Contractual Chains

Joel Watson*

May 2021†

Abstract

This paper develops a model of private bilateral 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. Of particular interest is whether such decentralized contracting can internalizes externalities that arise due to parties being unable to contract directly with others whose actions affect their payoffs. The paper investigates the prospect of 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/. 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 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) and Yale and the Center for Advanced Study in the Behavioral Sciences at Stanford for hospitality.

†A preliminary version of this paper was produced in 2011 and has been refined in fits and starts since. After many delays for various reasons, the author is now happy with the paper and eager to publish it.

1 Introduction

In many contractual settings there is multilateral productive interaction, contracting occurs before productive actions take place, and 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. These settings often feature what one may call *externalities due to lack of direct links (LDL)*, in which at least one agent is unable to contract directly with another agent whose productive action is an argument in her payoff function. A key issue is whether, through decentralized contracting, suitable contractual linkages between relationships will endogenously form to internalize LDL externalities via a *chain of bilateral contracts*.

In this paper, I develop a noncooperative game-theoretic model to study whether contractual chains can, in general, be expected to internalize LDL externalities. The model focuses on the following environment:

- Parties interact in the *contracting phase* followed by the *production phase*.
- The production phase is described by an n-player *underlying game* $\langle A, u \rangle$ that is commonly known to the players.
- Only bilateral contracting is possible. The pairs of players that can communicate and establish contracts are given exogenously by an undirected network of bilateral links, $L \subset N \times N$, where $N = \{1, 2, \dots, n\}$. Interaction in the contracting phase is private, so contracting between any pair of players is not observe by others.
- All productive actions are verifiable by everyone (*global verifiability*). Monetary transfers between the contracting parties are externally enforced, contracts may not impose monetary transfers on any third parties, and payoffs are linear in money.

Thus, a contract between players i and j is a function $m \colon A \to \mathbb{R}^n_0$ such that $m_k(a) = 0$ for all $k \notin \{i, j\}$ and every $a \in A$, where $\mathbb{R}^n_0$ denotes the vectors in $\mathbb{R}^n$ whose components sum to zero. 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 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. The pair of players (i, 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.

In this environment, 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. For instance, the contract between players i and i' could specify a transfer between them contingent on player k's action in the underlying game. The model rules out "contracts on contracts," such as if the contract between players i and i' could restrict the contract that players k and k' are allowed to create.

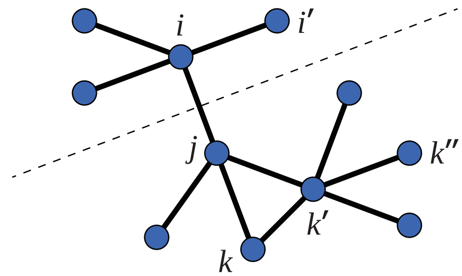

Figure 1: A network of contractual relationships.

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 private and voluntary decentralized contracting: First, each player receives messages from only those to whom she is linked in the network, and she does not observe messages exchanged between other players. Second, the contract formed between players i and j does not depend on the messages sent by, or received by, any other player k. Third, players have the option of rejecting contracts. I call a contracting institution *natural* if it satisfies these assumptions.[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 in the underlying game. In this exercise, the contracting institution is held fixed and we vary the underlying game and network.

Why is the possibility question interesting? First, LDL externalities exist in a plethora of economic settings, and they often traverse extensive networks and occur bidirectionally, so the question has practical significance.[2] Second, the possibility question does not have

[1]The third assumption ensures the definition of contract is conventional, in that it requires the consent of both parties. The second assumption embodies the principle that a contracting pair is free to form whatever contract is desired, uninhibited by others in the society, and it rules out contracts on contracts. One can also motivate the second assumption on the basis of a technological constraint, that messages involved in the formation of contracts are verifiable only locally. When contracting partners appear in court, the judge can observe their contract and the verifiable outcome of the underlying game, but the judge will not readily observe the contracts written in other contractual relationships and it may be prohibitively costly to gather such evidence.

[2]Examples include (i) collaboration agreements between firms on projects that rely on investments by their suppliers; (ii) data-transmission networks, where end users contract with local service providers and content

an obvious answer; addressing it requires a novel noncooperative modeling exercise with a number of subtleties. Third and most important, a positive general result would constitute an extended *Coase Theorem* that can serve as a useful benchmark for analysis of complex contractual settings.

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 in every productive setting. 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 $1-\varepsilon$. Furthermore, $\varepsilon = 0$ for many underlying games of interest.[3] Thus, there is a single contracting institution under which, in all productive settings with global verifiability, LDL externalities can be overcome by decentralized formation of contractual chains.

The proof of the Theorem entails an elaborate construction, but three essential economic elements are easy to describe. First, the contracting institution defined in the proof allows for *sequential contract formation*. Specifically, it facilitates creation 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* with penalties, whereby one player guarantees that specified 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. Third, the players agree on *cancellation penalties* that discourage them from cancelling tentative contracts except in particularly onerous situations. All three elements are realistic.[4]

The model developed herein is purposefully abstract in order to provide a general possibility result that may be a useful reference point for a variety of applications. Further, the model precisely accounts for fundamentals of contracting that may vary across applications,

providers but care also about the actions of "Tier-1" firms that transmit data between them; (iii) 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; (iv) sales of goods exhibiting network externalities, where each consumer cares about the other consumers' use of the seller's technology; (v) 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; and (vi) supply contracting in vertically integrated industries.

[3] 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), and is needed only to deal with a special case of underlying game. 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, which would be 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 (weak implementation is immediate without these constraints).

[4] Options to terminate are common in contracts across industries. Regarding assurance, parties typically bear responsibility for the work of their individual suppliers. Contracting partners sometimes even develop detailed criteria for the practices of each others' employees and suppliers. Such "talent management" is documented in the World Management Survey dataset, as discussed recently by Bernstein and Peterson (2020).

such as the extent of verifiability and the scope of external enforcement, which determine ways in which contracts can link relationships. Individual applications may require additional structure and may also depart from the benchmark setting that I focus on, with implications for efficiency and equilibrium contractual arrangements. For instance, an example presented at the end of this paper shows that inefficiency may be unavoidable if productive actions are only partially verifiable.

Related literature

As noted, the modeling exercise herein generalizes Coase's (1960) famous insight about how externalities can be internalized through contracting, regardless of the assignment of property rights. Coase's logic was put forth informally through a discussion of two-party examples and legal cases. It can be formalized by noting that for two-player settings of complete information, with full verifiability and enforcement, there exists a noncooperative bargaining game that is suitable for modeling contract negotiation and has an efficient equilibrium regardless of the economic parameters.

Ellingsen and Paltseva (2016) prove a Coase-style efficiency result for settings with any number of players. Their model has the same basic structure as mine: players interact in a contracting phase followed by an underlying game with full verifiability. The key difference is that Ellingsen and Paltseva examine centralized multilateral contracting, which for instance allows all of the players to join in a single contract.[5] In contrast, the setting that I analyze constrains the players to bilateral contracting and only with those to whom they are linked in the fixed network. Further, all contracting is private, so there is asymmetric information about the contracts that are formed. Thus, in relation to the Coase Theorem, my analysis explores whether decentralized bilateral contracting suffices to internalize externalities that extend across contractual relationships.

A variety of other papers develop game-theoretic models of multiple contractual relationships that share features with the present exercise; some are fully noncooperative models and the others are in the cooperative-theory tradition. Neither strand has examined the general question posed here regarding internalizing LDL externalities.[6] Using a fully noncooperative model with individual productive actions that fits the framework here, McAfee and Schwartz (1994) study private bilateral contracting between a monopoly supplier and multiple downstream firms. There are LDL externalities because the downstream firms are competitors in a market, although the authors restrict the contract between the supplier and a given down-

[5] In the contracting institution that Ellingsen and Paltseva (2016) examine, players simultaneously make public, multilateral offers and then each player accepts at most one of the contracts offered. Their modeling exercise builds on the model of Jackson and Wilkie (2005), which examines binding unilateral promises rather than contracts.

[6] An advantage of the noncooperative framework is that it allows for a precise categorization of externalities and feasible contractual linkages, on the basis of the enforcement technology and the specification of what is verifiable within and across relationships. A further distinction can be made between models that describe productive actions as taken by individual players and models that treat productive actions as essentially "public" (taken by a third party) and occurring automatically with contract formation. Individual-action modeling is required to understand the full extent to which a player's productive action can serve as an option (Watson 2007), especially as influenced by contracts with multiple partners.

stream firm to condition the transfer only on this downstream firm's orders. Other noncooperative models in the related literature focus on similar applications with specific networks and enforcement mechanisms, most without LDL externalities.[7]

On the cooperative-theory side, some models of bilateral contracting utilize the *Nash-in-Nash* solution, whereby for each relationship, the specified contract maximizes the Nash product holding fixed the contracts in all other relationships. Cremer and Riordan (1987) in this way examine vertical contracting with a single supplier and no LDL externalities. Horn and Wolinsky (1988) allow for LDL externalities but limit attention to linear contracts that condition a transfer from a firm on only the number of units delivered to this firm. Collard-Wexler et al. (2017) provide a result in the tradition of the "Nash program" that relates the Nash-in-Nash solution to an equilibrium of a fully noncooperative model of bargaining in a general public-action setting with no LDL exernalities.

The line on "matching with contracts" initiated by Hatfield and Milgrom (2005) studies stability concepts for models in which the fundamentals are an abstract set of feasible contracts available to each subset of players and payoffs as a function of the contracts chosen. Closest to my modeling exercise is the analysis of Rostek and Yoder (2019, 2020). Their model allows a player's payoff to depend on the contracts formed by others, but these are not LDL externalities because multilateral contracting is permitted. They focus, as other related entries do, on the existence of stable matchings and the characterization of stability conditions, highlighting specific settings such as where contracts are complementary. Additional discussion of this and other areas of the literature, along with notes about the relative advantages of noncooperative modeling, may be found in the Conclusion.

Overview

The next section provides stylized examples that are later used to illustrate key components of the analysis. Section 3 describes the general model. Section 4 presents the Theorem, describes the contracting institution used in the proof, and sketches elements of the equilibrium construction for the examples. The main parts of the proof are developed in Section 5, with a few details deferred to the Appendix. Section 6 discusses variations of the model and provides tangential results on network size, the impact of contracting constraints, the general existence of inefficient equilibria, and the problem of partially verifiable productive actions. The Conclusion offers additional comments, additional references, and notes on further steps in the research program.

[7] Segal's (1999) model of bilateral contracting between a principal and multiple agents effectively has only the principal taking an action in the underlying game, so there are no LDL externalities. See Galasso (2008) for a more recent study that looks at various bargaining protocols and provides additional references. Bernheim and Whinston's (1986a,b) common-agency framework is similar in this regard, as is Prat and Rustichini's (2003) setting of multiple agents.

2 Illustrative examples

This section describes simplistic examples to show that efficiency is hard, if not impossible, to achieve and which will later be used to illustrate components of the proof of the Theorem.

Collaboration agreement

Suppose that two information-technology or pharmaceutical firms, called players 2 and 3, seek to collaborate on an innovative "Project X." Player 2 has an existing relationship with a supplier called player 1, and success of Project X requires a specialized input from player 1. To create this specialize input, player 1 must divert resources from some of its other projects. Likewise, player 3 has an existing relationship with a supplier called player 4, whose specialized input will also be needed for Project X. Further, the suppliers may compete in an unrelated market; focusing their resources on Project X may reduces their ability to perform in this other market.

Imagine that the four parties cannot communicate and contract as a group, due to legal barriers, physical barriers, or transaction costs. Rather, only three bilateral contracts are possible: Player 2 can contract with its supplier, player 1; player 3 can contract with its supplier, player 4; and player 2 can contract with player 3. This contracting network is shown on the left side of Figure 2. There are then LDL externalities between several of the parties. Player 2's profit from the Project X depends on player 4's effort, but player 2 cannot contract directly with player 4. Likewise, player 3's profit depends on player 1's effort, but these firms are unable to contract together. Finally, players 1 and 4 care about each other's productive actions, but this pair cannot contract.

To manage Project X, players 2 and 3 will have to form a collaboration agreement—a contract that governs their interaction and may also contain provisions having to do with player 1's and player 4's productive actions.[8] Additionally, player 2 and player 1 will have to modify their supply agreement, and likewise for players 3 and 4.

Let us consider a case in which only players 1 and 4 have choices to make in the underlying game and they both have action spaces $\{0, 1, 2\}$. The actions of players 2 and 3 are fixed at $a_2 = a_3 = 1$. Let the payoff vector in the underlying game be given by the table on the right side of Figure 2. The efficient action profile is $a^* = (1, 1, 1, 1)$ and the Nash equilibrium of the underlying game is $\underline{a} = (0, 1, 1, 0)$. If a supplier deviates from the efficient action profile, it negatively affects payoffs of all the other players, so externalities extend throughout the network.

To generate intuition and calibrate expectations, consider how contracting may work in this example. Regardless of the contracting institution, efficiency requires transfers in the collaboration agreement between players 2 and 3 to depend on their suppliers' productive

[8] Collaboration agreements are common in the information-technology and pharmaceutical industries, and others. The U.S. Securities and Exchange Commission's Edgar Database of required SEC filings contains numerous collaboration agreements and other documents that reference them, although many details of the agreements are not available. A recent example in the pharmaceutical industry is a research collaboration agreement between Jounce Therapeutics and Celgene to design and test cancer therapies. An example in IT is a agreement between Bsquare and Amazon Web Services to collaborate on technology and standards for the "Internet of Things."

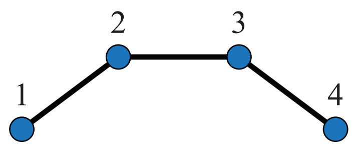

	$a_4 = 1$	$a_4 = 0$	$a_4 = 2$
$a_1 = 1$	2, 6, 6, 2	0, 2, 5, 4	0, 0, 5, 9
$a_1 = 0$	4, 5, 2, 0	4, 2, 2, 4	7, 2, 2, 2
$a_1 = 2$	9, 5, 0, 0	2, 2, 2, 7	2, 2, 2, 2

Figure 2: A collaboration-agreement example.

actions. To see this, suppose for a moment that there is an efficient equilibrium of the grand game in which the contract formed between players 2 and 3 specifies a transfer that does not depend on player 1's action a_1. In order for the equilibrium to be efficient, the contract formed between players 1 and 2 must specify a payment of at least 2 to player 1 conditional on $a_1 = 1$, for player 1 can guarantee a payoff of at least 4 by refusing to contract and by choosing $a_1 = 0$. But it is then apparent that player 2 has a profitable deviation. This firm can decline to contract with player 1 while forming the equilibrium contract with player 3. Player 1 will select either $a_1 = 0$ or $a_1 = 2$ in the underlying game ($a_1 = 1$ is dominated) and player 4 will select $a_4 = 1$ because, having not observed that player 2 deviated in the contracting phase with player 1, players 3 and 4 still believe that they are on the equilibrium path. Player 2's payoff increases by at least 1 when deviating in this way.

Therefore, to obtain an efficient outcome, it is essential for the players to form contracts that condition transfers in a given relationship on productive actions taken outside this relationship. Do the players have incentives to create such contracts in equilibrium and, further, in such a complementary form that would motivate them to choose a^* in the underlying game? A look at some contracting institutions suggests perhaps not.

Consider the two-round contracting institution studied by Ellingsen and Paltseva (2016). In the first round, players simultaneously offer contracts to their linked partners. That is, player 1 sends a contract offer to player 2, player 2 sends separate contract offers to players 1 and 3, and players 3 and 4 do similarly. Each player observes the offers made to her. For a given pair of linked players, the contracts that they offer to each other become available for an agreement. In the second round, simultaneously each player accepts at most one contract in each of her relationships, selecting from the available contracts. For a given linked pair, if a contract is accepted by both players then it goes into force; otherwise, they have the null contract specifying zero transfers. Players observe the acceptance choices only in their own relationships.

With this contracting institution, there is no efficient equilibrium of the grand game. To see why, suppose there is an efficient equilibrium and we will find a contradiction. The pairs $(1, 2)$ and $(3, 4)$ must form contracts that induce players 1 and 4 to select $a_1 = 1$ and $a_4 = 1$. Suppose player 1 were to deviate by declining to accept any contract with player 2 and then choosing $a_1 = 2$. This deviation is not observed by player 4, who still forms a contract with player 3 and selects $a_4 = 1$. So the deviation ensures player 1 of a payoff of at least 9, which is a lower bound on player 1's equilibrium payoff. This is incompatible with the fact that players 2 and 3 can each guarantee themselves a payoff of at least 2, player 4 can guarantee a payoff of at least 4, and the sum of payoffs is 16 in the efficient outcome.

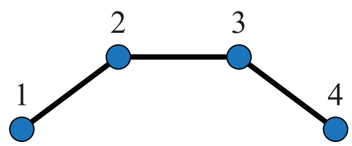

	$a_3 = 1$	$a_3 = 0$	$a_3 = 2$
$a_1 = 1$	0, 4, 4, 8	0, 0, 6, 7	0, 0, 6, 7
$a_1 = 0$	4, 3, 0, 2	4, 2, 4, 2	7, 2, 1, 1
$a_1 = 2$	9, 3, 0, 0	1, 2, 7, 1	1, 2, 1, 1

Figure 3: An example with a peripheral beneficiary.

The same logic holds for similar contracting institutions with more rounds. Consider an R-round process where R is an even number. In odd-numbered rounds, linked players who have not already formed a contract simultaneously send contract offers to each other. In even-numbered rounds, player simultaneously choose whether to accept contracts as described above. Given an efficient equilibrium, let $\overline{r}$ be the last round in which a contract is formed by the pair $(1, 2)$ or the pair $(3, 4)$. Player 1 or player 4 can profitably deviate in round $\overline{r}$ in the same manner just described, leading to a contradiction as before.

Collaboration and a peripheral beneficiary

Next consider a variant of the previous example with the same network but in which only players 1 and 3 take actions in the underlying game. The actions of players 2 and 4 are fixed at 1. Player 1 is a supplier for player 2, as before. Player 4 is now a beneficiary of successful collaboration between the others. Payoffs in the underlying game are given by the table on the right side of Figure 3. The efficient action profile is $a^* = (1, 1, 1, 1)$ and the Nash equilibrium of the underlying game is $\underline{a} = (0, 1, 0, 1)$.

This example adds an element to the incentive issues discussed in the context of the first example. Player 1 can guarantee herself a payoff of at least 4 by refusing to contract with player 2. Therefore achieving the efficient action profile a^* must involve a contracted transfer of at least 4 from player 2 to player 1 in equilibrium. Such a transfer implies that player 2's equilibrium payoff would be nonpositive unless this player receives a transfer from player 3. Because player 2 can guarantee a payoff of at least 2 by refusing to contract, the equilibrium contract for the pair $(2, 3)$ must specify a transfer to player 2 of at least 2 when a^* is chosen. Likewise, because player 3 can guarantee a payoff of at least 4 by refusing to contract, the equilibrium contract between players 3 and 4 must specify a transfer of at least 2 to player 3 when a^* is played.

In summary, the three bilateral contracts must work together to transfer some of player 4's benefit of the efficient productive action across the network to player 1. It remains to be seen whether there is an equilibrium of the grand game in which such contracts are written, and this depends on the contracting institution. It should be clear, for instance, that the two-round protocol described above would not support such an equilibrium. Player 4 would gain by deviating to not contract with player 3, while players 1 and 2 form their equilibrium contract and player 1 then chooses $a_1 = 4$.

3 The Model

This section describes the contractual setting and implementation concept.

3.1 Setting

There are n players who will interact in an underlying game $\langle A, u\rangle$, where $A = A_1 \times A_2 \times \cdots \times A_n$ is the space of action profiles and $u : A \to \mathbb{R}^n$ is the payoff function. Payoffs are in monetary units. Let $N = \{1, 2, \ldots, n\}$ denote the set of players. 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. The players commonly know the underlying game, but $\langle A, u\rangle$ is not observed by the external enforcer.

The external enforcer compels monetary transfers between contracting parties, as directed by the contracts they form. The outcome of the underlying game is fully verifiable, so contracts can specify transfers as a function of the action profile $a \in A$ that the players select. Contracts can be formed only in a restricted set of bilateral relationships given by a fixed network $L \subset N \times N$, meaning 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. Network L is undirected and thus symmetric, so $(i, j) \in L$ implies that $(j, i) \in L$. Assume that L is connected.

Although contracting partners can condition transfers between them on actions taken by third parties, their contract 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.[9] Thus, a contract between players i and j is a function $m : \mathcal{A} \to \mathbb{R}^n_0$, where $m_k(a) = 0$ for all $k \notin \{i, j\}$ and every $a \in \mathcal{A}$. Let C^{ij} be the set of of feasible contracts between players i and j. Denote by $\underline{m}$ the null contract that always specifies a transfer of zero.

Let $\mathcal{M}$ denote a set of contracts formed by the various contractual relationships, and 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 maximize $\sum_{i\in N} u_i(a)$, the players' joint value in the underlying game.

3.2 Contracting institution formalities

A *contracting institution* is an extensive game form with costless 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 $1, 2, \ldots, R$, where R is a finite integer. There is no discounting. In every round, each player i sends messages to every other player. Let λ^r_{ij} 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 $h^r_{ij} = (\lambda^1_{ij}, \lambda^2_{ij}, \ldots, \lambda^r_{ij})$ denote the sequence (or "history") of messages sent from player i to player j from round 1 to round r. Let $\underline{h}^r = (\underline{\lambda}, \underline{\lambda}, \ldots, \underline{\lambda})$ be the sequence of r null messages. A full history of messages through round r is given by $h^r = \{h^r_{ij}\}_{i\neq j}$. A corresponding history of messages between just players i and j is given by (h^r_{ij}, h^r_{ji}).

[9] Due 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").

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^1_{ij} 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 $\Lambda^r_{ij}(h^{r-1})$ denote the set of feasible messages from player i to player j in round r. The message sets are defined arbitrarily by the contracting institution and, for example, can include descriptions of contracts. Assume that the null message is an element of the feasible message set in every round. I also allow the contracting institution to specify a public randomization device following messages. Let ϕ denote this random draw.

Finally, 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. This mapping represents how the contracting institution interprets the sequence of messages, resulting in a contract that will be enforced. 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 $\mathcal{M} = \{\mu(i, j, h^R, \phi) \mid i < j\}$ 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 $\mathbb{R}^n_0$.

A contracting institution is called *natural* if the following assumptions hold: First, contracting in different relationships takes place independently. Second, players have the option of refusing to contract. Third, 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. Here are the formal descriptions of these assumptions:

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 $r \in \{2, 3, \dots, R\}$, $\Lambda^r_{ij}(h^{r-1})$ depends on h^{r-1} only through the component $(h^{r-1}_{ij}, h^{r-1}_{ji})$, so we can write $\Lambda^r_{ij}(h^{r-1}_{ij}, h^{r-1}_{ji})$. Second, the contract $\mu(i, j, h^R, \phi)$ depends on h^R only through the component (h^R_{ij}, h^R_{ji}), and we write $\mu(i, j, (h^R_{ij}, h^R_{ji}), \phi)$.

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,

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

for all h^R_{ij} and h^R_{ji}.

Part ii of the private-contracting assumption requires that, at the end of round r, player i's personal history is exactly $\{h^r_{ij}, h^r_{ji}\}_{j \neq i}$, 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 $\times_{j\neq i}\Lambda^r_{ij}$. To make precise part i of the private-contracting assumption, 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, $h^R_{ij} = \underline{h}^R$ and $h^R_{ji} = \underline{h}^R$, 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 $(\Lambda^r_{ij})_{i,j,r}$, 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.

3.3 The notion of implementation

For a given contracting institution, network L, and underlying game $\langle A, u\rangle$, the grand 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 $M(a)$, resulting in payoff vector $u(a) + M(a)$.

Note that, because of private contracting, the players have asymmetric information in the contracting phase and 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)—that is, probability distributions over strategy profiles—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.[10]

In the analysis that follows, I hold fixed n and 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 grand 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*

[10]It 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 then is worthwhile to analyze.

game $\langle A, u\rangle \in G$ and every connected network L, there is a sequential equilibrium of the grand 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.

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.*

At the risk of belaboring the point, 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.

4.1 Specifications of the PCA contracting institution

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, wherein 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.[11]

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^1_{ij} is the set of triples (c^0, c^1, p), where $c^0, c^1 \in C^{ij} \cap \overline{\mathcal{M}}$ and $p = (p^2, p^3, \ldots, p^R) \in \overline{Q}^{R-1}$. Here c^0 is the

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

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 $r \in \{2, 3, \ldots, R\}$, 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 $p = (0, 0, \ldots, 0)$.

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 let us say that these players "*made a provisional arrangement*" and (c^0, c^1) is their tentative contract. 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 each round $r \in \{2, 3, \ldots, R\}$, 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 messages are private and 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.

Note that 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 every underlying game $\langle A, u \rangle \in G$ and every connected network L, there is a sequential equilibrium of the grand 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 Illustrations of efficient equilibria

The proof of the Theorem involves intricate equilibrium constructions, and beliefs must be specified for the many information sets at which players have limited information about the actions of the others. This subsection illustrates the main components using the two examples discussed in Section 2. Assume $n = 4$ and suppose G is any set that contains the two underlying games described in Section 2. Fix the PCA contracting institution with $R = 3$ and with $\overline{\mathcal{M}}$ and $\overline{Q}$ assumed to contain the contracts and penalty numbers described below. These examples do not require the public randomization device, so we can assume the specifications described next are the same for $\phi = 1$ and $\phi = 0$.

Collaboration-agreement example

Consider the collaboration-agreement example with the underlying game shown in Figure 2. I next sketch an efficient equilibrium for the grand game by describing the actions and beliefs specified for key information sets, starting with those on the equilibrium path.

In round 1 of the contracting phase, players 2 and 3 are supposed to offer each other the contract $c^0 = c^1 = m^{23}$ defined by

$$m^{23}(1,1,1,1) = (0,0,0,0), \quad m^{23}(a_1,1,1,1) = (0,-4,4,0),$$
$$m^{23}(1,1,1,a_4) = (0,4,-4,0), \quad \text{and} \quad m^{23}(a_1,1,1,a_4) = (0,0,0,0)$$

for $a_1 \neq 1$ and $a_4 \neq 1$. In words, player 2 gives assurance that player 1 will select $a_1 = 1$, and player 3 gives assurance that player 4 will select $a_4 = 1$. A player pays an assurance penalty of 4 to the other if the supplier on the first player's side of the network fails to choose the efficient action in the production phase. Further, the players are supposed to offer each other the vector of cancellation penalties $p = (p^2, p^3) = (0, 4)$.

Players 1 and 2 are supposed to offer each other the contract m^{12} defined by

$$m^{12}(1,1,1,1) = (3,-3,0,0), \quad m^{12}(1,1,1,a_4) = (7,-7,0,0),$$
$$m^{12}(a_1,1,1,1) = (-5,5,0,0), \quad \text{and} \quad m^{12}(a_1,1,1,a_4) = (-1,1,0,0)$$

for $a_1 \neq 1$ and $a_4 \neq 1$. This contract specifies a base payment of 3 from player 2 to player 1. A deviation from $a_1 = 1$ requires player 1 to pay an assurance penalty of 8, whereas a deviation from $a_4 = 1$ requires player 2 to pay an assurance penalty of 4. These players are supposed to offer each other cancellation penalties of $p = (p^2, p^3) = (5, 6)$.

Likewise, players 3 and 4 are supposed to offer each other the contract m^{34} defined by

$$m^{34}(1,1,1,1) = (0,0,-3,3), \quad m^{34}(a_1,1,1,1) = (0,0,-7,7),$$
$$m^{34}(1,1,1,a_4) = (0,0,5,-5), \quad \text{and} \quad m^{34}(a_1,1,1,a_4) = (0,0,1,-1)$$

for $a_1 \neq 1$ and $a_4 \neq 1$, and the same vector of cancellation penalties $p = (p^2, p^3) = (5, 6)$.

If a given player i and her contracting partners behaved as specified in round 1, then assume this player believes no one deviated, and prescribe that this player send the null message to the others in round 2. The same is specified for the belief and behavior in round 3. Then, if player i detected no deviations in all three rounds of the contracting phase, this player believes that everyone contracted as prescribed (so contracts m^{12}, m^{23}, and m^{34} are in force) and that the other players will all choose action 1 in the production phase. In this contingency, player i is prescribed also to choose $a_i = 1$ in the production phase. On the equilibrium path, the payoff vector is $(5, 3, 3, 5)$, which for each player exceeds the payoff of the Nash equilibrium of the underlying game.

Clearly if player 1 reaches the production phase having not detected any deviations from the prescribed play, this player rationally chooses $a_1 = 1$ in response to the belief that player 4 will choose $a_4 = 1$. In this case, player 1 expects a payoff of $2 + 3 = 5$, whereas the best deviation, to $a_1 = 2$, would deliver a payoff of $9 - 5 = 4$. Note as well that, conditional on reaching round 2 or round 3 of the contracting phase without detecting a deviation, player 1 could not gain by canceling her contract with player 2. Such a cancellation would reduce her expected payoff from 5 to at most 4 because of the cancellation penalty. Player 4 has the same incentives.

Do players 1 and 4 want to behave in round 1 as prescribed? To address this question, we must incorporate the incentives and beliefs of the other players in later rounds and in the production phase. For example, whether player 1 would gain by deviating in round 1 (so that the null contract will be formed with player 2) depends on what such a deviation would cause player 4 to eventually choose in the production phase. If player 4 would still choose $a_4 = 1$, then player 1 would get a payoff of 9 by deviating in round 1 and then by choosing $a_1 = 2$ in the production phase. Critical to the equilibrium construction is that a deviation by player 1 in round 1 will lead to a sequence of contract cancellations, resulting in player 4 selecting $a_4 = 0$ in the production phase based on the belief that player 1 will choose the same action.

Here are the details: Suppose that player 1 were to deviate in round 1, so that the null contract will be formed with player 2. Upon observing the first-round deviation, player 2 believes that player 1 will choose $a_1 = 0$ in the production phase, which will put player 2 on the hook for the assurance penalty in the contract with player 3. Player 2 is better off canceling the contract with player 3 to avoid the assurance penalty, even if it changes her payoff in the underlying game, and she strictly prefers to cancel in round 2 when the cancellation penalty is 0 rather than wait until round 3 when it is 4. So the equilibrium strategy prescribes that, in this contingency, player 2 cancel m^{23} in round 2.

If player 3 observes player 2 cancel m^{23} in round 2, then player 3 believes that this was because player 1 deviated in round 1, leading player 3 to further believe that player 1 will select $a_1 = 0$ in the production phase. Now subject to the assurance penalty in the contract with player 4, player 3 has the incentive to cancel m^{34} in round 3, and this is what the equilibrium strategy prescribes. Finally, if player 4 learns that player 3 cancels m^{34} in round 3, then player 4 believes that player 1 deviated in round 1, that player 2 canceled m^{23} in round 2, and that player 1 will select $a_1 = 0$ in the production phase. Player 4 therefore rationally chooses $a_4 = 0$ in the production phase.

In summary, if player 1 deviates in round 1 to nullify the prescribed contract with player 2, a wave of cancellations ensues and spreads across the contractual network, ultimately coordinating the players on the Nash equilibrium of the underlying game. Player 1 does not gain with such a deviation. The incentives for player 4 are similar, as are the incentives for players 2 and 3. It is important to note that the wave of cancellations and play in the production phase are motivated by the players' beliefs off the equilibrium path, since players do not observe what is happening elsewhere in the network. Left out of this account are the specifications of beliefs and behavior at a variety of other off-path information sets, but the key elements of the equilibrium construction have been provided here.

Example of a peripheral beneficiary

In the equilibrium construction for the first example, players have the incentive to perpetuate a wave of contract cancellations because they prefer paying cancellation penalties rather than assurance penalties. For instance, if player 4 deviates in round 1 to form the null contract with player 3, then player 3 has the incentive to cancel the tentative contract with player 2 in round 2. Player 3 believes that player 4 will no longer choose her part of the efficient action profile in the underlying game, making player 3 liable for the assurance penalty in the contract with player 2.

However, in the case shown in Figure 3, player 4 has no action to take in the underlying game, and so player 3 faces no such liability. There is an efficient equilibrium, but the construction is a bit different than in the first example. It turns out that player 3 can be given the incentive to cancel the tentative contract with player 2 on the basis of a lost monetary transfer from player 4.

For instance, there is an efficient equilibrium in which in round 1 players 3 and 4 declare the contract m^{34} defined by $m^{34}(1,1,1,1) = (0,0,5,-5)$ and $m^{34}(a_1,1,a_3,1) = (0,0,-5,5)$ for $a_1 \neq 1$ and/or $a_3 \neq 1$, and they specify cancellation penalties of $p = (p^2,p^3) = (0,8)$. Contracts between the other players are similar to those described for the first example, except that they entail a transfer of money from player 3 to player 2 and from player 2 to player 1 when the efficient action profile is played in the production phase. If player 4 were to deviate in round 1 to form the null contract with player 3, it is in player 3's interest to cancel the tentative contract with player 2 in round 2, for otherwise player 3 would get a payoff below 4, which this player could guarantee by canceling.

5 Proof

The major parts of the proof are contained in this section. Subsection 5.1 proves a version of the Theorem for a large class of settings that I call the *Critical-Extremes Class*, which includes the collaboration example. For this class, an efficient outcome is implemented with certainty. Subsection 5.2 generalizes the analysis to cover all other settings, including underlying games with peripheral beneficiaries as in the second example, and completes the proof. Some of the details of the generalized analysis are shown in the Appendix.

The following additional notation is used below: 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. 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. Say that a set of contracts $\mathcal{M}$ *supports* a^* if a^* is a Nash equilibrium of the game $\langle A, u+M\rangle$.

5.1 The critical-extremes class

Whatever is the connected network L, we can find a minimally connected sub-network $K \subset L$, where for each pair of players i and j, there is exactly one sequence $\{k^t\}_{t=1}^{T} \subset N$

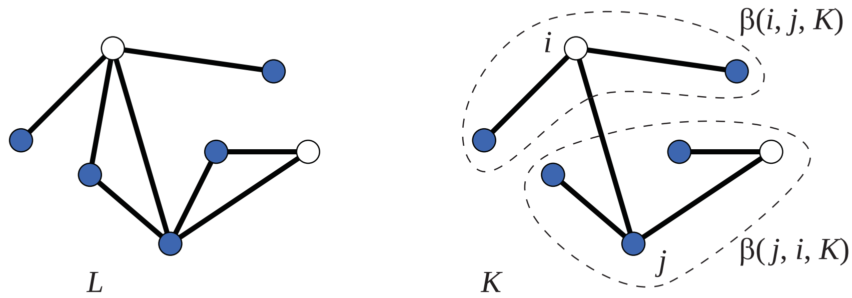

Figure 4: Construction of K and definition of β.

satisfying $k^1 = i$, $k^T = j$, and $(k^{t-1}, k^t) \in K$ for all $t = 2, 3, \ldots, T$. 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 sub-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. 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. Figure 4 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.

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 grand game in which the equilibrium outcome is efficient.*

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 subnetwork and actions profiles with the properties stated 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'. This is followed by the general technical development, including a precise specification of actions and beliefs at 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

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

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 $p^* = (p^{2*}, p^{3*}, \ldots, p^{R*})$ be any vector satisfying $p^{2*} > \gamma$ and $p^{r*} = p^{2*}(r-1)$ for all $r = 3, 4, \ldots, R$. 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

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

  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 $\#\{k \in \beta(j, i, K) \,|\, a_k \neq a_k^*\}$ 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, $p^{2*}, p^{3*}, \ldots, p^{R*}$, 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 $2, 3, \ldots, R$, 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 made no provisional arrangements and thus get the null contract.

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 $r \in \{1, 2, 3, \ldots, R\}$, 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 $2, 3, \ldots, r$. 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 $r \in \{1, 2, 3, \ldots, R\}$, 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 H1i, H2i, and so on) along with the actions specified at these information sets (labelled S1i, S2i, and so on). The first three entries refer to contingencies that will be on the equilibrium path and also those that deviate in only inessential ways.

H1i The null/empty history at the beginning of round 1 of the contracting phase.

S1i Player i sends the prescribed first-round messages to each of the other 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.

H2i For each $r \in \{1, 2, \ldots, R-1\}$, the personal histories through round r that essentially conform to the prescribed path.

S2i Player i sends the null message to all other players in round $r + 1$.

H3i The personal histories through round R of the contracting phase that essentially conform to the prescribed path.

S3i Player i selects action a_i^* in the production phase.

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

H4i For each $r \in \{1, 2, \ldots, R-1\}$ 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.

S4i Player i declares "cancel" to every contracting partner with whom the action "cancel" is available.

H5i The personal histories through round R such that for some player $j \in K^i$ and an integer $r \in \{1, 2, \ldots, R\}$ 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.

S5i Player i selects action $\underline{\alpha}_i$ in the production phase.

The inequalities in H4i and H5i 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.

H6i 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.

S6i 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 H1j-H6j, for every $j \neq i$. The belief at H1i is denoted B1i, the belief at H2i is denoted B2i, and so on. These are not the full appraisals at information sets H1j-H6j, but what is described here is sufficient for checking that the specified behavior S1j-S6j 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 S1j-S6j 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.

B1i At the beginning of the game (for personal history H1i), player i believes that all of the other players will behave according to strategy specifications S1-S6.

B2i-B3i For information sets in H2i and H3i, 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 $L^i \setminus K^i$.

B4i For information sets in H4i, 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 S2k-S6k. 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.

B5i For information sets in H5i, 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 B4i. 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 $L^i \backslash K^i$ 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 $k \in K^i \backslash \{j\}$. 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 S1i is optimal.

For information sets in H4i, 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 $a^*_{\beta(j,i,K)}$ 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 $r \in \{2, 3, \ldots, R\}$, 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 S4i is optimal.

Likewise, S5i and S6i are optimal because at information sets H5i and H6i 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 H1i for any player i. At this information set, let us assign probability ξ to each action that, for a single player $j \in L^i \backslash K^i$, has player i sending a message other than the null message to player j and has player i sending the prescribed message (given by S1i) 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 S1i, and assign the remaining probability mass to action S1i. At each information set in H2i-H6i, assign probability ξ^2 to every feasible action other than that prescribed by S2i-S6i; put the remaining probability on the action prescribed by S2i-S6i.

Let H denote the set of information sets in the grand 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 $\hat{H} = H \backslash \tilde{H}$. 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 grand game (the actions for information sets in $\tilde{H}$). For any player i and each information set $h \in \tilde{H}_i$, as $\xi \to 0$ 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 B1i-B6i. Likewise, $\sigma_{\tilde{H}}(\xi)$ converges to what is specified by S1-S6, which we can denote $\overline{\sigma}_{\tilde{H}}$. With rationality at these information sets already verified, we thus have for information sets $\tilde{H}$ 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 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 are specified as in the grand 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 find a strictly positive sequence $\{\xi_k\}$ converging to zero such that $\sigma_{\hat{H}}(\xi_k)$ converges to some profile $\overline{\sigma}_{\hat{H}}$.

Let us write $\sigma^k = (\sigma_{\tilde{H}}(\xi_k), \sigma_{\hat{H}}(\xi_k))$ and $\overline{\sigma} = (\overline{\sigma}_{\tilde{H}}, \overline{\sigma}_{\hat{H}})$. Clearly σ^k converges to $\overline{\sigma}$. 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 $\overline{Q}$ was assumed to contain the numbers $0, p^{2*}, p^{3*}, \ldots, p^{R*}$, and q. Otherwise, $\overline{\mathcal{M}}$ and $\overline{Q}$ were unconstrained except to be finite. So we can let $\mathcal{M}' \equiv \mathcal{M}^* \cup \{\underline{m}\}$ and $Q' \equiv \{0, p^{2*}, p^{3*}, \ldots, p^{R*}, q\}$, and the proof of Proposition 1 is complete.

5.2 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 deals with cases in which some players at the periphery of a network don't have actions in the underlying game but benefit from play of the efficient action profile, as in the second example. 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, $\underline{\alpha}_i(a_i^*) = 1$ 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.

Arbitrary underlying games

Returning to the analysis of arbitrary underlying games and networks, we have the following extension of Proposition 1 that is proved in Appendix A.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 grand game in which the equilibrium outcome is ε-efficient.*

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.

6 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.

6.1 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 the examples resented earlier, 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 Section 2, implying 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 grand 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* 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: By enforcing cancellation penalties, the PCA contracting institution restricts the players' ability to back out of contractual arrangements toward the end of the contracting phase. It turns out that a restriction of this sort 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 grand game has a^* played with a probability greater than $1 - \underline{\varepsilon}$.*

I turn next to folk-theorem considerations, where a characterization of multiple equilibria can be provided most easily for the critical-extremes class as a general version of the Theorem. 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 $\underline{\alpha}$ of $\langle A, u\rangle$, and any feasible payoff vector (incorporating transfers) $v \geq u(\underline{a})$, there is a sequential equilibrium of the grand 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.

6.2 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 $r = 1$. The other rounds are dates at which the players can communicate along edges of the network. Communication is private and locally verifiable in the sense that messages exchanged by a pair of players are available to the enforcer for evaluating the contract between these two players. The contracted transfers can be contingent on these messages, the random draw σ, and the action profile a. In this fashion, the contracts formed in the first round incorporate options triggered by messages sent in later rounds.

On a related note, in the proof of the Theorem, the exact penalties q and p were chosen for convenience to suffice for all of the contracting pairs. As a result, the chosen penalties

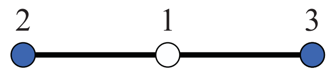

	$a_3 = 1$	$a_3 = 0$
$a_2 = 1$	6, −2, −2	3, −2, 0
$a_2 = 0$	3, 0, −2	0, 0, 0

Figure 5: A setting with partial verifiability.

were large. This is not necessary, for 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.

6.3 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. Consider a simple case of team production with partial output verification, where $n = 3$. The network and underlying game are shown in Figure 5.

Player 1 is the manager and players 2 and 3 are workers. Player 1 has no productive action in the underlying games, so $A_1 = \{1\}$. The other players have action spaces given by $A_2 = A_3 = \{0, 1\}$, where 1 stands for high effort and 0 represents low effort. Payoffs are given by $u_1(a) = 3(a_2 + a_3)$, $u_2(a) = -2a_2$, and $u_3(a) = -2a_3$. Partial verifiability of a is represented by the partition of A with these two elements: $\overline{\rho} = \{(1, 1)\}$ 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 or not. The efficient action profile is $(1, 1, 1)$

I claim that, regardless of the contracting institution (assumed natural), in every equilibrium of the grand game, action profile $(1, 0, 0)$ is played with probability 1 in the production phase. The analysis substantiating this claim is in Appendix A.2, 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 $(1, 1, 1)$ 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 $\overline{\rho}$ 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 $\overline{\rho}$ 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, because player 2 does not observe the deviation and still believes that her effort choice influences whether $\overline{\rho}$ or $\underline{\rho}$ is obtained. 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. The deviation is thus profitable, which means there is no equilibrium in which $(1, 1, 1)$ is played with certainty.

7 Conclusion

The modeling exercise herein offers a benchmark result on how LDL externalities can be internalized through endogenously formed chains of independent bilateral contracts, assuming connected networks and globally verifiable productive actions. I stress that this is a benchmark rather than a claim that efficiency will actually be achieved in any particular real setting. In this sense, the model can help categorize barriers to efficiency that carefully designed institutions may be able to overcome.

The model points to four distinct barriers to inefficiency in practice: (1) limited verifiability of productive actions, as demonstrated by the third example; (2) institutional rules or technologies that limit the feasible space of contracts; (3) suboptimal contracting institutions that, for instance, 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 related to exertion of bargaining power.

When evaluating the barriers to efficiency, it is sometimes helpful to categorize examples by prominent aspects of their networks and the structure of their underlying games. Figure 6 pictures four classes of networks. The networks shown would be suitable to model, from left to right, (i) vertical contracting with a single supplier, as well as common-agent or common-principal settings;[12] (ii) vertical contracting in a bipartite supply network and two-sided markets;[13] (iii) platforms and general intermediation networks;[14] and (iv) community interaction with an arbitrary contractual network.[15]

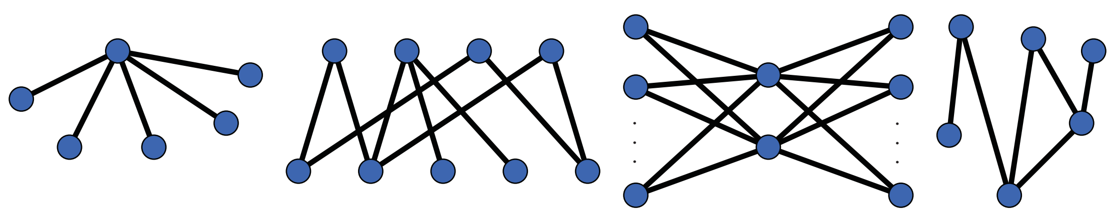

Figure 6: A few common networks.

[12]Bernhaim and Whinston (1986a,b), Segal (1999), and Galasso (2008) were noted in the Introduction. Martimort (2007) surveys the related literature.

[13]Kranton and Minehart (2001) initiated a line of research on buyer-seller networks; other work on such vertical contracting includes Elliott (2013) and Nocke and Rey (2018). Particularly relevant to the present modeling exercise is the analysis of collusion and competition with cross-licensing, such as in Jeon and Lefouili (2018, 2020) and Rey and Vergé (2019).

[14]Rochet and Tirole (2003, 2006) and Armstrong (2006) look at various pricing alternatives for intermediaries; Rysman (2009) provides an overview of this area of literature. Other entries 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.

[15]Regarding more general networks, an important but relatively undeveloped branch of contract theory views an organizational structure as a nexus of contracts (Jensen and Meckling 1976). Laffont and Martimort (1997) surveys this area of the literature and Cafeggi (2008) provides a legal perspective. Economides (1996) provides an overview of contractual issues in general networks, mostly notably data networks.

Other related technical work

I next add to the earlier discussion of related technical literature. First, it is worth expounding on the relative strengths of the noncooperative modeling approach taken herein, in comparison to the approach taken in other areas of the literature including cooperative matching theory and coalitional bargaining theory.[16] Two distinctions of my noncooperative approach are that (1) it accounts for productive interaction explicitly using noncooperative game theory and (2) it models contract formation noncooperatively as well. In contrast, cooperative matching and coalitional bargaining models typically account for productive interaction only as public-action implications of contracts; that is, they specify payoffs as a function of an abstract set of contracts that the players form. Further, in the case of cooperative matching, contracting is analyzed using a stability concept.

Regarding item 2, on modeling contract formation, the goal of the Nash program is to establish a mathematical equivalence between stability concepts and the outcomes of noncooperative bargaining games. As noted in the Introduction, some progress has been made in the matching-with-contracts context, but the program has not been advanced for settings with LDL externalities. Therefore, it is unclear whether any given stability condition would translate into equilibrium conditions in a noncooperative model of contracting.

Regarding item 1, on modeling productive interaction, two points are worth making. First, without an explicit account of productive actions, one could not distinguish various ways in which linkages may occur across contractual relationships, such as between "contracts on contracts" and contracting on only others' productive actions. The distinctions have practical importance, for these linkages differ in terms of expression, interpretation, enforcement, and verification requirements.

Second, the noncooperative approach provides a foundation for distinguishing types of externalities and understanding exactly what is required to internalize them. Consider, for instance, the collaboration-agreement example, where, among other things, player 4's productive action directly affects player 1's payoff in the underlying game. Compare this to a supply-chain setting in which player 4 may provide an intermediate good to player 3, who in turn may provide an intermediate good to player 1. For the latter setting, suppose player 1's payoff is a function of only the type and quantity of the intermediate good delivered by player 3, and player 3's cost of producing the good for player 1 depends on the intermediate good supplied by player 4. Thus, player 1 cares about player 4's productive action only to the extent that it affects the negotiated terms of her contract with player 3.

Because these two settings are distinguished by different production technologies, they are best described using a model that explicitly accounts for productive interaction, as accomplished herein by specifying the noncooperative underlying game. A modeling approach

[16] In coalitional bargaining models, centralized contracting is possible because the grand coalition can form a contract, 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. 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). One variation is a form of interconnected bilateral bargaining, such as the model of De Fontenay and Gans (2014), where disagreement in one bilateral relationship causes this relationship to sever and triggers a restart to negotiation in all the other relationships.

that abstracts from the underlying game by specifying payoffs as a function of an abstract set of contracts is not well suited to make the distinctions that these two examples illustrate. For instance, Fleiner et al. (2018), Fleiner et al. (2019), and others in the cooperative matching literature assume that a player's payoff depends on only the contracts this player signs, which does not allow for the externality in the collaboration-agreement example. Some matching models allow payoffs to be a function of the entire set of contracts formed, such as in Rostek and Yoder (2019, 2020) and Pycia and Yenmez (2019) for two-sided markets. These models can capture such an externality to some extent, but it is not clear how they could distinguish between, say, the collaboration-agreement and supply-chain examples without an explicit account of the production technology.[17]

In summary, the noncooperative approach, taken herein and in line with Jackson and Wilkie (2005) and Ellingsen and Paltseva (2016), has advantages that complement other approaches to the study of contractual networks.[18] The noncooperative approach provides a good foundation for exploring topics such as contracts on contracts and other methods of linking contractual relationships. Regarding contracts on contracts, Peters and Szentes (2008) tackle one of the key modeling components in their analysis of interactive promises. They examine settings 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 suggests that interactive contracts require the external enforcement system to develop a sophisticated language for the cross-referencing of promises.

Further directions

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 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.

[17] The supply chain example features what some call a pecuniary externality, which I would not classify as an externality generally but rather the downstream result of a possible contracting or market distortion. Other entries in the matching literature include Hatfield and Kominers (2012, 2015), Hatfield et al. (2013), Ostrovsky (2008), and Manea (2018).

[18] Additional related papers from the prior literature include Guttman (1978), Danziger and Schnytzer (1991), Guttman and Schnytzer (1992), Varian (1994), and Yamada (2003).

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.[19] Limited verifiability, in terms of both the example in Section 6.3 and differences between what is locally and globally verifiable, represent another line of potentially insightful inquiry. For settings in which LDL externalities cannot be internalized through contracting on 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. It would also be useful to examine special classes of underlying games and to categorize externalities along with the contractual and institutional elements needed to internalize them. Finally, it would be helpful to explore ways of extending the Nash program into the present setting.

A Appendix (intended to be Online Appendix)

This section contains the analysis deferred in Sections 5.2 and 6.3.

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

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

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.

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 $\overline{a} \in A$ such that $\overline{a}_i \neq a_i^*$ 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 $\hat{K}$.

[19] A 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.

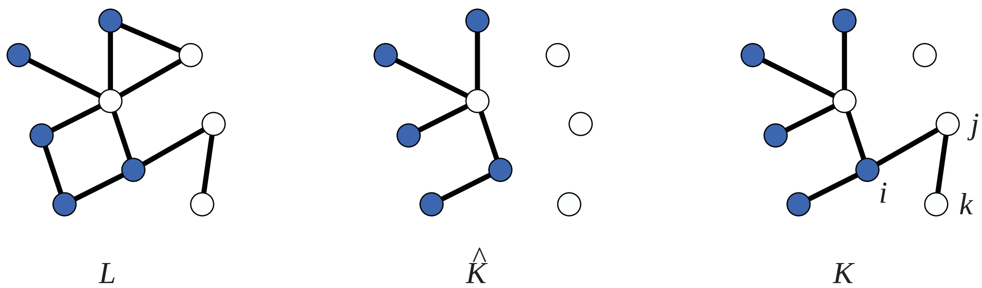

Figure 7: Selection of $\hat{K}$ and K.

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 $U_N(a^*) > U_N(\underline{a})$. 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.[20] The result is a network K with the desired properties.

Figure 7 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*.

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

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 $\overline{a}$ in the event that $\phi = 0$. 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 $\underline{\alpha}_i$ would not be revealed to have deviated from the specified behavior:

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

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 7. 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 7 for the players i, j, and k shown, we have $\tau(i,j) = 1$ and $\tau(j,k) = 2$.

Let $\overline{\tau} \equiv \max_{(i,j)\in K} \tau(i,j)$. Note that $\overline{\tau}$ 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 $p_{ij}^* = (p_{ij}^{2*}, p_{ij}^{3*}, \ldots, p_{ij}^{R*})$ with the following specifications. First, for any integer $r \geq 2$ such that $r < 2 + \overline{\tau} - \tau(i,j)$, let $p_{ij}^{r*} = 0$. Then for $r \geq 2 + \overline{\tau} - \tau(i,j)$, let $p_{ij}^{r*} = (1+\varepsilon)\gamma[r - 1 - \overline{\tau} + \tau(i,j))]$. 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 7. Pairs in the core have zero cancellation penalties through round $2 + \overline{\tau}$. 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 $\theta = \sum_{i<j}(1-\varepsilon)c^{1ij}(a^*) + \varepsilon c^{0ij}(\overline{a})$. 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 $\overline{a}$ and c^{1ij} assures a^*, and in these cases the transfers are the same. That is, for all $a \in A$,

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

  and

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

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

- The expected payoff vector of the prescribed actions exceeds the Nash equilibrium payoff vector for every player $i \in N(K)$. That is, $(1-\varepsilon)u_i(a^*) + \varepsilon\, u_i(\overline{a}) + \theta_i \geq u_i(\underline{\alpha})$ 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

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

  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 $2, 3, \ldots, R$, 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 $\overline{a}$.

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 made 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 7. 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_2, a_3) 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 $m_2^{12}(\overline{\rho})$ 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 $\zeta m_2^{12}(\overline{\rho}) + (1-\zeta) m_2^{12}(\underline{\rho}) - 2 \geq m_2^{12}(\underline{\rho})$, which simplifies to

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

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 $a = (1,1,1)$ 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.

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.

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.

Bernstein, L. and B. Peterson (2020), "Managerial Contracting: A Preliminary Study," preliminary draft, University of Chicago.

Bloch, F. and A. Gomes (2006), "Contracting with Externalities and Outside Options," *Journal of Economic Theory* 127 (1): 172–201.

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.

Coase, R. H. (1960), "The Problem of Social Cost," *Journal of Law and Economics* 3 (Oct.): 1–44.

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. Jankóz, A. Tamura, and A. Teytelboym (2018), "Trading Networks with Bilateral Contracts," unpublished manuscript.

Fleiner, T., R. Jagadeesan, Z. Jankóz, and A. Teytelboym (2019), "Trading Networks With Frictions," *Econometrica* 87 (5): 1633–1661.

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.

Hatfield, J. W. and P. R. Milgrom (2005), "Matching With Contracts," *American Economic Review* 95 (4): 913–935.

Hatfield, J. W. and S. D. Kominers (2012), "Matching in Networks with Bilateral Contracts," *American Economic Journal: Microeconomics* 4: 176–208.

Hatfield, J. W. and S. D. Kominers (2015), "Multilateral Matching," *Journal of Economic Theory* 156: 175–206.

Hatfield, J. W., S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp (2013), "Stability and Competitive Equilibrium in Trading Networks," *Journal of Political Economy* 121 (5): 966–1005.

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.

Jeon, D.-S. and Y. Lefouili (2018), "Cross-Licensing and Competition," *RAND Journal of Economics* 49 (3, Fall): 656–671.

Jeon, D.-S. and Y. Lefouili (2018), "Decentralizing Cooperation through Upstream Bilateral Agreements," TSE working paper.

Kranton, R., and D. Minehart (2001), "A Theory of Buyer-Seller Networks," *American Economic Review* 91 (3): 485–508.

Kreps, D.M. and R. Wilson (1982), "Sequential Equilibria," *Econometrica* 50 (4): 863–894.

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.

Manea, M. (2018), "Intermediation and Resale in Networks," *Journal of Political Economy* 126 (3): 1250–1301.

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.

Nocke, V. and P. Rey (2018), "Exclusive Dealing and Vertical Integration in Interlocking Relationships," *Journal of Economic Theory* 177: 183–221.

Ostrovsky, M. (2008), "Stability in Supply Chain Networks." *American Economic Review* 98 (3): 897–923.

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 (2019), "Matching with Externalities," unpublished manuscript.

Reisinger, M. (2014), "Two-Part Tariff Competition between Two-Sided Platforms," *European Economic Review* 68: 168–180.

Rey, P. and T. Vergé (2019), "Secret contracting in multilateral relations," TSE Working Paper 16–744.

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.

Rostek, M. and N. Yoder (2019), "Matching with Multilateral Contracts," manuscript in circulation.

Rostek, M. and N. Yoder (2020), "Matching with Complementary Contracts," *Econometrica* 88 (5): 1793–1827.

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.

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. (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