

Contractual Chains

Joel Watson*

June 2018[†]

Abstract

This paper develops a model of decentralized contracting, in which a network describes the pairs of players who can communicate and contract with each other. After contracting, the players interact in an underlying game with globally verifiable productive actions and externally enforced transfers. The model allows for a type of externality that the previous literature has not fully explored—in which parties are unable to contract directly with others whose actions affect their payoffs—and it covers a variety of applied settings such as sales of network-based goods, supply relations, two-sided markets, and data-transmission networks. The paper investigates the prospects for efficient outcomes under various contract-formation protocols (“contracting institutions”) and network structures. The main result is positive: There is a contracting institution that yields ε -efficient equilibria for any underlying game and connected network. A critical property is that the institution allows for sequential contract formation or cancellation. The equilibrium construction features *assurance contracts* and *cancellation penalties*.

*UC San Diego; <http://econ.ucsd.edu/~jwatson/>. Some of the work reported here was developed while the author visited Yale in 2010 and when the author was a fellow in residence at the Center for Advanced Study in the Behavioral Sciences at Stanford in 2012-2013. For their helpful comments and discussion, the author thanks Nageeb Ali, Gorm Grønnevet, Keri Hu, Natalia Lazzati, Marek Pycia, Maarten Pieter Schinkel, Joel Sobel, Alex Weiss, colleagues at Yale and UCSD, and seminar participants at Oxford, Washington University, UC Irvine, UBC, CalTech, Pittsburgh, Georgetown, the 2012 Econometric Society summer conference, UT Austin McCombs School, Columbia, Stanford, Mannheim, Paris School of Economics, Northwestern, Arizona, and UCLA. The author thanks the NSF for financial support (SES-1227527).

[†]First version: May 2011. Although this project has been presented widely, initial drafts of the paper were circulated only narrowly and the paper was not submitted for publication prior to June 2018. The explanation for the delay is complicated but includes that the author practiced insane perfectionism with this and other projects, and he took a long detour to develop a general definition of perfect Bayesian equilibrium which, while motivated by this project, in the end is not utilized here!

1 Introduction

In many contractual settings there is multilateral productive interaction and yet barriers prevent the parties from contracting in one large group. Instead, contracts are formed bilaterally in a decentralized manner, and only some pairs of agents have the opportunity to communicate and establish contracts. For example, the developer of a construction project may contract with a general building contractor, who separately establishes agreements with subcontractors for the provision of constituent components. This setting features an externality: The developer cares about the subcontractors' work but she does not contract directly with them. I call this an *externality due to lack of direct links* (an "LDL externality"), formally defined as a situation in which at least one agent is unable to contract directly with another whose productive action he cares about.

LDL externalities exist in a plethora of economic settings, and they often traverse extensive networks and occur bidirectionally. For example, the Internet comprises a variety of service providers, including a backbone of "Tier-1" firms, that collectively transmit data between content providers and end users. A given end user may not have the opportunity to contract with a Tier-1 firm, but she cares about whether this firm maintains its lines and gives priority to the content that she values. The end user contracts with her local Internet service provider, which itself may have an agreement with the Tier-1 firm. Other examples include the internal organization of a firm (where multiple workers have employment contracts with the firm but care about each others' productive actions and may not be able to contract with each other), sales of goods exhibiting network externalities (where each consumer cares about the other consumers' use of the seller's technology), and platforms that facilitate transactions between buyers and sellers (where agents on one side of the market care about whether agents on the other side make investments tied to a particular platform intermediary).

One might ask whether, in principle, LDL externalities can be overcome under private, decentralized contracting. The key issue is whether suitable contractual linkages between relationships will endogenously form in *chains of bilateral contracts*, motivating the parties to choose efficient productive actions. For instance, in the example of a construction project, contractual chains may run from the developer, through the general building contractor, to each subcontractor. Perhaps the developer and the general building contractor will agree to make monetary transfers contingent on the subcontractors' performance, and the general contractor will reach corresponding agreements with the individual subcontractors. The resulting collection of contracts may give the general contractor and subcontractors the incentive to perform at the efficient level.

In this paper, I study whether contractual chains can internalize LDL externalities in an environment with simultaneous productive actions preceded by bilateral contracting via exogenously given network links. All productive actions are verifiable by everyone ("global verifiability") but nothing else is observable or verifiable across contractual relationships. In this setting, contractual linkages can be made only by specifying transfers in one contractual relationship as a function of productive actions taken by agents in other relationships.




Figure 1: A network of contractual relationships.

For instance, in the data-network example, if an end user is able to contract with a particular local service provider, then their agreement can specify transfers between them as a function of a content provider’s productive actions, such as the entertainment it produces and distributes.

Rather than focus on any particular application, I propose an abstract game-theoretic model that is general in terms of the set of players, the specification of who can contract with whom, and the productive technology. Formally, I suppose that a set of n players will interact in an “underlying game” $\langle A, u \rangle$ with externally enforced monetary transfers and full global verifiability. Payoffs are linear in money. Before the underlying game is played, the players can establish contracts bilaterally. Only some pairs of players can communicate and establish contracts, and these pairs are given by an undirected network of bilateral links, $L \subset N \times N$, where $N = \{1, 2, \dots, n\}$. A contract between players i and j specifies transfers between them as a function of the outcome of the underlying game. That is, it is a function $m : A \rightarrow \mathbb{R}_0^n$ such that $m_k(a) = 0$ for all $k \notin \{i, j\}$ and every $a \in A$, where \mathbb{R}_0^n denotes the vectors in \mathbb{R}^n whose components sum to zero. I shall measure efficiency by the players’ joint values (the sum of their payoffs). Since this is meaningful only if it is feasible to transfer money across the network, I assume that L is connected.

A general illustration of network L is shown in Figure 1, where the edges of the graph denote the pairs of players who can communicate and contract. As an example of an LDL externality, player i ’s payoff in the underlying game may depend on player k ’s action a_k , and likewise player k ’s payoff may depend on player i ’s action a_i , and yet these players are unable to contract together. Players i and j can establish a contract, and so can the pair (j, k) , implying that a chain of contracts can in principle indirectly connect players i and k . Observe that a more distant LDL externality may exist between, say, players i' and k'' , and it is possible to indirectly connect them via a longer chain of contracts.

A novel aspect of the approach taken in this paper is that it emphasizes the role of the *contracting institution* that facilitates contract formation. Formally, a contracting institution is an extensive game form with payoff-irrelevant messages and with outcomes defined as externally enforced contracts. For a given contracting institution, the players will play a *grand game* in which they first interact as allowed in the contracting institution, then simultaneously select their actions in the underlying game, and finally receive payoffs including

the contracted transfers.

Critically, contracting institutions are restricted by the network of links and by the following assumptions that represent the notion of voluntary, decentralized contracting:

- **Private** – Each player receives messages from only those to whom he or she is linked in the network, and she does not observe messages exchanged between other players.
- **Independent** – The contract formed between players i and j does not depend on the messages sent by, or received by, any other player k .
- **Voluntary** – Players have the option of rejecting contracts.

I call a contracting institution *natural* if it satisfies these assumptions.


The natural contracting assumptions represent physical constraints and limitations of the external enforcement system that necessitate decentralized contracting rather than centralized planning. The third assumption ensures that the definition of “contract” is conventional, in that it requires the consent of both parties. This is a standard requirement for external enforcement in modern legal systems. The first assumption imposes the network limitation—that communication and contract formation take place only between linked players—and thus it describes real physical constraints on how the players can interact. The second assumption embodies the principle that any two contracting parties are free to form whatever contract they desire, uninhibited by others in the society. Further, it rules out “contracts on contracts,” whereby players in one relationship make their contractual specifications a function of the *contracts* formed by other relationships, not just a function of the productive action profile.¹

I focus on a narrow “possibility” question: Given that contracting must take place in a decentralized manner, is there a natural contracting institution under which efficient outcomes can be achieved for all underlying games and networks? If so, what are the key properties of the contracting institution? In technical terms, the first question asks whether there is a contracting institution that *implements efficient outcomes*, meaning that for *every* underlying game and *every* allowed network structure, there is a sequential equilibrium of the grand game in which an efficient action profile is played for sure in the underlying game. Note that in this exercise, the contracting institution is held fixed and we vary the underlying game and network.

¹Contracts on contracts have been studied to a limited extent. For instance, interactive promises are analyzed by Katz (2006) and Peters and Szentes (2012). Watson (2014) looks at interactive contracts in a setting similar to the one herein. I discuss this further in Sections 3, 4, and 6. One can also motivate the second assumption on the basis of limitations of the enforcement system, whereby messages involved in the formation of contracts are verifiable only locally. In reality, contracting partners rarely go to court. Instead, they voluntarily make the transfers that their contracts require, with the understanding that failure to comply would trigger a lawsuit, at which point the court would compel the required transfers. When a pair of contracting partners appears in court, the judge can observe their contract and the verifiable outcome of the underlying game, for evidence of these can be provided by the contracting parties. However, the judge will not readily observe the contracts written in other contractual relationships and it may be prohibitively costly for the parties to gather such evidence. It would then be feasible to enforce transfers as a function of only the outcome of the underlying game and not the messages sent in other contractual relationships.

To generate intuition and calibrate expectations, consider some simple logic and two examples. Note first that, because every player i is a member of at least one bilateral contractual relationship, there exists a contract that would force player i to select his part of an efficient action profile; the contract could specify a large punishment if player i fails to choose the prescribed action. However, such a contract may not arise in equilibrium.

Example 1 (unidirectional externality—subcontracting): Suppose $n = 3$. Player 1’s action space in the underlying game is $A_1 = \{b, c\}$, whereas players 2 and 3 have no actions. As a function of player 1’s action, payoffs are given by $u(b) = (0, 0, 4)$ and $u(c) = (1, 0, 0)$. The network is $L = \{(1, 2), (2, 1), (2, 3), (3, 2)\}$, so players 1 and 2 are linked and players 2 and 3 are linked.



In the figure above, which depicts the network, I indicate with a filled circle that player 1 has an action in the underlying game. The example illustrates a simple subcontracting relationship, where player 3 is the beneficiary, player 2 is the general contractor, and player 1 is the subcontractor who does the work that player 3 values.

In this example, the efficient outcome cannot be obtained with a natural contracting institution that has one-shot, simultaneous contracting. To see this, consider the contracting institution in which linked players form contracts simultaneously via the “Nash demand game.” That is, players 1 and 2 send contract demands to each other, and simultaneously players 2 and 3 do the same. If a pair of linked players name the same contract, then this contract goes into force; otherwise, they get the *null contract*, which specifies zero transfers always. Note that player 2 declares two contract demands, one for player 1 and one for player 3. Also note that, when the players select actions in the underlying game, they have private information about the contracts they formed earlier.

In order for player 1 to select action b in equilibrium, it must be that players 1 and 2 form a contract that pays player 1 at least 1 for choosing this action in the production phase. Because player 2 can guarantee himself a payoff of 0 by naming the null contract, the proposed equilibrium must also have players 2 and 3 form a contract that transfers at least 1 from player 3 to player 2 in the event that player 1 chooses b . Then player 2 is compensated by player 3 for the amount he must transfer to player 1 conditional on b being chosen. In the proposed equilibrium, players 1 and 2 name the same contract m while players 2 and 3 name the same contract m' , where $m(b) = (\kappa, -\kappa, 0)$ and $m'(b) = (0, \kappa', -\kappa')$, for some values κ and κ' satisfying $1 \leq \kappa \leq \kappa' \leq 4$. Also, it must be the case that $\kappa \geq 1 + m_1(c)$, so that b is supported in the production phase.

Unfortunately, this specification of behavior in the contracting phase is not consistent with individual incentives. If Player 3 deviates by naming the null contract in his relationship with player 2, then he will get a payoff of 4 rather than the payoff of $4 - \kappa'$ that he

would have received under the specified behavior. He gets 4 because his deviation does not disrupt the formation of contract m between players 1 and 2. At the productive phase, player 1 would not know or care that any such deviation occurred, and he would select b given the incentives that contract m provides.

The problem arises for any contracting institution in which player 3 has the last word on the formation of a contract with player 2. For instance, suppose that the contracting institution specifies the following order of moves: Player 2 offers a contract to player 1, who either accepts or rejects it; then player 2 offers a contract to player 3, who accepts or rejects. Again, there is no equilibrium in which player 1 chooses action b , for this would require the kind of contracts described above; but then player 3 could gain by rejecting player 2's offer.

In contrast, if the order of contracting were reversed, so that players 2 and 3 first negotiate, followed by players 1 and 2, then there is an efficient equilibrium. In this equilibrium, player 2 will not make a deal with player 1 unless he was previously successful in contracting with player 3. Thus, the first example illustrates that sequential contracting is essential for obtaining efficient outcomes, because it allows for this sort of conditioning by a player.

However, although the contracting institution just described works well for the externality of Example 1, where player 1's action of b benefits player 3, it would not deliver an efficient outcome if the externality ran the other way. The next example shows that the problem is ostensibly insurmountable in cases of complex externalities.

Example 2 (complex externality—teams, commons, platforms, etc.): Suppose $n = 4$ and $A_i = \{b, c\}$ for $i = 1, \dots, 4$. Payoffs in the underlying game are given by: $u(b, b, b, b) = (1, 1, 1, 1)$, $u(c, c, c, c) = (0, 0, 0, 0)$, and if at least one player selects c and another player selects b then those selecting c get 5 and those selecting b get -16 . The network is $L = \{(1, 2), (2, 1), (2, 3), (3, 2), (3, 4), (4, 3)\}$.



Consider a contracting institution with multiple rounds that gives players flexibility regarding when to establish contracts with each other. In each round, linked players who haven't already formed a contract simultaneously send contract offers to each other. If they name the same contract then this contract is formed and there is no further communication between them. Otherwise, they continue to communicate in the next round. Consider pure-strategy equilibria.

It is not difficult to confirm that an efficient equilibrium does not exist here. Supposing to the contrary, take an efficient equilibrium and note that players 1 and 4 must form contracts on the equilibrium path. Without loss of generality, assume that the contract between players 1 and 2 is formed in the same round or after players 3 and 4 form their contract. Let r be the round in which players 1 and 2 form their contract. If player 1 were to deviate by

naming the null contract in round r and the same in all later rounds, and then select c in the underlying game, then his payoff would be 5. This is because player 1's deviation would not be detected by player 4, whose contract with player 3 would be formed nonetheless, and so player 4 would still select b in the underlying game. We thus have a lower bound of 5 on player 1's equilibrium payoff, which contradicts that each player can guarantee herself at least zero.

Despite the negative theme of these examples, the message of this paper shall be positive. The sole theorem presented here establishes the existence of a natural contracting institution that virtually implements efficient outcomes. This means that for every underlying game, every connected network of contractual relationships, and every $\varepsilon > 0$, there is a sequential equilibrium of the grand game that achieves an efficient outcome with probability of at least $1 - \varepsilon$. Thus, there is a contracting institution under which, in all productive settings, LDL externalities can be overcome by decentralized formation of contractual chains. The key economic message is that verifiable productive actions are sufficient to establish the required linkages between relationships.

To preview the analysis, I next describe some of its main themes, a few technical notes, and an outline of the paper. Let us start on the technical side. Because sequential equilibrium is defined for only finite games, the analysis must be confined to a finite set of underlying games, but this arbitrarily chosen set is otherwise unrestricted. The “virtual” sense of implementation, whereby one achieves the implementation goal within ε for every positive ε , follows Matsushima (1988) and Abreu and Sen (1991). Implementation is “weak” in the sense that we are looking for the existence of an equilibrium with desired properties and are not insisting on uniqueness (an unachievable goal in settings like the one here). From a social-choice perspective, the natural contracting assumptions are the critical constraints that lead to the insights in this paper.²

The proof of the Theorem entails an elaborate construction, but its two key elements are intuitive and easy to describe. First, the contracting institution defined in the proof allows for sequential contract formation. Specifically, it facilitates formation of tentative contracts and then gives players the opportunity to cancel these arrangements over time, so that each player is able to unwind his provisional arrangement with one contracting partner in response to disruptions with another contracting partner. Second, the players coordinate on *assurance contracts*, whereby one player guarantees that specific third parties will select their part of an efficient action profile in the underlying game. For example, in the contract formed between players i and j in Figure 1, player i assures player j that all of the players on i 's side of the network relative to j (those located above the dashed line in the picture) will choose their part of the efficient profile. Likewise, player j assures player i that all of the players on j 's side of the network (those located below the dashed line) will do their part. A similar deal is reached between players j and k , as is the case with all other relationships that actively contract in equilibrium.³

²Weak implementation would be easy without these constraints, because the players commonly know the underlying game.

³Not all linked pairs actively contract in equilibrium—just enough to create a minimally connected net-

Equilibrium contracts specify *assurance penalties*, which motivate players to choose the efficient action profile and, importantly, also to cancel arrangements following out-of-equilibrium disruptions. For example, with reference to Figure 1, in player j 's tentative contract with player k , he assures player k that player i' will select a specified efficient action in the underlying game. But if the contracting between players i and j fails, it causes player j to fear that player i' will not select the efficient action and player j avoids assurance penalties by canceling his tentative contracts with player k and the others. Endogenously specified *cancellation penalties* motivate players not to cancel arrangements toward the end of the contracting phase, neutralizing the problematic forces highlighted in the two examples, but they aren't so great as to dissuade players from cancelling when contracting has been disrupted.

The proof starts by tackling a large class of settings called the *critical-extremes class*, where efficient and Nash-equilibrium actions in the underlying game differ for players at the extremity of the network. Example 2 shown above is in this category. For this class, the essential elements just described operate in the construction of equilibria that implement efficient outcomes with certainty—that is, with $\varepsilon = 0$ —so the allowance of “virtual” implementation is not needed. Broadening the analysis to include cases such as Example 1 shown above requires a variation of the equilibrium construction. In this case, players who do not have actions in the underlying game and are at the periphery of the network (such as player 3 in Example 1) establish “pay in” contracts with those who indirectly connect them to players who have critical actions in the underlying game. Again, implementation is with certainty.

Thus, virtual implementation plays a minor role in the modeling exercise and is needed only to deal with the case in which a player has more than one action in the underlying game and yet her efficient action coincides with her Nash-equilibrium action. For such a case, the equilibrium construction entails contracts that induce such a player to condition her action on the outcome of a public randomization device.

The next section describes the general model precisely. Section 3 briefly reviews some applications and related literature. Section 4 contains the Theorem and the main parts of the proof, including the complete argument for the critical-extremes class and a summary for the other cases. The remainder of the proof is contained in the Appendix. Section 5 provides tangential results and discusses variations of the model. The tangential results address how the required number of rounds in the contracting institution relates to the diameter of the network, the problem with letting players freely cancel contractual arrangements, and multiplicity of implemented payoffs. This section also shows that the Theorem does not hold in settings with partially verifiable productive actions. The Conclusion offers additional comments on the related literature, implications, and further steps in the research program.

work of active pairs. For example, in reference to Figure 1, only two of the pairs (j, k) , (j, k') , and (k, k') form contracts in equilibrium.

2 The Model

This section describes the contractual setting and implementation concept.

2.1 The Setting

There are n players who will interact in an underlying game $\langle A, u \rangle$, where $A = A_1 \times A_2 \times \dots \times A_n$ is the space of action profiles and $u: A \rightarrow \mathbb{R}^n$ is the payoff function. Payoffs are in monetary units. Let $N = \{1, 2, \dots, n\}$ denote the set of players. I assume that n is finite and that A is a subset of some finite set \mathcal{A} , which is interpreted as the set of all possible action profiles that the external enforcer can recognize. For any subset of players $J \subset N$, write $a_J \equiv (a_i)_{i \in J}$ as the vector of actions for these players. The players commonly know the action space and payoff function.

An external enforcer compels monetary transfers between the players, and he does so as directed by contracts that the players form. The enforcer does not observe the payoff function u or the feasible space of action profiles A . Otherwise, the contracting environment has *full verifiability* in that the enforcer observes the action profile $a \in A$ that the players select. Thus, transfers can be freely conditioned on the outcome of the underlying game.

Contracts can be formed only in a restricted set of bilateral relationships. This set of contractual relationships is given by a fixed network $L \subset N \times N$, with the interpretation that players i and j can form a contract if and only if $(i, j) \in L$. Contracting by larger groups of agents is not possible. For any player i , let

$$L^i \equiv \{j \in N \mid (i, j) \in L\}$$

be player i 's set of *contracting partners*, with whom he is linked in the network. Network L is undirected and thus symmetric, so $(i, j) \in L$ implies that $(j, i) \in L$. The network is generally not transitive; thus, player i may be able to contract with player j , and player j may be able to contract with player k , whereas player i is unable to contract with player k . Assume that L is connected.

Contracting partners can condition transfers between them on actions taken by third parties. For example, a contract between players 1 and 2 could specify that player 1 must pay an amount to player 2 in the event that player 3 selects a particular action in the underlying game. However, contracts may not impose transfers on third parties. For instance, the enforcer will not enforce a contract between players 1 and 2 that specifies a transfer to or from player 3.⁴ Thus, a contract between players i and j is a function $m: \mathcal{A} \rightarrow \mathbb{R}_0^n$, where $m_k(a) = 0$ for all $k \notin \{i, j\}$ and every $a \in \mathcal{A}$. Let C^{ij} be the set of feasible contracts between players i and j . Denote by \underline{m} the null contract that always specifies a transfer of zero.

⁴Due to full verifiability, the analysis here would not be affected if one allowed contracting partners to commit to make transfers to, but not from, third parties (“unbalanced transfers”).

Let \mathcal{M} denote a set of contracts formed by the various contractual relationships. Given such a set, let

$$M(a) \equiv \sum_{m \in \mathcal{M}} m(a).$$

Because of transferable utility and a connected network, an efficient action profile a^* must solve $\max_{a \in A} \sum_{i \in N} u_i(a)$; that is, it maximizes the players' joint value (the sum of the players' payoffs in the underlying game). I say that a set of contracts \mathcal{M} *supports* a^* if a^* is a Nash equilibrium of the game $\langle A, u + M \rangle$.

2.2 Contracting Institution Formalities

A *contracting institution* is an extensive game form with payoff-irrelevant messages that map to contracts enforced by the external enforcer. I constrain attention to a particular class of game forms in which the players send messages in discrete rounds $1, 2, \dots, R$, where R is a finite integer.⁵ There is no discounting. In every round, each player i sends messages to every other player. Let λ_{ij}^r denote the message that player i sends to player j in round r . Suppose that there is a special “null message” $\underline{\lambda}$ that we can interpret as silence.

Let $h_{ij}^r = (\lambda_{ij}^1, \lambda_{ij}^2, \dots, \lambda_{ij}^r)$ denote the sequence (or “history”) of messages sent from player i to player j from round 1 to round r . Let $\underline{h}^r = (\underline{\lambda}, \underline{\lambda}, \dots, \underline{\lambda})$ be the sequence of r null messages. A full history of messages through round r is given by $h^r = \{h_{ij}^r\}_{i \neq j}$. A corresponding history of messages between just players i and j is given by (h_{ij}^r, h_{ji}^r) .

The space of feasible messages in each round is allowed to be a function of the history of messages sent in earlier rounds. Let Λ_{ij}^1 denote the set of feasible messages from player i to player j in round 1. For $r \in \{2, 3, \dots, R\}$ and a history of messages h^{r-1} through round $r - 1$, let $\Lambda_{ij}^r(h^{r-1})$ denote the set of feasible messages from player i to player j in round r . Assume that the null message is an element of the feasible message set in every round. The message sets are defined arbitrarily by the contracting institution and, for example, can include descriptions of contracts.

I allow the contracting institution to specify a public randomization device following messages. Let ϕ denote this random draw.

The contracting institution specifies a mapping from the outcome of the contracting phase (the full history of communication through R , as well as the random draw ϕ) to the contracts formed between the various pairs of players. Thus, this mapping represents how the contracting institution interprets the sequence of messages. For any outcome of messages and random draw (h^R, ϕ) , let $\mu(i, j, h^R, \phi) \in C^{ij}$ denote the contract formed between players i and j . Because $\mu(i, j, h^R, \phi)$ and $\mu(j, i, h^R, \phi)$ would be the same contract, we can avoid the redundancy by defining $\mu(i, j, \cdot)$ only for $i < j$. Then, for a given sequence of messages and random draw (h^R, ϕ) , the set of contracts formed is

⁵A previous version of this paper focused on the setting of $R = \infty$, which makes for a simpler proof of the main result.

$\mathcal{M} = \{\mu(j, i, h^R, \phi) \mid i < j\}$ and the function describing the sum of contracted transfers is

$$M \equiv \sum_{i=1}^{n-1} \sum_{j>i} \mu(i, j, h^R, \phi).$$

Remember that M maps A to \mathbb{R}_0^n .

A contracting institution is called *natural* if the following assumptions hold: First, contracting is private in that (i) a given player can receive messages from only those to whom he or she is linked in the network, and (ii) a player cannot observe messages exchanged between other players. Second, contracting in different relationships takes place independently. Third, players have the option of refusing to contract. Here are the formal descriptions of these assumptions:

Part ii of the private contracting assumption requires that, at the end of round r , player i 's personal history is exactly $\{h_{ij}^r, h_{ji}^r\}_{j \neq i}$, the sequence of messages to and from this player. It follows that player i 's strategy in the contracting phase specifies, for each round r , a mapping from the feasible set of such personal histories to $\times_{j \neq i} \Lambda_{ij}^r$.

The assumption of independent contracting requires that, for any two players, their sets of feasible messages and the contract between them can depend on only the sequence of messages that they send to each other. That is, for any two players i and j , we have the following. First, for any $r \in \{2, 3, \dots, R\}$, $\Lambda_{ij}^r(h^{r-1})$ depends on h^{r-1} only through the component $(h_{ij}^{r-1}, h_{ji}^{r-1})$, so we can write $\Lambda_{ij}^r(h_{ij}^{r-1}, h_{ji}^{r-1})$. Second, the contract $\mu(i, j, h^R, \phi)$ depends on h^R only through the component (h_{ij}^R, h_{ji}^R) , and we write $\mu(i, j, (h_{ij}^R, h_{ji}^R), \phi)$.

To represent that players can refuse to contract, I assume that if player i always sends the null message to player j then the contract formed between these two players is exactly the null contract. That is,

$$\mu(i, j, (\underline{h}^R, h_{ji}^R), \phi) = \mu(i, j, (h_{ij}^R, \underline{h}^R), \phi) = \underline{m},$$

for all h_{ij}^R and h_{ji}^R .

To make precise the assumption that players may communicate only with players to whom they are linked (part i of private contracting), I suppose that the network L transforms the game form into an *effective game form* in which, for each pair of players i and j , if $(i, j) \neq L$ then these players are restricted to send each other the null message in each round. Thus, $h_{ij}^R = \underline{h}^R$ and $h_{ji}^R = \underline{h}^R$, for all pairs such that $(i, j) \neq L$. In this case, note that no information can be exchanged directly between players i and j , and their contract is null.

To summarize, a contracting institution is defined by a number of rounds R , message spaces $(\Lambda_{ij}^r)_{i,j,r}$, a public-randomization device, and functions $\mu(i, j, \cdot)$ for all $i < j$. It is called *natural* if it satisfies the assumptions detailed above.

2.3 The Notion of Implementation

For a given contracting institution, network L , and underlying game $\langle A, u \rangle$, the entire game between the players runs as follows:

1. **Contracting phase:** Players interact in the contracting institution to form contracts, resulting in the set \mathcal{M} .
2. **Production phase:** Players simultaneously select actions in the underlying game.
3. **External enforcement phase:** The enforcer observes the outcome a of the underlying game and compels the transfers $M(a)$. The payoff vector for the players is $u(a) + M(a)$.

Note that, because of private contracting, the players have asymmetric information at the time of productive interaction. For example, player i does not observe the contract formed between two other players j and k .

I analyze behavior using the concept of sequential equilibrium (Kreps and Wilson 1984), with beliefs at information sets expressed in terms of appraisals (Watson 2017) which is convenient for the kind of game studied here. To keep the game finite, as required to apply sequential equilibrium, I restrict attention to a finite subset of underlying games G . The set G is arbitrary, so it can be large and varied; the only restriction is that it is a finite set.⁶

In the analysis that follows, I hold fixed the number of players n and the set G , with the understanding that the grand space of action profiles \mathcal{A} is the union of actions profiles for the games in G . The primary objective is to evaluate the performance of a given contracting institution across various networks L and underlying games $\langle A, u \rangle \in G$. I use a standard notion of implementation, appropriately defined for the model developed here.

Definition 1: *Say that a given **contracting institution implements efficient outcomes** if for every underlying game $\langle A, u \rangle \in G$ and every connected network L , there is a sequential equilibrium of the entire game in which an efficient action profile a^* is played. Let us say that the **contracting institution implements ε -efficient outcomes** if for every underlying game $\langle A, u \rangle \in G$ and every connected network L , there is a sequential equilibrium of the entire game in which an efficient action profile a^* is played with probability at least $1 - \varepsilon$.*

Because G is finite, for small enough ε , implementation of ε -efficient outcomes implies an equilibrium joint value that is arbitrarily close to the maximal joint payoff, for every underlying game.

⁶It is possible to examine all underlying games by using a perfect Bayesian equilibrium concept, such as that defined in Watson (2017), but the characterization of some off-equilibrium-path contingencies turns out to be more complicated than it is worth.

3 Applications and Related Literature

This section briefly describes applications and connections to the related literature. To distinguish LDL externalities from other externalities that the literature has focused on, let us say that an *externality of non-contractibility* is present when a pair of linked players is not allowed to condition transfers on an action that is payoff relevant to them, perhaps one of their own actions or the action of a third player. This could be because the action is not verifiable or due to legal restrictions. Note that the model studied here has no externalities of noncontractibility.

Comparisons with some models in the related literature are difficult due to the way in which these models describe productive actions as essentially “public” (taken by a third party) and occurring automatically with contract formation. Consider a buyer and seller who may contract on a product to be created and delivered by the seller in exchange for a payment from the buyer. A public-action model might entail a noncooperative-game accounting of the individual actions that the parties take to form a contract, such as offer and acceptance choices, but would not include a noncooperative account of productive interaction. In such a public-action model, payoffs are specified as a function of the contractual provisions. For instance, if the parties agree that a particular good shall be produced and delivered, then production and delivery are assumed to occur as agreed and the payoffs correspond to the associated benefits and costs.

In contrast, the model herein includes a noncooperative account of the productive interaction, as specified by the underlying game. This allows us to distinguish between various ways in which contracts can link relationships, in particular payments conditioned on others’ verifiable productive actions as distinct from provisions conditioned on the specifications of other contracts (contracts on contracts). We can also clearly account for settings in which a player’s productive action is influenced by her contracts with multiple partners. Finally, accounting for individual productive actions is required to understand the full extent to which these actions can serve as options (Watson 2007).

Apart from variations in the modeling of productive actions, much of the related literature comprises special cases of the general model proposed here but focuses on externalities of non-contractibility. Quite a few related models do not address LDL externalities, in some cases for simplicity. Also, each of the models specifies a single contracting institution, typically a protocol with simultaneous commitments or offers. Holding aside common agency with complete information, which produces efficient equilibria, the literature typically finds that inefficient outcomes are inescapable even without externalities of non-contractibility.

For the applied settings described below, the analysis herein makes two contributions. The first is to illustrate the value of clearly accounting for the technology of production and contracting, including a specification of individual productive actions and verifiability. The second is to demonstrate that efficient outcomes may be realized in settings of “global verifiability” despite LDL externalities. Therefore, perhaps observed inefficiencies can be traced to verifiability constraints and barriers to contracting on verifiable actions.




Figure 2: Common-party vertical contracting, network-based goods.

Common-Party Vertical Contracting, Network-Based Goods (Star Networks)

In the case of vertical contracting with a common supplier, a monopoly seller of an intermediate good contracts separately with multiple downstream buyers. The underlying game accounts for actions parties take to produce, trade, and utilize the intermediate good. Cremer and Riordan (1987) study such a setting with a public-action model that includes idiosyncratic privately-observed productivity/demand shocks, but there are no LDL externalities because each buyer's payoff is assumed to not depend on the number of units delivered to other buyers. The solution concept is a stability notion requiring that the contract for each seller-buyer pair is selected according to the Nash bargaining solution, fixing the contracts with the other buyers—a solution usually called *Nash-in-Nash*.

In the model of Horn and Wolinsky (1988), there are two buyers and they simultaneously select quantities in the retail market (a Cournot-type underlying game). Because these firms care about each others' quantities, there is an LDL externality. Contracts are restricted to be linear and to not depend on the other downstream firm's quantity choice (neither directly nor indirectly as, say, through a profit-sharing arrangement), so a non-contractibility externality is assumed. McAfee and Schwartz (1994) examine a similar model with a simultaneous-ultimatum-offer contracting institution, where the seller offers a two-part tariff to each buyer and the payment is a function of only this buyer's quantity (a non-contractibility externality). Efficiency is realized in Cremer and Riordan's setting but generally not in Horn and Wolinsky's setting and McAfee and Schwartz's setting.⁷

Bernheim and Whinston's (1986a,b) *common agency* framework features a single agent, who has the only action in the underlying game, and multiple principals who simultaneously make commitments regarding how much money to transfer to the agent as a function of the agent's action. The agent does not have the option of declining the principals' offers, so in a sense the setting has unilateral commitments rather than contracting; however, the

⁷Segal's (1999) public-action model of contracting with externalities (see also Segal and Whinston 2003) seeks to synthesize and partly generalize these and other models. A principal and multiple agents are connected by a star network with the principal in the center. We can translate the model into one with individual trade actions, whereby the principal has the only action in the underlying game and it represents the number of units to trade with each agent. There are no LDL externalities in this representation. By assumption, the contract with one agent cannot specify transfers as a function of the action components for other agents, so externalities due to non-contractibility are assumed. Unless the principal can make the trade with one agent conditioned on communication with another (a type of multilateral contract), equilibria are typically inefficient. See Galasso (2008) for a more recent study that looks at various bargaining protocols and provides additional references. Moller (2007) performs an analysis of the common-principal problem with sequential contracting.




Figure 3: Supply relations, games played through agents.

transfers are constrained to be nonnegative, so the agent would never refuse. There are no LDL externalities and efficient equilibria always exist in settings of complete information.⁸

Another application is contracting for the development of parcels of land, whereby a number of land owners have individual actions that determine whether and how their parcels are developed. In some cases, another party seeks to acquire or coordinate development across all of the parcels, and this party can be viewed as seeking bilateral contracts with the land owners. Eminent domain may be modeled as altering the underlying game (a change in the default legal rules) or as forcing a particular centralized contract, the latter which would not necessarily be efficient because the local government may lack information about underlying-game payoffs. Similar applications include the organization of shopping centers, where the center’s operator contracts with individual retailers, and the organization of production teams, a setting described more generally below.

Supply Relations, Games Played Through Agents (Bipartite Networks)

In a bipartite supply network, there is a group of upstream firms and a group of downstream firms. Bilateral contracting occurs across the two groups but not within the groups. Actions in the underlying game may include production, trade, and utilization of intermediate goods. Collard-Wexler, Gowrisankaram, and Lee (2017) develop a public-action model in which the outcome of contract negotiation includes a binary outcome—agreement or no contract—and, if agreement, a monetary transfer. Utilities are a function of the set of agreements, plus the contracted transfers. For every agreement that a given player i cares about, this player can contract with someone who is a party to the agreement (because player i can contract with every player on the other side of the market). Thus, there are no LDL externalities. Transfers cannot be contingent on whether others have agreements, so players cannot internalize externalities.⁹

⁸More recent entries, such as Prat and Rustichini (1998) and Bergemann and Valimaki (2003), examine settings in which the principals make their commitments sequentially; as with the static case, there are efficient equilibria. The literature also looks at settings with incomplete information, where efficient outcomes are generally not attained. For a survey of the literature, see Martimort (2007).

⁹Collard-Wexler, Gowrisankaram, and Lee’s main exercise is to relate the Nash-in-Nash solution to the equilibria of a noncooperative contracting institution that has semi-independent bargaining by agents. In the bargaining protocol, players observe agreements as they are formed, so contracting is not fully private. Lee and Fong (2013) conduct a similar exercise for more general networks. The literature also includes models of supplier network formation and investments that take place before contracting to trade, such as Kranton and Minehart (2001) and more recently Elliott (2013).




Figure 4: Two-sided markets, platforms.

In another example of bipartite networks, Prat and Rustichini (2003) extend the common agency framework to have multiple agents. Only the agents have actions in the underlying game. The principals simultaneously make unilateral commitments regarding how much money to transfer to the agents as a function of the agents' actions. The authors examine the case in which a transfer to one agent can be conditioned on the actions of all of the agents, so there are no externalities of non-contractibility. There are also no LDL externalities because it is assumed that each agent does not care about the actions of the other agents. Contracting is "sparse" in the sense that agents cannot contract with one another and, likewise, principals cannot contract with one another. Still, efficient equilibria are not guaranteed, although they arise in some cases.

Two-Sided Markets, Platforms (Intermediation Networks)

In an intermediation network (also called a two-sided market), one or more platform intermediaries facilitate transactions between two classes of players on opposite sides of a market. For example, on one side are consumers, on the other side are merchants, and the network intermediary provides a payment-processing system. In a detailed model, actions in the underlying game would refer to whether the intermediary provides access to its system, consumers' and merchants' selection of payment methods for transactions, and investments associated with the intermediary's technology. The applied literature has typically focused on public-action models, which can be translated into underlying games in which the intermediary has the only productive action and so there would be no LDL externalities. Rochet and Tirole (2003, 2006) and Armstrong (2006) look at various pricing alternatives for platform intermediaries, including prices indexed by the number of participants served on the other side of the platform or the number of transactions processed. Market efficiency is not their main focus, but in some cases the preferred pricing formula can yield efficient outcomes.¹⁰

¹⁰Rysman (2009) provides an overview. Other recent papers include Weyl (2010) on pricing strategies and monopoly power, Lee (2014) on tipping points for platform adoption, Reisinger (2014) on two-part tariffs and equilibrium selection, Edelman and Wright (2015) on the microstructure of interaction between agents on the two sides of the market, and Hagiu and Wright (2015) on vertical integration.




Figure 5: Subcontracting, supply/distribution chains.

Subcontracting, Supply Chains, Distribution Chains (Tree networks)

The subcontracting application was described in the Introduction. Supply chains have a similar structure. There is also the case in which intermediate goods flow through tree networks in the opposite direction, which we can call distribution chains. Cremer and Riordan (1987) examine such a setting with multiple levels and no LDL externalities, finding that the Nash-in-Nash outcome is efficient.

More Complex Networks and Other Models

The category of more complex networks includes data networks and the Internet. In addition to the linkages described in the Introduction, in some cases consumers can contract directly with content providers even though they rely on service providers to transmit data between them (as shown in Figure 6). Another application is team production and organization of the firm, whereby the underlying game represents the productive interaction between managers and workers, and the network represents contractual limitations that arise due to legal constraints and/or transaction costs.¹¹ The special case of a star network captures the example of a team of workers who all have bilateral employment contracts with a manager. Commons problems (Lloyd 1833, Hardin 1968) feature a dense system of payoff interdependence with regard to productive actions such as extraction of an exhaustible common-property resource and, where bilateral contracting is possible, chains may facilitate internalization of the resulting LDL externalities. Other settings in this category include the network of international trade agreements, where a country’s policies and trade choices can affect welfare globally.¹²

There are some related papers in the more abstract theory literature. Fleinery, Jankó z, Tamura, and Teytelboym (2016) propose a public-action model of decentralized contracting in which each player’s payoff depend on the sets of contracts that this player signs (so there are no LDL externalities) and explore the relation between various stability concepts. Jackson and Wilkie (2005) study a model in which players can make unilateral commitments

¹¹An important but relatively undeveloped theory of the firm views an organizational structure as a nexus of contracts (Jensen and Meckling 1976). See Laffont and Martimort (1997) for a survey of the literature. Other organizational structures comprise networks of contracts but may not be regarded as traditional firms. Cafeggi (2008) provides one legal perspective.

¹²Economides (1996) provides an overview of contractual issues in general networks, mostly notably data networks. Bagwell, Staiger, and Yurukoglu (2017) analyze networks of international trade agreements with and without a most-favored-nations rule, using a public-action model without LDL externalities and the Nash-in-Nash concept.




Figure 6: Applications with more general contractual networks.

before playing an underlying game. The commitments are promises to make positive monetary transfers to other players as a function of the action profile in the underlying game. Every player can commit to make transfers to every other player, so the implied network is complete. The institution that determines promises is one-shot, simultaneous announcement of commitments. The authors characterize the equilibrium outcomes; they show that efficient outcomes can be supported in some cases but not generally.¹³

Further afield is the growing literature on coalitional bargaining. In these models, the grand coalition can form a contract (so centralized contracting is possible) but subgroups of players can shape the final agreement by first making agreements in their smaller coalitions. The incentives of coalitions to manipulate in this way sometimes precludes the attainment of an efficient outcome.¹⁴ One variation is a form of interconnected bilateral bargaining, where disagreement in one bilateral relationship causes this relationship to sever and triggers a restart to negotiation in all the other relationships.¹⁵ Finally, there is a small literature on matching with externalities; for instance, Pycia and Yenmez (2017) investigate how standard properties of stable matchings extend to settings with externalities.

¹³Ellingsen and Paltseva (2016) expand on Jackson and Wilkie’s (2005) model by allowing the players to form agreements rather than just make binding promises, and they also allow unilateral promises. They show that, regardless of the underlying game, there exist efficient equilibria. Related papers from the prior literature include Guttman (1978), Danziger and Schnytzer (1991), Guttman and Schnytzer (1992), Varian (1994), and Yamada (2003). Peters and Szentes (2008) examine a setting in which players can make unilateral commitments about how they will play in the underlying game, and each player’s promises can be conditioned on the promises of the others. The authors develop a mathematical apparatus to handle the infinite regress issue, and they prove a folk theorem (implying the existence of efficient equilibria). Their modeling exercise reveals that interactive contracts require the external enforcement system to develop a sophisticated language for the cross-referencing of promises.

¹⁴The typical model has a public, dynamic negotiation process. Payoffs are described in terms of a characteristic function over coalition structures or a state variable that can be changed over time by the active coalition. A representative sample of contributions is: Chatterjee et al. (1993), Seidmann and Winter (1998), Gomes (2005), Gomes and Jehiel (2005), Bloch and Gomes (2006), Hafalir (2007), and Hyndman and Ray (2007).

¹⁵De Fontenay and Gans (2014) present a model with only bilateral bargaining by players linked on a network. There are LDL externalities and non-contractibility externalities (linked pairs can contract on only their own “trade” and cannot condition transfers on actions taken in other relationships). Contracting takes place sequentially and privately, except that everyone knows when a link is broken. The authors isolate equilibria that produce the Myerson value of a characteristic function defined by bilaterally efficient actions.

4 Efficient Implementation

Here is the main result:

Theorem: Fix n and a finite set of underlying games G . Consider any number $\varepsilon > 0$. There exists a natural contracting institution that implements ε -efficient outcomes.

Note that the natural contracting institution identified here performs well for *all* connected networks and *all* underlying games in G . That is, with this single contracting institution, for every connected network and every underlying game, efficiency is approximately obtained in a sequential equilibrium.¹⁶

The proof features a contracting institution that I shall call the *PCA contracting institution*, where P stands for “provisional,” C for “cancellation options,” and A for “assurance.” This contracting institution has a two-part messaging structure, where tentative contracts are formed and then players have the option of canceling them. In the first round of messaging, contracting pairs engage in a Nash demand protocol that determines their *provisional arrangements*, which specify *tentative contracts* and *cancellation penalties*. In later rounds the players can unilaterally cancel their tentative contracts. Contracts that are not canceled will then be enforced. If a player is the first to cancel a contract, then this party must pay the specified cancellation penalty.¹⁷

The equilibrium construction features *assurance contracts*, where in the contract between two players i and j , player i guarantees that players on i 's side of the network will choose their part of the efficient action profile in the underlying game. A player k is on i 's side of the network relative to j if the path from player k to player j goes through player i . (This is in reference to a minimally connected subnetwork.) Assurance penalties are set high enough to motivate players to choose the efficient action profile. By choosing assurance penalties higher than the cancellation penalties, players are motivated to cancel tentative contracts in the event their arrangements with others are nullified. Cancellation penalties are still high enough to dissuade players from canceling arrangements toward the end of the contracting phase.

The major parts of the proof are contained in the remainder of this section. I begin in Subsection 4.1 with a precise description of the contracting institution. In Subsection 4.2 I describe a large class of settings that I call the *Critical-Extremes Class* and provide the proof for this class. This subsection also contains an informal description of the equilibrium construction for Example 2, which provides the basic idea for readers who prefer not to delve into the general details. Subsection 4.3 generalizes the analysis, gives an informal account of the equilibrium construction for Example 1, and completes the proof, with details shown in the Appendix.

¹⁶The value ε can be set arbitrarily small for the fixed set G . It is possible to strengthen the implementation criterion to virtual implementation, but this requires infinite message spaces and thus a PBE solution concept, as would be required to let G be the set of all underlying games.

¹⁷I thank Gorm Grønnevet for suggesting a version of the analysis in which cancellation penalties are used.

4.1 Specifications of the PCA Contracting Institution

In the PCA contracting institution, the players will be restricted to a finite set $\overline{\mathcal{M}}$ of feasible contracts and a finite set $\overline{Q} \subset \mathbb{R}$ of feasible penalties. The proof will indicate elements that $\overline{\mathcal{M}}$ and \overline{Q} are assumed to contain; it will be clear that any additional elements may be included in these sets without affecting the analysis. The null contract is an element of $\overline{\mathcal{M}}$ and $0 \in \overline{Q}$. Also taken as a parameter will be a number $\varepsilon \in (0, 1)$, assumed small, which will coincide with the value referenced in the statement of the Theorem.

The PCA contracting institution, with parameters $\overline{\mathcal{M}}$, \overline{Q} , and ε , can now be described. Let $R = n - 1$. (Any integer that is at least $n - 1$ will do.) The public randomization device is a binary random variable that takes value $\phi = 1$ with probability $1 - \varepsilon$ and value $\phi = 0$ with probability ε . Message spaces in the first round are such that the players name (i) feasible contracts for each of the two values of ϕ , and (ii) cancellation penalties for later rounds of the contracting phase. That is, for each pair of players (i, j) , message set Λ_{ij}^1 is the set of triples (c^0, c^1, p) , where $c^0, c^1 \in C^{ij} \cap \overline{\mathcal{M}}$ and $p = (p^2, p^3, \dots, p^R) \in \overline{Q}^{R-1}$. Here c^0 is the contract suggested for the $\phi = 0$ contingency, c^1 is suggested for the $\phi = 1$ contingency, and p is a vector of cancellation penalties. For each $r \in \{2, 3, \dots, R\}$, p^r is the cancellation penalty for round r . Identify the null message $\underline{\lambda}$ for round 1 as specifying $c^0 = \underline{m}$, $c^1 = \underline{m}$, and $p = (0, 0, \dots, 0)$.

If players i and j send the same message (c^0, c^1, p) in round 1, and if $c^0 \neq \underline{m}$ or $c^1 \neq \underline{m}$ or both, then call (c^0, c^1) their “provisional arrangement.” If the players send different messages in round 1, or if they both send (c^0, c^1, p) with $c^0 = c^1 = \underline{m}$, then say that the players “made no provisional arrangement.”

In rounds $r \in \{2, 3, \dots, R\}$, players can send the null message or declare “cancel,” and they are further restricted by the history as follows. If the pair (i, j) made no provisional arrangement in round 1, then they are restricted to send each other the null message in all other rounds. If players i and j made a provisional arrangement, then they each have the choice of “cancel” or the null message in later rounds, except that once “cancel” is declared by one or both of the players then these two players are restricted to the null message with each other for the rest of the contracting phase.

For each pair (i, j) , the function $\mu(i, j, \cdot)$ is defined in a straightforward way. If the players named different triples (c^0, c^1, p) in round 1, then they get the null contract \underline{m} . Next suppose players i and j named the same triple (c^0, c^1, p) in round 1 and they sent message $\underline{\lambda}$ in all later rounds. Then for random draw ϕ , the contract between them is defined to be $m \equiv c^\phi$. Finally, if the players named the same pair (c^0, c^1, p) in round 1 and one or both of them later sent the message “cancel,” then the contract between them is a constant transfer. For this case, let \hat{r} be the round in which the message “cancel” was sent. If player i alone sent the “cancel” message in round \hat{r} , then the constant transfer is $p^{\hat{r}}$ from player i to player j . If player j alone sent the “cancel” message in round \hat{r} , then the constant transfer is $p^{\hat{r}}$ in the other direction. If the players both canceled in round \hat{r} then the constant transfer is defined to be zero (the null contract \underline{m}).

The PCA contracting institution is natural. To see this, note that contracting is private

because we look at the effective game form that restricts messages to $\underline{\lambda}$ between pairs of players that are not in the network. Contracting is independent across relationships because the contract between players i and j is a function of only the messages that players i and j exchange. Finally, players can reject contracts (unilaterally force the null contract) by sending message $\underline{\lambda}$ in the first round.

To prove the Theorem, we must find a finite set of contracts $\overline{\mathcal{M}}$ and a finite set of penalty amounts \overline{Q} such that, for any underlying game $\langle A, u \rangle \in G$ and any connected network L , there is a sequential equilibrium of the entire game with the PCA contracting institution (specifying feasible contracts $\overline{\mathcal{M}}$ and feasible penalties \overline{Q}) in which an efficient action profile a^* is played with probability of at least $1 - \varepsilon$.

4.2 The Critical-Extremes Class

In this subsection I prove a restricted version of the Theorem, which covers a broad class of underlying games that includes many examples of interest and also illustrates the key elements of the theory. For this class, an efficient outcome is implemented with certainty.

Note that, whatever is the connected network L , we can find a minimally connected sub-network $K \subset L$, where each pair of players is connected (indirectly or directly) by exactly one path. This means that for each pair of players i and j , there is exactly one sequence $\{k^t\}_{t=1}^T \subset N$ with the following properties: $k^1 = i$, $k^T = j$, and $(k^{t-1}, k^t) \in K$ for all $t = 2, 3, \dots, T$. The sequence $\{k^t\}_{t=1}^T$ is called the *path from i to j* . A minimally connected subnetwork K exists by construction and is defined to be symmetric. For any such minimally connected network K , call a player i *extremal* if player i has exactly one contracting partner—that is, there is a single player $j \neq i$ for which $(i, j) \in K$.

Definition 2: *A given network L and underlying game $\langle A, u \rangle$ are said to be in the **Critical-Extremes Class** if $\langle A, u \rangle$ has a Nash equilibrium $\underline{\alpha} \in \Delta A$ and there is an efficient action profile a^* with the following property: L has a minimally connected subnetwork K such that $\underline{\alpha}_i(a_i^*) < 1$ for each player i who is extremal.*

Note that $\underline{\alpha}$ is a mixed action profile, with $\underline{\alpha}_i(a_i)$ giving the probability that player i puts on action a_i . Clearly, a pure-action equilibrium is a special case. In the critical-extremes class, each extremal player's Nash equilibrium action puts probability less than 1 on this player's component of the efficient action profile a^* . In the case of a pure-action equilibrium \underline{a} , we have $\underline{a}_i \neq a_i^*$ for each extremal player i .

Proposition 1: *Fix n and \mathcal{A} , and consider any network L and any n -player underlying game $\langle A, u \rangle$ with $A \subset \mathcal{A}$. Assume that L and $\langle A, u \rangle$ are in the critical-extremes class. There exists a finite set of contracts \mathcal{M}' and a finite set of penalties $Q' \subset \mathbb{R}$ such that the following holds for all finite sets of feasible contracts $\overline{\mathcal{M}}$ and penalties \overline{Q} such that $\mathcal{M}' \subset \overline{\mathcal{M}}$ and $Q' \subset \overline{Q}$. Let the contracting institution be the PCA contracting institution with feasible contracts $\overline{\mathcal{M}}$, feasible penalties \overline{Q} , and any ε . There is a sequential equilibrium of the entire game in which the equilibrium outcome is efficient.*




Figure 7: Construction of K and definition of β .

In this case, the efficient action profile has all of the extremal players selecting actions that differ from their underlying-game Nash equilibrium actions. Figure 7 illustrates L and K ; open nodes represent players for which $\underline{\alpha}_i(a_i^*) = 1$, which includes the case of having no action in the underlying game. The conditions of this theorem are trivially satisfied for any underlying game in which $\underline{\alpha}_i(a_i^*) = 0$ for all $i \in N$, so it doesn't matter which players are extremal.

The remainder of this subsection contains a proof of Proposition 1. Let us fix L and $\langle A, u \rangle$ in the critical-extremes class, and let K , $\underline{\alpha}$, and a^* be a minimally connected sub-network and actions profiles with the properties required in Definition 2.

There are several steps in the proof. I start by describing core elements of the equilibrium construction—namely the prescribed provisional arrangements for pairs in the network K , which also identifies elements needed for the sets \mathcal{M}' and Q' . Then, to provide intuition, I give an informal description of the equilibrium using Example 2 from the Introduction. This is followed by the general technical development, including a precise specification of actions and beliefs at some key information sets (personal histories) and then the analysis of the other information sets.

Prescribed provisional arrangements and cancellation penalties

For each $(i, j) \in K$, we can divide the set of players into two disjoint groups by relative proximity to players i and j on network K . Define:

$$\beta(i, j, K) \equiv \{k \in N \mid j \text{ is not on the path from } i \text{ to } k\}.$$

In words, $\beta(i, j, K)$ is the set of players that are on “ i 's side of network K ” relative to player j . Figure 1 illustrates β . Note that, for each $(i, j) \in K$, the sets $\beta(i, j, K)$ and $\beta(j, i, K)$ partition N .

Also, for each player i , define $K^i \equiv \{j \mid (i, j) \in K\}$ as player i 's *active contracting partners*, which is the set of players with which player i is supposed to establish non-null contracts. Let

$$\gamma \equiv \max_{\substack{a, a' \in A \\ i \in N}} [u_i(a) - u_i(a')],$$

let \hat{K} denote the set of extremal players (for whom $\underline{\alpha}_i(a_i^*) < 1$ by the definition of the critical-extremes class), and let

$$\omega \equiv \max_{i \in \hat{K}} \underline{\alpha}_i(a_i^*).$$

The number γ is the maximum payoff difference for the players in the underlying game and ω is an upper bound on the probability that, in the Nash equilibrium $\underline{\alpha}$ of the underlying game, any given extremal player selects his part of the efficient profile. Let $p^* = (p^{2*}, p^{3*}, \dots, p^{R*})$ be any vector satisfying $p^{2*} > \gamma$ and $p^{r*} = p^{2*}(r - 1)$ for all $r = 3, 4, \dots, R$. Let q be any number greater than $Rp^{2*}/(1 - \omega)$.

I'll next describe a set of contracts $\mathcal{M}^* = \{m^{ij}\}_{i < j}$, where m^{ij} denotes the contract for the pair (i, j) . In some places below I refer to the contract m^{ij} without specifying $i < j$; in the case of $i > j$, it is understood that $m^{ij} = m^{ji}$ but we do not include m^{ij} in \mathcal{M}^* to avoid double counting in the sum M . The contracts are specified to have the following properties:

- For each $(i, j) \notin K$, the contract is null ($m^{ij} = \underline{m}$).
- At the efficient action profile, the payoff vector with transfers exceeds the Nash equilibrium payoff vector: $u(a^*) + M(a^*) \geq u(\underline{\alpha})$.
- For each $(i, j) \in K$, the contract between i and j is an *assurance contract*, where

$$m_i^{ij}(a) = -m_j^{ij}(a) = m_i^{ij}(a^*) + q[\#\{k \in \beta(j, i, K) \mid a_k \neq a_k^*\} - \#\{k \in \beta(i, j, K) \mid a_k \neq a_k^*\}]$$

for all $a \in A$.

The second property simply requires finding transfers to satisfy the inequalities, and these obviously exist given that K is connected. The third property is achieved by construction. Note that $\#\{k \in \beta(j, i, K) \mid a_k \neq a_k^*\}$ is the number of players on j 's side of the subnetwork whose actions deviate from a^* . Likewise, $\#\{k \in \beta(i, j, K) \mid a_k \neq a_k^*\}$ is the number of players on i 's side of the subnetwork whose actions deviate from a^* .

The contract is null for any pair $(i, j) \notin L$, since this implies $(i, j) \notin K$. Thus, pairs who cannot contract have the null contract as required. Also, in general there will be pairs who can feasibly contract but who essentially do not because they coordinate on the null contract. These are pairs $(i, j) \in L \setminus K$.

Under the set of contracts just described, if a^* is played then the players get a payoff vector that exceeds that of the Nash equilibrium in the underlying game. In the contract for the pair $(i, j) \in K$, player i assures player j that all of the players on player i 's side of network K —that is, those in $\beta(i, j, K)$, including player i —will select their part of the efficient action profile. If any player in $\beta(i, j, K)$ were to deviate from $a_{\beta(i, j, K)}^*$ then player i is required to pay q to player j . This penalty is multiplied by the number of players on i 's side of the network (including i) who deviate. Because q exceeds the maximum payoff

difference in the underlying game, clearly a^* is a Nash equilibrium of $\langle A, u + M \rangle$, so the set of contracts supports profile a^* .

Let us assume that $\overline{\mathcal{M}}$ contains all of the contracts in \mathcal{M}^* and also the null contract, and let us assume that \overline{Q} contains the numbers $0, p^{2*}, p^{3*}, \dots, p^{R*}$, and q . Then the specified contracts are all feasible for the players, along with the null contract.

In the sequential equilibrium that I shall construct, on the equilibrium path the players behave as follows.

Prescribed path: For every $(i, j) \in K$, in round 1 of the contracting phase player i sends message (m^{ij}, m^{ij}, p^*) to player j . That is, player i suggests contract m^{ij} for both values of ϕ . For every $(i, j) \notin K$, in round 1 of the contracting phase player i sends the null message to player j . In rounds $2, 3, \dots, R$, each player i sends the null message to everyone else. In the production phase, the players select a^* .

For $(i, j) \in K$, let us call (m^{ij}, m^{ij}, p^*) the *prescribed first-round message from player i to player j* . Note that, in the prescribed path, players linked in K are supposed to make provisional arrangements that would lead to the set of contracts \mathcal{M}^* . The other contracting pairs are supposed to make no provisional arrangements and thus get the null contract.

Illustration: Example 2

To get a feel for the equilibrium construction to follow, consider Example 2 from the Introduction. Recall that players 1-4 are connected in sequence and that they all have actions in the underlying game. The efficient action profile is (b, b, b, b) and the Nash equilibrium of the underlying game is (c, c, c, c) . Suppose that the contracting phase has three rounds, which is all that is required for underlying games with four players. Provisional arrangements are made in round 1 and cancellations can occur in rounds 2 and 3.

Let the contracting partners coordinate on assurance contracts with an assurance penalty of $q = 14$ and cancellation penalties $p = (6, 12)$. For instance, players 2 and 3 coordinate on $c^0 = c^1 = m^{23}$ with $m^{23}(b, b, b, b) = (0, 0, 0, 0)$, $m^{23}(c, b, b, b) = (0, -14, 14, 0)$, $m^{23}(c, c, b, b) = (0, -28, 28, 0)$, \dots , $m^{23}(c, c, c, c) = (0, 0, 0, 0)$.

The specified contracts force all of the players to select b in the underlying game. Also, no player would benefit from being the first to cancel a prescribed provisional arrangement, because without any cancellations each player gets the payoff 1 and the cancellation penalty outweighs any gain in the underlying game. Further, players prefer not to deviate in the first round, assuming that it would lead to a wave of cancellations and then play of (c, c, c, c) in the underlying game.

For instance, suppose that player 4 deviates in the first round, so that no provisional arrangement is made with player 3 and the contract between them will be null. In the equilibrium construction, player 3 then believes that player 4 will select c in the underlying game, so player 3 will have to pay player 2 an assurance penalty of 14 unless he cancels his contract with player 2. The cancellation costs only 6, so player 3 cancels in round 2. Note

also that player 3 prefers not to wait until round 3 because the cancellation penalty would then double.¹⁸ Player 3's cancellation leads player 2 to conclude that player 4 deviated in the first round (even though player 2 did not observe the deviation) and player 2 then expects that players 3 and 4 will choose c in the underlying game. Player 2 optimally cancels with player 1 in the third round rather than be liable for assurance penalties. In the end, the players all expect each other to choose c , every player indeed selects c , and player 4's deviation only harms him.¹⁹

Key information sets and specified actions

Let us return to the general equilibrium construction. The core elements of the construction pertain to a collection of *key information sets* (*personal histories*) for each player. To describe these information sets, some additional terminology will be helpful. For $(i, j) \in K$, define $\delta(i, j, K)$ to be the length of the largest path between player i and players in the set $\beta(i, j, K)$, with $\delta(i, j, K) \equiv 0$ in the case of $\beta(i, j, K) = \{i\}$.

For any player i and any round $r \in \{1, 2, 3, \dots, R\}$, let us say that player i 's personal history through round r “*essentially conforms to the prescribed path*” if the following two conditions hold. First, for each $j \in K^i$, players i and j exchanged the prescribed first-round messages in round 1 and, if $r \geq 2$, these players sent the null message to each other in rounds $2, 3, \dots, r$. Second, for each $j \notin K^i$, players i and j made no provisional arrangement in round 1. Note that $j \notin K^i$ means the players were supposed to coordinate on the null message in round 1, leading to no provisional arrangement, but the condition allows for the players to have selected different messages in round 1, which also implies no provisional arrangement (so one or both players deviated in the first round in an inessential way).

For any player i and an active contracting partner $j \in K^i$, let us say that “*player j nullified the prescribed provisional arrangement with player i* ” if, in the first round of the contracting phase, player i sent the prescribed first-round message (m^{ij}, m^{ij}, p^*) to player j but player j did not send the prescribed first-round message to player i , and so the pair made no provisional arrangement. For any given personal history for player i through round $r \in \{1, 2, 3, \dots, R\}$, let us say that “*player j triggered vulnerability for player i in round r* ” if the following conditions hold for exactly one player $j \in K^i$. In the case of $r = 1$, player i 's interaction with the others essentially conformed to the prescribed path except that player j nullified the prescribed provisional arrangement with player i . In the case of $r > 1$, player i 's interaction through round r essentially conformed to the prescribed path except that, in round r , player j sent the message “cancel” to player i while player i sent the null message to player j .

Below is a list of the key information sets for player i (labelled H1 i , H2 i , and so on) along with the actions specified at these information sets (labelled S1 i , S2 i , and so on).

¹⁸The cancellation penalty rises from one round to the next by more than a player could gain by changing the action profile played in the underlying game, so the only rational move is to cancel at the first opportunity.

¹⁹In the general construction, the off-equilibrium-path beliefs are critical to support the wave of cancellations and subsequent play of the Nash equilibrium in the underlying game.

The first three entries refer to contingencies that will be on the equilibrium path and also those that deviate in only inessential ways.

- H1*i* The null/empty history at the beginning of round 1 of the contracting phase.
- S1*i* Player i sends the prescribed first-round messages to each of the players. That is, for each $j \in K^i$, player i sends message (m^{ij}, m^{ij}, p^*) to player j . For each $j \notin K^i$, player i sends the null message.
- H2*i* For each $r \in \{1, 2, \dots, R-1\}$, the personal histories through round r that essentially conform to the prescribed path.
- S2*i* Player i sends the null message to all other players in round $r + 1$.
- H3*i* The personal histories through round R of the contracting phase that essentially conform to the prescribed path.
- S3*i* Player i selects action a_i^* in the production phase.

Next we have information sets in which a player triggered vulnerability with player i .

- H4*i* For each $r \in \{1, 2, \dots, R-1\}$ and $j \in K^i$ satisfying $\delta(j, i, K) + 1 \geq r$, the personal histories through round r in which player j has triggered vulnerability for player i in round r .
- S4*i* Player i declares “cancel” to every contracting partner with whom the action “cancel” is available.
- H5*i* The personal histories through round R such that for some player $j \in K^i$ and an integer $r \in \{1, 2, \dots, R\}$ satisfying $\delta(j, i, K) + 1 \geq r$, the following conditions hold: In the sub-history through round r , player j triggered vulnerability for player i in round r . In the case of $r < R$, in round $r + 1$ player i sent the message “cancel” to every contracting partner with whom the action “cancel” was available, while the other players sent the null message to player i .
- S5*i* Player i selects action α_i in the production phase.

The inequalities in H4*i* and H5*i* require an explanation. Recall that $\delta(j, i, K)$ is the length of the largest path between player j and players in the set $\beta(j, i, K)$, so $\delta(j, i, K) + 1$ is the maximum number of rounds required for someone in set $\beta(j, i, K)$ to nullify a prescribed provisional arrangement and then a sequence of cancellations to flow all the way to player i . Because $R \geq n - 1$ and $\delta(i, j, K) + \delta(j, i, K) \leq n - 2$, the inequality $\delta(j, i, K) + 1 \geq r$ implies that $\delta(i, j, K) \leq R - r$, and so enough rounds remain for the sequence of cancellations to flow from player i to all of the players in $\beta(i, j, K)$ (that is, on player i 's side of the network).

Finally, we have an information set at the production phase for the case in which player i nullified all of her prescribed provisional arrangements in round 1.

- H6 i The personal histories through round R such that every player $j \in K^i$ sent the prescribed first-round message to player i , but player i formed no provisional arrangements. Thus, in particular, player i nullified all of her prescribed provisional arrangements. Note that player i then had no choices to make in rounds 2 through R .
- S6 i Player i selects action $\underline{\alpha}_i$ in the production phase.

Note that at all of the information sets described above, player i ignores any inessential deviations in round 1.

Partial appraisals at key information sets

I next summarize player i 's beliefs at all of the information sets addressed so far, about each player j 's behavior at information sets H1 j -H6 j , for every $j \neq i$. The belief at H1 i is denoted B1 i , the belief at H2 i is denoted B2 i , and so on. These are not the full appraisals at information sets H1 j -H6 j , but what is described here is sufficient for checking that the specified behavior S1 j -S6 j is sequentially rational. The full appraisals are constructed from a sequence of fully mixed behavior strategies, as required for a sequential equilibrium, and the construction will ensure that the partial appraisals described here are implied.

A theme will be that, absent a direct observation to the contrary, player i believes that every other player j will or has behaved as specified in S1 j -S6 j above. Note that player i could not observe a deviation in the messages exchanged between two other players. The only way that player i can detect a deviation is if, for some $j \in L^i$, player j deviates from the prescribed message to player i in round 1, or in the case of $j \in K^i$, player j sends the "cancel" message to player i in a later round.

- B1 i At the beginning of the game (for personal history H1 i), player i believes that all of the other players will behave according to strategy specifications S1-S6.
- B2 i -B3 i For information sets in H2 i and H3 i , player i continues to believe that all of the other players behave according to strategy specifications S1-S6, except for any observed inessential deviations from S1 between player i and players in the set $L^i \setminus K^i$.
- B4 i For information sets in H4 i , where some player $j \in K^i$ triggered vulnerability for player i in some round $r \leq \delta(j, i, K)$, player i believes the following: A player $k \in \beta(j, i, K)$ of distance $r - 1$ from player j in network K nullified all of his prescribed provisional arrangements in round 1 but otherwise behaves according to S2 k -S6 k . All other players behave according to S1-S6, except for any inessential deviations in round 1 that player i observed. Thus, player i believes that the players in $\beta(j, i, K)$ will select action profile $\underline{\alpha}_{\beta(j, i, K)}$ in the production phase.
- B5 i For information sets in H5 i , where some player $j \in K^i$ triggered vulnerability for player i in some round $r \leq \delta(j, i, K)$, and player i cancelled all of her remaining provisional arrangements in round $r + 1$, player i continues to believe what is described in B4 i . Thus, player i believes that the other players will select action profile $\underline{\alpha}_{-i}$ in the production phase.

B6i For information sets in H6i, player i believes that all of the other players behave according to strategy specifications S1-S6, except for any inessential deviations in round 1 that player i observed, and so they will select action profile $\underline{\alpha}_{-i}$ in the production phase.

Note why, at information sets H4i-H6i, player i believes that other players will select their part of the underlying-game Nash equilibrium $\underline{\alpha}$. As specified by S4, each player will extend a wave of cancellations that starts with a player nullifying prescribed provisional arrangements in the first round because, along the sequence, the players encounter nullifications early enough to satisfy the δ inequality. Each player in the sequence then reaches the production phase at an information set in H5 or H6.

Rationality at key information sets

Remember that, because the key information sets are a proper subset of all information sets for a player, specifications S1i-S6i only partially define player i 's strategy. Likewise, B1i-B6i only partially define player i 's appraisals at these information sets because, in particular, we have not specified what player i thinks the other players will do at information sets outside the set H1-H6. Nonetheless, we can easily verify that S1i-S6i are optimal actions for player i , regardless of what behavior is specified for player i at information sets outside the set S1i-S6i and regardless of what player i believes the other players will do outside of the set H1-H6.

To see this, first note that if the players behave as specified by S1-S6 then, by construction of the prescribed contracts, player i 's payoff is $u_i(a^*) + M_i(a^*)$, which exceeds $u_i(\underline{\alpha})$. At information sets in H3i, where play essentially conformed to the prescribed path through round R of the contracting phase, player i believes that contracts \mathcal{M}^* were formed and the other players will select a_{-i}^* , to which a_i^* is a best response given player i 's induced payoff function $u_i + M_i$. Here, recall that M_i incorporates assurance penalties, which player i must pay if she herself deviates from a_i^* . Thus, S3i is optimal.

Next consider information sets in H2i. If player i behaves as prescribed, then she will eventually obtain the payoff $u_i(a^*) + M_i(a^*) \geq u_i(\underline{\alpha})$. If she deviates by cancelling her provisional arrangement with one or more players then, whatever happens later in the game, her payoff must be less than $u_i(\underline{\alpha})$. This is because the maximum that player i could gain by altering play in the underlying game is γ , which is more than offset by payment of just one cancellation penalty. Further, since there are no loops in network K , player i 's cancellation with a player will not lead another player to eventually cancel with player i , so player i will not receive any cancellation penalties. For the same reason, player i will not receive any assurance penalties because, for each player $k \in K^i$ with whom player i retains the provisional arrangement, all of the players in $\beta(k, i, K)$ will play their part of a^* in the production phase.

At H1i, which is the beginning of the game, inessential deviations with players in the set $L^i \setminus K^i$ do not affect player i 's payoff. Nullifying the prescribed provisional arrangement with some player $j \in K^i$ (by not sending the prescribed round-1 message) constitutes

a more significant deviation. Given S4, such a deviation would lead to a wave of cancellations through $\beta(j, i, K)$ because the δ condition will hold in each pair of players along the sequence due to $R \geq n - 1$. From S5, each player $k \in \beta(j, i, K)$ will select $\underline{\alpha}_k$ in the production phase. The defining property of the critical extremes class then implies that $\underline{\alpha}_{\beta(j, i, K)}(a_{\beta(j, i, K)}^*) \leq \omega < 1$. So, with a probability of at least $1 - \omega$, player i would be on the hook for an assurance penalty with any contracting partner whose provisional arrangement was made and not cancelled.

In particular, suppose that player i nullifies with player $j \in K^i$ but forms the prescribed provisional arrangement with another player $k \in K^i \setminus \{j\}$. Player i believes that player k will not cancel his provisional arrangements with her. If player i does not cancel with player k , then player i expects the players in the set $\beta(k, i, K)$ to select $a_{\beta(k, i, K)}^*$ in the production phase, and so player i will not receive any assurance penalty from player k but with a probability of at least $1 - \omega$ will have to pay an assurance penalty of at least q to player k .

Because $q(1 - \omega) > \gamma$, the assurance penalty outweighs any gain in the underlying game. Thus, if player i were to nullify a prescribed provisional arrangement with one player j , then player i should do so with all of the players in K^i . This would lead everyone else to play $\underline{\alpha}_{-i}$ in the underlying game and player i optimally selects $\underline{\alpha}_i$ in response. In summary, by deviating in round 1 other than in an inessential way, player i 's resulting payoff cannot exceed $u_i(\underline{\alpha})$, so we know that S1*i* is optimal.

For information sets in H4*i*, a single player $j \in K^i$ has triggered vulnerability for player i . Player i believes that a wave of cancellations will occur through $\beta(j, i, K)$ by the end of the contracting phase and that these players will select $\underline{\alpha}_{\beta(j, i, K)}$ in the production phase. For each player in the set $k \in K^i \setminus \{j\}$, interaction between players i and k has occurred as prescribed, so player i believes that player k will not cancel his provisional arrangements with her. If player i does not cancel with player k , then player i expects the players in the set $\beta(k, i, K)$ to select $a_{\beta(k, i, K)}^*$ in the production phase. In this case, with a probability of at least $1 - \omega$ player i will have to pay an assurance penalty of at least q to player k , and player i will not receive any assurance penalty from him. We have selected the penalties so that $q(1 - \omega) > p^{r*} + \gamma$ for all $r \in \{2, 3, \dots, R\}$, so the cancellation penalty is lower than the expected assurance penalty and the difference exceeds any possible payoff gain in the underlying game, so it is optimal for player i to cancel her outstanding provisional arrangement with each $k \in K^i \setminus \{j\}$. Further, she strictly prefers to cancel in the current round rather than delay, because, round by round, the cancellation penalty rises by more than the maximum gain in the underlying game. We conclude that S4*i* is optimal.

Likewise, S5*i* and S6*i* are optimal because at information sets H5*i* and H6*i* player i believes that the other players will select $\underline{\alpha}_{-i}$ in the underlying game, and $\underline{\alpha}_i$ is a best response.

Remainder of the Equilibrium Construction

To complete the equilibrium construction, we must specify full appraisals and behavior for all information sets of the game, characterized by the limit of a suitably defined sequence of fully mixed behavior strategies. The first step is to specify a sequence of fully mixed actions at the information sets H1-H6. This will converge to the profile of actions S1-S6 and will imply, by the limit of conditional probabilities, the beliefs B1-B6. Rationality at H1-H6 is thus already shown. The second step is to establish that there is a convergent sequence of mixed actions at all of the other information sets and that the limit profile is rational with respect to the implied beliefs. This involves a familiar construction along the lines of what is used to prove the existence of a trembling-hand perfect equilibrium.

Here are the details of the first step. Let ξ be a small positive number. Consider information set H1*i* for any player *i*. At this information set, let us assign probability ξ to each action that, for a single player $j \in L^i \setminus K^i$, has player *i* sending a message other than the null message to player *j* and has player *i* sending the prescribed message (given by S1*i*) to all of the other players. Further, assign probability ξ to each action of player *i* that, for a single player $j \in K^i$, sends a message other than the prescribed message to player *j* (nullifying the prescribed provisional arrangement if player *j* sends the prescribed message) and sends the null message to every player $k \neq j$. Assign probability ξ^2 to all other actions except the prescribed action S1*i*, and assign the remaining probability mass to action S1*i*. At each information set in H2*i*-H6*i*, assign probability ξ^2 to every feasible action other than that prescribed by S2*i*-S6*i*; put the remaining probability on the action prescribed by S2*i*-S6*i*.

Let H denote the set of information sets in the entire game, let \tilde{H} denote the information sets H1-H6, let \tilde{H}_i denote the information sets in \tilde{H} that belong to player *i*, and let $\hat{H} = H \setminus \tilde{H}$. Let $\sigma_{\tilde{H}}(\xi)$ denote the profile of mixed actions defined in the previous paragraph, which gives part of a behavior strategy for the entire game (the actions for information sets in \tilde{H}). For any player *i* and each information set $h \in \tilde{H}_i$, as $\xi \rightarrow 0$ the probability distribution conditional on *h* regarding the behavior at information sets \tilde{H}_j for every player $j \neq i$ converge to the beliefs B1*i*-B6*i*. Likewise, $\sigma_{\tilde{H}}(\xi)$ converges to what is specified by S1-S6, which we can denote $\bar{\sigma}_{\tilde{H}}$. With rationality at these information sets already verified, we thus have for information sets \hat{H} the conditions necessary for a sequential equilibrium.

Proceeding to the second step, consider an artificial game played just on information sets \hat{H} , with behavior at the information sets \tilde{H} exogenously given by $\sigma_{\tilde{H}}(\xi)$, for a given ξ . Assume further that at every information set in this artificial game, the players are restricted to put probability of at least ξ on every action available. Payoffs for the artificial game are specified as in the entire game. Call this artificial game Γ^ξ . Because the set of behavior strategies in the artificial game is a compact and convex subset of a Euclidean space and payoffs are linear in the action probabilities, Γ^ξ has a Nash equilibrium $\sigma_{\hat{H}}(\xi)$. By compactness, we can then take a strictly positive sequence $\{\xi_k\}$ converging to zero such that $\sigma_{\hat{H}}(\xi_k)$ converges to some profile $\bar{\sigma}_{\hat{H}}$.

Let us write $\sigma^k = (\sigma_{\tilde{H}}(\xi_k), \sigma_{\hat{H}}(\xi_k))$ and $\bar{\sigma} = (\bar{\sigma}_{\tilde{H}}, \bar{\sigma}_{\hat{H}})$. Clearly σ^k converges to $\bar{\sigma}$. Further, since σ^k is a fully mixed behavior strategy for all *k*, the conditional-probability

formula applies at all information sets. So, for every information set $h \in H$, we have an implied mixed strategy conditional on reaching h . We can find a subsequence of $\{\sigma^k\}$ such that, for every information set h , the distribution conditional on h converges to some mixed strategy π^h . Note that π^h is an appraisal that gives player i 's belief and behavior at information set h . For $h \in \hat{H}$, by the construction of $\{\sigma^k\}$ and continuity of payoffs in action probabilities, π^h puts positive probability only on actions at h that are optimal in response to the mixed strategy profile for the other players. That is, π^h is rational at h . For $h \in \tilde{H}$, rationality was confirmed in the preceding subsections. Thus, the appraisal system $\{\pi^h\}_{h \in H}$ is sequentially rational and fully consistent by construction, and so it is a sequential equilibrium. By construction, it specifies the behavior and beliefs described in S1-S6 and B1-B6 at information sets H1-H6.

Note that, in this equilibrium construction, the set of feasible contracts $\overline{\mathcal{M}}$ was assumed to contain \mathcal{M}^* and also the null contract, and the set of feasible penalties $\overline{\mathcal{Q}}$ was assumed to contain the numbers $0, p^{2*}, p^{3*}, \dots, p^{R*}$, and q . Otherwise, $\overline{\mathcal{M}}$ and $\overline{\mathcal{Q}}$ were unconstrained except to be finite. So we can let $\mathcal{M}' \equiv \mathcal{M}^* \cup \{\underline{m}\}$ and $\mathcal{Q}' \equiv \{0, p^{2*}, p^{3*}, \dots, p^{R*}, q\}$, and the proof of Proposition 1 is complete.

4.3 Expanding and Completing the Analysis

A similar equilibrium construction can be done for underlying games and networks outside the critical-extremes class, but two additional elements are required. One is to deal with cases in which there are players at the periphery of a network who don't have actions in the underlying game but who benefit from play of the efficient action profile. The construction here requires "pay-in" rather than assurance contracts for peripheral players, but otherwise has the same elements as before and achieves implementation with certainty.

The second new element is required to take care of cases where, for some player i , $\alpha_i(a_i^*) = 1$ (the Nash equilibrium of the underlying game has player i selecting her part of the efficient profile for sure) and yet a_i^* is not a best response to a_{-i}^* in the underlying game. In this case, other players cannot use player i 's reversion to the Nash profile as a threat. Here we can take advantage of the public randomization device to create a suitable threat, and implementation is in the ε -efficient sense.

Illustration: Example 1

To get the flavor of the first new element (dealing with peripheral players), consider Example 1 from the Introduction. Recall that players 1, 2, and 3 are connected in a line and that only player 1 has an action in the underlying game. Players 2 and 3 are peripheral. Suppose that the contracting phase has two rounds (all that is required for underlying games with three players), so there is just one opportunity to cancel. Let players 1 and 2 coordinate on $c^0 = c^1 = m^{12}$ with $m^{12}(b) = (2, -2, 0)$ and $m^{12}(c) = (0, 0, 0)$, and $p = (0)$. Further, let players 2 and 3 coordinate on $c^0 = c^1 = m^{23}$ with $m^{23}(b) = (0, 3, -3,)$ and $m^{23}(c) = (0, 0, 0)$, and $p = (5)$. In words, players 1 and 2 make a provisional arrangement specifying that player 2 pays 2 if and only if player 1 chooses b, and there is no cancellation penalty.

Players 2 and 3 make a provisional arrangement specifying that player 3 pays 3 if and only if player 1 chooses b, and the cancellation penalty is 5.

With these provisional arrangements, the implied contracts give player 1 the incentive to choose b and the payoff will be (2, 1, 1). Neither player 1 nor player 2 would benefit by cancelling their provisional arrangement in the second round of the contracting phase, because then player 1 would select c and the payoff vector would be (1, 0, 0). For the same reason, these players would not want to deviate in the first round. Neither player 2 nor player 3 would benefit by cancelling unilaterally in the second round of the contracting phase, due to the cancellation penalty. Further, if player 3 were to nullify the specified arrangement with player 2 in the first round, it would give player 2 the incentive to cancel with player 1 in the second round, to player 3's disadvantage. Therefore, the players have the incentive to make and keep the prescribed provisional arrangements.

Arbitrary Underlying Games

Returning to the analysis of arbitrary underlying games and networks, we have the following version of Proposition 1:

Proposition 2: *Fix $\varepsilon > 0$, n , and \mathcal{A} , and consider any network L and any n -player underlying game $\langle A, u \rangle$ with $A \subset \mathcal{A}$. There exists a finite set of contracts \mathcal{M}' and a finite set of penalties $Q' \subset \mathbb{R}$ such that the following holds for all finite sets of feasible contracts $\overline{\mathcal{M}}$ and penalties \overline{Q} such that $\mathcal{M}' \subset \overline{\mathcal{M}}$ and $Q' \subset \overline{Q}$. Let the contracting institution be the PCA contracting institution with parameters $\overline{\mathcal{M}}$, \overline{Q} , and ε . There is a sequential equilibrium of the entire game in which the equilibrium outcome is ε -efficient.*

This proposition is proved in Appendix A.1.

Final Step to Complete the Proof of the Theorem

With Proposition 2, one more step proves the Theorem. Fix $\varepsilon > 0$, n , and a finite set G of underlying games. For each network L and underlying game $\langle A, u \rangle \in G$, the proposition identifies finite sets \mathcal{M}' and Q' . Let $\overline{\mathcal{M}}$ be the union of the sets \mathcal{M}' over all of the networks and underlying games, and let \overline{Q} be the corresponding union of the sets Q' . The sets $\overline{\mathcal{M}}$ and \overline{Q} are finite because G is finite and there are a finite number of networks relating the player set N . Proposition 2 then applies with these particular sets, for all networks and underlying games in G . Thus, the PCA contracting institution with parameters $\overline{\mathcal{M}}$, \overline{Q} , and ε implements ε -efficient outcomes.

5 Elaboration and Variations

This section provides some results that elaborate on the need for multiple contracting rounds and the achievable equilibrium values. These results are straightforward extensions of the analysis in the previous section and so are presented without formal proofs. I then discuss how one can reinterpret the contracting institution as a structure that facilitates option contracts, and I add a few words about required penalties. Finally, I provide an example to show that the positive message of the Theorem does not extend to settings with partial verifiability of productive actions.

5.1 Additional Results

Recall that a critical feature of the analysis is the sequential nature of contracting. The main result requires at least $n - 1$ rounds in the contracting phase. In fact, with unrestricted networks one cannot get away with fewer rounds. For instance, if we have a linear network as in Examples 1 and 2, it takes $n - 1$ rounds of communication for a disruption caused by a player at one end to transit (in whatever form of contract adjustments) to the other end. With fewer than $n - 1$ rounds in the contracting phase, we can construct counterexamples along the lines of what appears in the Introduction, which can be put in the form of the following result.

Result 1: *For any given $n \geq 3$, there exists an underlying game $\langle A, u \rangle$, a network L , and a value $\underline{\varepsilon} > 0$ such that, for every natural contracting institution with fewer than $n - 1$ contracting rounds, no sequential equilibrium of the entire game has a^* played with a probability greater than $1 - \underline{\varepsilon}$.*

If we restrict attention to networks with bounded diameter, then a shorter contracting phase will suffice. Let us say that *an institution implements ε -efficient outcomes on a class of networks (or on a class of underlying games)* if the implementation conditions hold for this class.

Result 2: *Consider the class of networks of diameter D or smaller. There exists a contracting institution with $D - 1$ rounds of contracting that implements ε -efficient outcomes.*

Next consider another prominent feature of the analysis: The need to restrict the players' ability to back out of contractual arrangements toward the end of the contracting phase. The PCA contracting institution restricts the players by enforcing cancellation penalties. Holding aside whether this is the most realistic form of restricting the players, it is clear that some sort of restriction is essential. Let us say that the *contracting institution allows player i to freely cancel contracts at R* if for every other player j and every sequence of messages $(h_{ij}^{r-1}, h_{ji}^{r-1})$, there is a message $\lambda^0 \in \Lambda_{ij}^R(h_{ij}^{r-1}, h_{ji}^{r-1})$ such that

$$\mu(i, j, (z_{ij}^{R-1} \lambda^0, z_{ji}^R), \sigma) = \underline{m}.$$

That is, if player i sends message λ^0 to player j at the end of the contracting phase, then their contract is null regardless of the messages they sent earlier.

Result 3: *Consider any contracting institution that allows some player i to freely cancel contracts at R . There exists an underlying game $\langle A, u \rangle$, a network L , and a value $\underline{\varepsilon} > 0$ such that no sequential equilibrium of the entire game has a^* played with a probability greater than $1 - \underline{\varepsilon}$.*

I turn next to folk-theorem considerations. Some models of interactive contracts (for instance, Peters and Szentes 2008) feature many different equilibria with a range of payoffs up to the efficient frontier. The theme extends to the setting here and is most easily expressed as a general version of the Theorem focusing on the critical-extremes class. For a given player set N , let \mathcal{L} denote the set of all connected networks.

Result 4: *Fix n and a finite set of underlying games G . Let Ω be the subset of $G \times \mathcal{L}$ that are in the critical-extremes class. Consider any number $\varepsilon > 0$. There exists a natural contracting institution with the following property. For any $(L, \langle A, u \rangle) \in \Omega$, any Nash equilibrium \underline{a} of $\langle A, u \rangle$, and any feasible payoff vector (incorporating transfers) $v \geq u(\underline{a})$, there is a sequential equilibrium of the entire game that yields a payoff vector that is within ε of v .*

This result is proved using a slight generalization of the steps to prove Proposition 2, where the public randomization device is used to randomize among action profiles in A to, along with transfers, generate an expected payoff near the desired v . The only complication with extending the result to settings outside the critical-extremes class is in how to deal with peripheral players. One can get a similar result holding fixed a group of peripheral players that are needed to pay in, as identified in the Appendix.

5.2 Notes About Option Contracts and Penalties

I next comment on the interpretation of the PCA contracting institution. As defined, the institution has R rounds of contracting and, following the exogenous random draw ϕ , the output of the institution for a pair of players is their “contract” m . A different, perhaps more realistic, interpretation is that the “contracting phase” comprises just the first round $r = 1$ and that the other rounds are simply dates at which the players can communicate along edges of the network.

To formalize this, one might want to expand the PCA institution so that contracts explicitly give the players options; they specify transfers as a function of the sequence of messages between the contracting parties, the random draw σ , and the action profile a . We would assume that messages are locally verifiable in the sense that messages between players i and j are available to the enforcer for evaluating the contract between these two players. To keep with what has been assumed for the PCA institution, we would assume that these messages could not be seen by players other than i and j . Such a formulation


would allow for additional options in the contracts compared to what is allowed in the PCA contracting institution, but the equilibrium construction used to prove the Theorem appears to extend to this setting.

Speaking of the equilibrium construction used in the proof, note that the penalties q and p there were chosen for convenience in the proof. In particular, it was convenient to use values of p and q that are sufficient for all of the contracting pairs, which are large. This is not necessary. One could find workable penalties for each relationship that match with the magnitude of the two players' possible deviation gains in the underlying game. It is not clear whether penalties that real courts would call excessive would be needed. Real courts are, for example, not as sensitive to probabilistic gains (requiring penalties to be scaled up) as the theory requires, but this is an issue that goes beyond the present modeling exercise.

5.3 An Example of Partial Verifiability

I finish this section with an example showing that the Theorem does not extend to settings with partial verifiability of productive actions. The example is a simple case of team production with partial output verification.

Example 3 (partial verifiability—team production): Suppose $n = 3$. Player 1 is the manager and players 2 and 3 are workers. Player 1 has no productive action. The other players have action spaces given by $A_i = \{0, 1\}$ for $i = 2, 3$, where 1 stands for high effort and 0 represents low effort. Payoffs are given by $u_1(a) = 3(a_2 + a_3)$, and for $i = 2, 3$, $u_i(a) = -2$ if $a_i = 1$ and $u_i(a) = 0$ if $a_i = 0$. The network is $L = \{(1, 2), (2, 1), (1, 3), (3, 1)\}$. Partial verifiability of a is given by the partition with these two elements: $\bar{\rho} = \{(1, 1)\}$ and $\underline{\rho} = \{(0, 1), (1, 0), (0, 0)\}$; that is, the enforcer can verify only whether the output $a_2 + a_3$ is 6.



In this example, the workers' cost of high effort is 2 and the manager obtains 3 for each unit of high effort exerted, so the efficient action profile is $(1, 1)$; that is, both workers exert high effort. A contract must be measurable with respect to the partition $\{\bar{\rho}, \underline{\rho}\}$ of A , so we can write the contracted transfers as a function of $\rho \in \{\underline{\rho}, \bar{\rho}\}$. Note that each player can guarantee himself/herself a payoff of at least 0 by declining to contract.

I claim that, regardless of the contracting institution (assumed natural), in every equilibrium of the entire game, the action profile $(0, 0)$ is played with probability 1. The analysis substantiating this claim is in the Appendix, but the basic logic is easy to describe here. Suppose that, for a given contracting institution, we seek to construct an equilibrium in which action profile $(1, 1)$ is played for sure. A key aspect of such an equilibrium is that, in the contract between players 1 and 2, the difference between the transfer to player 2 in the

event of $\bar{\rho}$ and the transfer in the event of $\underline{\rho}$ must be at least 2. Such a margin gives player 2 the incentive to select high effort because, with player 3 choosing high effort, player 2's effort choice determines whether $\bar{\rho}$ or $\underline{\rho}$ will be realized.

But if player 1 refuses to contract with player 3 while behaving with player 1 as the equilibrium dictates, then it would not affect player 2's choice of high effort (player 2 does not observe the deviation and still believes that her effort choice influences whether $\bar{\rho}$ or $\underline{\rho}$ is obtained) and yet $\underline{\rho}$ would be the outcome for sure. The deviation thus gives player 1 a gain of at least 2 in the interaction with player 2, whereas player 1 loses at most 1 in the interaction with player 3.²⁰ The deviation is thus profitable, which means there is no equilibrium in which (1, 1) is played with certainty. The Appendix extends the argument and provides more details.

6 Conclusion

The analysis presented here identifies a natural contracting institution that supports efficient outcomes through bilateral contracting in contractual settings with connected networks and globally verifiable productive actions. Thus, contractual chains arise endogenously and in a decentralized way to internalize all externalities that are due to lack of direct links. The modeling exercise herein points to three possible sources of inefficiency in practice: (1) limited verifiability of productive actions; (2) institutional rules or technologies that limit the feasible space of contracts; (3) suboptimal contracting institutions that, in particular, do not provide players with the opportunity to solidify or abandon contracts in sequence; and (4) coordination problems in equilibrium selection. The third item may also include problems that arise from opportunities for the players to exert bargaining power.

With these conclusions in mind, it may be useful to take a closer look at applications. Linkages across contractual relationships are common in the real world, and we should work to identify the features of real enforcement institutions that make the beneficial linkages possible. Further, for settings in which LDL externalities cannot be internalized through verifiable productive actions, it will be instructive to determine whether a deeper level of verifiability—such as that which would allow for contracts on contracts—can do the trick.

In essence, this paper has followed Hurwicz's (1994) prescription of incorporating "natural" constraints into problems of institutional design, in contrast to the perspective that posits a centralized policymaker who has complete control over the design of the game form in which economic agents will be engaged. Natural constraints include the nature of productive actions (as defined by an underlying game) and limitations on communication channels (as a contractual network may represent). As Jackson and Wilkie (2005) argue, Hurwicz's suggestion must be taken a step further since real mechanisms are not designed by an outsider. Rather, the players themselves determine the mechanism. Depending on

²⁰The surplus generated in the relationship between players 1 and 3 is 1.

the unit of analysis, some design elements are controlled by an “external planner” and others controlled by the players. In the model herein, the contracting institution is an object of external design, and it must obey the physical reality represented by the natural contracting assumptions. The contracts are the player-design element. These come together to determine the induced game between the players.

The model presented here leaves out some institutional constraints—for instance, those having to do with limits on the sophistication of the external enforcer. It would be useful to identify these constraints and examine how the design of the institution can restrict contracting in such a way as to improve the prospects of efficient outcomes.²¹ Limited verifiability, in terms of both the example in Section 5.3 and differences between what is locally and globally verifiable, represent another line of potentially insightful inquiry. It would also be useful to examine special classes of underlying games—for instance ones with “linear externalities” (as in Example 1) as opposed to “complex externalities” (as in Example 2). Finally, for applied work, it would be nice to explore the relation between equilibria of a noncooperative contracting institution (as developed here) and a cooperative-theory stability concept, as Collard-Wexler, Gowrisankaram, and Lee (2017) do in a more specialized environment with different informational assumptions.

A Appendix

This section contains the analysis deferred in Sections 4.3 and 5.3.

A.1 Analysis for Proposition 2

This section explains how to modify the proof of Proposition 1 in order to prove Proposition 2. First, without loss of generality, let us assume that ε is small enough so that, for every underlying game $\langle A, u \rangle \in G$ that does not have an efficient Nash equilibrium, there is a Nash equilibrium $\underline{\alpha}$ and an efficient action profile a^* such that

$$\sum_{i \in N} \left[(1 - \varepsilon) u_i(a^*) + \varepsilon \min_{a \in A} u_i(a) \right] \geq \sum_{i \in N} u_i(\underline{\alpha}). \quad (1)$$

In the equilibrium construction for any given underlying game $\langle A, u \rangle$, we shall utilize a Nash equilibrium profile $\underline{\alpha}$ that satisfies this inequality.

²¹A simple illustration along these lines is given by comparing the results of Jackson and Wilkie (2005) and Ellingsen and Paltseva (2016). One might ask if a legal system should enforce unilateral promises or just contracts. In the two-player setting, Ellingsen and Paltseva’s results suggest that the key is to enforce contracts, and then it does not matter whether promises are also enforced. But suppose promise-making and contracting are costly, and it is cheaper to make a promise than to form a contract. Then, it may be best to enforce only contracts in order to avoid the inefficiencies that arise when players only make strategic promises.

Consider any network L and any underlying game $\langle A, u \rangle \in G$. Let N^* be the set of players who have a nontrivial choice in the underlying game; that is, $i \in N^*$ if and only if A_i contains at least two actions. We can find an action profile $\bar{a} \in A$ such that $\bar{a}_i \neq a_i^*$ for all $i \in N^*$. For any $J \subset N$ and $\alpha \in \Delta A$, let $U_J(\alpha) \equiv \sum_{i \in J} u_i(\alpha)$.

Let $\hat{K} \subset L$ be an arbitrary subnetwork that minimally connects N^* . Some players outside of N^* may be included, as needed to indirectly connect the players in N^* , and these players each has two or more links. Let $N(\hat{K})$ denote the set of players connected by \hat{K} .

I next define a subnetwork of L , called K , which will designate the active contracting pairs who form non-null contracts. If $U_{N(\hat{K})}(a^*) \geq U_{N(\hat{K})}(\underline{\alpha})$ then let $K = \hat{K}$. In this case, the efficient profile a^* generates enough value to the players in $N(\hat{K})$ so that, by making transfers between them, they can all be made better off than if $\underline{\alpha}$ were to be played. If $U_{N(\hat{K})}(a^*) < U_{N(\hat{K})}(\underline{\alpha})$ then more players will have to be included to make the joint value exceed that of $\underline{\alpha}$. In this case, let K be the union of \hat{K} and additional paths in L , now connecting a set $N(K)$ of players, with $N(\hat{K}) \subset N(K)$. This can be done such that (i) K minimally connects $N(K)$ and (ii) K minimally achieves $U_{N(K)}(a^*) \geq U_{N(K)}(\underline{\alpha})$ in the sense that removing an extremal player (who is in $N(K) \setminus N(\hat{K})$ and is linked to just one other player) would reverse this inequality.

A simple algorithm suffices to deliver K . Starting with \hat{K} , we will add a link (i, j) for some player $i \in N(\hat{K})$ and some player $j \in N \setminus N(\hat{K})$ for which $(i, j) \in L$. Since L is connected, there is such a player j and the selection can be arbitrary. If $U_{\hat{K} \cup \{j\}}(a^*) \geq U_{\hat{K} \cup \{j\}}(\underline{\alpha})$ then the algorithm terminates. Otherwise, we continue by adding another player k who is linked via network L with a player in $\hat{K} \cup \{j\}$ but who is not yet included, again checking the joint value inequality. We continue in this way until the joint value inequality holds. The algorithm must reach its goal, because $U_N(a^*) > U_N(\underline{\alpha})$. Once this algorithm terminates, we conduct a paring routine in which any extremal player whose removal would not flip the joint value comparison is removed from the network.²² The result is a network K with the desired properties.

Figure 8 illustrates how the subnetworks \hat{K} and K relate to L . Filled nodes represent players in N^* (players active in the underlying game) and open nodes represent players in $N \setminus N^*$. Note that, unlike in the setting of critical extremes, the network of active contracting pairs K is generally not complete. Thus, there may be players who are supposed to have null contracts with everyone else. Also, there are multiple options for selecting \hat{K} and K ; just one is illustrated in the figure. Let us call the players in $N(K) \setminus N(\hat{K})$ *peripheral*.

Prescribed provisional arrangements and cancellation penalties

In the equilibrium to be constructed, pairs in K will form assurance contracts that specify action profile a^* in the event that $\phi = 1$ and action profile \bar{a} in the event that $\phi = 0$.

²²The paring routine is important because when adding a player, the joint value difference may rise more than it did in a previous round.




Figure 8: Selection of \hat{K} and K .

Recall that $\phi = 1$ occurs with probability $1 - \varepsilon$. As in the critical-extremes case, let ω be the maximal probability that a player $i \in N^*$ who plays \underline{a}_i would not be revealed to have deviated from the specified behavior:

$$\omega \equiv \max_{i \in N^*} \varepsilon \underline{\alpha}_i(\bar{a}_i) + (1 - \varepsilon) \underline{\alpha}_i(a_i^*).$$

Define γ as before:

$$\gamma \equiv \max_{\substack{a, a' \in A \\ i \in N}} [u_i(a) - u_i(a')].$$

The presence of peripheral players (those in $K \setminus \hat{K}$) necessitates a variation of the penalty structure utilized in the critical-extremes case, because peripheral players are not active in the underlying game. For instance, consider player i , j , and k shown in the right part of Figure 8. If player k were to nullify the prescribed provisional arrangement with player j , then player j does not fear paying an assurance penalty to player i because neither player j nor player k is active in the underlying game. In essence, then, the assurance contract between players i and j is a one-way commitment regarding a^* .

However, player j can still be given the incentive to cancel his provisional arrangement with player i in the event that player k nullified in the first round, because the prescribed contract between players j and k specifies a payment from k to j in the event that players in the set $\beta(j, k, K)$ behave as prescribed in the underlying game. Player j loses without the payment from player k , particularly because player j 's provisional arrangement with player i specifies a similar payment. That is, the assurance contracts for the (i, j) and (j, k) relationships are effectively *pay-in* contracts, whereby the peripheral players pay in exchange for action profile a^* . If players i and j coordinate on no cancellation penalty for round 2, then player j will cancel in this round if player k nullified in the first round.

So, in general, the equilibrium construction utilizes cancellation penalties that are sensitive to the distance players are from the “core group” $N(\hat{K})$. For every pair $(i, j) \in K$, let us define $\tau(i, j)$ to be the length of the path in network K between the pair (i, j) and the core group $N(\hat{K})$, defined for the player further away from the core group. If $i \in N(\hat{K})$ and $j \in N(\hat{K})$, then $\tau(i, j) = 0$. Otherwise, without loss of generality suppose $N(\hat{K})$ is on i 's side of the network K , so that $N(\hat{K}) \subset \beta(i, j, K)$; in this case, $\tau(i, j)$ is the number

of edges in K between player j and the closest player in $N(\hat{K})$. For instance, in Figure 8 for the players i, j , and k shown, we have $\tau(i, j) = 1$ and $\tau(j, k) = 2$.

Let $\bar{\tau} \equiv \max_{(i,j) \in K} \tau(i, j)$. Note that $\bar{\tau}$ is the greatest distance between players in $N(K)$ and the subset $N(\hat{K})$. For each pair $(i, j) \in K$, these two players are supposed to coordinate on a vector of cancellation penalties $p_{ij}^* = (p_{ij}^{2*}, p_{ij}^{3*}, \dots, p_{ij}^{R*})$ with the following specifications. First, for any integer $r \geq 2$ such that $r < 2 + \bar{\tau} - \tau(i, j)$, let $p_{ij}^{r*} = 0$. Then for $r \geq 2 + \bar{\tau} - \tau(i, j)$, let $p_{ij}^{r*} = (1 + \varepsilon)\gamma[r - 1 - \bar{\tau} + \tau(i, j)]$. Thus, the cancellation penalties start at zero, and then increase by more than γ for each successive round, ensuring that the difference is always larger than the maximum deviation gain in the underlying game. The positive penalty starts in round 2 for pairs at the end of peripheral paths, such as (j, k) in Figure 8. Pairs in the core have zero cancellation penalties through round $2 + \bar{\tau}$. Pairs in the periphery have cancellation penalties starting in intermediate rounds.

I next describe a set of contracts $\mathcal{M}^* = \{c^{0ij}, c^{1ij}\}_{i < j}$, where c^{0ij} denotes the contract for the pair (i, j) in the event that the random draw is $\phi = 0$ and c^{1ij} is the specified contract in the event that $\phi = 1$. The specifications are similar to what was done for the critical-extremes class, but there are extra conditions to take care of the peripheral relationships. Let q be any number greater than $R(1 + \varepsilon)\gamma/(1 - \omega)$. Also, since the effective contracts and actions in the underlying game will be a function of ϕ , it will be convenient to define $\theta = \sum_{i < j} (1 - \varepsilon)c^{1ij}(a^*) + \varepsilon c^{0ij}(\bar{a})$. We can find contracts with the following properties:

- For each $(i, j) \notin K$, $c^{0ij} = c^{1ij} = \underline{m}$.
- For each $(i, j) \in K$, c^{0ij} and c^{1ij} are assurance contracts with penalty q , where c^{0ij} assures \bar{a} and c^{1ij} assures a^* , and in these cases the transfers are the same. That is, for all $a \in A$,

$$c_i^{1ij}(a) = -c_j^{1ij}(a) = c_i^{1ij}(a^*) + q[\#\{k \in \beta(j, i, K) \mid a_k \neq a_k^*\} - \#\{k \in \beta(i, j, K) \mid a_k \neq a_k^*\}]$$

and

$$c_i^{0ij}(a) = -c_j^{0ij}(a) = c_i^{1ij}(a^*) + q[\#\{k \in \beta(j, i, K) \mid a_k \neq \bar{a}_k\} - \#\{k \in \beta(i, j, K) \mid a_k \neq \bar{a}_k\}],$$

which also implies that $\theta = \sum_{i < j} c^{1ij}(a^*) = \sum_{i < j} c^{0ij}(\bar{a})$.

- The expected payoff vector of the prescribed actions exceeds the Nash equilibrium payoff vector for every player $i \in N(K)$. That is, $(1 - \varepsilon)u_i(a^*) + \varepsilon u_i(\bar{a}) + \theta_i \geq u_i(\underline{\alpha})$ for every $i \in N(K)$.
- For $(i, j) \in K$ with $j \notin N(\hat{K})$ and $\hat{K} \subset \beta(i, j, K)$, so that j is a periphery player and the core group is on i 's side of the network K relative to j , we have

$$(1 - \varepsilon)u_i(a^*) + \varepsilon u_i(\bar{a}) + \theta - c_i^{1ij} < u_i(\underline{\alpha}). \quad (2)$$

That is, removing the prescribed contract for (i, j) would lead to a payoff for player i that is less than $u_i(\underline{\alpha})$.

To construct these contracts, one can start by finding transfers for the specified action profiles that satisfy the third and fourth conditions. Note that the fourth condition requires payments toward the core group for all pairs in the periphery, which can be arranged given that the peripheral players are needed to push the joint expected payoff above that of $\underline{\alpha}$. The second condition can then be provided because it binds only the transfers specified in the event that players deviate from the prescribed action profiles.

In the sequential equilibrium to be constructed, on the equilibrium path the players behave as follows.

Prescribed path: For every $(i, j) \in K$, in round 1 of the contracting phase player i sends message $(c^{0ij}, c^{1ij}, p_{ij}^*)$ to player j , so that contract c^{0ij} is suggested for $\phi = 0$ and c^{1ij} is suggested for $\phi = 1$. For every $(i, j) \notin K$, in round 1 of the contracting phase player i sends the null message to player j . In rounds 2, 3, \dots , R , each player sends the null message to everyone else. If $\phi = 1$ then the players select a^* in the production phase; if $\phi = 0$ then they select \bar{a} .

For $(i, j) \in K$, $(c^{0ij}, c^{1ij}, p_{ij}^*)$ is the *prescribed first-round message from player i to player j* . The other contracting pairs are supposed to make no provisional arrangements and thus get the null contract.

Overview of the equilibrium construction

From this point, the equilibrium construction works exactly as for the critical-extremes class. Players ignore inessential deviations in round 1. If a player j triggers vulnerability for player i in a round that is early enough to allow a wave of cancellations to reach all players, then player i cancels all outstanding contracts in the next round. Note that the only opportunity for an extremal player j in the periphery to nullify or cancel without penalty is by nullifying in round 1. The partner, player i , can freely cancel with his other partners in the next round, which will be optimal because otherwise player i 's expected payoff would fall below $u_i(\underline{\alpha})$ given Condition 2 above. The incentives work the same for other peripheral players, with ones closer to the core group being able to cancel freely for multiple rounds.

By construction, a player $i \in N(K)$ who nullifies or cancels provisional arrangements will get one of the following three outcomes. If the player nullifies all of his prescribed provisional arrangements, or cancels them all before cancellation penalties apply, then it triggers a wave of cancellations leading to play of $\underline{\alpha}$. If the player nullifies/cancels only some of his prescribed provisional arrangements, then he will either get a payoff below $u_i(\underline{\alpha})$ owing to Condition 2 or he will have to pay an assurance penalty that implies a lower payoff. Finally, if the player waits until cancellation penalties apply and then cancels

some provisional arrangements, he pays the cancellation penalties which also imply a lower payoff.

Consider how this works for the case shown in Figure 8. For players in $N(\hat{K})$, the options are exactly as in the critical extremes class, except that these players have no cancellation penalties until round 4. Any nullification or cancellation in the first three rounds within this core group triggers a wave of cancellations of all provisional arrangements for players in $N(\hat{K})$, including with any peripheral players, driven by beliefs to this effect. Therefore, no player in the core group wants to deviate from the prescribed path. Player k , at the end of the periphery chain, can nullify/cancel without penalty only by nullifying the prescribed provisional arrangement with player j in round 1. Player j then prefers to cancel in round 2 (with no penalty) the provisional arrangement forged with player i . Not doing so would result in a payoff below $u_j(\underline{\alpha})$ because player k would not be paying in. In round 3, player i cancels with his two contracting partners in \hat{K} for the same reason. Keeping just one of his provisional arrangements would be worse because it would lead to payment of an assurance penalty.

Constructing the appropriate sequence of fully-mixed behavior strategies to define beliefs and the complete equilibrium strategy profile works the same way as with the critical extremes class. As before, we put probability ε on perturbations in round 1 in which a player deviates in an inessential way with one other player but otherwise behaving as prescribed. Another perturbation, with the same probability, has a player nullifying with all partners. All other perturbations occur with probability ε^2 . For a player who has zero cancellation penalties in some rounds, assume that with probability ε in these rounds the player cancels all of his provisional arrangements. All other perturbations get probability ε^2 . This construction ensures that (i) a player who observes an inessential deviation will believe that the other players are otherwise playing as prescribed, (ii) a player i for whom vulnerability has been triggered will believe that it was started by a player who nullified or cancelled all of her prescribed provisional arrangements, triggering a wave of cancellations and play of $\underline{\alpha}$ on paths that do not involve player i . Thus, beliefs and behavior are as described for the critical-extremes class, with the additional possibility that a wave of cancellations was started in a round where the cancellation penalty is zero.

A.2 Analysis for the Partial Verifiability Example

For any given equilibrium, let f be the joint distribution of (a_1, a_2) on the equilibrium path. Consider that in some equilibrium contingency at the end of the contracting phase, player 2's contract with player 1 is m^{12} and player 2 is supposed to select high effort with positive probability. Let ζ be the probability that, in this contingency, player 2 thinks player 3 will select high effort. Noting that player 2 receives $m_2^{12}(\bar{\rho})$ if and only if both workers choose high effort, and otherwise player 2 receives $m_2^{12}(\underline{\rho})$, player 2's incentive condition requires $\zeta m_2^{12}(\bar{\rho}) + (1 - \zeta)m_2^{12}(\underline{\rho}) - 2 \geq m_2^{12}(\underline{\rho})$, which simplifies to

$$m_2^{12}(\bar{\rho}) - m_2^{12}(\underline{\rho}) \geq 2/\zeta.$$

That is, player 1 pays to player 2 a bonus of at least $2/\zeta$ from this contingency, in the event that both players 2 and 3 select high effort.

Let us integrate over the equilibrium paths in which both workers select high effort. Using Jensen's inequality with respect to the distribution of ζ , which has mean $\frac{f(1,1)}{[f(1,1)+f(1,0)]}$ over these paths, we find that player 1 pays to player 2 an expected bonus of at least

$$f(1,1) \cdot 2 \frac{f(1,1) + f(1,0)}{f(1,1)} = 2[f(1,1) + f(1,0)].$$

If player 1 were to deviate by refusing to contract with player 3 while still contracting with player 2 as specified by the equilibrium, then player 1 would save this expected bonus without changing player 2's action in the underlying game. There would be an associated loss in player 1's relationship with player 3 of no more than $f(1,1) + f(0,1)$, which is the expected surplus generated by player 3. In equilibrium, player 1 must be dissuaded from deviating and so we must have $2[f(1,1) + f(1,0)] \leq f(1,1) + f(0,1)$, which simplifies to $f(1,1) \leq f(0,1) - 2f(1,0)$. The same steps apply to player 1 considering whether to refuse to contract with player 3, which implies $f(1,1) \leq f(1,0) - 2f(0,1)$.

Summing the last two inequalities, we get $2f(1,1) \leq -f(1,0) - f(0,1)$, which cannot be satisfied if $f(1,1) > 0$, implying that $a = (1,1)$ occurs with zero probability. A further implication is that, if there is an equilibrium contingency in which a worker i is supposed to choose high effort with positive probability, then the other worker is sure to choose low effort and player i 's payment is not sensitive to this player's effort choice, which is contradicts rationality. Thus, workers select low effort for sure in equilibrium.²³

References

- Abreu, D. and A. Sen (1990), "Subgame Perfect Implementation: A Necessary and Almost Sufficient Condition," *Journal of Economic Theory* 50: 285–299.
- Armstrong, M. (2006), "Competition in Two-Sided Markets." *The RAND Journal of Economics* 37 (3): 668–691.
- Bagwell, K., R. W. Staiger, and A. Yurukoglu (2017), " 'Nash-in-Nash' Tariff Bargaining with and without MFN," unpublished manuscript.
- Bergemann, D. and J. Valimaki (2003), "Dynamic Common Agency," *Journal of Economic Theory* 111: 23–48.
- Bernheim, D. and M. Whinston (1986a), "Common Agency," *Econometrica* 54: 923–942.
- Bernheim, D. and M. Whinston (1986b), "Menu Auctions, Resource Allocations and Economic Influence," *Quarterly Journal of Economics* 101: 1–31.
- Bloch, F. and A. Gomes (2006), "Contracting with Externalities and Outside Options," *Journal of Economic Theory* 127 (1): 172–201.

²³Watson (2014) picks up on the general question of what forms of contractual linkages might be needed in settings like this, and it is a good topic for future research.

- Cafeggi, F. (2008), “Contractual Networks and the Small Business Act: Towards European Principles?” *European Review of Contract Law* 4 (4): 493–539.
- Chatterjee, K., B. Dutta, D. Ray, and K. Sengupta (1993), “A Noncooperative Theory of Coalitional Bargaining,” *Review of Economic Studies* 60: 463–477.
- Collard-Wexler, A., G. Gowrisankaram, and R. S. Lee (2017), “ ‘Nash-in-Nash’ Bargaining: A Microfoundation for Applied Work,” *Journal of Political Economy*, forthcoming.
- Cremer, J. and M. H. Riordan (1987), “On Governing Multilateral Transactions With Bilateral Contracts,” *The RAND Journal of Economics* 18 (3): 436–451.
- Danziger, L. and Schnytzer, A. (1991), “Implementing the Lindahl Voluntary-Exchange System,” *European Journal of Political Economy* 7: 55–64.
- De Fontenay, C.C. and J.S. Gans (2014), “Bilateral Bargaining with Externalities,” *The Journal of Industrial Economics* 62 (4): 756–788.
- Economides, N. (1996), “The Economics of Networks,” *International Journal of Industrial Organization* 14 (6): 673–699.
- Edelman, B. and J. Wright (2015), “Price Coherence and Excessive Intermediation,” *The Quarterly Journal of Economics* 130 (3): 1283–1328.
- Ellingsen, T. and E. Paltseva (2016), “Confining the Coase Theorem: Contracting, Ownership, and Free-Riding,” *The Review of Economic Studies* 83 (2): 547–586.
- Elliott, Matthew (2015), “Inefficiencies in Networked Markets,” *American Economic Journal: Microeconomics* 7 (4): 43–82.
- Fleiner, T., Z. Jankó, A. Tamura, and A. Teytelboym (2016), “Trading Networks with Bilateral Contracts,” unpublished manuscript.
- Galasso, A. (2008), “Coordination and Bargaining Power in Contracting with Externalities,” *Journal of Economic Theory* 143 (1): 558–570.
- Gomes, A. (2005), “Multilateral Contracting with Externalities,” *Econometrica* 73 (4): 1329–1350.
- Gomes, A. and Jehiel, P. (2005), “Dynamic Processes of Social and Economic Interactions: On the Persistence of Inefficiencies,” *Journal of Political Economy* 113: 626–667.
- Guttman, J.M. (1978), “Understanding Collective Action: Matching Behavior,” *American Economic Review* 68: 251–255.
- Guttman, J.M. and A. Schnytzer (1992), “A Solution to the Externality Problem Using Strategic Matching,” *Social Choice and Welfare* 9: 73–88.
- Hafalir, I.E. (2007), “Efficiency in Coalition Games with Externalities,” *Games and Economic Behavior* 61: 242–258.
- Hagiu, A. and J. Wright (2015), “Multi-Sided Platforms,” *International Journal of Industrial Organization* 43: 162–174.
- Hardin, G. (1968), “The Tragedy of the Commons,” *Science* 162: 1243–1248.
- Horn, H. and A. Wolinsky (1988), “Bilateral Monopolies and Incentives for Merger,” *The RAND Journal of Economics* 19 (3): 408–419.

- Hurwicz, L. (1994), “Economic Design, Adjustment Processes, Mechanisms, and Institutions,” *Economic Design* 1: 1–14.
- Hyndman, K. and D. Ray (2007), “Coalition Formation with Binding Agreements,” *Review of Economic Studies* 74: 1125–1147.
- Jackson, M.O. and S. Wilkie (2005), “Endogenous Games and Mechanisms: Side Payments Among Players,” *Review of Economic Studies* 72 (2): 543–566.
- Jensen, M.C. and W.H. Meckling (1976), “Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure,” *Journal of Financial Economics* 3 (4): 305–360.
- Katz, M.L. (2006), “Observable Contracts as Commitments: Interdependent Contracts and Moral Hazard,” *Journal of Economics and Management Strategy* 15(3): 685–706.
- Kranton, R., and D. Minehart (2001), “A Theory of Buyer-Seller Networks,” *American Economic Review* 91 (3): 485–508.
- Laffont, J.-J. and D. Martimort (1997), “The Firm as a Multicontract Organization,” *Journal of Economics and Management Strategy* 6 (2): 201–234.
- Lee, R. (2014), “Competing Platforms,” *Journal of Economics and Management Strategy* 23 (3): 507–526.
- Lee, R. S. and K. Fong (2013), “Markov-Perfect Network Formation: An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks,” unpublished manuscript.
- Lloyd, W.F. (1833), *Two Lectures on the Checks to Population*, Oxford: Oxford University Press.
- Martimort, D. (2007), “Multi-Contracting Mechanism Design,” in *Advances in Economic Theory: Proceedings of the World Congress of the Econometric Society*, edited by Blundell, Newey and Persson, Cambridge University Press.
- Matsushima, H. (1988), “A New Approach to the Implementation Problem,” *Journal of Economic Theory* 45: 128–144.
- McAfee, R. P. and M. Schwartz (1994), “Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity,” *The American Economic Review* 84 (1): 210–230.
- Peters, M. and B. Szentes (2012), “Definable and Contractible Contracts,” *Econometrica* 80 (1): 363–411.
- Prat, A. and A. Rustichini (2003), “Games Played Through Agents,” *Econometrica* 71 (4): 989–1026.
- Pycia, M. and M. B. Yenmez (2017), “Matching with Externalities,” unpublished manuscript.
- Reisinger, M. (2014), “Two-Part Tariff Competition between Two-Sided Platforms,” *European Economic Review* 68: 168–180.
- Rochet, J.C. and J. Tirole (2003), “Platform Competition in Two-Sided Markets,” *Journal of the European Economic Association* 1 (4): 990–1029.
- Rochet, J.C. and J. Tirole (2006), “Two-Sided Markets: A Progress Report,” *The RAND Journal of Economics* 37 (3): 645–667.

- Rysman, M. (2009), “The Economics of Two-Sided Markets,” *Journal of Economic Perspectives* 23 (3): 125–143.
- Segal, I. (1999), “Contracting with Externalities,” *Quarterly Journal of Economics* 114 (2): 337–88.
- Segal, I. and M. Whinston (2003), “Robust Predictions for Bilateral Contracting with Externalities,” *Econometrica* 71 (3): 757–791.
- Seidmann, D.J. and E. Winter (1998), “Exploring Gains from Trade in Multilateral Bargaining: A Theory of Gradual Coalition Formation,” *Review of Economic Studies* 65: 793–815.
- Varian, H. (1994b), “A Solution to the Problem of Externalities when Agents are Well-Informed,” *American Economic Review* 84: 1278–1293.
- Yamada, A. (2003), “Efficient Equilibrium Side Contracts,” *Economics Bulletin* 3 (6): 1-7.
- Watson, J. (2014), “Benefits of Contractual Linkages,” unpublished draft.
- Watson, J. (2007), “Contract, Mechanism Design, and Technological Detail,” *Econometrica* 75: 55–81.
- Watson, J. (2017), “A General, Practicable Definition of Perfect Bayesian Equilibrium,” unpublished draft.
- Weyl, G. (2010), “A Price Theory of Multi-Sided Platforms,” *American Economic Review* 100 (September): 1642–1672.