_1 > \underline{u}_1 + \underline{u}_2$, then there is an $\eta^* > 0$, such that for every breakdown probability $\eta < \eta^*$, there exists a unique strategy profile σ^+ in $\mathcal{R}((\Gamma^n)_{n\geq 0})$. It consists of the following: both players propose π_1 in each period, and accept (and only accept) π_1. The payoff of player i in the game Γ is $U_i(\sigma;\Gamma) = v_i(\pi_1 - \underline{u}_1 - \underline{u}_2) + \underline{u}_i$ $(i = 1, 2)$.*

- *If $\pi_1 < \underline{u}_1 + \underline{u}_2$, then for every breakdown probability $\eta \in (0, 1]$, there exists a strategy profile σ^- in $\mathcal{R}((\Gamma^n)_{n\geq 0})$: both players propose π_1 in each period, and reject all proposals (including π_1). Every strategy profile in $\mathcal{R}((\Gamma^n)_{n\geq 0})$ induces the same outcome as σ^-, in which alternating offers in every period are rejected and the game ends when bargaining breaks down, and both players get their respective outside option values $\underline{u}_i$ $(i = 1, 2)$.*

- *In both cases, for each sequence $\sigma^n \in \mathcal{T}[\sigma; (\Gamma^n)_{n\geq 0}]$, the payoff of Player i under σ^n in the game Γ^n satisfies $U_i(\sigma^n;\Gamma^n) \to U_i(\sigma;\Gamma)$, where σ is σ^+ or σ^-, depending on whether $\pi_1 > \underline{u}_1 + \underline{u}_2$ or $\pi_1 < \underline{u}_1 + \underline{u}_2$.*

One can easily adapt the proof of Theorem 3 and lemma 12 to show the lemma above. The proof is indeed much simpler, as the set of possible agreements in each Rubinstein's game Γ^n is finite rather than continuum. Moreover, because of the discrete nature of action sets, one obtains uniqueness in the strategy profile $\sigma \in \mathcal{R}((\Gamma^n)_{n\geq 0})$ when $\pi > \underline{u}_1 + \underline{u}_2$, a stronger result than uniqueness in outcome. We omit the proof to avoid repeating the same argument that we provided earlier.

Now we are ready to prove Theorem 1.

Proof of Theorem 1. We proceed in four steps in each of the existence and uniqueness proofs:

Existence: We need to show that the strategy profile $x^{*\eta}$ is a restricted equilibrium for the three-node bargaining game $\Gamma_{\text{3-node}}(\eta, S, C, f, s^0)$. Recall that $\sigma^{*\eta}$ is defined in Section 8.2.

To this end, we construct a restricted trembling sequence $(\sigma^n)_{n\geq 0}$, by "pasting" together sequences of restricted trembling equilibria in the Rubinstein's alternating offer game across every stage.

Step 1: We first consider the negotiation between Nodes 1 and 2 in Stage aa over $y_1^b\left(s_1^b, s_2^{3b}\right)$ - the payment from Node 2 to Node 1 associated with some choice of treatments $\left(s_1^b, s_2^{3b}\right) \in \mathcal{S}_{1,2}^B$ for the contingency $B_{2,3}$, if Nodes 1 and 2 agreed on the pair of treatments $\left(s_1^b, s_2^{3b}\right)$ in Stage a. The negotiation process is identical to that of the Rubinstein's alternating-offers game $\Gamma\left(\pi_1^b\left(s_1^b, s_2^{3b}\right), 0, 0\right)$, where[17]

$$\pi_1^b\left(s_1^b, s_2^{3b}\right) = f_1\left(s_1^b, s_2^{3b}\right) + f_2\left(s_1^b, s_2^{3b}, s_3^0\right). \tag{69}$$

For simplicity, we will write π_1^b for $\pi_1^b\left(s_1^b, s_2^{3b}\right)$. Therefore we need to show that the strategy profile $\sigma^{*\eta}_{aa,\left(s_1^b, s_2^{3b}\right)}$ - the restriction of $\sigma^{*\eta}$ to Stage aa and the choice of breakdown treatments $\left(s_1^b, s_2^{3b}\right)$ at the contingency $B_{2,3}$ - is a restricted equilibrium in the game $\Gamma\left(\pi_1^b, 0, 0\right)$, and construct a restricted trembling sequence $\left(\sigma^n_{aa,\left(s_1^b, s_2^{3b}\right)}\right)_{n\geq 0}$ for $\sigma^{*\eta}_{aa,\left(s_1^b, s_2^{3b}\right)}$.

If $\pi_1^b \leq 0$, then by Lemma 12, $\sigma^*\left(\pi_1^b, 0, 0\right)$ is a restricted equilibrium for the game $\Gamma\left(\pi_1^b, 0, 0\right)$. Recall that the strategy profile $\sigma^*\left(\pi_1^b, 0, 0\right)$ is defined in Lemma 12. There exists a restricted trembling sequence $\sigma^n_{aa,\left(s_1^b, s_2^{3b}\right)}$ for $\sigma^*\left(\pi_1^b, 0, 0\right)$ such that

$$U_i\left(\sigma^n_{aa,\left(s_1^b, s_2^{3b}\right)}; \Gamma\left(\pi_1^b, 0, 0\right)\right) \to 0, \qquad i = 1, 2.$$

Following equations (10a) and (10b), we note that $\sigma^*\left(\pi_1^b, 0, 0\right)$ consists of Node 1 asking for the maximum payment $\bar{y}_1^B\left(s_1^b, s_2^{3b}\right)$ from Node 2, and Node 2 offering the minimum payment $\underline{y}_1^B\left(s_1^b, s_2^{3b}\right)$ to Node 1 in each round. Node 1 accepts offers $y_1^b\left(s_1^b, s_2^{3b}\right)$ that are strictly larger than $\underline{y}_1^B\left(s_1^b, s_2^{3b}\right)$, while Node 2 rejects all offers. Therefore $\sigma^*\left(\pi_1^b, 0, 0\right) = \sigma^{*\eta}_{aa,\left(s_1^b, s_2^{3b}\right)}$.

If $\pi_1^b > 0$, then by Corollary 8, $\sigma^*\left(\pi_1^b, 0, 0\right)$ is a restricted equilibrium for the game $\Gamma\left(\pi_1^b, 0, 0\right)$. Recall that the strategy profile $\sigma^*\left(\pi_1^b, 0, 0\right)$ is defined in Corollary 8. There

[17] This is so because breakdown leads to $f_1\left(s_1^0, s_2^0\right) = f_2\left(s_1^0, s_2^0, s_3^0\right) = 0$ in payoff for both Nodes 1 and 2, and agreement leads to total a total payoff of $f_1\left(s_1^b, s_2^{3b}\right) + f_2\left(s_1^b, s_2^{3b}, s_3^0\right)$ for the two nodes to share.

exists a restricted trembling sequence $\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)}$ for $\sigma^*\left(\pi_1^b,0,0\right)$ such that

$$U_i\left(\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)};\Gamma\left(\pi_1^b,0,0\right)\right)\to v_i\pi_1^b,\qquad i=1,2.$$

Following equations (21) and (22), we note that $\sigma^*\left(\pi_1^b,0,0\right)$ consists of Node 1 asking for $y_1^{B\eta}\left(s_1^b,s_2^{3b}\right)$ from Node 2, and Node 2 offering $\tilde{y}_1^{B\eta}\left(s_1^b,s_2^{3b}\right)$ to Node 1; Node 1 accepts offers that are at least $\tilde{y}_1^{B\eta}\left(s_1^b,s_2^{3b}\right)$, while Node 2 accepts offers that are at most $y_1^{B\eta}\left(s_1^b,s_2^{3b}\right)$. Again we have $\sigma^*\left(\pi_1^b,0,0\right)=\sigma^{*\eta}_{aa,\left(s_1^b,s_2^{3b}\right)}$.

Combining the two cases of $\pi_1^b\leq 0$ and $\pi_1^b>0$, we conclude that $\sigma^{*\eta}_{aa,\left(s_1^b,s_2^{3b}\right)}$ is a restricted equilibrium in the game $\Gamma\left(\pi_1^b,0,0\right)$, and we have a restricted trembling sequence $\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)}$ for $\sigma^{*\eta}_{aa,\left(s_1^b,s_2^{3b}\right)}$ such that

$$U_i\left(\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)};\Gamma\left(\pi_1^b,0,0\right)\right)\to v_i\cdot\left(\pi_1^b\right)^+,\qquad i=1,2. \tag{70}$$

Likewise for the negotiation between Nodes 2 and 3 in Stage bb over $y_3^b\left(s_2^{1b},s_3^b\right)$, the bargaining process is identical to that of the game $\Gamma\left(\pi_3^b\left(s_2^{1b},s_3^b\right),0,0\right)$, where

$$\pi_3^b\left(s_2^{1b},s_3^b\right):=f_2\left(s_1^0,s_2^{1b},s_3^b\right)+f_3\left(s_2^{1b},s_3^b\right) \tag{71}$$

The strategy profile $\sigma^{*\eta}_{bb,\left(s_2^{1b},s_3^b\right)}$ is a restricted equilibrium for the game $\Gamma\left(\pi_1^b,0,0\right)$, and we have a restricted trembling sequence $\sigma^n_{bb,\left(s_2^{1b},s_3^b\right)}$ for $\sigma^{*\eta}_{bb,\left(s_2^{1b},s_3^b\right)}$ such that

$$U_i\left(\sigma^n_{bb,\left(s_2^{1b},s_3^b\right)};\Gamma\left(\pi_3^b,0,0\right)\right)\to v_i\cdot\left(\pi_3^b\right)^+,\qquad i=2,3, \tag{72}$$

where $v_3:=v_1=\dfrac{1}{2-\eta}$.

Step 2: We next consider the negotiation between Nodes 1 and 2 in Stage a over $\left(s_1^b,s_2^{3b}\right)\in\mathcal{S}^B_{1,2}$ - the choice of treatments for the contingency $B_{2,3}$. Given that Nodes 1 and 2 will follow the strategy profile $\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)}$ in Stage aa, implementation of any pair of treatments $\left(s_1^b,s_2^{3b}\right)$ leads to a payoff of $v_i^n\left(s_1^b,s_2^{3b}\right):=U_i\left(\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)};\Gamma\left(\pi_1^b,0,0\right)\right)$ for Node i $(i=1,2)$, where by

equation (70),

$$v_i^n\left(s_1^b, s_2^{3b}\right) \xrightarrow{n\to\infty} v_i \cdot \left(\pi_1^b\left(s_1^b, s_2^{3b}\right)\right)^+, \qquad i = 1, 2.$$

On the other hand, a breakdown of the negotiation over $\left(s_1^b, s_2^{3b}\right) \in \mathcal{S}_{1,2}^B$ would give both nodes their respective outside option values 0. Thus we are in the setting of Lemma 13, where $l = \left|\mathcal{S}_{1,2}^B\right|$, $\Pi = \left\{\pi_1^b\left(s_1^b, s_2^{3b}\right) : \left(s_1^b, s_2^{3b}\right) \in \mathcal{S}_{1,2}^B\right\}$, $v_{ij}^n = v_i^n\left(s_1^b, s_2^{3b}\right)$, and $\underline{u}_i^n = 0$, for $i = 1, 2$, $j = 1, \cdots, l$, $n \geq 0$. For simplicity of notation, let

$$\Gamma_{a,B_{2,3}}^n := \Gamma\left[\left(v_{ij}^n\right)_{\substack{i=1,2\\ j=1,\cdots,l}}, \underline{u}_1^n = 0, \underline{u}_2^n = 0\right]$$

and

$$\Gamma_{a,B_{2,3}} = \Gamma\left[\left(v_{ij}\right)_{\substack{i=1,2\\ j=1,\cdots,l}}, \underline{u}_1 = 0, \underline{u}_2 = 0\right].$$

We need to show that the strategy profile $\sigma_{a,B_{2,3}}^{*\eta}$ is in the set $\mathcal{R}\left(\left(\Gamma_{a,B_{2,3}}^n\right)_{n\geq 0}\right)$, and construct a sequence $\left(\sigma_{a,B_{2,3}}^n\right)_{n\geq 0} \in \mathcal{T}\left[\sigma_{a,B_{2,3}}^{*\eta}; \left(\Gamma_{a,B_{2,3}}^n\right)_{n\geq 0}\right]$. Since

$$\underset{\left(s_1^b, s_2^{3b}\right)\in\mathcal{S}_{1,2}^B}{\operatorname{argmax}} \ \pi_1^b\left(s_1^b, s_2^{3b}\right) = \left(s_1^B, s_2^{3B}\right)$$

and $\pi_1^b\left(s_1^B, s_2^{3B}\right) > \pi_1^b\left(s_1^0, s_2^0\right) = 0$. Hence by Lemma 13, there exists a strategy profile in $\mathcal{R}((\Gamma_{a,B_{2,3}}^n)_{n\geq 0})$, which is that both nodes propose $\left(s_1^B, s_2^{3B}\right)$ in each period, and accept (and only accept) $\left(s_1^B, s_2^{3B}\right)$. This strategy profile is identical to the candidate strategy profile $\sigma_{a,B_{2,3}}^{*\eta}$. Under $\sigma_{a,B_{2,3}}^{*\eta}$, the payoff of node i in $\Gamma_{a,B_{2,3}}$ is $U_i\left(\sigma_{a,B_{2,3}}^{*\eta}; \Gamma_{a,B_{2,3}}\right) = v_i \cdot \pi_1^b\left(s_1^B, s_2^{3B}\right)$ $(i = 1, 2)$. Note that by equation (18), the payoff of Node 2

$$v_2 \cdot \pi_1^b\left(s_1^B, s_2^{1B}\right) = \underline{u}_{23}^\eta \tag{73}$$

where $\underline{u}_{23}^\eta$ is, we recall, the outside option value of Node 2 in its bargaining with Node 3. Also by Lemma 13, there exists a sequence $\sigma_{a,B_{2,3}}^n \in \mathcal{T}\left(\sigma_{a,B_{2,3}}^{*\eta}; \left(\Gamma_{a,B_{2,3}}^n\right)_{n\geq 0}\right)$, such that

$$U_i\left(\sigma_{a,B_{2,3}}^n; \Gamma_{a,B_{2,3}}^n\right) \to U_i\left(\sigma_{a,B_{2,3}}^{*\eta}; \Gamma_{a,B_{2,3}}\right) = v_i \cdot \pi_1^b\left(s_1^B, s_2^{3B}\right) \tag{74}$$

The same analysis applies to the negotiation between Nodes 2 and 3 in Stage b over $\left(s_2^{1b}, s_3^b\right) \in \mathcal{S}_{2,3}^B$. We define $\sigma_{b,B_{1,2}}^n$ in a similar way.

For every n, the strategy profiles $\left(\sigma_{a,B_{2,3}}^n, \sigma_{b,B_{1,2}}^n, \left[\sigma_{aa,\left(s_1^b, s_2^{3b}\right)}^n\right]_{\left(s_1^b, s_2^{3b}\right) \in \mathcal{S}_{1,2}^B}, \left[\sigma_{bb,\left(s_2^{1b}, s_3^b\right)}^n\right]_{\left(s_2^{1b}, s_3^b\right) \in \mathcal{S}_{2,3}^B}\right)$ determine the strategies for all three nodes at the breakdown contingency in the four stages of bargaining game. Let $\sigma_{\{B_{1,2}, B_{2,3}\}}^n$ denote this combined strategy profile.

Remark 5. The restriction of the strategy profile $\sigma^{*\eta}$ to the breakdown contingencies $B_{1,2}$ and $B_{2,3}$ fixes the outside option values in the three-node bargaining game $\Gamma_{\text{3-node}}(\eta, S, C, f, s^0)$, for all the bilateral bargaining problems in all stages and across all contingencies. By equation (73), $\left(\underline{u}_{ij}^\eta\right)$ is the outside option value of Node i in its bargaining against Node j. Along the converging sequence of trembling equilibria $\sigma_{\{B_{1,2}, B_{2,3}\}}^n$, the corresponding outside option values $\underline{u}_{ij}^n$ converge to $\underline{u}_{ij}^\eta$ by equation (74).

Step 3: Now consider the negotiation between Nodes 1 and 2 in Stage aa over $y_1(s_1; s_2, s_3)$ - the payment from Node 2 to Node 1 associated with some choice of treatment $s_1 \in C_1(s_2)$ for the contingency $(s_2, s_3) \in \mathcal{S}_{2,3}$. Following Remark 5, we already know that the outside option values of Nodes 1 and 2 are $\underline{u}_{12}^n$ and $\underline{u}_{21}^n$ respectively along the sequence of $\mu_{\{B_{1,2}, B_{2,3}\}}^n$, and

$$\underline{u}_{12}^n \to \underline{u}_{12}^\eta; \qquad \underline{u}_{21}^n \to \underline{u}_{21}^\eta. \tag{75}$$

For each $(s_1, s_2, s_3) \in \mathcal{S}$, let

$$\pi_1(s_1; s_2, s_3) = f_1(s_1, s_2) + f_2(s_1, s_2, s_3) - y_3^\eta(s_2, s_3), \tag{76}$$

where y_3^η is defined in equation (17).

Claim 1. We claim that $\pi_1(s_1; s_2, s_3)$ is the total payoff to be shared by Nodes 1 and 2 if the two nodes agreed on the treatment s_1 for the contingency (s_2, s_3) in Stage a.

This is the case if and only if the expected payment from Node 2 to Node 3, conditional on the two nodes agreeing on (s_2, s_3) in Stage b and Stage bb, does not depend on n and

is equal to $y_3^\eta(s_2, s_3)$. When we later complete the construction of the sequence σ^n, we will verify that this claim is correct. Note that $\pi_1(s_1; s_2, s_3)$ does not depend on n.

Let $\Gamma^n_{aa,(s_1;s_2,s_3)} = \Gamma\left(\pi_1(s_1; s_2, s_3), \underline{u}^n_{12}, \underline{u}^n_{21}\right)$. If we assume that Claim 1 holds, then we need to show that $\sigma^{*\eta}_{aa,(s_1;s_2,s_3)} \in \mathcal{R}\left[\left(\Gamma^n_{aa,(s_1;s_2,s_3)}\right)_{n\geq 0}\right]$, and construct a sequence $\left(\sigma^n_{aa,(s_1;s_2,s_3)}\right)_{n\geq 0} \in \mathcal{T}\left[x^{*\eta}_{aa,(s_1;s_2,s_3)}; (\Gamma^n_{aa,(s_1;s_2,s_3)})_{n\geq 0}\right]$. Following the same argument as in *Step 1*, Lemma 11 and Lemma 12 imply that

$$\sigma^{*\eta}_{aa,(s_1;s_2,s_3)} \in \mathcal{R}\left[\left(\Gamma^n_{aa,(s_1;s_2,s_3)}\right)_{n\geq 0}\right], \tag{77}$$

and there exists a sequence $\left(\sigma^n_{aa,(s_1;s_2,s_3)}\right)_{n\geq 0} \in \mathcal{T}\left[\sigma^{*\eta}_{aa,(s_1;s_2,s_3)}, \left(\Gamma^n_{aa,(s_1;s_2,s_3)}\right)_{n\geq 0}\right]$, such that

$$U_i\left(\sigma^n_{aa,(s_1;s_2,s_3)}; \Gamma^n_{aa,(s_1;s_2,s_3)}\right) \xrightarrow{n\to\infty} v_i \cdot (\pi_1(s_1; s_2, s_3) - \underline{u}^\eta_{12} - \underline{u}^\eta_{21})^+ + \underline{u}^\eta_{i(3-i)}, \qquad i = 1, 2. \tag{78}$$

$$U_1\left(\sigma^n_{aa,(s_1;s_2,s_3)}; \Gamma^n_{aa,(s_1;s_2,s_3)} \mid \text{agreement}\right) = v_1 \cdot (\pi_1(s_1; s_2, s_3) - \underline{u}^\eta_{12} - \underline{u}^\eta_{21})^+ + \underline{u}^\eta_{12}. \tag{79}$$

We next consider the negotiation between Nodes 1 and 2 in Stage a over $s_1 \in C_1(s_2)$ - the choice of treatment for the contingency (s_2, s_3). Given that Nodes 1 and 2 will play $\left(\sigma^n_{aa,(s_1;s_2,s_3)}\right)$ in Stage aa, implementation of any choice of treatment s_1 for the contingency (s_2, s_3) leads to a payoff of $v^n_i(s_1; s_2, s_3) := U_i\left(\sigma^n_{aa,(s_1;s_2,s_3)}; \Gamma^n_{aa,(s_1;s_2,s_3)}\right)$ for Node i $(i = 1, 2)$, where by equation (78),

$$v^n_i(s_1; s_2, s_3) \xrightarrow{n\to\infty} v_i \cdot (\pi_1(s_1; s_2, s_3) - \underline{u}^\eta_{12} - \underline{u}^\eta_{21})^+ + \underline{u}^\eta_{i(3-i)}, \qquad i = 1, 2.$$

Thus we are in the setting of Lemma 13, where $l = |C_1(s_2)|$, $\Pi = \{\pi_1(s_1; s_2, s_3) : s_1 \in C_1(s_2)\}$, $v^n_{ij} = v^n_i(s_1; s_2, s_3)$, and $\underline{u}^n_i = \underline{u}^n_{i(3-i)}$, for $i = 1, 2$, $j = 1, \cdots, l$, $n \geq 0$. Let

$$\Gamma^n_{a,(s_2,s_3)} := \Gamma\left[\left(v^n_{ij}\right)_{\substack{i=1,2\\ j=1,\cdots,l}}, \underline{u}^n_1 = \underline{u}^n_{12}, \underline{u}^n_2 = \underline{u}^n_{21}\right]$$

and

$$\Gamma_{a,(s_2,s_3)} = \Gamma\left[\left(v_{ij}\right)_{\substack{i=1,2\\ j=1,\cdots,l}}, \underline{u}_1 = \underline{u}^\eta_{12}, \underline{u}_2 = \underline{u}^\eta_{21}\right].$$

Following the same argument as in *Step 2*, Lemma 13 implies that

$$\sigma^{*\eta}_{a,(s_2,s_3)} \in \mathcal{R}\left[\left(\Gamma^n_{a,(s_2,s_3)}\right)_{n\geq 0}\right] \tag{80}$$

and there exists a sequence $\left(\sigma^n_{a,(s_2,s_3)}\right)_{n\geq 0} \in \mathcal{T}\left[\sigma^{*\eta}_{a,(s_2,s_3)}; \left(\Gamma^n_{a,(s_2,s_3)}\right)_{n\geq 0}\right]$ such that

$$\begin{aligned}
& U_1\left(\sigma^n_{\{a,aa\},(s_2,s_3)}; \Gamma^n_{a,(s_2,s_3)} \mid \text{agreement}\right) \\
\xrightarrow{n\to 0} \quad & v_1 \cdot [\pi_1(s^*_1(s_2,s_3); s_2, s_3) - \underline{u}^\eta_{12} - \underline{u}^\eta_{21}] + \underline{u}^\eta_{12} \qquad \text{by equation (79)} \\
= \quad & f_1(s^*(s_2,s_3), s_2) + y^\eta_1(s^*_1(s_2,s_3); s_2, s_3) \qquad \text{by equations (16) and (76)}
\end{aligned} \tag{81}$$

Recall that s^*_1 is defined in Section 8.1.

Step 4: We next consider the negotiation between Nodes 2 and 3 in Stage bb over $y_3(s_2, s_3)$ - the payment from Node 2 to Node 3 associated with some choice of treatment $(s_2, s_3) \in \mathcal{S}_{2,3}$. Following Remark 5, we already know that the outside option values of Nodes 2 and 3 are $\underline{u}^n_{23}$ and $\underline{u}^n_{32}$ respectively along the sequence of $\sigma^n_{\{B_{1,2},B_{2,3}\}}$, and

$$\underline{u}^n_{23} \to \underline{u}^\eta_{23}; \qquad \underline{u}^n_{32} \to \underline{u}^\eta_{32}. \tag{82}$$

Given that Nodes 1 and 2 play $\sigma^n_{\{a,aa\}}$ in Stage a and Stage aa, the total payoff $\pi^n_3(s_2, s_3)$ to be shared by Nodes 2 and 3 satisfies

$$\begin{aligned}
& \lim_{n\to\infty} \pi^n_3(s_2,s_3) = U(s^*_1(s_2,s_3), s_2, s_3) - \lim_{n\to\infty} U_1\left(\sigma^n_{\{a,aa\},(s_2,s_3)}; \Gamma^n_{a,(s_2,s_3)} \mid \text{agreement}\right) \\
= \quad & U(s^*_1(s_2,s_3), s_2, s_3) - [f_1(s^*(s_2,s_3), s_2) + y^\eta_1(s^*_1(s_2,s_3); s_2, s_3)] \\
= \quad & \underbrace{(2-\eta)\left[f_3(s_2,s_3) + y^\eta_3(s_2,s_3)\right] - (1-\eta)\underline{u}^\eta_{32} + \underline{u}^\eta_{23}}_{:=\pi_3(s_2,s_3)}
\end{aligned}$$

The first equality follows from equation (81), and the second from equation (17). Let $\Gamma^n_{bb,(s_2,s_3)} = \Gamma\left(\pi^n_3(s_2,s_3), \underline{u}^n_{32}, \underline{u}^n_{23}\right)$. Following the same argument as in *Step 1*, Lemma 11 and Lemma 12 imply that

$$\sigma^{*\eta}_{bb,(s_2,s_3)} \in \mathcal{R}\left[\left(\Gamma^n_{bb,(s_2,s_3)}\right)_{n\geq 0}\right], \tag{83}$$

and there exists a sequence $\left(\sigma^n_{bb,(s_2,s_3)}\right)_{n\geq 0} \in \mathcal{T}\left[\sigma^{*\eta}_{bb,(s_2,s_3)}; \left(\Gamma^n_{bb,(s_2,s_3)}\right)_{n\geq 0}\right]$, such that

$$U_i\left(\sigma^n_{bb,(s_2,s_3)}; \Gamma^n_{bb,(s_2,s_3)}\right) \xrightarrow{n\to\infty} v_i \cdot (\pi_3(s_2,s_3) - \underline{u}^\eta_{32} - \underline{u}^\eta_{23})^+ + \underline{u}^\eta_{i(5-i)}, \qquad i = 2,3. \tag{84}$$

$$\begin{aligned} U_3\left(\sigma^n_{bb,(s_2,s_3)}, \Gamma^n_{bb,(s_2,s_3)} \mid \text{agreement}\right) &= v_1 \cdot (\pi_3(s_2,s_3) - \underline{u}^\eta_{32} - \underline{u}^\eta_{23})^+ + \underline{u}^\eta_{32} \\ &= f_3(s_2,s_3) + y^\eta_3(s_2,s_3) \end{aligned} \tag{85}$$

Equation (85) shows that the expected payment from Node 2 to Node 3, conditional on the two nodes agreeing on (s_2, s_3) in Stage b and Stage bb, does not depend on n and is equal to $y^\eta_3(s_2,s_3)$. Hence we have verified Claim 1.

Finally we consider the negotiation between Nodes 2 and 3 in Stage b over $(s_2,s_3) \in \mathcal{S}_{2,3}$. Given that Nodes 2 and 3 will play $\left(\mu^n_{bb,(s_2,s_3)}\right)$ in Stage bb, implementation of any choice of treatments (s_2,s_3) leads to a payoff of $v^n_i(s_2,s_3) := U_i\left(\sigma^n_{bb,(s_2,s_3)}; \Gamma^n_{bb,(s_2,s_3)}\right)$ for Node i $(i = 2,3)$, where by equation (84),

$$v^n_i(s_2,s_3) \xrightarrow{n\to\infty} v_i \cdot (\pi_3(s_2,s_3) - \underline{u}^\eta_{32} - \underline{u}^\eta_{23})^+ + \underline{u}^\eta_{i(5-i)}, \qquad i = 1,3.$$

Thus we are in the setting of Lemma 13, where $l = |\mathcal{S}_{2,3}|$, $\Pi = \{\pi_3(s_2,s_3) : (s_2,s_3) \in \mathcal{S}_{2,3}\}$, $v^n_{ij} = v^n_i(s_2,s_3)$, and $\underline{u}^n_i = \underline{u}^n_{i(5-i)}$, for $i = 2,3$, $j = 1,\cdots,l$, $n \geq 0$. Let

$$\Gamma^n_{b,A_{1,2}} := \Gamma\left[\left(v^n_{ij}\right)_{\substack{i=1,2 \\ j=1,\cdots,l}}, \underline{u}^n_1 = 0, \underline{u}^n_2 = 0\right]$$

and

$$\Gamma_{b,A_{1,2}} = \Gamma\left[\left(v_{ij}\right)_{\substack{i=1,2 \\ j=1,\cdots,l}}, \underline{u}_1 = 0, \underline{u}_2 = 0\right].$$

where $A_{1,2}$ denotes the contingency that Nodes 1 and 2 reach an agreement. Following the same argument as in *Step 2*, Lemma 13 implies that

$$\sigma^{*\eta}_{b,A_{1,2}} \in \mathcal{R}\left[\left(\Gamma^n_{b,A_{1,2}}\right)_{n\geq 0}\right] \tag{86}$$

By combining equations (77), (80), (83) and (86), we conclude that $\sigma^{*\eta}$ is a restricted

equilibrium in the three-node bargaining game $\Gamma_{\text{3-node}}(\eta, S, C, f, s^0)$. This completes the proof of the existence part.

Uniqueness: We will proceed in four steps, as in the proof of the existence part. Suppose σ is a restricted equilibrium of the three-node contracting game $\Gamma_{\text{3-node}}(\eta, S, C, f, s^0)$, and $(\sigma^n)_{n\geq 0} \in \mathcal{T}[\sigma, \lambda; \Gamma_{\text{3-node}}(\eta, S, C, f, s^0)]$ is a restricted trembling sequence for σ, for some $\lambda > 0$ sufficiently small. We need to show that the strategy profile σ induces the same deterministic outcome $\left[s^{**}, \left(y_1^\eta(s^{**}), y_3^\eta(s_2^{**}, s_3^{**})\right)\right]$ as $\sigma^{*\eta}$. Instead of using the existence results provided by Corollary 8 and Lemmas 11 to 13, we will use their uniqueness counterparts. There is one issue that we will need to deal with: when applying Lemmas 11 to 13, we need to first verify the convergence of total payoff π^n of a bargaining pair. We can guarantee the convergence of this value by taking subsequences.

One will see that *Step 1* and *Step 2* in the uniqueness part are very similar to that of the existence part, since the total payoff and the outside option values of any pair of nodes' bargaining problem are constant contingent on breakdown of the other pair of nodes.

Step 1: Consider the negotiation between Nodes 1 and 2 in Stage aa over the $y_1^b\left(s_1^b, s_2^{3b}\right)$. The negotiation process is identical to that of the Rubinstein's alternating-offers game $\Gamma\left(\pi_1^b\left(s_1^b, s_2^{3b}\right), 0, 0\right)$, where π_1^b is defined in (69). As σ is a restricted equilibrium of the three-node contracting game $\Gamma_{\text{3-node}}(\eta, S, C, f, s^0)$, then its restriction $\sigma_{aa,\left(s_1^b, s_2^{3b}\right)}$ to Stage aa and the choice of breakdown treatments $\left(s_1^b, s_2^{3b}\right)$ at the contingency $B_{2,3}$ is a restricted equilibrium in the game $\Gamma\left(\pi_1^b, 0, 0\right)$. The sequence $\left(\sigma^n_{aa,\left(s_1^b, s_2^{3b}\right)}\right)_{n\geq 0}$ is a restricted trembling sequence for $\sigma_{aa,\left(s_1^b, s_2^{3b}\right)}$ in the game $\Gamma\left(\pi_1^b, 0, 0\right)$. Corollary 8 and lemma 12 imply that $\sigma_{aa,\left(s_1^b, s_2^{3b}\right)}$ induces the same payoff to Nodes 1 and 2 as $\sigma^{*\eta}_{aa,\left(s_1^b, s_2^{3b}\right)}$ in the game $\Gamma\left(\pi_1^b, 0, 0\right)$. That is,

$$U_i\left(\sigma_{aa,\left(s_1^b, s_2^{3b}\right)}; \Gamma\left(\pi_1^b, 0, 0\right)\right) = U_i\left(\sigma^{*\eta}_{aa,\left(s_1^b, s_2^{3b}\right)}; \Gamma\left(\pi_1^b, 0, 0\right)\right) = v_i\left(\pi_1^b\right)^+, \qquad i = 1, 2.$$

Along the restricted trembling sequence $\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)}$,

$$U_i\left(\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)};\Gamma\left(\pi_1^b,0,0\right)\right)\to U_i\left(\sigma_{aa,\left(s_1^b,s_2^{3b}\right)};\Gamma\left(\pi_1^b,0,0\right)\right)=v_i\left(\pi_1^b\right)^+,\qquad i=1,2.$$

Similar results can be obtained for the negotiation between Nodes 2 and 3 in Stage bb over $y_3^b\left(s_2^{1b},s_3^b\right)$. Letting π_3^b be given by (71), we have,

$$U_i\left(\sigma_{bb,\left(s_2^{1b},s_3^b\right)};\Gamma\left(\pi_3^b,0,0\right)\right)=v_i\left(\pi_3^b\right)^+,\qquad i=2,3.$$

Along the restricted trembling sequence $\sigma^n_{bb,\left(s_2^{1b},s_3^b\right)}$,

$$U_i\left(\sigma^n_{bb,\left(s_2^{1b},s_3^b\right)};\Gamma\left(\pi_3^b,0,0\right)\right)=v_i\left(\pi_3^b\right)^+,\qquad i=2,3. \tag{87}$$

Step 2: We next consider the negotiation between Nodes 1 and 2 in Stage a over $\left(s_1^b,s_2^{3b}\right)\in\mathcal{S}^B_{1,2}$ - the choice of treatments for the contingency $B_{2,3}$. Given that Nodes 1 and 2 will follow the strategy profile $\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)}$ in Stage aa, implementation of any pair of treatments $\left(s_1^b,s_2^{3b}\right)$ leads to a payoff of $v_i^n\left(s_1^b,s_2^{3b}\right):=U_i\left(\sigma^n_{aa,\left(s_1^b,s_2^{3b}\right)};\Gamma\left(\pi_1^b,0,0\right)\right)$ for Node i $(i=1,2)$, where by equation (87),

$$v_i^n\left(s_1^b,s_2^{3b}\right)\xrightarrow{n\to\infty}v_i\cdot\left(\pi_1^b\left(s_1^b,s_2^{3b}\right)\right)^+,\qquad i=1,2.$$

On the other hand, a breakdown of the negotiation over $\left(s_1^b,s_2^{3b}\right)\in\mathcal{S}^B_{1,2}$ would give both nodes their respective outside option values 0. Thus we are in the setting of Lemma 13, where $l=\left|\mathcal{S}^B_{1,2}\right|$, $\Pi=\left\{\pi_1^b\left(s_1^b,s_2^{3b}\right):\left(s_1^b,s_2^{3b}\right)\in\mathcal{S}^B_{1,2}\right\}$, $v_{ij}^n=v_i^n\left(s_1^b,s_2^{3b}\right)$, and $\underline{u}_i^n=0$, for $i=1,2$, $j=1,\cdots,l$, $n\geq 0$. For simplicity of notation, let

$$\Gamma^n_{a,B_{2,3}}:=\Gamma\left[\left(v_{ij}^n\right)_{\substack{i=1,2\\j=1,\cdots,l}},\underline{u}_1^n=0,\underline{u}_2^n=0\right]$$

and

$$\Gamma_{a,B_{2,3}}=\Gamma\left[\left(v_{ij}\right)_{\substack{i=1,2\\j=1,\cdots,l}},\underline{u}_1=0,\underline{u}_2=0\right].$$

Then $\sigma_{a,B_{2,3}} \in \mathcal{R}\left(\left(\Gamma^n_{a,B_{2,3}}\right)_{n\geq 0}\right)$, and $\left(\sigma^n_{a,B_{2,3}}\right)_{n\geq 0} \in \mathcal{T}\left[\sigma^{*\eta}_{a,B_{2,3}}; \left(\Gamma^n_{a,B_{2,3}}\right)_{n\geq 0}\right]$. Since

$$\underset{\left(s_1^b, s_2^{3b}\right)\in \mathcal{S}^B_{1,2}}{\operatorname{argmax}} \ \pi_1^b\left(s_1^b, s_2^{3b}\right) = \left(s_1^B, s_2^{3B}\right)$$

and $\pi_1^b\left(s_1^B, s_2^{3B}\right) > \pi_1^b\left(s_1^0, s_2^0\right) = 0$.

Hence by Lemma 13, $\sigma_{a,B_{2,3}}$ is such that both nodes propose $\left(s_1^B, s_2^{3B}\right)$ in each period, and accept (and only accept) $\left(s_1^B, s_2^{3B}\right)$. Thus $\sigma_{a,B_{2,3}} = \sigma^{*\eta}_{a,B_{2,3}}$. Applying Lemma 13 again, we obtain

$$U_2\left(\sigma^n_{a,B_{2,3}}; \Gamma^n_{a,B_{2,3}}\right) \to U_2\left(\sigma_{a,B_{2,3}}; \Gamma_{a,B_{2,3}}\right) = U_2\left(\sigma^{*\eta}_{a,B_{2,3}}; \Gamma_{a,B_{2,3}}\right) = \underline{u}^\eta_{23}. \tag{88}$$

The same analysis applies to the negotiation between Nodes 2 and 3 in Stage b over $\left(s_2^{1b}, s_3^b\right) \in \mathcal{S}^B_{2,3}$.

Following Remark 5 and (88), we have

$$\underline{u}^n_{ij} \to \underline{u}^\eta_{ij} \tag{89}$$

where $\underline{u}^n_{ij}$ are the outside option values determined by the breakdown strategies $\left(\sigma^n_{\{B_{1,2},B_{2,3}\}}\right)_{n\geq 0}$.

Step 3: In Step 3, we first consider the negotiation between Nodes 2 and 3 in Stage bb over $y_3(s_2, s_3)$. Under σ^n, suppose $\pi^n_3(s_2, s_3)$ is the total payoff to be shared by Nodes 2 and 3 in Stage bb, if the two nodes agreed on the treatments (s_2, s_3) in Stage b. For every $n \geq 0$, since the payoff of Node 1 under σ^n is at least his outside option value $\underline{u}^\eta_{12} = 0$, we have

$$\pi^n_3(s_2, s_3) \leq U\left(s_1^*(s_2, s_3), s_2, s_3\right).$$

Since $\left(\pi^n_3(s_2, s_3)\right)_{n\geq 0}$ is a bounded real sequence, there exists a converging subsequence:

$$\pi^{\varphi(n)}_3(s_2, s_3) \xrightarrow{n\to\infty} \pi_3(s_2, s_3), \tag{90}$$

for some $\pi_3(s_2, s_3) \leq U\left(s_1^*(s_2, s_3), s_2, s_3\right)$. Let $y_{3|\text{agreement}}\left(\sigma^n_{bb,(s_2,s_3)}\right)$ be the expected payment from Node 2 to Node 3 under σ^n, conditional on the two nodes agreeing on (s_2, s_3) after

bargaining in Stage b and Stage bb. By only considering the subsequence $\varphi(n)$, we have a sequence of converging total payoffs $\left(\pi_3^{\varphi(n)}(s_2,s_3)\right)_{n\geq 0}$ and two sequences of converging outside option values $(\underline{u}_{23}^n)_{n\geq 0}$ and $(\underline{u}_{32}^n)_{n\geq 0}$ (by equation (89)) for the bargaining problem between Nodes 2 and 3.

- If $\pi_3(s_2,s_3) > \underline{u}_{23}^{\eta}$, then we are in the setting of Lemma 11. We thus have

$$\begin{aligned}\lim_{n\geq 0} y_{3|\text{agreement}}\left(\sigma_{bb,(s_2,s_3)}^{\varphi(n)}\right) &= v_3[\pi_3(s_2,s_3) - \underline{u}_{23}^{\eta}] - f_3(s_2,s_3)\\ &\leq v_3\left[U\left(s_1^*(s_2,s_3),s_2,s_3\right) - \underline{u}_{23}^{\eta}\right] - f_3(s_2,s_3)\end{aligned}$$

  Recall that $v_3 = 1/(2-\eta)$.

- If $\pi_3(s_2,s_3) \leq \underline{u}_{23}^{\eta}$, then we are in the setting of Lemma 12. We thus have

$$\limsup_{n\to\infty} y_{3|\text{agreement}}\left(\sigma_{bb,(s_2,s_3)}^{\varphi(n)}\right) \leq -f_3(s_2,s_3) + 2\lambda.$$

Since the subsequence $\left(y_{3|\text{agreement}}\left(\sigma_{bb,(s_2,s_3)}^{\varphi(n)}\right)\right)_{n\geq 0}$ is bounded, there further exists a converging subsequence:

$$y_{3|\text{agreement}}\left(\sigma_{bb,(s_2,s_3)}^{\phi(n)}\right) \xrightarrow{n\to\infty} y_{3|\text{agreement}}(s_2,s_3) \tag{91}$$

for some $y_{3|\text{agreement}}(s_2,s_3)$. Here $(\phi(n))_{n\geq 0}$ is a subsequence of the sequence $(\varphi(n))_{n\geq 0}$. Combining the two inequalities above, we have

$$y_{3|\text{agreement}}(s_2,s_3) \leq v_3\, U\left(s_1^*(s_2,s_3),s_2,s_3\right) - f_3(s_2,s_3) + 2\lambda. \tag{92}$$

Now we consider the negotiation between Nodes 1 and 2 in Stage aa over the payment $y_1(s_1;s_2,s_3)$, and then the negotiation between the two nodes in Stage a over the choice of treatment $s_1(s_2,s_3) \in C_1(s_2)$ for the contingency (s_2,s_3). Since we fix a contingency $(s_2,s_3) \in \mathcal{S}_{2,3}$, thus we can, for simplicity, write s_1 for $s_1(s_2,s_3)$.

Under σ^n, let $\pi_1^n(s_1;s_2,s_3)$ be the total payoff to be shared by Nodes 1 and 2 if the two nodes agreed on the treatment s_1 for the contingency (s_2,s_3) in Stage a. Along the

subsequence $\sigma^{\phi(n)}$, we have

$$\begin{aligned}\pi_1^{\phi(n)}(s_1;s_2,s_3) &= f_1(s_1,s_2)+f_2(s_1,s_2,s_3)-y_{3|\text{agreement}}\left(\sigma^{\phi(n)}_{bb,(s_2,s_3)}\right)\\ &\to \underbrace{f_1(s_1,s_2)+f_2(s_1,s_2,s_3)-y_{3|\text{agreement}}(s_2,s_3)}_{:=\pi_1(s_1;s_2,s_3)} \qquad \text{by (91).}\end{aligned} \tag{93}$$

By only considering the subsequence $\phi(n)$, we have a sequence of converging total payoffs $\left(\pi_1^{\phi(n)}(s_1;s_2,s_3)\right)_{n\geq 0}$ and two sequences of converging outside option values $\left(\underline{u}_{12}^{\phi(n)}\right)_{n\geq 0}$ and $\left(\underline{u}_{21}^{\phi(n)}\right)_{n\geq 0}$ for the bargaining problem between Nodes 1 and 2. Let $\Gamma^{\phi(n)}_{aa,(s_1;s_2,s_3)} = \Gamma\left(\pi_1^{\phi(n)}(s_1;s_2,s_3),\underline{u}_{12}^{\phi(n)},\underline{u}_{21}^{\phi(n)}\right)$. By Lemmas 11 and 12, we have

$$\begin{aligned}&U_i\left(\sigma^{\phi(n)}_{aa,(s_1;s_2,s_3)};\Gamma^{\phi(n)}_{aa,(s_1;s_2,s_3)}\right)\xrightarrow{n\to\infty} v_i\cdot(\pi_1(s_1;s_2,s_3)-\underline{u}^{\eta}_{12}-\underline{u}^{\eta}_{21})^{+}+\underline{u}^{\eta}_{i(3-i)}, \qquad i=1,2.\\ &U_1\left(\sigma^{\phi(n)}_{aa,(s_1;s_2,s_3)};\Gamma^{\phi(n)}_{aa,(s_1;s_2,s_3)}\;\middle|\;\text{agreement}\right)=v_1\cdot(\pi_1(s_1;s_2,s_3)-\underline{u}^{\eta}_{12}-\underline{u}^{\eta}_{21})^{+}+\underline{u}^{\eta}_{12}.\end{aligned} \tag{94}$$

On the other hand, (92) implies that

$$\pi_1(s_1^*(s_2,s_3);s_2,s_3)\geq v_2\,U(s_1^*(s_2,s_3),s_2,s_3)-2\lambda.$$

Since we have the assumption (2) that

$$U(s_1^*(s_2,s_3),s_2,s_3)>U(s_1^0,s_2^{1B},s_3^B),$$

and $\lambda>0$ can be chosen to be aribitrarily close to 0, then by choosing a λ sufficiently small, we have

$$\pi_1(s_1^*(s_2,s_3);s_2,s_3)>v_2\,U(s_1^0,s_2^{1B},s_3^B)=\underline{u}^{\eta}_{21}. \tag{95}$$

We next consider the negotiation between Nodes 1 and 2 in Stage a over $s_1\in C_1(s_2)$ - the choice of treatment for the contingency (s_2,s_3) - along the subsequence $\phi(n)$. Given that Nodes 1 and 2 will play $\sigma^{\phi(n)}_{aa,(s_1;s_2,s_3)}$ in Stage aa, implementation of any choice of treatment s_1 for the contingency (s_2,s_3) leads to a payoff of $v_i^n(s_1;s_2,s_3):=U_i\left(\sigma^{\phi(n)}_{aa,(s_1;s_2,s_3)};\Gamma^{\phi(n)}_{aa,(s_1;s_2,s_3)}\right)$ for Node i $(i=1,2)$. Thus we are in the setting of Lemma 13, where $l=|C_1(s_2)|$, $\Pi=\{\pi_1(s_1;s_2,s_3):s_1\in C_1(s_2)\}$, $v_{ij}^n=v_i^n(s_1;s_2,s_3)$, and $\underline{u}_i^n=\underline{u}^{\phi(n)}_{i(3-i)}$, for $i=1,2$, $j=1,\cdots,l$,

$n \geq 0$, and the maximum total payoff

$$\pi_1(s_1^*(s_2, s_3); s_2, s_3) > \underline{u}_{21}^{\eta}.$$

Let

$$\Gamma_{a,(s_2,s_3)}^n := \Gamma\left[\left(v_{ij}^n\right)_{\substack{i=1,2 \\ j=1,\cdots,l}}, \underline{u}_1^n = \underline{u}_{12}^{\phi(n)}, \underline{u}_2^n = \underline{u}_{21}^{\phi(n)}\right]$$

and

$$\Gamma_{a,(s_2,s_3)} = \Gamma\left[\left(v_{ij}\right)_{\substack{i=1,2 \\ j=1,\cdots,l}}, \underline{u}_1 = \underline{u}_{12}^{\eta}, \underline{u}_2 = \underline{u}_{21}^{\eta}\right].$$

As $\sigma_{a,(s_2,s_3)}^{\phi(n)}$ converges weak*, pointwise, to $\sigma_{a,(s_2,s_3)}$, and

$$\left(\sigma_{a,(s_2,s_3)}^{\phi(n)}\right)_{n\geq 0} \in \mathcal{T}\left[\sigma_{a,(s_2,s_3)}^{*\eta}; \left(\Gamma_{a,(s_2,s_3)}^n\right)_{n\geq 0}\right],$$

Lemma 13 implies that $\sigma_{a,(s_2,s_3)}$ consists of Nodes 1 and 2 offer each other the contingent treatment contract $s_1^*(s_2, s_3)$ at each period. Both nodes accept this offer $s_1^*(s_2, s_3)$ and reject all other offers $s_1 \in C_1(s_2)\backslash\{s_1^*(s_2, s_3)\}$. It then follows that from (94) that

$$U_1\left(\sigma_{\{a,aa\},(s_2,s_3)}^{\phi(n)}; \Gamma_{a,(s_2,s_3)}^{\phi(n)} \mid \text{agreement}\right) \xrightarrow{n\to 0} \quad v_1 \cdot [\pi_1(s_1^*(s_2, s_3); s_2, s_3) - \underline{u}_{21}^{\eta}] \tag{96}$$

Step 4: We next consider the negotiation between Nodes 2 and 3 in Stage bb over $y_3(s_2, s_3)$ - the payment from Node 2 to Node 3 associated with some choice of treatment $(s_2, s_3) \in \mathcal{S}_{2,3}$. Following Remark 5, we already know that the outside option values of Nodes 2 and 3 are $\underline{u}_{23}^n$ and $\underline{u}_{32}^n$ respectively along the sequence of $\sigma_{\{B_{1,2},B_{2,3}\}}^n$, and

$$\underline{u}_{23}^n \to \underline{u}_{23}^{\eta}; \qquad \underline{u}_{32}^n \to \underline{u}_{32}^{\eta}. \tag{97}$$

Along the subsequence $\phi(n)$, given that Nodes 1 and 2 play $\sigma_{\{a,aa\}}^{\phi(n)}$ in Stage a and Stage

aa, the total payoff $\pi_3^{\phi(n)}(s_2,s_3)$ to be shared by Nodes 2 and 3 satisfies

$$\begin{aligned}
&\lim_{n\to\infty}\pi_3^{\phi(n)}(s_2,s_3) = U(s_1^*(s_2,s_3),s_2,s_3) - \lim_{n\to\infty} U_1\left(\sigma^{\phi(n)}_{\{a,aa\},(s_2,s_3)};\Gamma^{\phi(n)}_{a,(s_2,s_3)} \mid \text{agreement}\right)\\
=\ & \underbrace{U(s_1^*(s_2,s_3),s_2,s_3) - v_1\cdot[\pi_1(s_1^*(s_2,s_3);s_2,s_3) - \underline{u}_{21}^{\eta}]}_{:=\pi_3(s_2,s_3)} \qquad \text{by (96)}.
\end{aligned}$$

Let $y_1(s_1^*(s_2,s_3),s_2,s_3) = v_1\cdot[\pi_1(s_1^*(s_2,s_3);s_2,s_3) - \underline{u}_{21}^{\eta}] - f_1(s_1^*(s_2,s_3),s_2)$, then we have

$$\begin{aligned}
&f_2(s_1(s_2,s_3),s_2,s_3) - y_1(s_1^*(s_2,s_3),s_2,s_3) - y_{3|\text{agreement}}(s_2,s_3) - \underline{u}_{21}^{\eta}\\
&\qquad = (1-\eta)\left[f_1(s_1^*(s_2,s_3),s_2) + y_1(s_1^*(s_2,s_3),s_2,s_3) - \underline{u}_{12}^{\eta}\right].
\end{aligned} \tag{98}$$

$$\pi_3(s_2,s_3) = f_2(s_1^*(s_2,s_3),s_2,s_3) + f_3(s_2,s_3) - y_1(s_1^*(s_2,s_3),s_2,s_3). \tag{99}$$

Letting $\Gamma^{\phi(n)}_{bb,(s_2,s_3)} = \Gamma\left(\pi_3^{\phi(n)}(s_2,s_3),\underline{u}_{32}^{\phi(n)},\underline{u}_{23}^{\phi(n)}\right)$, then

$$\sigma_{bb,(s_2,s_3)} \in \mathcal{R}\left[\left(\Gamma^{\phi(n)}_{bb,(s_2,s_3)}\right)_{n\geq 0}\right], \tag{100}$$

and $\left(\sigma^{\phi(n)}_{bb,(s_2,s_3)}\right)_{n\geq 0} \in \mathcal{T}\left[\sigma_{bb,(s_2,s_3)};\left(\Gamma^{\phi(n)}_{bb,(s_2,s_3)}\right)_{n\geq 0}\right]$. Lemmas 11 and 12 imply that

$$U_i\left(\sigma^{\phi(n)}_{bb,(s_2,s_3)};\Gamma^{\phi(n)}_{bb,(s_2,s_3)}\right) \xrightarrow{n\to\infty} v_i\cdot(\pi_3(s_2,s_3) - \underline{u}_{23}^{\eta})^+ + \underline{u}_{i(5-i)}^{\eta}, \qquad i=2,3, \tag{101}$$

$$U_3\left(\sigma^{\phi(n)}_{bb,(s_2,s_3)},\Gamma^{\phi(n)}_{bb,(s_2,s_3)} \mid \text{agreement}\right) \xrightarrow{n\to\infty} v_1\cdot(\pi_3(s_2,s_3) - \underline{u}_{23}^{\eta})^+. \tag{102}$$

By the definition of $y_{3|\text{agreement}}(s_2,s_3)$, we know that

$$U_3\left(\sigma^{\phi(n)}_{bb,(s_2,s_3)},\Gamma^{\phi(n)}_{bb,(s_2,s_3)} \mid \text{agreement}\right) \xrightarrow{n\to\infty} f_3(s_2,s_3) + y_{3|\text{agreement}}(s_2,s_3). \tag{103}$$

We obtain, from (99) and (102), a linear equation of $y_{3|\text{agreement}}(s_2,s_3)$, when $\pi_3(s_2,s_3)\geq \underline{u}_{23}^{\eta}$:

$$\begin{aligned}
&f_2(s_1^*(s_2,s_3),s_2,s_3) - y_1(s_1^*(s_2,s_3),s_2,s_3) - y_{3|\text{agreement}}(s_2,s_3) - \underline{u}_{23}^{\eta}\\
&= (1-\eta)\left[f_3(s_2,s_3) + y_{3|\text{agreement}}(s_2,s_3) - \underline{u}_{32}^{\eta}\right].
\end{aligned} \tag{104}$$

Equations (98) and (104) form a linear system of two equations with two unknowns

$$\left(y_1(s_1^*(s_2,s_3),s_2,s_3), y_{3|\text{agreement}}(s_2,s_3)\right).$$

The equations are exactly the same as the ones for

$$(y_1^\eta(s_1^*(s_2,s_3);s_2,s_3), y_3^\eta(s_2,s_3))$$

given by equations (16) and (17). Hence when $\pi_3(s_2,s_3) \geq \underline{u}_{23}^\eta$,

$$y_1(s_1^*(s_2,s_3),s_2,s_3) = y_1^\eta(s_1^*(s_2,s_3),s_2,s_3) \tag{105a}$$

$$y_{3|\text{agreement}}(s_2,s_3) = y_3^\eta(s_2,s_3) \tag{105b}$$

Then it follows from (99) that

$$\begin{aligned}\pi_3(s_2,s_3) &= U(s_1^*(s_2,s_3),s_2,s_3) - v_1 \cdot [\pi_1(s_1^*(s_2,s_3);s_2,s_3) - \underline{u}_{21}^\eta] \\ &= \frac{2-\eta}{3-\eta} U(s_1^*(s_2,s_3),s_2,s_3) + c,\end{aligned} \tag{106}$$

where c is a constant that only depends on the outside option values $(\underline{u}_{21}^\eta, \underline{u}_{23}^\eta)$, but not on (s_2,s_3).

On the other hand, we show that

$$\pi_3(s_2^{**},s_3^{**}) > \underline{u}_{23}^\eta.$$

Recall that $(s_1^{**},s_2^{**},s_3^{**})$ is the socially efficient vector of treatments. If $\pi_3(s_2^{**},s_3^{**}) \leq \underline{u}_{23}^\eta$, then by (102) and (103), we have

$$f_3(s_2^{**},s_3^{**}) + y_{3|\text{agreement}}(s_2^{**},s_3^{**}) = 0.$$

Since $\pi_1(s_1^{**},s_2^{**},s_3^{**}) = f_1(s_1^{**},s_2^{**}) + f_2(s_1^{**},s_2^{**},s_3^{**}) - y_{3|\text{agreement}}(s_2^{**},s_3^{**})$, we have

$$\pi_1(s_1^{**},s_2^{**},s_3^{**}) = U(s_1^{**},s_2^{**},s_3^{**}).$$

It then follows from (99) that

$$\begin{aligned}\pi_3(s_2^{**}, s_3^{**}) = &\ U(s_1^{**}, s_2^{**}, s_3^{**}) - v_1 \cdot [\pi_1(s_1^*(s_2, s_3); s_2, s_3) - \underline{u}_{21}^{\eta}] \\ = &\ v_2\, U(s_1^{**}, s_2^{**}, s_3^{**}) + v_1\, \underline{u}_{21}^{\eta} \\ > &\ v_2\, U(s_1^{B}, s_2^{3B}, s_3^{0}) \\ = &\ \underline{u}_{23}^{\eta}.\end{aligned}$$

This leads to a contradiction. Therefore $\pi_3(s_2^{**}, s_3^{**}) > \underline{u}_{23}^{\eta}$.

Finally we consider the negotiation between Nodes 2 and 3 in Stage b over $(s_2, s_3) \in \mathcal{S}_{2,3}$. Given that Nodes 2 and 3 will play $\sigma_{bb,(s_2,s_3)}^{\phi(n)}$ in Stage bb, implementation of any choice of treatments (s_2, s_3) leads to a payoff of $v_i^n(s_2, s_3) := U_i\left(\sigma_{bb,(s_2,s_3)}^{\phi(n)}; \Gamma_{bb,(s_2,s_3)}^{\phi(n)}\right)$ for Node i $(i = 2, 3)$, where by equation (84),

$$v_i^n(s_2, s_3) \xrightarrow{n\to\infty} v_i \cdot (\pi_3(s_2, s_3) - \underline{u}_{32}^{\eta} - \underline{u}_{23}^{\eta})^+ + \underline{u}_{i(5-i)}^{\eta}, \qquad i = 1, 3.$$

Thus we are in the setting of Lemma 13, where $l = |\mathcal{S}_{2,3}|$, $\Pi = \{\pi_3(s_2, s_3) : (s_2, s_3) \in \mathcal{S}_{2,3}\}$, $v_{ij}^n = v_i^n(s_2, s_3)$, and $\underline{u}_i^n = \underline{u}_{i(5-i)}^{\phi(n)}$, for $i = 2, 3$, $j = 1, \cdots, l$, $n \geq 0$, and the maximum total payoff

$$\begin{aligned}\max_{(s_2,s_3)\in\mathcal{S}_{23}} \pi_3(s_2, s_3) = &\ \frac{2-\eta}{3-\eta} \max_{(s_2,s_3)\in\mathcal{S}_{23}} U(s_1^*(s_2, s_3), s_2, s_3) + c \\ = &\ \frac{2-\eta}{3-\eta}\, U(s_1^{**}, s_2^{**}, s_3^{**}) + c \\ = &\ \pi_3(s_2^{**}, s_3^{**}) > \underline{u}_{23}^{\eta}.\end{aligned}$$

Let

$$\Gamma_{b,A_{1,2}}^n := \Gamma\left[\left(v_{ij}^n\right)_{\substack{i=1,2 \\ j=1,\cdots,l}}, \underline{u}_1^n = 0, \underline{u}_2^n = 0\right]$$

and

$$\Gamma_{b,A_{1,2}} = \Gamma\left[\left(v_{ij}\right)_{\substack{i=1,2 \\ j=1,\cdots,l}}, \underline{u}_1 = 0, \underline{u}_2 = 0\right].$$

where $A_{1,2}$ denotes the contingency that Nodes 1 and 2 reach an agreement. As $\sigma_{b,A_{1,2}}^{\phi(n)}$ converges weak*, pointwise, to $\sigma_{b,A_{1,2}}$, and

$$\left(\sigma_{b,A_{1,2}}^{\phi(n)}\right)_{n\geq 0} \in \mathcal{T}\left[\sigma_{b,A_{1,2}}^{*\eta}; \left(\Gamma_{b,A_{1,2})}^n\right)_{n\geq 0}\right],$$

Lemma 13 implies that $\sigma_{b,A_{1,2}}$ consists of Nodes 2 and 3 offer each other the socially efficient treatments (s_2^{**}, s_3^{**}) at each period. Both nodes accept this offer (s_2^{**}, s_3^{**}) and reject all other offers $(s_2, s_3) \in \mathcal{S}_{2,3} \backslash \{(s_2^{**}, s_3^{**})\}$.

Combining the four steps above, we conclude that under the strategy profile σ induces the same deterministic outcome as $\sigma^{*\eta}$, in which with probability 1, $\left[s^{**}, \left(y_1^{\eta}(s^{**}), y_3^{\eta}(s_2^{**}, s_3^{**})\right)\right]$ is immediately implemented as an accepted offer in period 1 of the respective bargaining stages. This establishes uniqueness of Theorem 1.

□

C Foundations for an Axiomatic Solution

Here, we extend the axioms of Nash (1950) to a network setting. Under the stated axioms for a solution to a *network bilateral bargaining problem (with exogeneous outside option values)*, the unique outcome is socially efficient. Our axioms are preliminary and may be adjusted in subsequent versions.

C.1 Network Bargaining Problem with Exogenous Outside Values

The players of a network bilateral bargaining problem are the nodes of an undirected graph $G = (V, E)$, defined by a finite set V of at least two nodes, and by a set $E \subset V \times V$ of edges, which are the pairs of "directly connected" nodes. The network G is assumed to contain no circles. [18] The set of nearest neighbors of node i is denoted by $\mathcal{N}_i(G) = \{j \in V : (i, j) \in E\}$. The action of node i is chosen from a finite non-empty set S_i. If i is an isolated node (that is, if $\mathcal{N}_i(G) = \emptyset$), then its action set S_i is assumed to be a singleton. As we shall see, this is without loss of generality. We let $S = \prod_{i \in V} S_i$. Node i has a utility function $f_i : S_i \times \widehat{S}_i \to \mathbb{R}$, where $\widehat{S}_i = \prod_{j \in \mathcal{N}_i(G)} S_j$. We let $f = (f_i)_{i \in V}$ denote the collection of utility functions of all nodes. For simplicity, we often abuse the notation by writing $f_i(s)$ in place of $f_i\left(s_i, s_{\mathcal{N}_i(G)}\right)$.

[18] An ordered list $(i_1, \ldots, i_m)$ of nodes in V is a circle if i_p and i_{p+1} are directly connected for all p, where by convention, $i_{m+1} = i_1$.

For any pair $(i,j) \in E$ of direct counterparties, given some action s_j of Node j, the action of Node i is restricted to $C_{ij}(s_j)$, for a given correspondence C_{ij} from S_j into the non-empty subsets of S_i. Without loss of generality, we assume that C_{ij} and C_{ji} are consistent, in that $s_i \in C_{ij}(s_j)$ if and only if $s_j \in C_{ji}(s_i)$. A pair of actions $(s_i, s_j) \in S_i \times S_j$ is said to be *compatible* if and only if $s_i \in C_{ij}(s_j)$. With a slight abuse of notation, we sometimes write $(s_i, s_j) \in C_{ij}$ for $s_i \in C_{ij}(s_j)$. We let $C = (C_{ij})_{(i,j)\in E}$ and write $s \in C$ if $s_i \in C_{ij}(s_j)$ for every pair $(i,j) \in E$. We say that C is feasible if there is some s in C. Any such s is likewise called feasible.

The total social welfare of a feasible action vector s is

$$U_\sigma(s) = \sum_{i \in V} f_i(s)$$

We assume that there is a unique action vector $s^{**}(\sigma)$ that is *socially efficient*. That is, the maximum of the social welfare $U_\sigma(s)$ is uniquely achieved at $s^{**}(\sigma)$,

$$s^{**}(\sigma) = \underset{s \in C}{\operatorname{argmax}}\ U_\sigma(s).$$

The uniqueness of the maximizer is a generic property of the utility functions $(f_i)_{i\in V}$ that are non-degenerate.[19]

We consider a setting in which each pair of directly connected nodes bargain over their actions and the payment between them. For each such pair $(i,j) \in E$, the associated outside option value of i is some given parameter $\underline{u}_{ij}$. We later endogenize these outside option values. For condensed notation, we let $\underline{u} = \left(\underline{u}_{ij}\right)_{(i,j)\in E}$ denote the collection of outside options.

A *network bilateral bargaining problem (with exogenous outside options)* is defined as

$$\sigma = (G, S, C, f, \underline{u}),$$

[19] A utility function f_i is *degenerate* if the dependence of f_i on at least one of its arguments is superficial. Letting the set of collections f of non-degenerate utility functions be endowed with Lebesgue measure, then for almost every f, the social welfare $\sum_{i\in V} f_i(s)$ admits a unique maximizer.

with the property that C is feasible and $\underline{u}$ satisfies a condition for outside option values, to be provided. An outcome (s, y) of σ consists of some feasible action vector $s \in C$ and some payments $y \in \mathbb{R}^E$ assigning to each edge (i, j) a payment y_{ij} from i to j, with $y_{ij} = -y_{ji}$. The total utility for node i is

$$u_i(s, y) = f_i\left(s_i, s_{\mathcal{N}_i(G)}\right) - \sum_{j \in \mathcal{N}_i(G)} y_{ij}.$$

A *solution mapping* for network bargaining problems (with exogenous outside option values) is some function F that maps such a problem to an associated outcome.

The solution outcome $(s^*(\sigma), y^*(\sigma))$ is anticipated to be socially efficient ($s^*(\sigma) = s^{**}(\sigma)$), and the payment vector y^* satisfies

$$u_i(s^{**}(\sigma), y^*(\sigma)) - \underline{u}_{ij} = u_j(s^{**}(\sigma), y^*(\sigma)) - \underline{u}_{ji}, \qquad \forall (i, j) \in E.$$

That is, nodes i and j has equitable gain from trade relative to their respective outside option values. The equations above constitute a fully determined linear system in the payments $\left(y^*_{ij}(\sigma)\right)_{(i,j)\in E}$. The payment vector $y^*(\sigma)$ is thus uniquely determined by the linear system.

The outside option values $(\underline{u}_{ij})_{(i,j)\in E}$ are assumed to satisfy the following condition

$$u_i(s^{**}(\sigma), y^*(\sigma)) - \underline{u}_{ij} = u_j(s^{**}(\sigma), y^*(\sigma)) - \underline{u}_{ji} \geq 0, \qquad \forall (i, j) \in E.$$

We assume this condition for every NBBP with exogenous outside option values. Once we introduce endogenous outside options through an overlaying model structure, we will demonstrate that this condition is automatically satisfied, without any additional assumption.

Such a network bargaining problem features $|E|$ pairs of bilateral bargaining problems, where $|E|$ is the number of edges in the network. Given its decentralized nature, a network bilateral bargaining problem (with exogenous outside option values) has the following intrinsic characteristics regarding its information structure, which are implicit in its definition.

(i) Each node has complete information over the primitives $(G, S, C, f, \underline{u})$ of the problem, but observes only the actions of its direct counterparties. In this sense, the infor-

mation structure of a network bilateral bargaining problem is similar to that of a non-cooperative game with complete but imperfect information.

(ii) If node j is a direct counterparty to both nodes i and k, then j is able to inform k of the action of i when j bargains with k. Likewise, j can provide information about the action of k to node i when j bargains with i.

(iii) This possibility of communication is common knowledge.

Property *(i)* is a key aspect where a network bilateral bargaining game conceptually differs from a cooperative game.[20] In a cooperative game, groups of players ("coalitions") may enforce cooperative behavior. In this sense, a network bilateral bargaining game is, by nature, not cooperative, as it lacks a direct mechanism by which coalitions may enforce coordinated behavior on the members of the coalition. On the other hand, properties *(ii), (iii)* provide a possibility for coordination among different bargaining pairs. Whether this possibility comes to fruition, however, depends on the incentives of internal nodes to make these communications truthfully.

We propose a solution concept for of network bilateral bargaining problems (with exogenous outside option vallues) that respects these characteristics. In particular, our solution concept uses the possibility of communication as a coordination device. We will show that internal nodes indeed have incentive to make these communications truthfully in the setting of our solution concept.

C.2 Axioms

We define a set of axioms for network bargaining game solutions. We show that there exists a unique solution mapping F^* satisfying these axioms, and $F^*(\sigma)$ is socially efficient for each problem σ.

Axiom I: Consistency with Nash Bargaining.

[20] Many classic cooperative games are studied by, for examples, Shapley (1953) and Myerson (1977a,b,c).

We say that F is *consistent with Nash Bargaining* if for any two-node network bilateral bargaining problem $\sigma = (G, S, C, f, \underline{u})$ (with exogenous outside options), the outcome $F(\sigma) = (s(\sigma), y(\sigma))$ is that implied by the Nash Bargaining Solution. That is,

$$s(\sigma) = \underset{s \in C}{\operatorname{argmax}} \{f_1(s) + f_2(s)\}$$
$$y^*_{ij}(\sigma) = \frac{1}{2}\left([f_i(s(\sigma)) - \underline{u}_i] - [f_j(s(\sigma)) - \underline{u}_j]\right).$$

We shall treat consistency with Nash Bargaining as an axiom, although it can be based on more primitive underlying axioms, such as the the original four axioms of Nash (1950), or the axioms of fairness and Pareto optimality of Myerson (1977b).

Axiom II: Multilateral Stability

Our next axiom, *Multilateral Stability*, provides a notion of stability of solutions. The spirit of the axiom, to be given a precise definition, is that if the bargaining outcomes for some subset of edges are "frozen" at those prescribed by a multilaterally stable solution mapping F, then applying the solution mapping F to the NBBP induced on the residual network will not change the bargaining outcomes on the residual network. One may also interpret multilateral stability as an axiom of self-consistency, in that a solution mapping F must be consistent with its own prediction when applied to a problem induced on a subnetwork.

Formally, for a solution mapping F, a given network bilateral bargaining problem $\sigma = (G, S, C, f, \underline{u})$ with exogenous outside options, and a subset $E' \subseteq E$ of edges, we let $\sigma_{\text{sub}}(F, \sigma, E')$ be the network bilateral bargaining problem induced by solution mapping F on the subnetwork with edge set E'. Letting $F(\sigma) = (s(\sigma), y(\sigma))$, then the primitives of $\sigma_{\text{sub}}(F, \sigma, E')$ are

given by

$$G' = (V, E'), \qquad \underline{u}'_{ij} = \underline{u}_{ij},$$

$$S'_i = \{s_i(\sigma)\} \text{ if } (i,j) \in E \backslash E' \text{ for some } j \in V, \qquad S'_i = S_i \text{ otherwise,}$$

$$C_{ij} = \left\{(s_i, s_j) \in S'_i \times S'_j \,:\, (s_i, s_j) \in C_{ij}\right\},$$

$$f'_i : S'_i \times \prod_{(i,j)\in E'} S'_j \ni s \quad \mapsto \quad f_i(s) - \sum_{(i,j)\in E\backslash E'} y_{ij}(\sigma) \in \mathbb{R}.$$

We say that a solution mapping F satisfies *multilateral stability* if, for any given network bilateral bargaining problem σ (with exogenous outside option values) and any given subset $E' \subseteq E$ of edges, when letting σ' denote the subproblem $\sigma_{\text{sub}}(F, \sigma, E')$, we have

$$F(\sigma') = F(\sigma)_{G'}, \tag{107}$$

where $F(\sigma)_{G'}$ denotes the restriction of the outcome $F(\sigma)$ to the subnetwork G'. When applied to the special case $E' = \emptyset$, (107) places no restriction on the solution mapping F.

Axiom III: Independence of Irrelevant Actions

We first introduce the notion of a *reduced network bilateral bargaining problem.*

Definition 5. If $\sigma = (G, S, C, f, \underline{u})$ and $\sigma' = (G, S', C', f', \underline{u})$ are two network bilateral bargaining problems (with exogenous outside options) with the same underlying graph G and the same outside option values $\underline{u}$. We say that σ' is a *reduced network bilateral bargaining problem* of σ if $S' \subseteq S$, $C' \subseteq C$ and f' is the restriction of f to S'.[21]

Our third axiom is an adaptation to this network setting of Nash's Independence of Irrelevant Alternatives. We say that F satisfies *Independence of Irrelevant Actions* if it respects the following property. We suppose that $\sigma' = (G, S', C', f', \underline{u})$ is a reduced network

[21] That is, the utility function f'_i is the restriction of f_i to the set $S'_i \times \widehat{S}'_i$, for all $i \in V$.

bilateral bargaining problem of $\sigma = (G, S, C, f, \underline{u})$. Letting $F(\sigma) = (s(\sigma), y(\sigma))$, if $s(\sigma)$ is in C' (that is, if the actions associated with the outcome $F(\sigma)$ are feasible for the reduced problem σ'), then $F(\sigma') = F(\sigma)$.

The idea of the axiom is as follows. Suppose that the agents in a given network find a solution $(s(\sigma), y(\sigma))$ to their network bilateral bargaining problem. Then, for some reason, they realize that some of their feasible actions are no longer valid choices, although $s(\sigma)$ itself remains feasible. Under Independence of Irrelevant Alternatives, taking away some actions that they would not have chosen anyway has no effect on their solution outcome: they continue to choose $(s(\sigma), y(\sigma))$.

The axiom of Independence of Irrelevant Actions is compelling only in the setting of exogenous outside options. With endogenous outside options, the removal of some "non-equilibrium" actions may influence the outside options. We will address this issue by an overlaying model structure once we introduce endogenous outside options.

Axiom IV: Bilateral Optimality

For the axiom of Bilateral Optimality, we only consider networks that are connected.[22] We will see that this restriction provides the weakest form of the axiom, yet without losing any strength. The first three axioms have not made use of the possibility of communication between pairs of connected nodes. The last axiom, *Bilateral Optimality*, relies critically on that *(i)* information is able to propagate in the network through a series of local communication between pairs of connected nodes, and that *(ii)* network players has the right incentive to make these communication truthfully. In other words, Bilateral Optimality may not be applied to any network bilateral bargaining problems. We first motivate the conditions under which Bilateral Optimality is applicable, then provide a formal definition of the axiom.

Information propagation in a network bilateral bargaining problem.

[22] A network G is connected if any given pair of nodes are path-connected.

We always proceed under the premise that the actions of a pair of connected nodes are directly observable to each other. We now consider how information about actions can be propagated more widely in a network through only a series of local communication between pairs of directly connected nodes. We will define a notion of information transfer along a path in a graph G, which is defined as an ordered list $(i_1, i_2, \ldots, i_m)$ of distinct nodes with the property that i_h is directly connected to i_{h+1} for all $h < m$. Two nodes are said to be *path-connected* if they are elements of the same path.

Given a network bargaining problem $\sigma = (G, S, C, f, \underline{u})$, we fix some pair (a, b) of directly connected nodes, along with a compatible pair (s_a, s_b) of respective actions for these two nodes. We call $\alpha = (a, s_a, b, s_b)$ a "message." Our objective is to consider the ability to transfer such a message from node to node along a path in the graph G. Freezing the actions s_a and s_b for a and b, and leaving the problem otherwise minimally affected, a restricted network bargaining problem $\sigma(\alpha) = [G, S^\alpha, C^\alpha, f, \underline{u}]$ is induced by $\alpha = (a, s_a, b, s_b)$ in the obvious way.[23]

Definition 6. Given a solution mapping F, a network bilateral bargaining problem σ with a connected graph and a message of the form $\alpha = (a, s_a, b, s_b)$, letting $F(\sigma(\alpha)) = (s(\alpha), y(\alpha))$, we say that a path (i, j, k) can transfer the message α if

$$(s_k(\alpha_{ij}), y_{kj}(\alpha_{ij})) = (s_k(\alpha), y_{kj}(\alpha)), \tag{108}$$

where $\alpha_{ij} = (i, s_i(\alpha), j, s_j(\alpha))$.

In order to motivate this notion of information transfer, we let the actions of nodes a and b be (s_a, s_b) and $\alpha = (a, s_a, b, s_b)$. We *suppose* that the information contained in the message α is available to nodes i and j, and that nodes i and j choose actions $(s_i(\alpha), s_j(\alpha))$ based on the message α. If Node k were to observe α and adhere to the solution mapping F, then the action of k and her payment to j would be given by $(s_j(\alpha), y_{kj}(\alpha))$. However, Node k cannot directly observe α (unless $k = a$ or b). Instead, node k directly observes the action

[23] That is, the action set S^α is constructed from S by replacing S_a with $\{s_a\}$ and S_b with $\{s_b\}$, the pairwise compatibility correspondences and the utility functions are the restrictions of C and f to the new action set S^α. The network and outside option values of $\sigma(\alpha)$ remain the same as those of σ.

of j and "listens" to her report regarding the action of i . If j truthfully reports the action $s_i(\alpha)$ of i to k , then the action of k and her payment to j would be $(s_k(\alpha_{ij}), y_{kj}(\alpha_{ij}))$. If (108) holds, node k thus indirectly learns any payoff relevant information about α. In this sense, node j can "virtually" transfer the information about α to k.

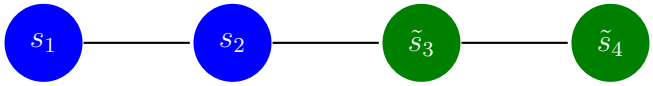

Figure 7 – A 4-node path illustration of information transfer

To give a concrete example, we consider the 4-node path illustrated in Figure 7. We suppose that nodes 1 and 2 have chosen actions (s_1, s_2), while nodes 3 and 4 have yet to choose their actions. Node 3 can observe s_2 but not s_1. Node 2, however, has the ability to inform Node 3 of s_1 when bargaining with Node 3. Hence, upon truthful report by Node 2, the information about $\alpha = (1, s_1, 2, s_2)$ can be transferred to Node 3 via the path $(1, 2, 3)$. Likewise, Node 4 is able to observe the action of Node 3, and Node 3 can provide information about s_2 when bargaining with Node 4. Now, we suppose that Node 3 has incentive to make truthful communication to 4 , and to choose action $s_3(\alpha)$ based on the message α. Then Node 4 observes the action $s_3(\alpha)$ of Node 3 and learns the action s_2 of Node 2. From the perspective of Node 4, having these information is equivalent to knowing α if equation (108) holds for $i = 2$, $j = 3$ and $k = 4$. In this way, Node 4 would effectively learn s_1. In other words, information about the actions of nodes 1 and 2 reaches Node 4 by "flowing" through the network through only a series of local communication.

More generally, if a network bilateral bargaining problem admits information propagation via a solution mapping F (a property that we will next define), then payoff-relevant information about the actions of any pair of direct counterparties can "flow" to the entire network through only a series of local communication, *provided* that the utilities and solution mapping F supply the internal nodes of the network with the "right" incentives to carry out truthful communication and act accordingly. We will discuss the strategic aspect of the internal nodes shortly.

Definition 7. A network bargaining problem σ with graph G admits *information propagation* via a bargaining solution mapping F if any path of the form (i, j, k) in G can transfer any message of the form (a, s_a, b, s_b).

If σ admits information propagation via F, a simple induction implies that for any compatible actions (s_a, s_b) of directly connected nodes (a, b), and for any node k that is connected by some path to a or b, the message $\alpha = (a, s_a, b, s_b)$ can be recursively transferred to Node k through that path via the solution mapping F. For a path of length $m \geq 2$, message α experiences $m-1$ consecutive transfers by the internal nodes along the path before reaching node k. Figure 7 corresponds to the case where $m = 3$.

The next proposition gives an equivalent characterization of this property regarding information propagation, which is based on a "global" property of the bargaining problem σ. This global property of σ is easier to be verified than its local version (108). Given a message α, we write $F(\sigma(\alpha)) = (s(\alpha), y(\alpha))$.

Definition 8 (Stable Actions)**.** Given a solution mapping F and an network bargaining problem σ, a feasible action vector s is *stable* with respect to (F, σ) if, for any directly connected pair (i, j) of nodes, when letting $\alpha = (i, s_i, j, s_j)$, we have

$$s(\alpha) = s.$$

That is, s is stable with respect to (F, σ) if, when applying the solution mapping F to the version of σ obtained by freezing the actions of nodes i and j, we obtain the same action vector s.

Proposition 19. *Given a solution mapping F that satisfies the axiom of Multilateral Stability, a network bargaining problem σ admits information propagation via F if and only if the following holds: for any pair (i, j) of directly connected nodes, and for any pair of compatible actions (s_i, s_j), when letting $\alpha = (i, s_i, j, s_j)$, the vector $s(\alpha)$ of actions is stable with respect to (F, σ).*

The type of "frozen behavior" that is considered here differs in two respects from that

associated with the axiom of Multilateral Stability. Information propagation is applied in each instance by freezing only the actions of a single pair of nodes, whereas for Multilateral Stability, the actions and the payments of a general subgraph are frozen. Second, for information propagation, the frozen behavior need not be that determined by the solution mapping F, whereas for Multilateral Stability, the frozen behavior is prescribed by F. Although the two notions of stability are mildly related, they serve rather different purposes in characterizing a solution.

In the notion of information propagation, all nodes are *assumed* to *(i)* choose actions as prescribed by the solution mapping F, and *(ii)* carry out truthful communication with its direct counterpatires. Any strategic deviation by nodes is ruled out by assumption. Next, we will give conditions on utility functions under which it is incentive compatible for nodes to act as such.

Incentive Compatibility.

We fix a network bilateral bargaining problem $\sigma = (G, S, C, f, \underline{u})$ and a solution mapping F. Given a node i, since the graph G is connected and contains no circle, G can be viewed as a tree with root node i. Suppose that node j is an immediate offspring of i in the tree. We write $\mathcal{D}_i(j)$ for the set of all descendants of j (including node j itself) in this tree.

Definition 9. A given solution mapping F is *incentive compatible* for a given network bilateral bargaining problem $\sigma = (G, S, C, f, \underline{u})$ (with exogenous outside option values) if the following two conditions are satisfied:

(i) The problem σ admits information propagation via F.

(ii) For every nodes i and $j' \in \mathcal{N}_i(G)$, every action s_i of node i, and every actions $(s_j)_{j \in \mathcal{N}_i(G)}$ of the nearest neighbors of i that are compatible with s_i, letting $\alpha_{i,j} = (i, s_i, j, s_j)$,

$$U\left(s_i, \left(s_k\left(\alpha_{i,j}\right)\right)_{j \in \mathcal{N}_i(G), k \in \mathcal{D}_i(j)}\right) - \sum_{j \in \mathcal{N}_i(G)} \sum_{k \in \mathcal{D}_i(j)} u_k\left(F\left(\sigma\left(\alpha_{i,j}\right)\right)\right) \leq u_i\left(F\left(\sigma\left(\alpha_{i,j'}\right)\right)\right) \tag{109}$$

To motivate this notion of incentive compatibility, we fix a node i, and view the graph G as a tree with root node i. Suppose the action of node i is s_i. We consider whether node i has incentive to coordinate the actions $(s_j)_{j\in\mathcal{N}_i(G)}$ of its nearest neighbors according to the bargaining solution mapping F. Given a nearest neighbor $j'\in\mathcal{N}_i(G)$ of i, if Node i truthfully communicate the action $s_{j'}$ of j' to its other nearest neighbors, then the net payoff of Node i is $u_i\left(F\left(\sigma\left(\alpha_{i,j'}\right)\right)\right)$. This is the benchmark payoff of node i that we will compare against when evaluating alternative choices of i. If Node i misguides its direct counterparties in a way that results in their actions being $(s_j)_{j\in\mathcal{N}_i(G)}$, then the net payoff of Node i is

$$u_i\left(s_i,(s_j)_{j\in\mathcal{N}_i(G)}\right)-\sum_{j\in\mathcal{N}_i(G)}y_{ij}(\alpha_{i,j}). \tag{110}$$

In particular, this payoff only depends on the actions $\left(s_i,(s_j)_{j\in\mathcal{N}_i(G)}\right)$ of Node i and its direct countarparties $j\in\mathcal{N}_i(G)$, and not on the bargaining outcomes further down the tree. This plays a key role in simplifying the calculation of the payoff of node i.

Suppose that every node $j\in\mathcal{N}_i(G)$ truthfully communicates the action s_i of Node i to its immediate offsprings, who then carry out truthful communication with their immediate offsprings, etcetera, then the message (i,s_i,j,s_j) would be transferred to all descendants of node j. The utility of some descendant node $k\in\mathcal{D}_i(j)$ is thus $u_k\left(F\left(\sigma\left(\alpha_{i,j}\right)\right)\right)$, and the total utility of all nodes in the network is

$$U\left(s_i,\left(s_k\left(\alpha_{i,j}\right)\right)_{j\in\mathcal{N}_i(G),k\in\mathcal{D}_i(j)}\right).$$

Thus, the net payoff of node i is given by the left hand side of (109). The two payoffs of node i, given by (110) and the left hand side of (109) respectively, must be equal. If this payoff is less than the benchmark payoff $u_i\left(F\left(\sigma\left(\alpha_{i,j'}\right)\right)\right)$ (that is, if inequality (109) holds), then node i has incentive to truthfully communicate the action $s_{j'}$ of j' to its other direct counterparties.

Roughly speaking, if a solution mapping F is incentive compatible for a network bilateral bargaining problem, then every node, given its information, has the incentive to carry out

truthful communication and to choose actions as prescirbed by F. Incentive compatibility implies that if a pair (i, j) of directly connected nodes is committed to some respective actions (s_i, s_j), then the behavior of the entire network would be given by $F(\sigma(s_i, s_j))$. This is true even when (s_i, s_j) differ from those prescribed by the solution mapping F. By way of comparison, when thinking of F as an "equilibrium" solution concept, Multilateral Stability is concerned with behavior "in equilibrium," whereas Incentive Compatibility can apply to "out-of-equilibrium" behavior.

The next proposition gives a sufficient condition for incentive compatibility, which is easier to be verified than the original definition.

Proposition 20. *A solution mapping F is incentive compatible for a network bilateral bargaining problem $\sigma = (G, S, C, f, \underline{u})$ (with exogenous outside option values) if*

(i) the problem σ admits information propagation via F, and

(ii) for any feasible action vector $s \in C$, associating a message $\alpha_k = (i_k, s_{i_k}, j_k, s_{j_k})$ (for some $(i_k, j_k) \in E$) to every node k, we have

$$\sum_{k \in V} [u_k(F(\sigma(\alpha_k))) - f_i(s)] \geq 0.$$

Bilateral Optimality.

Loosely speaking, a solution mapping F satisfies *Bilateral Optimality* if, whenever two directly connected nodes assume that other nodes will react to their chosen actions as specified by F, it is optimal for these two nodes to themselves choose outcomes implied by F. More precisely, solution mapping F satisfies Bilateral Optimality if, for any network bilateral bargaining problem $\sigma = (G, S, C, f, \underline{u})$ (with exogenous outside option values) for which F is incentive compatible, and for any directly connected nodes i and j, letting

$F(\sigma) = (s(\sigma), y(\sigma))$, the actions $(s_i(\sigma), s_j(\sigma))$ solve

$$\max_{(s_i,s_j)\in C_{ij}} \Big\{ f_i\Big(s_i, s_j, (s_k(i,s_i,j,s_j))_{k\in\mathcal{N}_i(G)\setminus\{j\}}\Big) - \sum_{k\in\mathcal{N}_i(G)\setminus\{j\}} y_{ik}(i,s_i,j,s_j) + f_j\Big(s_j, s_i, (s_k(i,s_i,j,s_j))_{k\in\mathcal{N}_j(G)\setminus\{i\}}\Big) - \sum_{k\in\mathcal{N}_j(G)\setminus\{i\}} y_{jk}(i,s_i,j,s_j)\Big\}.$$

That is, the actions $(s_i(\sigma), s_j(\sigma))$ maximize the total utility of nodes i and j associated with the common conjecture that the remaining network $G_{\text{sub}} = (V, E\setminus\{(i,j)\})$ will achieve its own solution outcome $F(\sigma(i,j,s_i,s_j))$, given the actions (s_i, s_j) agreed by nodes i and j.

The Axiomatic Solution

Theorem 4. *There is a unique solution mapping F^* satisfying Axioms I-IV. For each NBBP $\sigma = (G, S, C, f, \underline{u})$, the solution outcome $F^*(\sigma) = (s^*(\sigma), y^*(\sigma))$ is socially efficient (that is, $s^*(\sigma) = s^{**}(\sigma)$), and the payment vector $y^*(\sigma)$ solves the system of linear equations*

$$u_i(s^{**}(\sigma), y^*(\sigma)) - \underline{u}_{ij} = u_j(s^{**}(\sigma), y^*(\sigma)) - \underline{u}_{ji}, \quad (i,j) \in E. \tag{111}$$

The system (111) of payment equations states that the solution outcome $F^*(\sigma)$ provides for an equal sharing of the surplus between each pair of directly connected nodes, relative to their respective outside options.

C.3 Endogenous Outside Option Values

The assumption of exogenous outside option values $\underline{u}$ is a useful modeling technique, in that it allows us to establish the four natural axioms that uniquely determine the solution mapping F^*. Without this assumption, Independence of Irrelevant Actions is no longer an compelling axiom, since the removal of some "non-equilibrium" actions may influence the "threat point" of various contracting pairs, and thus their bargaining outcomes.

In this section, we consider a more natural model of network bilateral bargaining problems, in which the outside option values are endogenously determined by the set of available actions of each node in the event that the node fails to reach an agreement with some of its direct counterparties. A *network bilateral bargaining problem with endogenous outside option values* is defined as

$$\gamma = (G, S, C, f, \underline{S}),$$

where G, S, C and f are the same model components as for a network bilateral bargaining problem with exogenous outside option values, and $\underline{S}$ consists of subsets $\underline{S}_i(N_i)$ of S_i, for each node i and each subset N_i of nearest neighbors of i. The set $\underline{S}_i(N_i)$ consists of actions that are available to i if i fails to reach an agreement with all its nearest neighbors in the set N_i. If disagreements are reached between node i and all its direct counterpatires (that is, if $N_i = \mathcal{N}_i(G)$), then by convention, the $S_i^0(N_i)$ is a singleton. Naturally, no payment is made between a pair of disagreeing nodes. When applying this setting to the special case of a two-node network, a unique "conflict outcome" is thus implemented if the two nodes fail to reach an agreement, as in Nash (1950).

All disagreement events are assumed to be observed by the entire network, and thus become common knowledge *prior* to any agreement is made. This assumption is identical to the possibility of renegotiation in the bargaining protocol of Stole and Zwiebel (1996). In contrast, the actions and payment negotiated by a given pair of direct counterparties are not observable to other nodes in the network, an assumption that remains the same as for a network bilateral problem with exogenous outside option values. In other words, whether or not a given pair of nodes reaches an agreement is common knowledge, the detailed terms of the agreed-upon contract (if there is one) is nevertheless not observable to other nodes.

An *outcome*, in the case of endogenous outside option values, remains the same as that for the case of exogenous outside option values. That is, an outcome (s, y) specifies a feasible action vector and a payment for each edge. A *solution mapping* is some function Φ that maps an network bilateral bargaining problem with endogenous outside options to an outcome.

Given a network bilateral bargaining problem $\gamma = ((V, E), S, C, f, \underline{S})$ with endogenous

outside option values, our approach is to transform γ into an NBBP with exogenous outside options, and then deduce the solution outcome using the unique solution mapping F^* that satisfies the four axioms. Given a pair $(i,j) \in E$ of directly connected nodes. In the event that the pair fails to reach an agreement, we obtain a network bargaining problem γ_{ij} (with endogenous outside option values) that is induced from γ in the obvious way.[24] In particular, there are $|E|-1$ edges in the network of γ_{ij}, as the edge (i,j) is removed. When nodes i and j carry out bilateral bargaining, their outside option values $\left(\underline{u}_{ij}, \underline{u}_{ji}\right)$ are given by their respective payoffs in the event of a disagreement between the pair. Therefore,

$$\underline{u}_{ij} = u_i\left(\Phi\left(\gamma_{ij}\right)\right), \qquad \underline{u}_{ji} = u_j\left(\Phi\left(\gamma_{ij}\right)\right). \tag{112}$$

We let $s^{**}(\gamma)$ be the socially efficient action vector for γ, and $y^*(\underline{u})$ be the unique payment vector that solves the linear system

$$u_i(s^{**}(\gamma), y^*(\underline{u})) - \underline{u}_{ij} = u_j(s^{**}(\gamma), y^*(\underline{u})) - \underline{u}_{ji}, \qquad \forall (i,j) \in E.$$

If the outside option values $\underline{u}$, derived from equation (112), satisfy

$$u_i(s^{**}(\gamma), y^*(\underline{u})) - \underline{u}_{ij} = u_j(s^{**}(\gamma), y^*(\underline{u})) - \underline{u}_{ji} \geq 0, \qquad \forall (i,j) \in E, \tag{113}$$

we say that $\underline{u}$ satisfies *the condition for outside option values*. If this is the case, then $\sigma(\gamma) = (G, S, C, f, \underline{u})$ is a well defined NBBP (with exogenous outside option values), and the problem γ, which has endogenous outside option values, can be equivalent transformed to $\sigma(\gamma)$. Hence, the solution outcome for γ must be given by

$$\Phi(\gamma) = F^*(\sigma(\gamma)).$$

The next theorem establishes results on existence, uniqueness and efficiency of the solution

[24] That is, the network of γ_{ij} is given by $G' = (V, E \backslash \{(i,j)\})$, the action set S' is constructed from S by replacing S_i with $\underline{S}_i(\{j\})$ and S_j with $\underline{S}_j(\{i\})$, the pairwise compatibility correspondences and the utility functions are the respective restrictions of C and f to the new action set S', and the disagreement action sets are constructed from $\underline{S}$ by replacing $\underline{S}_i(N_i)$ with $\underline{S}_i(N_i \cup \{j\})$ and $\underline{S}_j(N_j)$ with $\underline{S}_j(N_j \cup \{i\})$, for every subsets $N_i \subseteq \mathcal{N}_i(G')$ and $N_j \subseteq \mathcal{N}_j(G')$.

to network bilateral bargaining problems with endogenous outside option values.

Theorem 5. *(i) There is a unique solution mapping Φ^* such that, for every network bilateral bargaining problem $\gamma = [(V, E), S, C, f, \underline{S}]$ (with endogenous outside option values), letting*

$$\underline{u}^*_{ij}(\gamma) = u_i\left(\Phi^*\left(\gamma_{ij}\right)\right), \qquad \underline{u}^*_{ji}(\gamma) = u_j\left(\Phi^*\left(\gamma_{ij}\right)\right). \tag{114}$$

for every pair $(i, j) \in E$, and

$$\sigma(\gamma) = (G, S, C, f, \underline{u}^*),$$

if $\underline{u}^*(\gamma)$ satisfies the condition for outside option values, then

$$\Phi^*(\gamma) = F^*(\sigma(\gamma)). \tag{115}$$

(ii) For any given network bilateral bargaining problem γ (with endogenous outside option values), the condition for outside option values holds for $\underline{u}^*(\gamma)$ derived from equation (114). (iii) For every network bilateral bargaining problem γ with endogenous outside option values, we have $\Phi^*(\gamma) = F^*(\sigma(\gamma))$. In particular, $\Phi^*(\gamma)$ is socially efficient.

The proof of Theorem 5, given in Appendix D.6, relies on a useful intermediate result regarding how the hold-up power of various nodes differentiates in a network. We next define this hold-up power and establish this intermediate result.

We fix a network bilateral bargaining problem $\gamma = (G, S, C, f, \underline{S})$. Again without loss of generality, we assume that the underlying graph $G = (V, E)$ is connected. Given a pair $(i, j) \in E$ of direct counterparties, if one removes the edge (i, j), the graph $G' = (V, E\backslash\{(i, j)\})$ can be viewed as two disjoint trees with roots i and j respectively. Any given node $k \in V$ either belongs to the tree rooted at i or at j. If k is in the tree rooted at i, and the length of the shortest path connecting from i to k is d, we say that k is a d-to-i node. A d-to-j node is similarly defined. We now suppose that k is a 1-to-i node (that is, k is a direct counterparty of i and $k \neq j$), and consider the bilateral bargaining problem between

i and k. The gap between the outside option values of the two nodes is

$$\underline{u}^*_{ki}(\gamma) - \underline{u}^*_{ik}(\gamma). \tag{116}$$

In the event that the pair (i, j) fails to reach an agreement, node k is able to hold up node i, and this gap becomes

$$\underline{u}^*_{ki}(\gamma_{ij}) - \underline{u}^*_{ik}(\gamma_{ij}). \tag{117}$$

The second outside option value gap (117) is wider than the first one (116), as

$$\begin{aligned}(117) - (116) &= \underline{u}^*_{ki}(\gamma_{ij}) - \underline{u}^*_{ik}(\gamma_{ij}) - [\underline{u}^*_{ki}(\gamma) - \underline{u}^*_{ik}(\gamma)] \\ &= u_k(\Phi^*((\gamma_{ik})_{ij})) - u_i(\Phi^*((\gamma_{ik})_{ij})) - [u_k(\Phi^*(\gamma_{ik})) - u_i(\Phi^*(\gamma_{ik}))] \\ &= u_i(\Phi^*(\gamma_{ik})) - u_i(\Phi^*((\gamma_{ik})_{ij})) \geq 0,\end{aligned} \tag{118}$$

where $(\gamma_{ik})_{ij}$ is the network bilateral bargaining problem (with endogenous outside option values) that is induced from γ after both pairs (i, j) and (i, k) fail to reach an agreement. The last inequality follows from the induction hypothesis, since the underlying graph of γ_{ik} has $|E|-1$ edges. The last equality follows from the fact that node k receives the same net payoff in γ_{ik} or $(\gamma_{ik})_{ij}$. This is because node k is not connected to i or j in the network bilateral bargaining problem γ_{ik}, the disagreement event between (i, j) is thus payoff irrelevant for node k.

This difference (118) between the two outside option value gaps is the *hold-power of k against i in the event of a disagreement between (i, j)*. The higher this difference is, the more the disagreement event hurts the bargaining position of i relative to k in her bilateral bargaining against k, and thus to a greater extent node k is able to hold up node i. For simplicity of terminology, we sometimes simply refer to this difference as the *hold-up power of k against i*, without specifying the disagreeing pair (i, j).

If k is a d-to-i node, where the distance d is strictly larger than 1, letting $(\ell_0, \ldots, \ell_d)$ be

the unique path[25] connecting node k to i, we define the hold-up power of k against i as

$$\sum_{a=0}^{d-1} \Big(\quad \big[\underline{u}^*_{\ell_a \ell_{a+1}}(\gamma_{ij}) - \underline{u}^*_{\ell_{a+1} \ell_a}(\gamma_{ij})\big] \\ - \big[\underline{u}^*_{\ell_a \ell_{a+1}}(\gamma) - \underline{u}^*_{\ell_{a+1} \ell_a}(\gamma)\big] \Big)$$

That is, the hold-up power of k against i is the the sum of the differences between the pairwise outside option value gaps along the path $(\ell_0, \ldots, \ell_d)$ connecting node k to i.

Proposition 21. *Given a pair of direct counterparties (i, j) and a d-to-i node k for some integer $d \geq 1$, The hold-up power of k against i (in the event of a disagreement between (i, j)) is non-negative.*

D Proofs for Appendix C

D.1 Determining the Payment Vector

The following proposition determines the payments using Multilateral Stability.

Proposition 22. *If a solution mapping F satisfies Multilateral Stability, then for each network bilateral bargaining problem $\sigma = (G, S, C, f, \underline{u})$, the solution outcome $F(\sigma)$ satisfies*

$$u_i(F(\sigma)) - \underline{u}_{ij} = u_j(F(\sigma)) - \underline{u}_{ji}, \quad \forall\, (i, j) \in E. \tag{119}$$

Proof. For any given pair (i, j) of connected nodes, we consider the network bilateral bargaining problem $\sigma' \equiv \sigma_{\text{sub}}(F, \sigma, \{(i, j)\})$ induced by F on the subnetwork with edge set $\{(i, j)\}$. Then σ' is simply a two-person bargaining problem with complete information between i and j, whose outside options are $\underline{u}_{ij}$ and $\underline{u}_{ji}$ respectively. It follows from Axiom I that $F(\sigma')$ is given by the Nash Bargaining Solution. Since i and j have transferable utility, the two nodes must have equal gain from trading. Thus, equation (119) must hold. □

We let V_i be the subset of nodes that are path-connected to i, and $F(\sigma) = (s(\sigma), y(\sigma))$.

[25] In particular, $\ell_0 = k, \ell_d = i$.

Proposition 23. *The system* (119) *of linear equations, together with*

$$\sum_{j \in V_i} u_j(F(\sigma)) = \sum_{j \in V_i} f_j(s(\sigma)), \tag{120}$$

uniquely determines the net payoff $u_i(F(\sigma))$ *of each node* i*,*

$$u_i(F(\sigma)) = c_i(\underline{u}) + \frac{1}{|V_i|} \sum_{j \in V_i} f_j(s(\sigma)),$$

where the constant $c_i(\underline{u})$ *depends only on the outside options* $\{\underline{u}_j : j \in V_i\}$*, and* $\sum_{j \in V_i} c_j(\underline{u}) = 0$*. The payment vector* $y(\sigma)$ *is uniquely determined by* (119)*.*

Proof. We fix a node $i \in V$. For every node j that is path-connected with i, letting $(\ell_0, \ldots, \ell_d)$ be unique shortest path connecting from j to i, we have

$$u_j(F(\sigma)) = u_i(F(\sigma)) + \sum_{a=0}^{d-1} \left(\underline{u}_{\ell_a \, \ell_{a+1}} - \underline{u}_{\ell_{a+1} \, \ell_a} \right).$$

That is, we can write $u_j(F(\sigma))$ as the sum of $u_i(F(\sigma))$ and a constant that only depends on the outside option values $\underline{u}$. This is true for every node $j \in V_i$. Combining this observation with equation (120), we can thus uniquely solve for $u_i(F(\sigma))$ as the sum of $u_i(F(\sigma))$ and a constant that only depends on $\underline{u}$.

Since the underlying graph G contains no circles, it is a disjoint union of trees. One can direct each such tree, by saying that $J(i)$ is the unique *parent* of node i (and i is an *immediately offspring* of $J(i)$) if i and $J(i)$ are directly connected and if the path from $J(i)$ to the unique root node is shorter than the path from i to the root node. One can inductively compute the payments, beginning with those from the *leave* nodes (the nodes with no offspring nodes). Each leave node i pays to her parent $J(i)$ the amount

$$y_{i,J(i)}(\sigma) = f_i(s(\sigma)) - u_i(F(\sigma)).$$

If node i is neither a leave node nor the root node, her payment to her parent node is then

$$y_{i,J(i)} = f_i(s(\sigma)) - u_i(F(\sigma)) + \sum_{j \in K(i)} y_{ji}(\sigma),$$

where $K(i)$ is the set of immediate offspring of node i, and $\{y_{ji}(\sigma) : j \in K(i)\}$ have been determined by induction. The induction ends with the root node, which has no parent. □

D.2 Proof of Proposition 19

We fix a network bilateral bargaining problem σ (with exogenous outside option values) that admits information propagation via a given solution mapping F, and fix a pair (a, b) of direct counterparties, along with some pair of compatible actions (s_a, s_b). We let $\alpha = (a, s_a, b, s_b)$. For any path of the form (i, j, k), we have

$$F_{jk}\left(\sigma\left(i, s_i(\alpha), j, s_j(\alpha)\right)\right) = (s_k(\alpha), y_{kj}(\alpha)).$$

In particular, letting $\alpha_{ij} = (i, s_i(\alpha), j, s_j(\alpha))$, we have

$$s_k(\alpha_{ij}) = s_k(\alpha).$$

This is true for every k that is a 1-to-j node. We suppose that the equality above holds for every k that is a d'-to-j node, for every integer $d' < d$ with some integer $d \geq 2$. Given a d-to-j node k, letting $(\ell_0, \ldots, \ell_d)$ be the unique shortest path connecting from j to k, we have

$$s_k(\alpha_{\,\ell_{d-2}\,\ell_{d-1}}) = s_k(\alpha). \tag{121}$$

Since ℓ_{d-2} and ℓ_{d-1} are $(d-2)$-to-j and $(d-1)$-to-j nodes respectively, the induction hypothesis implies

$$s_{\,\ell_{d-2}}(\alpha) = s_{\,\ell_{d-2}}(\alpha_{ij}), \qquad s_{\,\ell_{d-1}}(\alpha) = s_{\,\ell_{d-1}}(\alpha_{ij}).$$

Substituting the two equalities above into (121), we obtain

$$s_k(\ell_{d-2}, s_{\,\ell_{d-2}}(\alpha_{ij}), \ell_{d-1}, s_{\,\ell_{d-1}}(\alpha_{ij})) = s_k(\alpha).$$

On the other hand, since the path $(\ell_{d-2}, \ell_{d-1}, k)$ can transfer the message α_{ij}, it follows that

$$s_k(\ell_{d-2}, s_{\ell_{d-2}}(\alpha_{ij}), \ell_{d-1}, s_{\ell_{d-1}}(\alpha_{ij})) = s_k(\alpha_{ij}).$$

Hence, $s_k(\alpha) = s_k(\alpha_{ij})$ for every node k that is a $(d+1)$-to-j node. It follows from induction and symmetry that this holds for every node k in the graph G. Since (i, j) can be any pair of direct counterparites, the vector $s(\alpha)$ of actions is thus stable with respect to (F, σ).

We now establish the converse. We fix a solution mapping F that satisfies Multilateral Stability and a network bilateral bargaining problem σ (with exogenous outside option values) such that the stability property stated in Proposition 19 holds. We also fix a pair (a, b) of direct counterparties, along with some pair of compatible actions (s_a, s_b). We let $\alpha = (a, s_a, b, s_b)$. For any path of the form (i, j, k), we let $\alpha_{ij} = (i, s_i(\alpha), j, s_j(\alpha))$. Since the action vector $s(\alpha)$ is stable with respect to (F, σ), it follows that $s(\alpha_{ij}) = s(\alpha)$. Since F satisfies Multilateral Stability, it follows from Proposition 23 that $y(\alpha_{ij}) = y_\alpha$. Hence, the path (i, j, k) can transfer the message α.

D.3 Proof of Proposition 20

Proposition 20 follows from the observation that condition *(ii)* in Definition 9 is a special case of condition *(ii)* in Proposition 20.

D.4 Proof of Theorem 4

We first prove uniqueness, then existence.

Uniqueness: We first set up the following definition.

Definition 10 (Quasi-Maximizing Actions)**.** For an NBBP $\sigma = ((V, E), S, C, f, \underline{u})$, a feasible action vector $s^* \in C$ *quasi-maximizes* social welfare for σ if and only if, for any pair (i, j) of directly connected nodes, s^* solves

$$\max_{s \in C(i,j,s_i^*,s_j^*)} U_\sigma(s),$$

where

$$C\left(i,j,s_i^*,s_j^*\right) = \{s \in C : s_i = s_i^*, s_j = s_j^*).$$

We let $S^*(\sigma)$ denote the set of action vectors that quasi-maximize the social welfare for σ.

Lemma 14. *Suppose $\sigma = ((V,E),S,C,f,\underline{u})$ and $\sigma' = ((V,E),S',C',f',\underline{u})$ are such that σ' is a reduced NBBP of σ. If $s^* \in S^*(\sigma) \cap C'$, then $s^* \in S^*(\sigma')$.*

Proof. Fix some $(i,j) \in E$. Because $C' \subseteq C$ and $s^* \in C'$, we know that $s^* \in C'\left(i,j,s_i^*,s_j^*\right)$. Thus s^* solves

$$\max_{s \in C'\left(i,j,s_i^*,s_j^*\right)} U_\sigma(s).$$

As $U_{\sigma'}$ and U_σ agree on S', s^* thus solves

$$\max_{s \in C'\left(i,j,s_i^*,s_j^*\right)} U_{\sigma'}(s).$$

Because (i,j) was arbitrary, it follows that $s^* \in S^*(\sigma')$. □

For a positive integer k, we let Σ_k denote the collection of network bilateral bargaining problems with exogenous outside option values, whose graph (V,E) has $|E| = k$ edges. Suppose F^* is a solution satisfying Axioms I-IV. Letting $F^*(\sigma) = (s^*(\sigma), y^*(\sigma)))$, the uniqueness part is summarized as the following claim, which we will establish by an induction over the number $|E|$ of edges.

Claim 2. For every NBBP $\sigma = [(V,E),S,C,f,\underline{u}]$ with exogenous option values, $s^*(\sigma) = s^{**}(\sigma)$ and $y^*(\sigma)$ solves (111).

When $|E| = 1$, we are in the situation of a two-person bargaining problem. Suppose $\sigma \in \Sigma_1$ is such a problem. Then Axiom I implies $s^*(\sigma) = s^{**}(\sigma)$ and the equation (111) of equitable split of trading gain.

Suppose that Claim 2 holds for $|E| = p-1$. We show that it holds for $|E| = p$. We fix an arbitrary NBBP $\sigma = [(V,E),S,C,f,\underline{u}]$ in Σ_p.

Step 1: The action vector $s^*(\sigma)$ quasi-maximizes the social welfare for σ.

For any fixed pair (i,j) of directly connected nodes. We consider the NBBP $\sigma' \equiv \sigma_{\text{sub}}(F^*, \sigma, E')$ induced on the edge set $E' = E \backslash \{i,j\}$. Letting $G' = (V, E')$, since F^* satisfies Multilateral Stability, we have $F^*(\sigma') = F^*(\sigma)_{G'}$, which implies that

$$s^*(\sigma') = s^*(\sigma).$$

Since the edge set of the subproblem σ' is E', we have $\sigma' \in \Sigma_{p-1}$. By the induction hypothesis, the action vector $s^*(\sigma')$ of the solution outcome $F^*(\sigma')$ is given by

$$s^*(\sigma') = s^{**}(\sigma') = \underset{s \in C\left(i,j,s_i^*(\sigma), s_j^*(\sigma)\right)}{\operatorname{argmax}} U_{\sigma'}(s).$$

Since $s^*(\sigma) = s^*(\sigma')$, and $U_\sigma(s) = U_{\sigma'}(s)$ for every $s \in C\left(i,j,s_i^*(\sigma), s_j^*(\sigma)\right)$, we have

$$s^*(\sigma) = \underset{s \in C\left(i,j,s_i^*(\sigma), s_j^*(\sigma)\right)}{\operatorname{argmax}} U_{\sigma}(s).$$

This is true for all $(i,j) \in E$, therefore, $s^*(\sigma)$ quasi-maximizes social welfare for σ. That is, $s^*(\sigma) \in S^*(\sigma)$.

We summarize the result of *Step 1* with the following lemma, which we will later use in other parts of the proof.

Lemma 15. *Assuming the induction hypothesis that Claim 2 holds for $|E| = p-1$, Multilateral Stability implies that for any NBBP $\sigma \in \Sigma_p$, $s^*(\sigma) \in S^*(\sigma)$.*

Step 2: Application of Independence of Irrelevant Actions.

We let $S_i^*(\sigma)$ be the projection of the quasi-maximizing action set $S^*(\sigma)$ onto S_i, and $S_{ij}^*(\sigma)$ denote the projection of $S^*(\sigma)$ onto $S_i \times S_j$. We let $S_{\text{rec}}^*(\sigma) = \prod_{i \in V} S_i^*(\sigma)$ be the rectangle hull of $S^*(\sigma)$, and $f^*(\sigma) = f_{|S_{\text{rec}}^*(\sigma)}$ denote the restriction of f to $S_{\text{rec}}^*(\sigma)$. For any $(i,j) \in E$, we let $C_{ij}^*(\sigma) = S_{ij}^*(\sigma)$, and

$$C^*(\sigma) = \left\{ s \in S : (s_i, s_j) \in C_{ij}^*(\sigma),\ \forall (i,j) \in E \right\}.$$

Therefore, $S^*(\sigma) \subseteq C^*(\sigma)$, ensuring that $C^*(\sigma)$ is feasible. We define the NBBP

$$\sigma^* = \left[G, S^*_{\text{rec}}(\sigma), C^*(\sigma), f^*(\sigma), \underline{u}\right].$$

Comparing σ^* to the original NBBP σ, we have $S^*_{\text{rec}}(\sigma) \subseteq S$ and $C^*(\sigma) \subseteq C$. Thus, σ^* is a reduced NBBP of σ. Because $s^*(\sigma) \in S^*(\sigma)$ and $S^*(\sigma) \subseteq C^*(\sigma)$, we have

$$s^*(\sigma) \in C^*(\sigma).$$

Because F^* satisfies Independence of Irrelevant Actions, we have $F^*(\sigma) = F^*(\sigma^*)$. Hence, it suffices to characterize the solution outcome $F^*(\sigma^*)$ for the reduced NBBP σ^*.

Remark 6. Because $S^*(\sigma) \subseteq C^*(\sigma)$ and the socially optimal actions $s^{**}(\sigma)$ quasi-maximize U_σ, we see that $s^{**}(\sigma) \in C^*(\sigma)$. Thus,

$$s^{**}(\sigma^*) = \underset{s \in C^*(\sigma)}{\operatorname{argmax}}\, U_{\sigma^*}(s) = \underset{s \in C^*(\sigma)}{\operatorname{argmax}}\, U_\sigma(s) = s^{**}(\sigma).$$

Step 3: The solution F^* is incentive compatible for the NBBP σ^*.

First, we characterize the set of actions that are stable with respect to (F^*, σ). A feasible action vector $s \in C$ is stable with respect to (F^*, σ) if, for any directly connected pair (i, j) of nodes, when letting $\alpha = (i, s_i, j, s_j)$, we have $s^*(\alpha) = s$, where $s^*(\alpha)$ denotes $s^*(\sigma(\alpha))$. As $\sigma(\alpha) \in \Sigma_p$, Lemma 15 applied to the NBBP $\sigma(\alpha)$ implies that $s^*(\alpha)$ quasi-maximizes social welfare for $\sigma(\alpha)$. In particular, if we fix the actions of (i, j) to be (s_i, s_j), the remaining components of $s^*(\alpha)$ jointly maximize the total social welfare $U_{\sigma(\alpha)}$ of $\sigma(\alpha)$. As $U_{\sigma(\alpha)}$ and U_σ are equal on the set $C(i, j, s_i, s_j)$, we have

$$s^*(\alpha) = \underset{\tilde{s} \in C(i,j,s_i,s_j)}{\operatorname{argmax}}\, U_\sigma(\tilde{s}).$$

Therefore, a feasible action vector s is stable with respect to (F^*, σ) if and only if we have, for all $(i, j) \in E$,

$$s = \underset{\tilde{s} \in C(i,j,s_i,s_j)}{\operatorname{argmax}}\, U_\sigma(\tilde{s}),$$

which is equivalent to $s \in S^*(\sigma)$.

We summarize the result of this first part as follows.

Lemma 16. *Assuming the induction hypothesis that Claim 2 holds for $|E| = p-1$. For any NBBP $\sigma \in \Sigma_p$, Multilateral Stability implies that an action vector s is stable with respect to (F^*, σ) if and only if $s \in S^*(\sigma)$.*

Second, we show that the solution F^* admits information propagation for the NBBP σ^*. We fix a pair $(i,j) \in E$ and some actions $\left(s_i', s_j'\right) \in C_{ij}^*(\sigma)$. Letting $\alpha' = \left(i, s_i', j, s_j'\right)$, we need to show that $s^*\left(\sigma^*\left(\alpha'\right)\right)$ is stable with respect to (F^*, σ^*). As $\sigma^* \in \Sigma_p$, Lemma 16 implies that the stability of $s^*\left(\sigma^*\left(\alpha'\right)\right)$ with respect to (F^*, σ^*) is equivalent to

$$s^*\left(\sigma^*\left(\alpha'\right)\right) \in S^*(\sigma^*). \tag{122}$$

As $\sigma^*\left(\alpha'\right) \in \Sigma_p$, Lemma 15 implies that $s^*\left(\sigma^*\left(\alpha'\right)\right) \in S^*\left(\sigma^*\left(\alpha'\right)\right)$. In particular,

$$s^*\left(\sigma^*\left(\alpha'\right)\right) = \underset{s \in C^*\left(i,j,s_i',s_j'\right)}{\operatorname{argmax}} U_{\sigma^*}(s). \tag{123}$$

Since $\left(s_i', s_j'\right) \in C_{ij}^*(\sigma)$, then by the definition of $C_{ij}^*(\sigma)$, there exists some action vector $\bar{s} \in S^*(\sigma)$ such that $\bar{s}_i = s_i'$ and $\bar{s}_j = s_j'$. Thus, $\bar{s} \in C^*\left(i, j, s_i', s_j'\right)$ and

$$\bar{s} = \underset{s \in C\left(i,j,s_i',s_j'\right)}{\operatorname{argmax}} U_{\sigma}(s) = \underset{s \in C^*\left(i,j,s_i',s_j'\right)}{\operatorname{argmax}} U_{\sigma}(s) = \underset{s \in C^*\left(i,j,s_i',s_j'\right)}{\operatorname{argmax}} U_{\sigma^*}(s). \tag{124}$$

The last equality holds because U_{σ^*} and U_{σ} are equal on $C^*(\sigma)$. Comparing the two maximization problems in (123) and (124), we have $\bar{s} = s^*\left(\sigma^*\left(\alpha'\right)\right)$. Since σ^* is a reduced NBBP of σ, and $\bar{s} \in S^*(\sigma) \subseteq C^*(\sigma)$, Lemma 14 implies that $\bar{s} \in S^*(\sigma^*)$. Combining this fact with $s^*\left(\sigma^*\left(\alpha'\right)\right) = \bar{s}$, we have $s^*\left(\sigma^*\left(\alpha'\right)\right) \in S^*(\sigma^*)$. This establishes (122). The solution F^* therefore admits information propagation for the NBBP σ^*.

Lastly, we will show that condition *(ii)* in Proposition 20 holds for (F^*, σ^*) to complete the proof that F^* is incentive compatible for σ^*.

We fix some $s \in C^*(\sigma)$, and associate to every node $k \in V$ a message $\alpha_k = (i_k, s_{i_k}, j_k, s_{j_k})$

as in the statement of Proposition 20. Lemma 15 implies that $s^*(\sigma^*(\alpha_k)) \in S^*(\sigma^*(\alpha_k))$, Thus,

$$U_{\sigma^*}(s) \leq \max_{\tilde{s} \in C^*(\alpha_k)} U_{\sigma^*}(\tilde{s}) = U_{\sigma^*}(s^*(\sigma^*(\alpha_k))). \tag{125}$$

Summing up the inequalities (125) over $k \in V$, we obtain, by Proposition 23,

$$|V| \, U_{\sigma^*}(s) \leq \sum_{k \in V} [U_{\sigma^*}(s^*(\sigma^*(\alpha_k))) + |V| \, c_k(\underline{u})] = |V| \sum_{k \in V} u_k(F^*(\sigma^*(\alpha_k))).$$

Dividing both sides by $|V|$ leads to the desired inequality:

$$\sum_{k \in V} [u_k(F(\sigma^*(\alpha_k))) - f_i(s)] \geq 0.$$

Therefore, the solution F^* is incentive compatible for the NBBP σ^*.

Step 4: Apply Bilateral Optimality to (F^*, σ^*) to determine the action vector $s^*(\sigma^*)$.

As the solution F^* is incentive compatible for the NBBP σ^*, Bilateral Optimality implies that that for any directly connected pair $(i,j) \in E$, the actions $(s^*(\sigma^*)_i, s^*(\sigma^*)_j)$ solve

$$\max_{(s_i,s_j) \in C^*_{ij}} \left\{ u_i(F^*(\sigma^*(i,s_i,j,s_j))) + u_j(F^*(\sigma^*(i,s_i,j,s_j))) \right\}. \tag{126}$$

It follows from Proposition 22 that

$$\begin{aligned} & u_i(F^*(\sigma^*(i,s_i,j,s_j))) + u_j(F^*(\sigma^*(i,s_i,j,s_j))) \\ & \qquad = \frac{2}{|V|} U_{\sigma^*}(s^*(\sigma^*(i,s_i,j,s_j))) + c_i(\underline{u}) + c_j(\underline{u}) \\ & \qquad = \max_{\tilde{s} \in C^*(i,j,s_i,s_j)} \frac{2}{|V|} U_{\sigma^*}(\tilde{s}) + c_i(\underline{u}) + c_j(\underline{u}). \end{aligned}$$

Thus the maximization problem (126) is equivalent to maximize $U_{\sigma_i^*}(s)$ over $s \in C^*$, which is solved by $s = s^{**}(\sigma^*)$. Therefore, the maximum in (126) is achieved by $\left(s_i^{**}(\sigma^*), s_j^{**}(\sigma^*)\right)$. As F^* satisfies Bilateral Optimality, we have

$$\left(s_i^*(\sigma^*), s_j^*(\sigma^*)\right) = \left(s_i^{**}(\sigma^*), s_j^{**}(\sigma^*)\right) = \left(s_i^{**}(\sigma), s_j^{**}(\sigma)\right). \tag{127}$$

The second equality above follows from Remark 6. Since equation (127) holds for every pair $(i,j) \in E$, thus, $s^*(\sigma^*) = s^{**}(\sigma)$. Since Independence of Irrelevant Actions implies that $s^*(\sigma) = s^*(\sigma^*)$ (this is shown in *Step 2*), we finally obtain

$$s^*(\sigma) = s^{**}(\sigma).$$

That is, the action vector $s^*(\sigma)$ of $F^*(\sigma)$ is socially optimal for σ.

Step 5: Apply Multilateral Stability again to determine the payment vector $y^*(\sigma)$.

It follows from Proposition 22 that the payment vector $y^*(\sigma)$ satisfies

$$u_i(s^{**}(\sigma), y^*(\sigma)) - \underline{u}_{ij} = u_j(s^{**}(\sigma), y^*(\sigma)) - \underline{u}_{ji}, \quad (i,j) \in E.$$

This establishes that Claim 2 holds for $|E| = p$.

By induction, we conclude that Claim 2 holds, completing the uniqueness proof.

Existence

We show that the solution F^* determined by $F^*(\sigma) = (s^{**}(\sigma), y^*(\sigma))$, where the payment vector solves (111), satisfies Axiom I-IV.

Axiom I: For any two-player NBBP $\sigma = (G, S, C, f, \underline{u})$, the solution outcome $F^*(\sigma) = (s^{**}(\sigma), y^*(\sigma))$ is the same as that implied by the Nash Bargaining Solution. Therefore, F^* satisfies Axiom I.

Axiom II: For any NBBP $[(V,E), S, C, f, \underline{u}]$ and any subset $E' \subseteq E$ of edges, we let $\sigma' = \sigma_{\text{sub}}(F, \sigma, E')$ and $G' = (V, E')$. The action vector of the solution outcome $F^*(\sigma')$ is given by $s^{**}(\sigma')$, which is equal to $s^{**}(\sigma)$ since $s^{**}(\sigma)$ is a feasible action vector for σ'.

The system (119) of payment equations for σ' says that under both outcomes $F^*(\sigma')$ and $F^*(\sigma)_{G'}$, each pair of nodes equally share the surplus relative to their respective outside options. Thus, the payment vectors $y^*(\sigma')$ and $y^*(\sigma)_{G'}$ are the same. Therefore,

$$F^*(\sigma') = F^*(\sigma)_{G'}.$$

That is, the solution F^* satisfies axiom of Multilateral Stability.

Axiom III: Given two network bilateral bargaining problems $\sigma = (G, S, C, f, \underline{u})$ and $\sigma' = (G, S', C', f', \underline{u})$ such that σ' is a reduced NBBP of σ, if $s^{**}(\sigma) \in C'$, then

$$s^{**}(\sigma') = \underset{s \in C'}{\operatorname{argmax}}\, U_{\sigma'}(s) = \underset{s \in C'}{\operatorname{argmax}}\, U_{\sigma}(s) = s^{**}(\sigma).$$

The payment vectors $y^*(\sigma)$ and $y^*(\sigma')$ are determined by the same system of payment equations (119). Therefore, $F^*(\sigma) = F^*(\sigma')$, establishing the axiom of Independence of Irrelevant Actions.

Axiom IV: Given an NBBP $\sigma = [(V, E), S, C, f, \underline{u}]$ and a pair $(i, j) \in E$ of directly connected nodes, we consider the maximization problem

$$\max_{(s_i, s_j) \in C_{ij}} \left\{ u_i \left(F^* \left(\sigma \left(i, s_i, j, s_j \right) \right) \right) + u_j \left(F^* \left(\sigma \left(i, s_i, j, s_j \right) \right) \right) \right\} \tag{128}$$

The same argument as in *Step 4* of the uniquenesss proof implies that (128) is equivalent to

$$\max_{(s_i, s_j)} \quad \max_{\tilde{s} \in C(i, j, s_i, s_j)} \quad U_{\sigma}(\tilde{s}),$$

The maximization problem above is equivalent to

$$\max_{s \in C}\, U_{\sigma}(s).$$

which is solved by $s = s^{**}(\sigma)$. Hence, (128) is solved by $(s_i, s_j) = \left(s_i^{**}(\sigma), s_j^{**}(\sigma) \right)$. Thus, the solution F^* satisfies the axiom of Bilateral Optimality.

D.5 Proof of Proposition 21

When $d = 1$, it follows from (118) that the hold-up power of k against i is non-negative.

For any $d \geq 2$, letting $(\ell_0, \ldots, \ell_d)$ be the unique path connecting node k to i, we have, for every integer a between 0 and $d - 1$,

$$\underline{u}^*_{\ell_{a+1}\, \ell_a}(\gamma_{ij}) = u_{\ell_{a+1}} \left(\Phi^* \left((\gamma_{\ell_{a+1}\, \ell_a})_{ij} \right) \right) = u_{\ell_{a+1}} \left(\Phi^* \left(\gamma_{\ell_{a+1}\, \ell_a} \right) \right) = \underline{u}^*_{\ell_{a+1}\, \ell_a}(\gamma).$$

This is because node ℓ_a is not connected to i or j in the network bilateral bargaining problem $\gamma_{\,\ell_{a+1}\,\ell_a}$, the disagreement event between (i,j) is thus payoff irrelevant for node ℓ_a. The hold-up power of k against i can thus be simplified to be

$$\sum_{a=0}^{d-1} \left[\underline{u}^*_{\ell_a\,\ell_{a+1}}(\gamma_{ij}) - \underline{u}^*_{\ell_a\,\ell_{a+1}}(\gamma)\right]$$

This hold-up power being non-negative is equivalent to

$$\underline{u}^*_{\ell_1\,k}(\gamma_{ij}) - \underline{u}^*_{\ell_1\,k}(\gamma) \;\leq\; \sum_{a=1}^{d-1} \left[\underline{u}^*_{\ell_a\,\ell_{a+1}}(\gamma_{ij}) - \underline{u}^*_{\ell_a\,\ell_{a+1}}(\gamma)\right] \tag{129}$$

That is, the *relative hold-up power (against i)* of ℓ_1 with respect to k is weakly less than the hold-up power of ℓ_1 against i. This relative hold-up power of ℓ_1 with respect to k is equal to

$$\underline{u}^*_{\ell_1\,k}(\gamma_{ij}) - \underline{u}^*_{\ell_1\,k}(\gamma) = u_{\,\ell_1}\left(\Phi^*\left(\left(\gamma_{\ell_1\,k}\right)_{ij}\right)\right) - u_{\,\ell_1}\left(\Phi^*\left(\gamma_{\ell_1\,k}\right)\right).$$

Since the network bilateral bargaining problem $\gamma_{\ell_1\,k}$ (with endogenous outside option values) has $|E|-1$ edges, the induction hypothesis implies that $\underline{u}^*(\gamma_{\ell_1\,k})$ satisfies the condition for outside option values. Therefore, one has $u_p(\Phi^*(\gamma_{\ell_1\,k})) = u_p(F^*(\sigma(\gamma_{\ell_1\,k})))$ for every node p. In particular, for every pair $(p,q) \in E\backslash\{\ell_1\,k\}$,

$$u_p(\Phi^*(\gamma_{\ell_1\,k})) - u_q(\Phi^*(\gamma_{\ell_1\,k})) = \underline{u}^*_p(\gamma_{\ell_1\,k}) - \underline{u}^*_q(\gamma_{\ell_1\,k})$$

By adding up the equalities above along the path $(\ell_1,\ldots,\ell_d)$, we obtain

$$u_{\,\ell_1}(\Phi^*(\gamma_{\ell_1\,k})) - u_i(\Phi^*(\gamma_{\ell_1\,k})) = \sum_{a=1}^{d-1} \left[\underline{u}^*_{\ell_a\,\ell_{a+1}}(\gamma_{\ell_1\,k}) - \underline{u}^*_{\ell_{a+1}\,\ell_a}(\gamma_{\ell_1\,k})\right]$$

Likewise, the induction hypothesis for $(\gamma_{\ell_1\,k})_{ij}$ implies

$$u_{\,\ell_1}\left(\Phi^*\left(\left(\gamma_{\ell_1\,k}\right)_{ij}\right)\right) - u_i\left(\Phi^*\left(\left(\gamma_{\ell_1\,k}\right)_{ij}\right)\right) = \sum_{a=1}^{d-1} \left[\underline{u}^*_{\ell_a\,\ell_{a+1}}\left(\left(\gamma_{\ell_1\,k}\right)_{ij}\right) - \underline{u}^*_{\ell_{a+1}\,\ell_a}\left(\left(\gamma_{\ell_1\,k}\right)_{ij}\right)\right]$$

Subtracting the two equations above, we obtain

$$
\begin{aligned}
& u_{\ell_1}\left(\Phi^*\left((\gamma_{\ell_1 k})_{ij}\right)\right) - u_{\ell_1}\left(\Phi^*\left(\gamma_{\ell_1 k}\right)\right) + \left[u_i(\Phi^*(\gamma_{\ell_1 k})) - u_i\left(\Phi^*\left((\gamma_{\ell_1 k})_{ij}\right)\right)\right] \\
&= \sum_{a=1}^{d-1} \left(\left[\underline{u}^*_{\ell_a \ell_{a+1}}\left((\gamma_{\ell_1 k})_{ij}\right) - \underline{u}^*_{\ell_{a+1} \ell_a}\left((\gamma_{\ell_1 k})_{ij}\right)\right] - \left[\underline{u}^*_{\ell_a \ell_{a+1}}(\gamma_{\ell_1 k}) - \underline{u}^*_{\ell_{a+1} \ell_a}(\gamma_{\ell_1 k})\right] \right) \\
&= \sum_{a=1}^{d-1} \left[\underline{u}^*_{\ell_{a+1} \ell_a}(\gamma_{\ell_1 k}) - \underline{u}^*_{\ell_{a+1} \ell_a}\left((\gamma_{\ell_1 k})_{ij}\right)\right] \\
&= \sum_{a=1}^{d-1} \left[\underline{u}^*_{\ell_{a+1} \ell_a}(\gamma) - \underline{u}^*_{\ell_{a+1} \ell_a}\left(\gamma_{ij}\right)\right],
\end{aligned}
\tag{130}
$$

Since the network bilateral bargaining problem $\gamma_{\ell_1 k}$ (with endogenous outside option values) has $|E| - 1$ edges, the induction hypothesis implies that $\underline{u}^*(\gamma_{\ell_1 k})$ satisfies the condition for outside option values. In particular,

$$
u_i(\Phi^*(\gamma_{\ell_1 k})) = u_i(F^*(\sigma(\gamma_{\ell_1 k}))) \geq \underline{u}^*_{ij}(\gamma_{\ell_1 k})) = u_i\left(\Phi^*\left((\gamma_{\ell_1 k})_{ij}\right)\right)
$$

Combing the inequality above with equation (130), we obtain the desired inequality (129).

D.6 Proof of Theorem 5

We first establish the uniqueness part of the theorem. Suppose that Φ^* is a solution mapping that satisfies the condition stated in *(i)* of Theorem 5. We determine the solution outcome $\Phi^*(\gamma)$ inductively over the number of edges of the problem γ. If γ is a two-node bilateral bargaining problem, then the outside option values of the two nodes are $\underline{u}^*_1 = u_1(\underline{s}_1, \underline{s}_2)$ and $\underline{u}^*_2 = u_2(\underline{s}_1, \underline{s}_2)$, where $(\underline{s}_1, \underline{s}_2)$ is the unique exogenously given conflict point of the two nodes. It is clear that $(\underline{u}_1, \underline{u}_2)$ satisfy the condition for outside option values. Thus, the solution outcome $\Phi^*(\gamma)$ is simply given by the Nash Bargaining Solution $F^*(\sigma)$.

We suppose that for every network bilateral bargaining problem γ (with endogenous outside option values) whose number of edges is less than or equal to $|E| - 1$, the solution outcome $\Phi^*(\gamma)$ is uniquely determined, and $\underline{u}^*(\gamma)$ satisfies the condition for outside option values. This is our *induction hypothesis*. We now consider a network bilateral bargaining

problem $\gamma = [(V, E), S, C, f, \underline{S}]$ (with endogenous outside option values) whose number of edges is equal to $|E|$. We will show that $\underline{u}^*(\gamma)$ satisfies the condition for outside option values.

For any given node $i \in V$, we let

$$u_i^*(\gamma) = u_i(s^{**}(\gamma)) - y_{ij}^* \left(\underline{u}^*(\gamma)\right),$$

which is the candidate net payoff of node i in the network bilateral bargaining problem γ (with endogenous outside option values), *if* $\underline{u}^*(\gamma)$ satisfies the condition for outside option values. We then fix another arbitrary node k (not necessarily a direct counterparty of i), letting $(\ell_0, \ldots, \ell_d)$ be a path connecting from k to i, we have

$$u_k^*(\gamma) - u_i^*(\gamma) = \sum_{a=0}^{d-1} [\underline{u}^*_{\ell_a \ell_{a+1}}(\gamma) - \underline{u}^*_{\ell_{a+1} \ell_a}(\gamma)]. \tag{131}$$

Now we fix a node j that is a direct counterparty of node i. Since the network bilateral bargaining problem γ_{ij} (with endogenous outside option values) has $|E| - 1$ edges, the induction hypothesis implies that for any given node k,

$$u_k^*(\gamma_{ij}) - u_i^*(\gamma_{ij}) = \sum_{a=0}^{d-1} [\underline{u}^*_{\ell_a \ell_{a+1}}(\gamma_{ij}) - \underline{u}^*_{\ell_{a+1} \ell_a}(\gamma_{ij})]. \tag{132}$$

Subtracting equation (131) from equation (132), we obtain

$$\begin{aligned} &[u_k^*(\gamma_{ij}) - u_k^*(\gamma)] - [u_i^*(\gamma_{ij}) - u_i^*(\gamma)] \\ &= \sum_{a=0}^{d-1} \left(\left[\underline{u}^*_{\ell_a \ell_{a+1}}(\gamma_{ij}) - \underline{u}^*_{\ell_{a+1} \ell_a}(\gamma_{ij})\right] - \left[\underline{u}^*_{\ell_a \ell_{a+1}}(\gamma) - \underline{u}^*_{\ell_{a+1} \ell_a}(\gamma)\right] \right) \end{aligned}$$

The right hand side of the equation above is the hold-up power of node k against i (in the event of a disagreement between (i, j)). Thus, it is non-negative by Proposition 21. If

$$u_i^*(\gamma_{ij}) > u_i^*(\gamma),$$

then it must be that

$$u_k^*(\gamma_{ij}) > u_k^*(\gamma).$$

This is true for every node $k \in V$, which implies

$$\sum_{k \in V} u_k^*(\gamma_{ij}) > \sum_{k \in V} u_k^*(\gamma) = U(s^{**}(\gamma)).$$

This contradicts the optimality of the socially efficient action vector $s^{**}(\gamma)$. Hence, we have

$$u_i^*(\gamma_{ij}) \leq u_i^*(\gamma),$$

for every pair $(i, j) \in E$, which is precisely the condition for outside option values.

The existence result immediately follows from the observation that the solution mapping Φ^* constructed inductively in the uniqueness proof satisfies the condition in Theorem 5.

E Axioms for Nash Bargaining Solutions

This appendix provides simple axioms for the Nash Bargaining Solution for network bilateral bargaining problems (NBBPs) with completely connected graphs. For our main results, we apply these axioms in the case of a network consisting of two connected nodes.

The feasible utility set of a NBBP $\sigma = ((V, E), S, C, f, \underline{u})$ is

$$\mathcal{U}(\sigma) = \{(u_i(s, y))_{i \in V} : (s, y) \in \Omega[G, S, C]\} \cup \{\underline{u}\}.$$

The permutation $\sigma_P = ((V^P, E^P), S^P, C^P, f^P, \underline{u}^P)$ of σ is defined by some bijection $P : V \to$

V as:

$$\begin{aligned} V^P &= \{P(1), \ldots, P(|V|)\} \\ E^P &= \{(P(i), P(j) : (i,j) \in E\} \\ S^P_{P(i)} &= S_i \\ C^P_{P(i)} &= C_i \\ f^P_{P(i)} &= f_i \\ \underline{u}^P_{P(i)} &= \underline{u}_i. \end{aligned}$$

A solution F for NBBPs is said to be:

- *Utility focussed* if, for any NBBPs σ and σ' with the same feasible utility sets, we have $u(F(\sigma)) = u(F(\sigma'))$.
- *Pareto optimal* if, for each NBBP σ, the solution outcome $F(\sigma)$ is Pareto Optimal for σ.
- *Symmetric* if for any NBBP $\sigma = ((V,E), S, C, f, \underline{u})$ and any permutation σ_P of σ by P, we have
$$u_{P(i)}(F(\sigma_P) = u_i(F(\sigma)), \quad i \in V.$$
- *Translation preserving* if for any NNBPs σ and σ' with
$$\mathcal{U}(\sigma') = \mathcal{U}(\sigma) + \{v\},$$
for some v, we have $u(F(\sigma')) = u(F(\sigma)) + v$.

The Nash Bargaining Solution F^N for an NBBP $\sigma = ((V,E), S, C, f, \underline{u})$ whose graph (V,E) is completely connected is defined by $F^N(\sigma) = (s^*, y^*)$, where

$$s^* = \operatorname*{argmax}_{s \in C} \sum_{i \in V} f_i(s) \tag{133}$$

$$y^*_{ij} = \frac{1}{2}\left([f_i(s^*) - \underline{u}_i] - [f_j(s^*) - \underline{u}_j]\right), \quad (i,j) \in E. \tag{134}$$

Proposition 24. *When restricted to NBBPs whose graphs are completely connected, a solution F is utility focussed, Pareto optimal, symmetric, and translation preserving if and only if it is the Nash Bargaining Solution.*

Proof. The Nash Bargaining Solution is, by simple inspection, utility focussed, Pareto optimal, symmetric, and translation preserving.

Conversely, suppose a solution F is utility focussed, Pareto optimal, symmetric, and translation preserving. Consider an NBBP $\sigma = ((V, E), S, C, f, \underline{u})$ for which (V, E) is completely connected. By translation preservation, we can assume without loss of generality that $\underline{u} = 0$. By Pareto optimality and the fact (V, E) is completely connected, $u(F(\sigma))$ is an element of the symmetric hyperplane $H = \{v : \sum_i v_i = v^*\}$, where

$$v^* = \max_{s \in C} \sum_{i \in V} f_i(s).$$

Moreover, H is contained by the feasible set $\mathcal{U}(\sigma)$. By the symmetry of F, it follows that $u_i(F(\sigma)) = u_j(F(\sigma))$ for all i and j, implying that $F(\sigma) = F^N(\sigma)$. □

F Independence of Strategically Irrelevant Information

This appendix motivates and then defines a restriction on perturbed games, by which tremble probabilities cannot depend on strategically irrelevant information.

F.1 Motivation

In order to motivate the idea of independence of strategically irrelevant information, consider the 3-person extensive-form game depicted in Figure 8. At all but one of the terminal nodes, the three players, Column, Row, and Box, have equal payoffs.

The strategy profiles [*Left,(Up,Up),West*], marked in red, and [*Right,(Up,Up),West*], marked in blue, are the two Nash equilibria of the game. Both of these Nash equilibria are in fact extensive form trembling hand perfect in the sense of Selten (1975). However, [*Left,(Up,Up),West*] is the only reasonable equilibrium for this game, in the following sense.

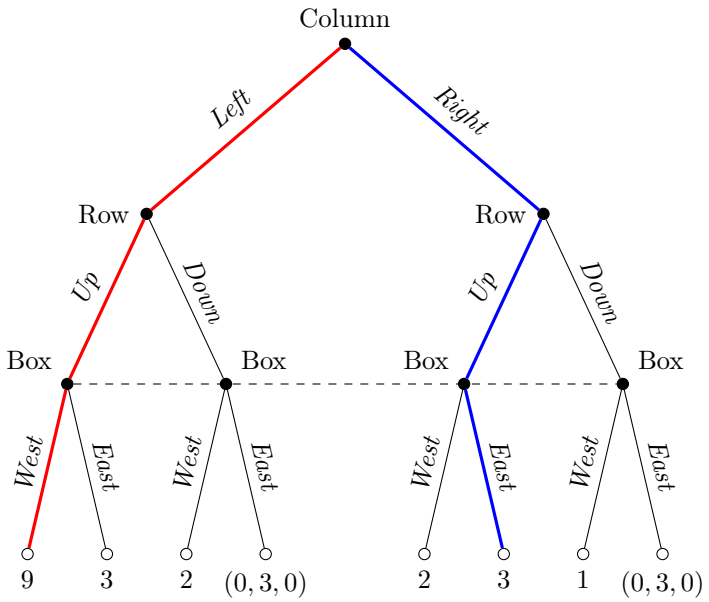

Figure 8 – a three player game in extensive form

The strategy *(Up, Up)* weakly dominates all of Row's other strategies. If one eliminates the weakly dominated strategies from Row's strategy space, then for Column, the strategy *Left* weakly dominates *Right*. If Column plays *Left*, then Box's best response is to choose *West*.

On the other hand, [*Right,(Up, Up),East*] can be sustained as an extensive form trembling hand perfect equilibria in the following way. Consider the strategy profile by which Column plays *Left* with probability ϵ, Box chooses *West* with probability ϵ^2, and Row deviates to *Down* with probability ϵ following *Left*, while he deviates with probability ϵ^2 following *Right*. It is easy to see that against this strategy profile, *(Up, Up)* is optimal for Row and *East* is optimal for Box. To see that *Right* is optimal for Column requires a little bit of calculation. The key here is that Row deviates to *Down* more often following *Right* than following *Left*. Because Column worries about a mistake by Row, and does not care much about the ϵ^2 probability that Box chooses *West*, Column is strictly better off playing *Right*, because Row deviates with lower likelihood there. As ϵ goes to 0, one obtains the limiting strategy profile [*Right,(Up, Up),East*]. Furthermore, [*Right,(Up, Up),East*] is also an extended proper equilibrium (Milgrom and Mollner, 2016), and thus a proper equilibrium (Myerson, 1977a),

that is sustained by the same sequence of trembling equilibria. This is so because deviating to *Down* is indeed more costly for Row when he faces *Right* than when he faces *Left*, given that *Right* is on the equilibrium path.

However, this is an "unreasonable" tremble for Row, since, when following *Left* or *Right*, Row faces two games that are equivalent to each other, for whatever choice Box makes. If Box plays *West*, Row faces a one-person decision problem in which *(i)* choosing *Up* gives him a payoff of 9, whereas choosing *Down* gives him only 2 at the information set following *Left*, and *(ii)* choosing *Up* pays him 3, while choosing *Down* pays only 0 at the information set following *Right*. These two problems are essentially the same for Row. Likewise if Box plays *East*, Row's problems are also identical irrespective of the information set in which he finds himself. He therefore has no reason to tremble more often in one information set than in the other. In other words, the information that Row obtains by observing Column's choice is irrelevant for his own play. Minimum tremble probabilities should be independent of irrelevant information. If one requires that in a perturbed game,[26] the minimum probabilities are the same for Row in the two information sets. That is, if $\epsilon(\textit{Left}) = \epsilon(\textit{Right})$ and $\chi(\textit{Left}) = \chi(\textit{Right})$, then it is strictly better for Column to play *Left* than *Right*.

More generally, a player might be facing a multiplayer non-cooperative game at an information set, rather than a single-person decision problem. We now formalize our independence requirement for abstract extensive form games in which both the strategy spaces and the time horizon can be finite or infinite. We first specify the sense in which two games are "equivalent," and then provide a formal definition of this notion of "irrelevant information."

F.2 Strategic Equivalence between Multistage Games

We fix a multistage game $\Gamma = (N, A, \mathcal{I}, \Theta, p, \zeta, u)$. A *pure strategy* of a player i is a behavioral strategy $\sigma_i = (\sigma_{i0}, \sigma_{i1}, \ldots)$ with the property that, for all t and any I_{it}, the probability measure $\sigma_{it}(I_{it})$ assigns all probability mass to a single point x_{it} in A_{it}. We use $x_i = (x_{i0}, x_{i1}, \ldots)$ to denote such a pure strategy. By letting X_{it} denote the set of measurable

[26] The perturbation here is in the sense of Selten (1975).

functions from $\mathcal{I}_{it}$ to A_{it}, we can thus equate $X_i = \prod_{t\geq 0} X_{it}$ with the set of pure strategies of player i. A *pure strategy profile* is thus an element of $X = \prod_{i\in N} X_i$. The associated normal form is $((X_i)_{i\in N}, (U_i)_{i\in N})$.

The reduced normal form[27] of a multistage game is obtained from the normal form by deleting from X_i, for each player i, any pure strategy x_i that is "redundant," in the sense that there exists a behavioral strategy $\sigma_i \neq x_i$ such that $U_j(x_i, x_{-i}) = U_j(\sigma_i, x_{-i})$ for all $j \in N$ and all $x_{-i} \in X_{-i}$.

The normal form of a game, however, does not provide enough information to detect all such redundant strategies. Consider, for example, the game depicted in Figure 9. Whether Player 1 chooses A or B is irrelevant for the play of the game, but the normal form illustrated in Table 1 does not allow us to detect this redundancy.

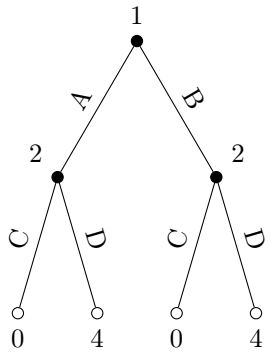

Figure 9 – Extensive form

		P1 A	P1 B
P2	C, C	0	0
	C, D	0	4
	D, C	4	0
	D, D	4	4

Table 1 – Normal form

Given this, we now propose a method for detecting redundant strategies for Γ. For $K \subseteq N \times \mathbb{Z}^+$, we apply the subscript "K" to denote the projection that maps A to $A_K = \prod_{(i,t)\in K} A_{it}$, and likewise maps a to $a_K = (a_{it})_{(i,t)\in K}$. Let Z_t be the set of measurable functions from H_t to Θ. We can view an element of z_t of Z_t as a "choice of Nature" in period t.

Consider some some disjoint subsets T and T' of $\mathbb{Z}^+$ (sets of times) with the property

[27] Kohlberg and Mertens (1986) argued that the reduced normal form captures all information about a game that is relevant to decision making.

that

$$t \notin T \cup T' \implies t \geq \sup T'. \tag{135}$$

For disjoint subsets K and K' of $N \times \mathbb{Z}^+$ and some $z_T \in Z_T$, $\theta_{T'} \in \Theta_{T'}$, $x_K \in X_K$, and $a_{K'} \in A_{K'}$, we let

$$\Gamma\left[(z_T, x_K), (\theta_{T'}, a_{K'}), T, K, T', K'\right] \tag{136}$$

be the multistage game obtained from Γ by fixing:

- For each $t \in T$, the choice function of nature in period t to be z_t.
- For each $(i,t) \in K$, the strategy of player i in period t to be x_{it}.
- For each $t \in T'$, the state of nature in period t to be θ_t.
- For each $(i,t) \in K'$, the action of player i in period t to be a_{it}.

If one or more of the index sets T, K, T', and K' in (136) are empty, we will omit them from the notation. Given an index set $K \subseteq N \times \mathbb{Z}^+$, we say that the strategies in K are *irrelevant* for the play of the game Γ if all of the multistage games in

$$\{\Gamma\left[x_K, K\right] : x_K \in X_K\} \tag{137}$$

have the same reduced normal form. In this case, we denote by $\Gamma(K)$ the common reduced normal form of the multistage games $\{\Gamma\left[x_K, K\right] : x_K \in X_K\}$. (We note that $\Gamma(K)$ may contain players with a singleton strategy set.) In Figure 9, the strategy of Player 1 in the first step of the game is irrelevant.

Given two multistage games Γ_1 and Γ_2 with respective player sets N_1 and N_2, we say that Γ_1 and Γ_2 are *strategically equivalent* if there exist $K_1 \subset N_1 \times \mathbb{Z}^+$ and $K_2 \subset N_2 \times \mathbb{Z}^+$ such that *(i)* the strategies in K_1 and K_2 are irrelevant for the play of Γ_1 and Γ_2 respectively, and *(ii)* $\Gamma_1(K_1)$ and $\Gamma_2(K_2)$ can be obtained from each other by adding or deleting players with a singleton strategy set, by performing increasing affine transformations of players' payoff functions, and by relabeling players and strategies[28].

[28]Suppose $\Gamma_1(K_1) = ((X_i)_{i\in N_1}, (u_i)_{i\in N_1})$ and $\Gamma_2(K_2) = ((Y_i)_{i\in N_2}, (v_i)_{i\in N_2})$, and suppose players with a

Remark 7. The notion of strategic equivalence here is more general than that obtained by comparing reduced normal forms, because it removes redundant strategies in the extensive forms that cannot be detected in the normal forms. As a special case, two multistage games with the same reduced normal forms are strategically equivalent, in the sense defined above. As a different example, if Γ is a multistage game such that the strategies associated with some index set K of players and times are irrelevant for the play of the game Γ, then Γ is strategically equivalent to $\Gamma[x_K, K]$ for every $x_K \in X_K$.

Consider two multistage games $\Gamma_\ell = (N_\ell, A_\ell, \mathcal{I}_\ell, \Theta_\ell, p_\ell, \zeta_\ell, u_\ell)$, $\ell \in \{1, 2\}$. For some $(i, t) \in N_1 \times \mathbb{Z}^+$, $(j, \tau) \in N_2 \times \mathbb{Z}^+$, some information set $I_1 \in \mathcal{I}_{1it}$ of Γ_1, some $I_2 \in \mathcal{I}_{2j\tau}$ of Γ_2, and a homeomorphism g between the action set A_{1it} in Γ_1 and $A_{2j\tau}$ in Γ_2, we say that Γ_1 at I_1 is *strategically isomorphic* to Γ_2 at I_2 through the homeomorphism g if the following conditions hold:

(i) Γ_1 and Γ_2 are strategically equivalent.

(ii) In the equivalence of Γ_1 and Γ_2 defined above, player i in Γ_1 is identified[29], through relabeling of players, with player j in Γ_2.

(iii) The homeomorphism g between A_{1it} and $A_{2j\tau}$ is defined by relabeling of strategies after which Γ_1 and Γ_2 are strategically equivalent.[30]

F.3 Strategically Irrelevant Information

We fix a multistage game $\Gamma = (N, A, \mathcal{I}, \Theta, p, \zeta, u)$. Given an information set $I_{it} \in \mathcal{I}_{it}$ of player i in period t, $\zeta^{-1}(I_{it}) \subseteq \Theta_{\leq t} \times A_{<t}$ can be viewed as the subset of partial histories

singleton strategy set have been deleted from both normal form games. Then what we mean here by relabeling players and strategies is that there exists a bijection π between N_1 and N_2 (which relabels players), and for every $i \in N_1$, a bijection f_i between X_i and $Y_{\pi(i)}$ ((which relabels strategies)), such that $u_i((x_i)_{i \in N_1}) = a_i v_{\pi(i)}((f_i(Y_{\pi(i)})_{i \in N_1}) + b_i$ with $a_i > 0$.

[29] This means, in the notation of Footnote 28, that neither player i nor player j have singleton strategy sets in (Γ_1, K_1) and (Γ_2, K_2) respectively, and that $\pi(i) = j$.

[30] This means, in the notation of Footnote 28, that $g(a_1) = a_2$ if and only if the subset $X_i(I_1, a_1) = \{x_i \in X_i : x_{it}(I_1) = a_1\}$ of pure strategies of player i matches $Y_j(I_2, a_2) = \{y_j \in Y_j : y_{j\tau}(I_2) = a_2\}$ through the bijection f_i.

that "reach" the information set I_{it}. We can represent $\zeta^{-1}(I_{it})$ as a set of the form

$$\zeta^{-1}(I_{it}) = H_{it}(I_{it}) \times \Theta_{O^c_{1it}(I_{it})} \times A_{O^c_{2it}(I_{it})}$$

where $O_{1it}(I_{it}) \subseteq \{0, \ldots, t\}$, $O_{2it}(I_{it}) \subseteq N \times \{0, \ldots, t-1\}$, and $H_{it}(I_{it}) \subsetneq \Theta_{O_{1it}(I_{it})} \times A_{O_{2it}(I_{it})}$. (As usual, S^c denotes the complement of a set S.) We obtain uniqueness for this representation by always taking O_{1it} and O_{2it} to be minimal, in the sense there there does not exist a strict subset of O_{1it} nor a strict subset of O_{2it} that also admits such a representation. (We allow O_{1it} and O_{2it} to be empty.) This representation means that player i learns from the information set I_{it} precisely the fact that

$$\left((\theta_s)_{s \in O_{1it}(I_{it})},\, (a_{j\tau})_{(j,\tau) \in O_{2it}(I_{it})}\right) \in H_{it}(I_{it}),$$

whereas player i learns nothing at all about the complementary elements of the partial history. The minimality of the index sets O_{1it} and O_{2it} implies that $H_{it}(I_{it})$, whenever it exists, must be a strict subset, as indicated, thus providing nontrivial information in the form of a binding restriction on the partial history.

We further define $M_{it}(I_{it}) \subseteq N \times \{0, \ldots, t-1\}$ by letting $(j, \tau) \in M_{it}(I_{it})$ if and only if for every (θ, a) and $(\tilde{\theta}, \tilde{a})$ such that $\zeta_{it}(\theta_{\leq t}, a_{<t}) = I_{it}$ and $\zeta_{it}(\tilde{\theta}_{\leq t}, \tilde{a}_{<t}) \neq I_{it}$, we have

$$\zeta_{j\tau}\left(\theta_{\leq \tau}, a_{<\tau}\right) \neq \zeta_{j\tau}\left(\tilde{\theta}_{\leq \tau}, \tilde{a}_{<\tau}\right).$$

This means that for every $(j, \tau) \in M_{it}(I_{it})$, by period τ player j has learned whether the information set I_{it} has been reached.

Proposition 25. *The index sets $O_{2it}(I_{it})$ and $M_{it}(I_{it})$ are disjoint.*

Proof. If $(j, \tau) \in O_{2it}(I_{it})$ then there exists $\bar{a}_{j\tau} \in A_{j\tau}$ such that if the action of player j action in period τ is $\bar{a}_{j\tau}$, the information set I_{it} is not reached. Formally, $a_{j\tau} = \bar{a}_{j\tau}$ implies $(\theta_{\leq t}, a_{<t}) \notin \zeta_{it}^{-1}(I_{it})$ for every complete history (θ, a). Now let (θ, a) be a complete history under which I_{it} is reached. That is, $(\theta_{\leq t}, a_{<t}) \in \zeta_{it}^{-1}(I_{it})$. Let $\tilde{a} = (a_{-j\tau}, \bar{a}_{j\tau})$ (that is, $\tilde{a}$ is obtained from a by replacing $a_{j\tau}$ by $\bar{a}_{j\tau}$). Then we have $(\theta_{\leq t}, \tilde{a}_{<t}) \notin \zeta_{it}^{-1}(I_{it})$. On the other

hand, $\zeta_{j\tau}(\theta_{\leq\tau}, \tilde{a}_{<\tau}) = \zeta_{j\tau}(\theta_{\leq\tau}, a_{<\tau})$. This shows that player j in period τ cannot tell whether the information set I_{it} has been reached. Thus $(j,\tau) \notin M_{it}(I_{it})$. □

Proposition 26. *Suppose $(j,\tau) \in O_{2it}(I_{it})$ and $(j',\tau') \in M_{it}(I_{it})$. Then, for every information set $I_{j\tau} \in \mathcal{I}_{j\tau}$ of player j in period τ, we have $(j',\tau') \notin O_{2j\tau}(I_{j\tau})$. (This means that player j in period τ cannot observe the action of player j' in period τ'.)*

Proof. If $(j',\tau') \in O_{2j\tau}(I_{j\tau})$ for some $I_{j\tau} \in \mathcal{I}_{j\tau}$, then it must be that $\tau' < \tau$. Then the same argument used in the proof of Proposition 25, after replacing (j,τ) by (j',τ'), implies that $(j',\tau') \notin M_{it}(I_{it})$. This leads to a contradiction, completing the proof. □

Definition 11. For any times t and τ and players i and j, we say that two information sets I_{it} and $I_{j\tau}$ *contain the same strategic information* if there exist subsets K_i and K_j of $N \times \mathbb{Z}^+$, and a homeomorphism g between the action sets A_{it} and $A_{j\tau}$, such that:

(i) $K_i \cap O_{2it}(I_{it}) = \emptyset$ and $K_j \cap O_{2j\tau}(I_{j\tau}) = \emptyset$.

(ii) $O_{2it}(I_{it}) \cup K_i \cup M_{it}(I_{it}) = N \times \mathbb{Z}^+$ and $O_{2j\tau}(I_{j\tau}) \cup K_j \cup M_{j\tau}(I_{j\tau}) = N \times \mathbb{Z}^+$.

(iii) Letting $T_i = \{0,\ldots,t\}\backslash O_{1it}(I_{it})$ and $T_j = \{0,\ldots,\tau\}\backslash O_{2j\tau}(I_{j\tau})$, for every $z \in Z$, $x \in X$, $h_{it} \in H_{it}(I_{it})$, and $h_{j\tau} \in H_{j\tau}(I_{j\tau})$, at least one of the following conditions applies:

(a) The game $\Gamma_{it} \equiv \Gamma\left[(z_T, x_{K_i}), h_{it}, (T_i, K_i), (O_{1it}(I_{it}), O_{2it}(I_{it}))\right]$ at I_{it} is strategically isomorphic to the game $\Gamma_{j\tau} \equiv \Gamma\left[(z_{T_j}, x_{K_j}), h_{j\tau}, (T_j, K_j), (O_{1j\tau}(I_{j\tau}), O_{2j\tau}(I_{j\tau}))\right]$ at $I_{j\tau}$ through the homeomorphism g.

(b) The strategy of player i in period t is irrelevant in Γ_{it} and the strategy of player j in period τ is irrelevant in $\Gamma_{j\tau}$.

We now give an interpretation of this notion of two information sets containing the same strategic information. Suppose that player i finds himself in some information set I_{it} at time t. This informs player i of something specific about Nature's states $(\theta_s)_{s\in O_{1it}(I_{it})}$, and about the actions $(a_{j\tau})_{(j,\tau)\in O_{2it}(I_{it})}$.

For each $(j,\tau) \notin M_{it}(I_{it})$, however, player j, when in period τ, cannot tell whether the information set I_{it} has been reached. In this case, the decision problem of player j in period τ should not be embedded into the "subgame" faced by player i at the information set I_{it}, because doing so would violate the information structure of the original game. That is, if $(j,\tau) \notin M_{it}(I_{it})$ and if the decision problem of player j in period τ is embedded into the "subgame" faced by player i at I_{it}, then player j naturally knows that the information set I_{it} has been reached. He is not supposed to learn this information by period τ in the original game. We thus fix the strategy of player j in period τ, for all (j,τ) in some index set $K \subseteq N \times \mathbb{Z}^+$. The construction implies that $O_{2it}(I_{it}) \cup K \cup M_{it}(I_{it}) = N \times \mathbb{Z}^+$, which guarantees that the set K is large enough to include all indices that are not in $M_{it}(I_{it})$. Fixing Nature's choice function in the periods of some index set T guarantees that the multistage game $\Gamma\big[(z_T, x_K), h_{it}, (T,K), (O_{1it}(I_{it}), O_{2it}(I_{it}))\big]$ is well defined. Thus "containing the same strategic information" means that, in all conceivable circumstances, player i in period t as well as player j in period τ are either indifferent in their choice of strategies, or, in case they are not indifferent, they face games that are strategically isomorphic, from the vantage points of their respective information sets.

Given a multistage game $\Gamma = (N, A, \mathcal{I}, \Theta, p, \zeta, u)$, we say that a perturbed game $\widehat{\Gamma} = (\Gamma, \epsilon, \chi)$ of Γ *respects independence of strategically irrelevant information* if $\epsilon_{it}(I_{it}) = \epsilon_{j\tau}(I_{j\tau})$ and $\chi_{j\tau}(I_{j\tau}) = \chi_{it}(I_{it}) \circ g$ whenever I_{it} and $I_{j\tau}$ contain the same strategic information for players i and j, respectively, where g is the homeomorphism between A_{it} and $A_{j\tau}$ defined in Definition 11.

References

Abreu, D., and M. Manea, 2012, "Bargaining and Efficiency in Networks," *Journal of Economic Theory*, 147(1), 43–70.

Aumann, R., 1964, "Mixed and Behavoir Strategies in Infinite Extensive Games," *Advances in Game Theory*, pp. pp. 627–650.

Bertsekas, D., 2012, *Dynamic Programming and Optimal Control, Volume II*. Athena Scientific, 4th edn.

Bertsekas, D. P., and S. E. Shreve, 1978, *Stochastic Optimal Control: The Discrete Time Case*, vol. 139 of *Mathematics in Science and Engineering*. Academic Press, New York.

Binmore, K., A. Rubinstein, and A. Wolinsky, 1986, "The Nash Bargaining Solution in Economic Modelling," *The Rand Journal of Economics*, 17(2), pp. 176–188.

Björnerstedt, J., and J. Stennek, 2007, "Bilateral Oligopoly and The Efficiency of Intermediate Goods Markets," *International Journal of Industrial Organization*, 25(5), 884–907.

Blackwell, D., 1965, "Discounted Dynamic Programming," *Annals of Mathematical Statistics*, 36(1), 226–235.

Bolton, P., and M. Oehmke, 2015, "Should Derivatives Be Privileged in Bankruptcy?," *Journal of Finance*, 70, 2353–2394.

Calvo-Armengol, A., 1999, "A Note on Three-Player Noncooperative Bargaining with Restricted Pairwise Meetings," *Economics Letters*, 65(1), 47–54.

Carbonell-Nicolau, O., 2011a, "Perfect and Limit Admissible Perfect Equilibria in Discontinuous Games," *Journal of Mathematical Economics*, 47(4-5), 531–540.

——— , 2011b, "The Existence of Perfect Equilibrium in Discontinuous Games," *Games*, 2(3), 235–256.

Condorelli, D., and A. Galeotti, 2016, forthcoming, "Bilateral Trading in Networks," *Review of Economic Studies*.

David, A., and A. Lehar, 2014, "Why Are Banks Highly Interconnected," .

de Fontenay, C., and J. S. Gans, 2013, "Bilateral Bargaining with Externalities," working paper, Melbourne Business School Discussion Paper No. 2004-32.

Drew, F., and J. Tirole, 1991, *Game Theory*. Cambridge, Massachusetts: MIT Press.

Duffie, D., 2012, *Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets*. Princeton University Press, Princeton :.

Duffie, D., and D. Skeel, 2012, "A Dialogue on the Costs and Benefits of Automatic Stays for Derivatives and Repurchase Agreements," in *Bankruptcy Not Bailout: A Special Chapter 14*, ed. by K. E. Scott, and J. B. Taylor. Hoover Press.

Dunford, N., and J. T. Schwartz, 1988, *Linear Operators Part I: General Theory*. New York, Interscience Publishers.

El Karoui, N., and X. Tan, 2013, "Capacities, Measurable Selection and Dynamic Programming Part I: Abstract Framework," *ArXiv e-prints*.

Elliott, M., and F. Nava, 2015, "Decentralized Bargaining: Efficiency and the Core," working paper, Cambridge University.

Gale, D. M., and S. Kariv, 2007, "Financial Networks," *American Economic Review*, 97(2), 99–103.

Hart, O., J. Tirole, D. W. Carlton, and O. E. Williamson, 1990, "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity. Microeconomics*, 1990, pp. 205–286.

Hatfield, J. W., S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp, 2013, "Stability and Competitive Equilibrium in Trading Networks," *Journal of Political Economy*, 121(5), 966–1005.

Jackson, M., and A. Wolinsky, 1996a, "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 71, 44–74.

Jackson, M. O., 2008, *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA.

Jackson, M. O., T. Rodriguez-Barraquer, and X. Tan, 2012, "Epsilon-Equilibria of Perturbed Games," *Games and Economic Behavior*, 75(1), 198–216.

Jackson, M. O., and A. Wolinsky, 1996b, "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 71(1), 44–74.

Kallianpur, G., and J. Xiong, 1995, "Stochastic Differential Equations in Infinite-Dimensional Spaces: An Extended Version of the Lectures Delivered by Gopinath Kallianpur as Part of the 1993 Barrett Lectures at the University of Tennessee, Knoxville, TN.," Institute for Mathemaical Statistics.

Kohlberg, E., and J.-F. Mertens, 1986, "On the Strategic Stability of Equilibria," *Econometrica*, 54(5), pp. 1003–1037.

Krishna, V., and R. Serrano, 1996, "Multilateral Bargaining," *The Review of Economic Studies*, 63(1), 61–80.

Kuhn, H. W., 1953, "Extensive Games and the Problem of Information," *Contributions to the Theory of Games*, 2, 193–216.

Lehar, A., 2014, "Restructuring Failure and Optimal Capital Structure," University of Calgary.

Lensberg, T., 1988, "Stability and the Nash Solution," *Journal of Economic Theory*, 45(2), 330–341.

Lou, Y., and C. A. Otto, 2015, "Debt Dispersion, Creditor Conflicts, and Covenants in Corporate Loans," HEC, Paris.

Manea, M., 2016, "Intermediation and Resale in Networks," MIT, Department of Economics.

Milgrom, P., and J. Mollner, 2016, "Extended Proper Equilibrium," Department of Economics, Stanford University.

Myerson, R., 1977a, "Graphs and Cooperation in Games," *Mathematics of Operations Research*, 2, 225–229.

Myerson, R., and P. Reny, 2015, "Sequential Equilibrium Distributions in Multi-Stage Games with Infinite Sets of Types and Actions (Draft notes)," Department of Economics, University of Chicago.

Myerson, R. B., 1977b, "Graphs and Cooperation in Games," *Mathematics of Operations Research*, 2(3), pp. 225–229.

Myerson, R. B., 1977c, "Values of Games in Partition Function Form," *International Journal of Game Theory*, 6(1), 23–31.

Nash, J., 1950, "The Bargaining Problem," *Econometrica*, 18, 155–162.

Navarro, N., 2007, "Fair Allocation in Networks with Externalities," *Games and Economic Behavior*, 58(2), 354–364.

Navarro, N., and A. Perea, 2013, "A Simple Bargaining Procedure for the Myerson Value," *The B.E. Journal of Theoretical Economics*, 13, 131–150.

Parthasarathy, K. R., 1972, *Probability Measures on Metric Spaces*. Academic Press.

Reny, P., 1999, "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games," *Econometrica*, 67, 1029–1056.

Roe, M., 2013, "Derivatives Markets in Bankruptcy," Harvard Law School, Working Paper.

Rubinstein, A., 1982, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, pp. 97–109.

Segal, I., 1999, "Contracting with Externalities," *The Quarterly Journal of Economics*, 114(2), 337–388.

Selten, R., 1975, "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4:1-2, 25–55.

Shapley, L. S., 1953, "A Value for n-Person Games," *Contributions to the theory of games*, 2, 307–317.

Simon, L., and M. Stinchcombe, 1995, "Equilibrium Refinement for infinite Normal-Form Games," *Econometrica*, 63, 1421–1443.

Stole, L., and J. Zwiebel, 1996, "Intra-Firm Bargaining under Non-Binding Contracts," *Review of Economic Studies*, 63(3), 375–410.

Watson, J., 2016, "Contractual Chains: Decentralized Contracting with Externalities," working paper, University of California, San Diego.