

Contractual Chains

Joel Watson*

October 2023

Abstract

This paper develops a model of private bilateral contracting, in which an exogenous network determines the pairs of players who can communicate and contract with each other. After contracting, the players interact in an underlying game with globally verifiable productive actions and externally enforced transfers. The paper investigates whether such decentralized contracting can internalize externalities that arise due to parties being unable to contract directly with others whose productive actions affect their payoffs. The contract-formation protocol, called the “contracting institution,” is treated as a design element. The main result is positive: There is a contracting institution that yields efficient equilibria for any underlying game and connected network. A critical property is that the institution allows for sequential contract formation or revision. 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 the editor and referees, Nageeb Ali, Gorm Grønnevet, Keri Hu, Natalia Lazzati, Kelvin Leong, 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, UCLA, and Wisconsin. 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. The author thanks JD Watson for encouragement.

1 Introduction

In many contractual settings, there is multilateral productive interaction but barriers prevent the parties from contracting all together. Instead, contracting is possible only in certain small groups that are exogenously specified. These settings often feature *externalities due to lack of direct links (LDL)*, in which agents are unable to contract directly with others whose productive actions they care about. A fundamental issue is whether LDL externalities can be internalized through such decentralized contracting, leading to efficient outcomes.

In this paper, I develop a noncooperative game-theoretic model to study the efficiency issue. The model has the following structure:

- Parties interact in the *contracting phase* followed by the *production phase*, the latter a simultaneous-move *underlying game* that is commonly known to the players.
- Only bilateral contracting is possible. An exogenous network describes the pairs of players that can communicate and establish contracts. Contracting is private and independent across these contractual relationships.
- All productive actions are verifiable by everyone, so there is *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.

An illustration of the set of players and network is shown in Figure 1, where each node is a player and 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 productive action, and likewise player k 's payoff may depend on player i 's productive action, and yet these players are unable to contract together. The pair (i, j) can establish a contract, and so can the pair (j, k) , implying that a chain of contractual relationships can, in principle, arise endogenously to 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 contractual relationships.

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 j 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 j could restrict the contract that players j and k are allowed to create.

To study the prospect of efficient outcomes without arbitrarily specifying the noncooperative protocol for interaction in the contracting phase (and its idiosyncratic constraints), I take the novel approach of treating this protocol, which I call the *contracting institution*, as a design element. Formally, a contracting institution is an extensive game form in which the players freely send messages that determine their externally enforced contracts. For a given contracting institution, the players will play a *grand game* in which they first interact in the contracting institution, then simultaneously select their actions in the underlying game, and finally receive payoffs including the contracted transfers.




Figure 1: A network of contractual relationships.

Critically, the contracting institution is restricted by the network of links and by assumptions that represent the notion of private, independent, and voluntary 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 any pair of players depends on only the messages they exchange, not on messages sent or received by other players. Thus, third parties cannot dictate the terms of a contract, and contracts on contracts are not feasible.¹ Third, players can reject contracts, which ensures the definition of contract is conventional in that the consent of both parties is required.

I focus on a “possibility” question: Is there a contracting institution that *implements efficient outcomes*, meaning that, fixing the contracting institution, for every underlying game and every connected network, the grand game has an efficient sequential equilibrium? If so, then we can say that under the right conditions for contracting, in all productive settings with global verifiability, LDL externalities can be overcome by decentralized formation of contractual chains.

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

The Theorem presented here answers the possibility question in the affirmative, showing that LDL externalities can generally be internalized. The proof entails an elaborate equi-

¹The second assumption can also be motivated on the basis of contracts being verifiable only locally (when enforcing a contract between two players, the court does not observe contracts written by others).

²Examples 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 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.

librium construction with subtleties, but three essential economic elements can be easily described. First, the contracting institution allows for *sequential contract formation*. It has multiple rounds in which contractual arrangements can be made and adjusted. In equilibrium, a player's behavior with one partner will be sensitive to her experiences with other partners. Second, the players coordinate on *assurance contracts* with penalties, whereby players guarantee that specified third parties will select their part of an efficient action profile in the underlying game. Third, the players agree on *cancellation penalties* that discourage them from cancelling tentative contracts except in particularly onerous situations.³

While the Theorem presents as an encouraging result about attaining efficiency, its more practical use may be as a reference point for applications. The analysis shows that efficient contracting relies on having the right kind of contracting institution as well as players coordinating on a socially desired equilibrium, conditions that some real settings may lack. More generally, by precisely accounting for the contracting institution and enforcement technology, the modeling framework helps classify methods of establishing contractual linkages across relationships. The framework can easily be modified to examine variations in the fundamentals of contracting, such as the extent of verifiability and the scope of external enforcement. For instance, an example presented in the Appendix 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) insight about how externalities can be circumvented 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 game of contract negotiation that 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 allows all of the players to join in a single contract, and there are no LDL externalities.⁴

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 others are in the cooperative-theory tradition. Neither strand has examined the general question posed here regarding internalizing LDL externalities.⁵ Using a fully noncooperative

³Options to terminate are common in contracts across industries. Assurance arrangements are also common, particularly with respect to the performance of subcontractors. Contracting partners sometimes 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).

⁴In 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.

⁵An 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

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 downstream firm to condition the transfer on only 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.⁶

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 externalities.

The line on “matching with contracts” initiated by Hatfield and Milgrom (2005) studies stability concepts for models in which the fundamentals are a feasible contracts available to subsets of players and payoffs as a function of the contracts chosen. Closest to my modeling exercise is the model of Rostek and Yoder (2020, 2022), which allows for multilateral contracts and LDL externalities directly via contacts. As with cooperative theory generally, their model abstracts from the details of production and enforcement technology. They focus on the existence of stable matchings and the characterization of stability conditions, and do not address conditions for efficiency. Thus, the objectives pursued and the methods developed herein are complementary to the objectives and methods of cooperative matching theory. Additional discussion of this and other areas of the literature, along with notes about the relative advantages of noncooperative modeling, may be found in Section 5.

Overview

The general model is developed in the next section. Section 3 uses simple examples to discuss barriers to efficient contracting. Section 4 presents the Theorem and describes the contracting institution and a variety of technical elements used in the proof, the constructive part of which appears in Appendix A. Section 5 discusses the modeling approach, additional references, tangential results, and further steps in the research program. Appendix B contains further discussion and the formalities of results stated informally in Sections 3 and 5.

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.

⁶Segal'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. Galasso (2008) 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 The Model

2.1 Setting

The set of players is $N = \{1, 2, \dots, n\}$ for some positive integer n . These players will engage in productive interaction described by a commonly known, simultaneous-move *underlying game* $\langle A, u \rangle$, where $A = A_1 \times A_2 \times \dots \times A_n$ is the space of action profiles and $u : A \rightarrow \mathbb{R}^n$ is the payoff function. Payoffs are in monetary units. A set G comprises the universe of underlying games. Let $\mathcal{A} \equiv \cup\{A \mid \langle A, u \rangle \in G\}$.

Before choosing actions in the underlying game, players have the opportunity to communicate and form contracts that direct an external enforcer to compel monetary transfers between them. This interaction is restricted to a set of bilateral relationships given by a fixed undirected and irreflexive network $L \subset N \times N$, meaning players i and j can communicate if and only if $(i, j) \in L$. Contracting by larger groups of agents is not possible. Contracting takes place via a protocol that I call the *contracting institution*, described formally below.

The outcome of the underlying game is fully verifiable. Therefore contracts can specify transfers as a function of the entire action profile $a \in \mathcal{A}$ that is played.⁷

Definition 1: The **contract** for a pair of players (i, j) is a mapping $m^{ij} : \mathcal{A} \rightarrow \mathbb{R}_0^n(i, j)$, where $\mathbb{R}_0^n(i, j) \equiv \{t \in \mathbb{R}^n \mid t_i + t_j = 0 \text{ and } t_k = 0 \text{ for } k \neq i, j\}$ is the set of n -vectors describing transfers between i and j .

Realistically, contracting partners can condition transfers between them on the productive actions taken by third parties in the underlying game, but their contract may not impose transfers on third parties.

The *grand game* comprises interaction in the contracting phase followed by choices in the underlying game, which I will also call the *production phase*. Grand-game payoffs are given by the vector $u(a) + M(a)$, where $M(a) \equiv \sum_{i,j \in N; i < j} m^{ij}(a)$. Restricting to $i < j$ in this expression avoids double counting, since m^{ij} and m^{ji} refer to the same contract. Note that M maps feasible outcomes of the underlying games, \mathcal{A} , to the set of *balanced transfers* $\mathbb{R}_0^n \equiv \{t \in \mathbb{R}^n \mid t_1 + t_2 + \dots + t_n = 0\}$.

2.2 Contracting institution and design problem

The contracting institution comprises a communication protocol and a description of how the players' messages shall be interpreted by the external enforcer as contracts formed between pairs of players. It can be described formally as an extensive game form with costless messages that map to a contract m^{ij} for every pair of players (i, j) .

I consider the problem of *designing* a contracting institution that, once fixed, must serve for every underlying game in G and every network in a given set \mathcal{L} . The welfare goal is to achieve efficient outcomes, which means that for any given underlying game $\langle A, u \rangle$ and network $L \in \mathcal{L}$, there is an efficient equilibrium of the grand game. In the case of L being

⁷Although the external enforcer can recognize all elements of \mathcal{A} , she does not observe which underlying game is played and therefore cannot paternalistically impose transfers to induce behavior in furtherance of any particular welfare objective.

connected (the setting we will focus on), due to transferable utility and balanced transfers, the equilibrium outcome is efficient if and only if the action profile played in the underlying game maximizes $\sum_{i \in N} u_i(a)$.

The design problem is constrained in two ways. First, the enforcement system allows only for contracting that is voluntary and independent across relationships. Second, the institution is limited to private contracting between only the pairs of players linked in the network. A novel aspect of the latter constraint is that it varies with the network.

I limit attention to the class of game forms in which the players simultaneously send messages to each other in discrete rounds $\underline{r}, \underline{r} + 1, \underline{r} + 2, \dots, \bar{r}$, where \underline{r} and \bar{r} are arbitrary integers.⁸ There is no discounting. A public random draw ϕ occurs after round \bar{r} , and contracts can be conditioned on ϕ . Let Φ denote the space of public draws.

To represent that contracting takes place privately in bilateral relationships, the contracting institution is assumed to have the following structure: For each $r \in \{\underline{r}, \dots, \bar{r}\}$ and $i \in N$, player i 's action in round r of the contracting phase is a vector of messages $d_i^r = (\lambda_{ij}^r)_{j \neq i}$, where λ_{ij}^r is the message player i sends to player j . Each player observes only the messages she sent to and received from the other individuals, not any messages exchanged between other players. Denote by $h_{ij} = (\lambda_{ij}^{\underline{r}}, \lambda_{ij}^{\underline{r}+1}, \dots, \lambda_{ij}^{\bar{r}})$ the sequence of messages from player i to player j , and let $h_{ij}^r = (\lambda_{ij}^{\underline{r}}, \dots, \lambda_{ij}^r)$ denote the sequence through any given round r .

To represent that contracting is independent across relationships, we require two things. First, the set of feasible messages that player i can send to player j in round r depends only on the messages exchanged earlier between players i and j , and so can be written as $\Lambda_{ij}^{\underline{r}}$ for the first round and $\Lambda_{ij}^r(h_{ij}^{r-1}, h_{ji}^{r-1})$ for $r > \underline{r}$. Assume that a special *null message* $\underline{\lambda}$, meaning silence, is always feasible.

Second, the recognized contract between players i and j is a function of only the sequence of messages sent between them as well as the random draw ϕ . Thus, letting H_{ij} denote the feasible sequences (h_{ij}, h_{ji}) of messages between players i and j , their contract m^{ij} is the output of some function $\mu^{ij}: H_{ij} \times \Phi \rightarrow \mathcal{M}^{ij}$, where $\mathcal{M}^{ij} \equiv \{m^{ij}: \mathcal{A} \rightarrow \mathbb{R}_0^n(i, j)\}$ denotes the set of feasible contracts for players i and j . Also, since μ^{ij} and μ^{ji} are the same contract, $\mu^{ij}(h_{ij}, h_{ji}, \phi) = \mu^{ji}(h_{ji}, h_{ij}, \phi)$ is required.

To represent that contracting is voluntary, assume that player i can decline to contract with player j by sending the null message $\underline{\lambda}$ to player j in every round. That is, we have $\mu^{ij}((\underline{\lambda}, \underline{\lambda}, \dots, \underline{\lambda}), h_{ji}, \phi) = \underline{m}$ for all h_{ji} , where \underline{m} is the *null contract* that specifies $\underline{m}(a) = 0$ for every $a \in \mathcal{A}$.

To review, a contracting institution specifies integers \underline{r} and \bar{r} ; the space of the public random draw Φ ; the public draw distribution; message spaces Λ_{ij}^r for all $i, j \in N, i \neq j$, and $r \in \{\underline{r}, \dots, \bar{r}\}$; and the function μ^{ij} for each pair of players $i, j \in N$, such that μ^{ij} satisfies the assumption regarding voluntary contracting.

Finally, to represent how the contracting institution is constrained by the network L , we layer on the assumption that, for every pair $(i, j) \notin L$, players i and j are restricted to send each other the null message. In other words, pairs of players that are not linked are unable to communicate directly, and their contract will be null.

⁸We could normalize \underline{r} up front, but the general numbering will be convenient for organizing components of the contracting institution in the proof of the main result.

2.3 Equilibrium concept and implementation

Because each player does not observe messages sent between pairs of other players in the contracting phase, there is a great deal of asymmetric information in the grand game. For instance, at the end of the contracting phase, each player knows only the contracts he created with other individuals; he does not observe the contracts formed in other relationships.

To impose the stringent requirement of full consistency for belief updating on the plethora of information sets in the grand game, I analyze behavior using the concept of sequential equilibrium (Kreps and Wilson 1984). Beliefs at information sets are expressed in terms of appraisals (Watson 2017)—probability distributions over strategy profiles—which is convenient for the kind of game studied here. To keep the grand game finite, as required to apply sequential equilibrium, I assume that \mathcal{A} is finite and look only at finite contracting institutions. I assume also that the universe of underlying games G is finite.⁹

Definition 2: Fix n , G , and \mathcal{L} . A given contracting institution is said to **implement efficient outcomes** if for every underlying game $\langle A, u \rangle \in G$ and every network $L \in \mathcal{L}$, there is a sequential equilibrium of the grand game in which the outcome is efficient.

Our organizing question, on whether LDL externalities can be internalized through rational decentralized contracting, can be viewed as a policy problem. We have a setting in which the external enforcement technology can verify messages sent in the contracting phase, the public draw, and the outcome of the underlying game. The enforcer does not observe which underlying game is played or the network. The contracting institution is constrained to allow for only private, independent, and voluntary contracting. Achieving efficient implementation would allow us to conclude that, with the right kind of contracting institution, LDL externalities can always be internalized, whatever are the underlying game and network.

3 Barriers Illustrated by Simple Examples

In this section, I describe simple examples to demonstrate barriers to efficient contracting. It is worth noting at the outset that, regardless of the underlying game and network, if every player has at least one link, then there exists a set of feasible contracts such that $\langle A, u + M \rangle$ has an efficient Nash equilibrium a^* . Thus, the main issue is whether the players would have the incentive to form such contracts and then choose a^* .

3.1 Disconnected network

Consider a setting with two information-technology firms, called players 2 and 3, whose operations have potential synergies. Player 2 has an existing relationship with a supplier called player 1, and player 3 has an existing relationship with a supplier called player 4. Players 2 and 3 would jointly benefit if their suppliers create specialized inputs in service of the synergy, but this would require the suppliers to divert resources from other projects and

⁹One can allow G to be infinite by imposing bounds on underlying game elements, but extending the analysis in this way does not generate further insights.




Figure 2: Example of two firms with suppliers, disconnected network.

reduce their ability to compete in an unrelated market. For simplicity, suppose only players 1 and 4 have choices to make in the underlying game and they both have action space $\{0, 1\}$. 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.

Imagine that only two pairs of players can communicate and contract, due to legal or physical barriers: Player 2 can contract with its supplier, and player 3 can contract with its supplier. The contracting network is shown on the left side of Figure 2. Note that every player is in at least one contractual relationship, but the network is discontinuous.

In this example, efficiency does not equate to maximizing the sum of the player's payoffs, since transfers cannot be made between the two disconnected components of the network. Yet any outcome in which action profile $(0, 1, 1, 0)$ is played in the production phase is inefficient, because we can find contracts m^{12} and m^{34} that, along with the choice of action profile $(1, 1, 1, 1)$, would give every player a strictly higher payoff in the grand game.¹⁰ Unfortunately, regardless of the contracting institution, action profile $(0, 1, 1, 0)$ is chosen with certainty in every sequential equilibrium of the grand game.

To see why, consider the incentives of players 1 and 2. Because there is no communication between them and the other players, whether they deviate from equilibrium in the contracting phase will not affect player 4's choice of a_4 in the production phase. Under the null contract with player 2, player 1 rationally must choose $a_1 = 0$ in the productive phase, for it dominates $a_1 = 1$ in the underlying game. Since the joint payoff for players 1 and 2 is strictly higher with $a_1 = 0$ than with $a_1 = 1$, at least one of these players strictly prefers to deviate from an equilibrium that would have player 1 choose $a_1 = 1$ with positive probability, by being silent throughout the contracting phase to get the null contract.

The following, more general, statement is Result 1 in Appendix B.1: *For any given $n \geq 4$, there exists an underlying game and a network in which every player has a link and, regardless of the contracting institution, there is no sequential equilibrium of the grand game in which an efficient action profile of the underlying game is played with positive probability.* Hereinafter, I limit attention to settings in which network L is connected.

3.2 Collaboration agreement

Next consider a case with four players in the same roles as in the previous example, but suppose that players 2 and 3 can communicate and contract, in addition to each contracting with her supplier. The contracting network is shown on the left side of Figure 3. In this example,

¹⁰The joint value for players 1 and 2 is 8 with action profile $(1, 1, 1, 1)$ and 6 with action profile $(0, 1, 1, 0)$, and likewise for players 3 and 4.




Figure 3: A collaboration-agreement example.

players 2 and 3 may seek to exploit their operational synergies by forming a *collaboration agreement*—a contract that governs their interaction and may also contain provisions having to do with their suppliers’ productive actions.¹¹

As before, suppose that only players 1 and 4 have choices to make in the underlying game. Assume they both have action space $\{0, 1, 2\}$. The actions of players 2 and 3 are fixed at $a_2 = a_3 = 1$. Payoffs in the underlying game are given by the table in Figure 3.

Because the network is connected, efficiency requires play of an action profile in the underlying game that maximizes the sum of the players’ payoffs. The efficient action profile is $a^* = (1, 1, 1, 1)$ and the Nash equilibrium of the underlying game is $\underline{a} = (0, 1, 1, 0)$. Since each supplier’s productive action affects the payoff of every other player, LDL externalities extend throughout the network.

Let us explore what might be needed for efficient contracting. First note that, regardless of the contracting institution, efficiency requires transfers in the collaboration agreement between players 2 and 3 to depend on their suppliers’ productive actions. Suppose, to the contrary, that the contract between players 2 and 3 does not condition their transfer on player 1’s action a_1 . The contract still 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 then player 2 can strictly gain by declining 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), with private contracting required here. In the first round, players simultaneously offer con-

¹¹Collaboration agreements are common in high-tech industries. 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” technology standards.

tracts separately to each of their linked partners. In the second round, players simultaneously choose at most one contract to accept in each of their relationships, selecting between the contracts offered by the two linked players. If in a given relationship, the same contract is accepted by both players, then this contract goes into force; otherwise, they have the null contract specifying zero transfers.

With this contracting institution, the grand game has no efficient equilibrium. To see why, suppose there is an efficient equilibrium, and we will find a contradiction. Pairs (1, 2) and (3, 4) must form contracts that induce players 1 and 4 to select $a_1 = a_4 = 1$. Suppose player 1 were to deviate in the second round by declining to accept any contract with player 2 and then choose $a_1 = 2$. This deviation is not observed by player 4, who still forms a contract with player 3 and selects $a_4 = 1$. The deviation ensures player 1 of a payoff of at least 9, which becomes 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.¹²

The same logic holds for similar contracting institutions with more rounds, with different pairs of players having to contract at different dates. Say that a given contracting institution exhibits *dated commitment* if, for every pair (i, j) , there is a round \hat{r}^{ij} such that these players can communicate through round \hat{r}^{ij} only, and their contract is null if either sends message λ to the other in round \hat{r}^{ij} . A special case of dated commitment, with $\hat{r}^{ij} = \bar{r}$, allows players to unwind their contractual commitments unilaterally at the end of the contracting phase. Result 2 in Appendix B.1 states that, *for any given $n \geq 4$, there exists an underlying game and a connected network such that, for every contracting institution exhibiting dated commitment, sequential equilibrium outcomes of the grand game are bounded away from efficient.*

Hence, to implement efficient outcomes, a contracting institution must give each linked pair of players the opportunity to make a contractual commitment and then continue to communicate and possibly modify the contract. Further, enough time is needed for players to make adjustments in response to their experiences in other relationships, as stated as Result 3 in Appendix B.1: *For any given $n \geq 3$, there exists an underlying game and a connected network such that, for every contracting institution with strictly fewer than $n - 1$ contracting rounds, sequential equilibrium outcomes of the grand game are bounded away from efficient.*

3.3 Collaboration and a peripheral beneficiary

The next example adds an element to the incentive issues discussed in the previous subsection: contracting with a beneficiary at the periphery of the network who is not active in the underlying game. Consider a variant of the collaboration-agreement example with the same connected 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 4. The efficient action profile is $a^* = (1, 1, 1, 1)$ and the Nash equilibrium of the underlying game is $\underline{a} = (0, 1, 0, 1)$.

¹²If player 2 refuses to contract with players 1 and 3, then player 1 could not rationally choose $a_1 = 1$ (it is dominated in the underlying game) and player 2's payoff is at least 2. The same reasoning holds for player 3.




Figure 4: An example with a peripheral beneficiary.

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, the equilibrium contract for the pair (3, 4) must specify a transfer of at least 2 to player 3 when a^* is played. For efficiency the players must have the incentive to establish contracts with these properties and which motivate players 1 and 3 to choose the high action in the production phase. It remains to be seen whether, depending on the contracting institution, there is an equilibrium of the grand game in which such contracts are written.

4 Efficient Implementation

The examples and results presented in the previous section suggest that, to achieve efficient implementation through decentralized, private contracting, a successful contracting institution must facilitate sequential contracting, in some manner encouraging players to initiate their contractual commitments early in the contracting phase while also allowing them to later adjust the contracts with some partners in response to their experience with other partners. Moreover, it must be flexible, giving the players sufficient scope to handle any network and underlying game. The implementation problem is therefore nontrivial. Despite the hurdles, the answer to our possibility question is positive. The main result is stated next.

Theorem: *Take as given any integer $n \geq 2$ and any finite set G of finite n -player underlying games. Let \mathcal{L} be the set of all connected networks in $N \times N$. There exists a contracting institution (representing private, independent, and voluntary contracting) that implements efficient outcomes.*

The proof of the Theorem has two parts. The first is to identify a contracting institution that will implement efficient outcomes. This is done in Section 4.1, except for specifying a probability number ε and set of feasible contracts $\overline{\mathcal{M}}^{ij}$ for every pair of players (i, j) . The second part is to construct an efficient equilibrium of the grand game for every connected network and underlying game. This part involves organizing classes of personal histories, specifying strategies and beliefs, and checking sequential rationality. Subsections 4.2-4.4 develop key elements of the equilibrium construction and also derive ε and the contracts $\overline{\mathcal{M}}^{ij}$.

is assumed to contain. Subsection 4.5 provides an overview of the construction, explaining what happens on the equilibrium path and after two sample deviations. The precise details of the construction are laid out in Appendix A.2. Appendix A.1 provides proofs of lemmas presented in subsections 4.2 and 4.3.

The equilibrium construction is formidable because of the large number of information sets in the grand game and because players have distinctly different information about actions taken previously. Further, the construction must be done generically, requiring numerous organizational steps. In fact, a full equilibrium construction is not undertaken. Rather, existence is established using a novel *partial-construction* method, specifying strategies and beliefs for a subset of information sets. An existence result, reported in Watson (2023), then guarantees that the partial construction extends to a fully specified sequential equilibrium.

Before proceeding to the technical details, it may be helpful to preview some features of the proof. In equilibrium, the players are endogenously partitioned into a cluster of *active* players whose incentives in the production phase require shaping by external enforcement (the *core group*) and the set of peripheral *passive* players, together forming a minimally connected subnetwork. The players coordinate to make contractual arrangements sequentially, starting in the periphery and working toward the core. Not all linked pairs form contracts. Players also have the opportunity to send cancellation messages in later rounds.

Contractual arrangements provide for primary contracts that turn out to be in force on the equilibrium path, as well as secondary contracts triggered by unilateral cancellations. Primary contracts work together to force play of an efficient action profile a^* in the underlying game, and the primary contract chosen by a pair (i, j) has *assurance penalties* by which player i guarantees that others on her side of the network will do their part. Cancellation either discharges a player from her contractual obligations or forces play of a particular action in the underlying game, depending on the random draw ϕ , resulting in an active player departing from a^* in the underlying game with positive probability under equilibrium beliefs.

Thus, if player j cancels or otherwise disrupts contracting with player i , then player i expects players on j 's side of the network to depart from a^* , making player i liable to pay assurance penalties in her other contractual arrangements, which she then has the incentive to cancel because cancellation penalties are small compared to assurance penalties. Escalating cancellation penalties deter late cancellations.

Finally, no player wants to be the first to disrupt a contractual relationship because doing so leads to a wave of cancellations and eventual play of an inferior action profile in the underlying game. In fact, after such a deviation, in a high-probability realization of ϕ , in the production phase players end up coordinating on a Nash equilibrium of the underlying game. By construction, this makes any deviating player worse off.

4.1 Featured contracting institution

The proof features what I shall call the *SCO contracting institution*, where SCO stands for Sequential Contracting and Options. The institution is defined as follows.

Let $\underline{r} = 1 - n$ and $\bar{r} = n - 2$, so that there are $2n - 2$ rounds of messages. Set $\Phi = \{0, 1, \dots, n\}$. The probability of $\phi = 0$ is set to $1 - n\varepsilon$, and for each $i \in N$ the probability of $\phi = i$ is set to ε , where ε is specified in Subsection 4.2 below. In rounds \underline{r}

through 0, each pair of contracting partners engages in a recurring Nash-demand protocol to determine what I call their *conditional arrangement*, which specifies their contract for the underlying game as a function of ϕ and whether either player cancels in rounds 1 through \bar{r} .

Consider any pair of players (i, j) . The feasible conditional arrangements for this pair, denoted by C^{ij} , is the set of functions mapping $\{0, (1, i), \dots, (\bar{r}, i), (1, j), \dots, (\bar{r}, j)\} \times \Phi$ to $\overline{\mathcal{M}}^{ij}$, where $\overline{\mathcal{M}}^{ij}$ is a finite subset of \mathcal{M}^{ij} that is assumed to contain the contracts identified in Subsection 4.3 below and includes the null contract, but is otherwise arbitrary. For a given sequence (h_{ij}^r, h_{ji}^r) , if there is a round $\ell \leq r$ in which $\lambda_{ij}^\ell = \lambda_{ji}^\ell = c^{ij}$ for some $c^{ij} \in C^{ij}$, then let us say (h_{ij}^r, h_{ji}^r) *records that players i and j made conditional arrangement c^{ij}* .

Here is an inductive definition of the set of feasible messages from player i to player j in each round: In round \underline{r} the set is defined as $\Lambda_{ij}^{\underline{r}} \equiv C^{ij} \cup \{\underline{\lambda}\}$. For $r \in \{\underline{r} + 1, \dots, 0\}$, if $(h_{ij}^{r-1}, h_{ji}^{r-1})$ records that players i and j made a conditional arrangement in an earlier round, then $\Lambda_{ij}^r(h_{ij}^{r-1}, h_{ji}^{r-1}) \equiv \{\underline{\lambda}\}$. Otherwise, $\Lambda_{ij}^r(h_{ij}^{r-1}, h_{ji}^{r-1}) \equiv C^{ij} \cup \{\underline{\lambda}\}$. That is, in words, once these players have made a conditional arrangement, then they are restricted to silence with each other until round 1.

For $r \in \{1, \dots, \bar{r}\}$, if $(h_{ij}^{r-1}, h_{ji}^{r-1})$ records that players i and j made a conditional arrangement earlier, and if $\lambda_{ij}^\ell = \lambda_{ji}^\ell = \underline{\lambda}$ for $\ell \in \{1, \dots, r-1\}$, then $\Lambda_{ij}^r(h_{ij}^{r-1}, h_{ji}^{r-1}) \equiv \{\text{“cancel”}, \underline{\lambda}\}$. Otherwise, $\Lambda_{ij}^r(h_{ij}^{r-1}, h_{ji}^{r-1}) \equiv \{\underline{\lambda}\}$. That is, if players i and j did not make a conditional arrangement, then they are restricted to silence in rounds 1 through \bar{r} . If they made a conditional arrangement, then they each have the option of sending the cancel message to the other, until one or both of them do so.

For any given sequence (h_{ij}, h_{ji}) of messages between players i and j through round \bar{r} , if there is a round ℓ at which $\lambda_{ij}^\ell = \text{“cancel”}$ and $\lambda_{ji}^\ell = \underline{\lambda}$, then let us say that (h_{ij}^r, h_{ji}^r) *records player i cancelling with player j in round ℓ* .

The function μ^{ij} is defined next. Consider any $c^{ij} \in C^{ij}$ and any sequence (h_{ij}, h_{ji}) . If (h_{ij}, h_{ji}) records that players i and j made conditional arrangement c^{ij} and does not record either player cancelling, then let $\mu^{ij}(h_{ij}, h_{ji}, \phi) \equiv c^{ij}(0, \phi)$ for every $\phi \in \Phi$. If (h_{ij}, h_{ji}) records that players i and j made conditional arrangement c^{ij} and records player i cancelling in some round ℓ , then define $\mu^{ij}(h_{ij}, h_{ji}, \phi) \equiv c^{ij}((\ell, i), \phi)$ for every $\phi \in \Phi$. Note that in the case of cancellation, the resulting contract can depend on the identity of the cancelling player and the round in which it occurred, as specified by the conditional arrangement.

4.2 Organizing elements of underlying games and networks

To specify ε along with various elements that will appear in the equilibrium constructions, a few special actions must be identified for each underlying game in G . Also, for every combination of a connected network and underlying game, a special subnetwork and payoffs will be identified. This subsection presents the relevant definitions. I use the following standard notation: For a subset of players $J \subset N$, $a_J \equiv (a_i)_{i \in J}$ denotes the vector of actions for these players, and $-i \equiv N \setminus \{i\}$. Every network in $N \times N$ considered hereinafter is assumed to be undirected and irreflexive, as we have assumed for L .




Figure 5: Generic example, active and passive players.

Focal elements for a given underlying game

For each underlying game $\langle A, u \rangle \in G$, we must define several elements, denoted by a^* , $\underline{\alpha}$, \underline{N} , \overline{N} , a^i , and \hat{a}_i^j for $i, j \in N$. They are defined as follows.

Let $a^* \in A$ be any efficient action profile, which maximizes the joint value $\sum_{i \in N} u_i(a)$, and let $\underline{\alpha} \in \Delta A$ be any Nash equilibrium in the underlying game. Call the underlying game *nontrivial* if its associated profile $\underline{\alpha}$ is inefficient, and *trivial* if it is efficient. We can restrict attention to the nontrivial case, for the trivial case is easy to handle as explained later.

For the given underlying game, define the sets of *passive players* \underline{N} and active players \overline{N} as follows. For any $J \subset N$, let $\Omega(J)$ contain each player i for whom a_i^* is a weakly dominant action in the underlying game, conditional on every player $j \in J \setminus \{i\}$ choosing a_j^* :

$$\Omega(J) \equiv \{i \in N \mid u_i(a_i^*, a_{-i}) \geq u_i(a'_i, a_{-i}) \text{ for every } a'_i \in A_i \\ \text{and } a_{-i} \in A_{-i} \text{ satisfying } a_{J \setminus \{i\}} = a_{J \setminus \{i\}}^*\}.$$

Then \underline{N} is defined as the largest set satisfying $\underline{N} = \Omega(\underline{N})$, and $\overline{N} \equiv N \setminus \underline{N}$.¹³ Figure 5 illustrates how, for a given underlying game, the set of players may be partitioned into the active and passive subsets, shown for the generic example of a network pictured in Figure 1 in the Introduction. In this example, there are four active players, represented by filled nodes in the right diagram; the passive players are depicted by open nodes.

In the underlying game, each passive player i optimally chooses a_i^* if she believes that the other passive players also select their efficient actions, regardless of what the active players select. For every active player i , we can find a profile $a^i \in A$ such that $a_{\underline{N}}^i = a_{\underline{N}}^*$, $a_i^i \neq a_i^*$, and a_i^i is a best response to a_{-i}^i in the underlying game. For each passive player i , let $a^i \equiv a^*$. Finally, for every $i, j \in N$, let \hat{a}_i^j be any best response for player i to a_{-i}^j in the underlying game, with $\hat{a}_i^i \equiv a_i^i$ specified in the case of $j = i$ and $\hat{a}_i^j = a_i^*$ in the case of i passive.

For example, consider the four-player setting shown in Figure 6, where players 2 and 3 are restricted to actions $a_2 = a_3 = 1$ in the underlying game and the payoffs, as a function of player 1's and player 4's actions, are shown in the table on the right. For this underlying

¹³ Ω is monotone, so one can calculate \underline{N} inductively by $\Omega^1 \equiv \Omega(\emptyset)$, $\Omega^{\ell+1} \equiv \Omega(\Omega^\ell)$, and $\underline{N} \equiv \Omega^n$.




Figure 6: Example of focal elements for an underlying game.

game, players 2 and 3 are trivially passive, players 1 and 4 are active, $a^* = (0, 1, 1, 1)$, $\underline{\alpha}$ is the pure-strategy profile $(0, 1, 1, 0)$, and $a^1 = (2, 1, 1, 1)$.

Keep in mind that a^* , $\underline{\alpha}$, \underline{N} , \overline{N} , $\{a^i\}_{i \in N}$, and $\{\hat{a}_i^j\}_{i, j \in N}$ all depend on the underlying game. For ease of notation, this dependence will not be made explicit hereinafter.

Global parameters

I next define two global parameters. Let γ be an arbitrary number satisfying $\gamma > 2|u_i(a)|$ for every $i \in N$, $a \in A$, and $\langle A, u \rangle \in G$. Let ε be any strictly positive number satisfying

$$\sum_{i \in N} u_i(a^*) > \sum_{i \in N} [(1 - n\varepsilon)u_i(\underline{\alpha}) + n\varepsilon\gamma], \quad (1)$$

and also $\varepsilon < 1 - \underline{\alpha}_i(a_i^*)$ for each player i for whom $\underline{\alpha}_i(a_i^*) < 1$, for every nontrivial underlying game in G . These numbers will be used to define penalties in the equilibrium contracts.¹⁴

Definitions pertaining to subnetworks

For any network $K \subset N \times N$, let the set of players with links be given by

$$N^K \equiv \{i \mid (i, j) \in K \text{ for some } j\}.$$

For players i and j , a *path from i to j* , if one exists, is given by a sequence $\{k^t\}_{t=1}^T \subset N$ that is distinct (no player appears multiple times) and satisfies $k^1 = i$, $k^T = j$, and $(k^{t-1}, k^t) \in K$ for all $t = 2, 3, \dots, T$. Players k^1, k^2, \dots, k^T are then said to be on this path from i to j , and T is called the path length.

A network K is called *minimally connected* (a tree) if for $i, j \in N^K$, there is exactly one path from i to j . Then for each $(i, j) \in K$ we can define:

$$\beta(i, j, K) \equiv \{k \in N^K \mid i \text{ is on the path from } j \text{ to } k\}.$$

In words, $\beta(i, j, K)$ is the set of players that are on “ i ’s side of network K ” relative to player j , and this includes player i . Note that minimally connected does not imply that connected for N , because $N^K \neq N$ is allowed.

For a given underlying game and any given minimally connected network K satisfying $\overline{N} \subset N^K$, let \hat{N}^K denote the set of *core players*, defined as the set of all active players and

¹⁴The numbers γ and ε exist because \mathcal{A} , G , and n are finite.




Figure 7: Generic example, minimally connected subnetwork K linking all active players.

those passive players that reside between active players in the network. Call each player $i \in N^K \setminus \hat{N}^K$ *peripheral*. Let us also define $\hat{\beta}(j, i, K) \equiv \beta(j, i, K) \cap \hat{N}^K$.

Figure 7 illustrates a minimally connected network K containing the active players, for our running generic example of a network and underlying game. The left diagram repeats the illustration of active and passive players in Figure 5. The right diagram shows a minimally connected subnetwork K linking players numbered 1-11, including the active players 2, 5, 7, and 8. In this example, the core group is $\hat{N}^K = \{2, 5, 6, 7, 8\}$. Players 1, 3, 4, 9, 10, and 11 are peripheral. Note that, for instance, $\beta(5, 6, K) = \{1, 2, 4, 5\}$ and $\hat{\beta}(6, 5, K) = \{6, 7, 8\}$.

Focal elements for a given underlying game and network

For each combination of an underlying game $\langle A, u \rangle \in G$ and allowed network L , we must define several elements that will be used in the equilibrium constructions: a special subnetwork K , a profile \underline{a}^{ik} for every $i \in \hat{N}^K$ and $k \in N$, and a value \underline{w}_i for every $i \in N^K$. In the equilibrium constructions, \underline{w}_i will be player i 's expected value of refusing to contract with everyone. In the case of more than two core players, it relates to a collection of action profiles that will be relevant in off-equilibrium path contingencies in which some players are induced to select their part of action profile a^k identified earlier, for each k .

Definition 3: Take as given an underlying game $\langle A, u \rangle \in G$ and a minimally connected network K satisfying $\bar{N} \subset N^K$. For every $i \in \hat{N}^K$ and $k \in N$, the ***ik*-default profile** $\underline{a}^{ik}(A, u, K)$ is constructed as follows:

- For each j such that $(i, j) \in K$ and $|\hat{\beta}(j, i, K)| > 1$, specify $\underline{a}_{j'}^{ik}(A, u, K) = a_{j'}^k$ for every $j' \in \hat{\beta}(j, i, K)$.
- For j such that $(i, j) \in K$ and $|\hat{\beta}(j, i, K)| = 1$, specify $\underline{a}_j^{ik}(A, u, K) = \hat{a}_j^k$.
- For $j \notin \hat{N}^K$, specify $\underline{a}_j^{ik}(A, u, K) = a_j^*$, completing the description of $\underline{a}_{-i}^{ik}(A, u, K)$.
- Finally, let $\underline{a}_i^{ik}(A, u, K)$ be any best response for player i to \underline{a}_{-i}^{ik} in the underlying game, subject to $\underline{a}_i^{ik}(A, u, K) \equiv \hat{a}_i^k$ in every case in which \hat{a}_i^k is a best response.¹⁵

¹⁵Thus $\underline{a}_i^{ik} = a_i^*$ is specified in the case of $i \in \underline{N}$. Further, $\underline{a}_i^{ik} = \hat{a}_i^k$ if there is a player j for which

For example, consider Figure 7 with $i = 5$ and $k = 8$. Then $\underline{a}_6^{ik}(A, u, K) = a_6^k = a_6^*$, $\underline{a}_7^{ik}(A, u, K) = a_7^k$, and $\underline{a}_8^{ik}(A, u, K) = a_8^k$, because $\hat{\beta}(6, 5, K) = \{6, 7, 8\}$ contains more than one active player. Likewise, $\underline{a}_2^{ik}(A, u, K) = \hat{a}_2^k$ because $\hat{\beta}(2, 5, K) = \{2\}$ contains exactly one active player. Every other player $j \neq i$ is passive and has $\underline{a}_j^{ik}(A, u, K) = a_j^*$.

In some of the equilibrium constructions, where K is the set of pairs that are supposed to establish non-null contracts, if a player $i \in \hat{N}^K$ deviates by refusing to contract with everyone, then later in the event of $\phi = k$, player i will believe that the other players will choose profile $\underline{a}_{-i}^{ik}(A, u, K)$ in the underlying game. Player i will choose $\underline{a}_i^{ik}(A, u, K)$ to best respond. If player i is passive, his choice will be a_i^* .

Definition 4: Take as given an underlying game $\langle A, u \rangle \in G$ and a minimally connected network K satisfying $\bar{N} \subset N^K$. For every $i \in N^K$, let **default payoff** $\underline{w}_i(A, u, K)$ be defined as follows:

- If $|\hat{N}^K| \leq 2$, then set $\underline{w}_i(A, u, K) \equiv u_i(\underline{\alpha})$.
- If $|\hat{N}^K| > 2$ and $i \in \hat{N}^K$, set $\underline{w}_i(A, u, K) \equiv (1 - n\varepsilon)u_i(\underline{\alpha}) + \varepsilon \sum_{k \in N} u_i(\underline{a}^{ik}(A, u, K))$.
- If $|\hat{N}^K| > 2$ and $i \notin \hat{N}^K$, set $\underline{w}_i(A, u, K) \equiv (1 - n\varepsilon)u_i(\underline{\alpha}) + \varepsilon \sum_{k \in N} u_i(\underline{a}^{\bar{k}k}(A, u, K))$, where player \bar{k} is the closest active player to player i .

Note that each profile $\underline{a}^{ik}(A, u, K)$ has all peripheral players and those outside the network choosing their part of action profile a^* . This implies that $\underline{w}_i(A, u, K)$ depends on K only through the paths between core players. For instance, enlarging K by inserting a link to an additional peripheral player does not change $\underline{w}_i(A, u, K)$ for all others in the network. The next definition will help identify the pairs of players that establish non-null contracts in the equilibrium constructions.

Definition 5: For a given underlying game $\langle A, u \rangle \in G$ and network K , call K **adequate** if every active player is in N^K , K is minimally connected, and

$$\sum_{i \in N^K} u_i(a^*) > \sum_{i \in N^K} \underline{w}_i(A, u, K). \quad (2)$$

Call K **essential** if it is adequate and no proper subset of K is also adequate.

Lemma 1: Take as given any underlying game in G and any connected network L . There exists a network $K \subset L$ that is essential.

For each underlying game in G and connected network L , let us select an essential subnetwork and refer to it as K , now fixed in relation to $\langle A, u \rangle$ and L . Likewise, let us write \underline{a}^{ik} and \underline{w}_i as the corresponding ik -default profile and default payoff for player i , suppressing the dependence on $\langle A, u \rangle$ and L . Lemma 1 guarantees existence; if there is more than one essential subnetwork, K is an arbitrary selection from the set of essential subnetworks.

$(i, j) \in K$ and $\hat{\beta}(j, i, K)$ contains all active players except player i (which means all other active players choose their part of a^k). Also, recall that $\hat{a}_i^i = a_i^i$ in the case of $k = i$.

Summary

To summarize this subsection, we have defined global parameters γ and ε , the latter which is part of the specification of the SCO contracting institution. For each underlying game, we have identified action profiles and individual actions a^* , \underline{a} , \underline{N} , \overline{N} , a^i , and \hat{a}_i^j for $i, j \in N$. Further, for every underlying game and connected network, we have identified an essential subnetwork K and an ik -default profile \underline{a}^{ik} for every $i \in \hat{N}^K$ and $k \in N$, and we have defined default value \underline{w}_i for every $i \in N^K$. All of these selections shall be fixed throughout the analysis hereinafter; their dependence on $\langle A, u \rangle$, L , and K will not be made explicit in the notation. Likewise, let us write $\beta(j, i)$ and $\hat{\beta}(j, i)$, dropping K as an explicit argument.

4.3 Feasible and featured contracts

I next specify contracts that will be featured in the equilibrium constructions and that $\overline{\mathcal{M}}$ is assumed to contain. An \tilde{a} -forcing contract imposes a large penalty on a contracting partner who deviates from \tilde{a} in the production phase. An \tilde{a} -assurance contract goes further by requiring a contracting partner to pay a large penalty for every deviation from \tilde{a} that takes place on this player's side of the network. For convenience, penalties will be sufficiently large for use with all of the underlying games in G and all networks. One of the penalties is

$$\psi \equiv \gamma \max\{(n-1), 1/\varepsilon\},$$

and others are multiples of γ . Denote by e^{ij} the vector in $\mathbb{R}_0^n(i, j)$ giving -1 to player i and 1 to player j .

The following definitions and constructive elements all are specific to a given underlying game $\langle A, u \rangle \in G$, connected network L , and their associated essential subnetwork K .

Definition 6: Consider any $(i, j) \in K$. The \tilde{a} -forcing contract with baseline transfer $\tau \in \mathbb{R}_0^n(i, j)$ is given by:

$$m(a) = \begin{cases} \tau + e^{ij}\psi & \text{if } a_i \neq \tilde{a}_i \text{ and } a_j = \tilde{a}_j \\ \tau + e^{ji}\psi & \text{if } a_i = \tilde{a}_i \text{ and } a_j \neq \tilde{a}_j \\ \tau & \text{otherwise} \end{cases}.$$

The \tilde{a} -assurance contract with baseline transfer $\tau \in \mathbb{R}_0^n(i, j)$ specifies:

$$m(a) = \tau + e^{ij}\psi \#\{k \in \beta(i, j) \mid a_k \neq \tilde{a}_k\} + e^{ji}\psi \#\{k \in \beta(j, i) \mid a_k \neq \tilde{a}_k\}.$$

In the equilibrium constructions, contracting partners will make conditional arrangements that commit them to a^* -assurance contracts if neither sends the cancel message to the other in rounds 1 through \bar{r} of the contracting phase. Various forcing contracts will be triggered by the cancel message, depending on who sends it and in what round. The next lemma identifies the assurance contracts that the players will coordinate on.

Define for each $i \in N^K$ the *periphery index* for this player, denoted by $\rho(i)$, as follows: In the case of i being peripheral, let $\rho(i)$ be the length of the path from i to the closest core player. For example, in Figure 7, we have $\rho(5) = 0$ because player 5 is a core player,

$\rho(3) = 1$ because player 3 is one branch away from nearest core player 6, and $\rho(11) = 2$ because player 11 is two branches away from nearest core player 8.

Lemma 2: *Take as given a nontrivial underlying game $\langle A, u \rangle \in G$ and its essential network K . There exists a collection of contracts $\{\check{m}^{ij}\}_{(i,j) \in K}$, with $\check{m}^{ij} = \check{m}^{ji}$ for $i \neq j$, such that the following conditions hold, where $\check{M} \equiv \sum_{(i,j) \in K, i < j} \check{m}^{ij}$.*

- a) *For each pair $(i, j) \in K$, $\check{m}^{ij} = \check{m}^{ji}$ is an a^* -assurance contract.*
- b) *For each player $i \in N^K$, $u_i(a^*) + \check{M}_i(a^*) > \underline{w}_i$.*
- c) *For each pair $(i, j) \in K$ satisfying $\rho(j) = \rho(i) + 1$, $u_i(a^*) + \check{M}_i(a^*) - \check{m}_i^{ij}(a^*) < \underline{w}_i$.*

A significant aspect of the first condition is the following for every player i in the production phase: Suppose that player i believes that the relationships in K have established these assurance contracts, the other relationships have the null contract, and a^*_{-i} will be chosen by the other players in the underlying game. Then player i rationally selects a^*_i . The second condition ensures that this outcome, with a^* played in the production phase, gives player i a higher payoff than her default payoff. The third condition states that the inequality is reversed for a player who, all else held fixed, would lose the contracted transfer from a peripheral partner who is further from the core group.

4.4 Target conditional arrangements

In the sequential equilibrium to be constructed for any given underlying game and network, a pair of players (i, j) will make a conditional arrangement, and therefore form a contract, if and only if $(i, j) \in K$. I denote by \check{c}^{ij} the conditional arrangement that they will coordinate on. The next definition identifies these equilibrium conditional arrangements. For accounting purposes, \check{c}^{ij} and \check{c}^{ji} refer to the same contract, so statements about players i and j as pair (i, j) also apply as pair (j, i) . Recall that $h_{ij} = (\lambda_{ij}^r, \dots, \lambda_{ij}^{\bar{r}})$ denotes the sequence of messages sent from player i to player j in the contracting phase.

Definition 7: *The set of target conditional arrangements, denoted by $(\check{c}^{ij})_{(i,j) \in K}$, is defined as follows.*

- $\check{c}^{ij}(0, \phi) = \check{m}^{ij}$ for every $\phi \in \Phi$.
- If $\rho(i) = \rho(j) = 0$, then $\check{c}^{ij}((r, i), 0) = \underline{m} + r\gamma e^{ij}$ and, for every $\phi \in \{1, 2, \dots, n\}$, $\check{c}^{ij}((r, i), \phi)$ is the a^ϕ -forcing contract with baseline transfer $r\gamma e^{ij}$.
- If $\rho(i) = \rho(j) - 1$, then $\check{c}^{ij}((r, i), \phi) = \underline{m} + (r - 1)\gamma e^{ij}$ for every $\phi \in \Phi$.
- If $\rho(i) = \rho(j) + 1$, then $\check{c}^{ij}((r, i), \phi) = \underline{m} + r\gamma e^{ij}$ for every $\phi \in \Phi$.

The meaning of the target provisional arrangements is straightforward. Every contracting pair arranges to form an a^* -assurance contract if neither cancels. If a core player cancels with another core player, then a penalty is paid by the first player, the players are otherwise released from their obligations in the high-probability event of $\phi = 0$, and they get the a^k -forcing contract in the low-probability event that $\phi \in N$. These provisions ensure that these

players depart from action profile a^* with positive probability following a cancellation. The cancellation penalty increases in r . If a cancellation occurs in a pair that includes a peripheral player, then a penalty is paid and the players are otherwise released from their obligations regardless of ϕ ; further, the penalty is zero in round 1 for a cancellation made “outward,” away from the core group.¹⁶

4.5 Overview of the equilibrium construction

The task from here is to show that an efficient sequential equilibrium of the grand game exists, for any given underlying game and connected network. This is accomplished in four steps. First, I define a subset of information sets in the grand game, denoted by Ξ , that includes the personal histories that will be on the equilibrium path as well as a number of critical off-path information sets. Second, I partially construct the sequential equilibrium by specifying the players’ beliefs at, and prescribed actions for, the information sets in Ξ . Each belief is also partial, describing what the player thinks has or will happen at only the information sets in Ξ . Third, I verify that the prescribed actions at the information sets in Ξ are optimal given the beliefs, regardless of the actions taken at the other information sets. Fourth, I apply the theorem of Watson (2023) to establish the existence of a sequential equilibrium that in an appropriate sense agrees with the partial construction.

The full-blown equilibrium construction is presented in Appendix A.2. In this subsection I provide intuition by highlighting aspects of the construction. Specifically, I describe the equilibrium path of play, as well as paths induced by two sample deviations, using as an illustration the generic example discussed before and reproduced in Figure 8 on page 22. Recall that the figure shows, for a given underlying game and network L , the role of each player as either active or passive and an adequate subnetwork K .

Play on the equilibrium path

On the equilibrium path, conditional arrangements are formed sequentially, starting at the extreme of network K and working inward to the core group. Specifically, pair $(i, j) \in K$ forms its conditional arrangement in round $-\max\{\rho(i), \rho(j)\}$. There are no cancellations, and each player selects her part of a^* in the underlying game.

For example, in a case illustrated by Figure 8, the following occurs on the equilibrium path, where unspecified actions are prescribed to be the default message λ (silence):

- In rounds $1 - n$ through -3 , all players are silent.
- In round -2 , players 9 and 11 send message $\check{c}^{9,11}$ to each other, and players 9 and 10 send message $\check{c}^{9,10}$ to each other, forming these conditional arrangements. Note that in these relationships, the outer peripheral player has periphery index 2.

¹⁶I thank Gorm Grønnevet for suggesting a version of the analysis in which cancellation penalties are used in the equilibrium construction. The statements in the definition above cover all of the possibilities for contracting partners and messages in rounds 1 through \bar{r} , so the target provisional arrangements are well defined. On network locations, the only possibilities are $\rho(i) = \rho(j) + 1$, $\rho(i) = \rho(j) - 1$, and $\rho(i) = \rho(j) = 0$.

- In round -1 , players 1 and 2 send message $\check{c}^{1,2}$ to each other, players 4 and 5 send message $\check{c}^{4,5}$ to each other, players 3 and 6 send message $\check{c}^{3,6}$ to each other, and players 8 and 9 send message $\check{c}^{8,9}$ to each other, forming these conditional arrangements. In these relationships, the outer peripheral player has periphery index 1.
- In round 0, players 2 and 5 send message $\check{c}^{2,5}$ to each other, players 5 and 6 send message $\check{c}^{5,6}$ to each other, players 6 and 7 send message $\check{c}^{6,7}$ to each other, and players 6 and 8 send message $\check{c}^{6,8}$ to each other, forming these conditional arrangements. These are relationships between core players, with periphery index 0.
- In rounds 1 through $n - 2$, players are silent, so there are no cancellations.
- In the production phase, a^* is played regardless of the random draw ϕ .

It should be clear that if play in the contracting phase proceeds as just described, then in the production phase each player i has the incentive to select a_i^* if she believes that the others will choose a_{-i}^* . Passive players have this incentive based on believing other passive players act the same way. Active players are bound by assurance contracts that penalize them heavily if they would deviate from a^* . Therefore, on-path incentives in the production phase are set, subject to working out the details of the beliefs.

Likewise, if upon reaching some round $r \geq 1$ the personal history of a player $i \in N^K$ is exactly as expected on the equilibrium path, then player i prefers not to cancel any contracts. For example, suppose players i and j formed their target conditional arrangement, and $\rho(i) \geq \rho(j)$. If player i were to cancel the contract formed with player j , then she must pay a cancellation penalty of at least γ , which exceeds any gain in the underlying game that the cancellation might induce. In the case of $\rho(i) < \rho(j)$, player i could cancel for free in round 1, but it will turn out that this also is of no benefit.

First sample deviation

Next I describe two examples of equilibrium play in the continuation after a unilateral deviation. In the first scenario, play occurs as on the equilibrium path until round 0, at which point player 5 sends message $\underline{\lambda}$ to every other player, effectively declining to form the target conditional arrangements with players 2 and 6. All other players choose their equilibrium actions in round 0. Here is what happens from round 1 in the continuation of the game:

Player 6, having observed the failure of player 5 to make the conditional arrangement with her, will believe that this was the only deviation. Thus, player 6 believes that player 5 formed his target conditional arrangements with players 2 and 4. Player 6 further believes that player 5 will cancel with players 2 and 4 in round 1, and that in the production phase player 5 will therefore be forced to select $a_5^5 \neq a_5^*$ in the event of $\phi = 5$, putting player 6 on the hook for the huge assurance penalty in her contracts with players 3, 7 and 8. Because the cancellation penalty is much lower than the assurance penalty, and because cancellation penalties increase over the rounds, in round 1 player 6 cancels with players 3, 7 and 8. Then player 8 is induced to have a similar belief and cancels with player 9 in round 2, and player 9 cancels with players 10 and 11 in round 3.

Likewise, following player 5's initial deviation, player 2 believes that player 5 formed his target conditional arrangements with players 4 and 6, will cancel these in round 1, and in




Figure 8: Generic example, adequate subnetwork K , equilibrium construction.

the production phase will be forced to select $a_5^5 \neq a_5^*$ when $\phi = 5$. Liable for an assurance penalty in her contract with player 1, in round 1 player 2 cancels with player 1.

Thus, player 5's initial deviation leads to a wave of cancellations though the network, resulting in cancelled or null contracts in all relationships. Although the players observe different things and have different beliefs about what happened in the contracting phase, in the production phase the players in N^K all think (correctly) that every relationship has either a cancelled or null contract. Recall that cancellations lead to contracts that specify only transfers that are constant in a in the high-probability event that $\phi = 0$. In the production phase in event $\phi = 0$, every player i believes that $\underline{\alpha}_{-i}$ will be chosen by the other players and player i rationally responds by choosing $\underline{\alpha}_i$. We see that by deviating in round 0, player 5's expected payoff becomes \underline{w}_5 , which from Lemma 2(b) is strictly less than what player 5 would get by adhering to the prescribed path.

Second sample deviation

In the second scenario, play occurs as on the equilibrium path until round -2 , at which point player 10 sends message $\underline{\lambda}$ to every other player, effectively declining to form the target conditional arrangement with player 9. All other players choose their equilibrium actions in round -2 , implying that players 9 and 11 formed their target conditional arrangement. Here is what happens from round -1 in the continuation of the game:

Having made the target conditional arrangement with player 11 but not with player 10, player 9 realizes that by ignoring player 10's deviation (forming the conditional arrangement with player 8 in round -1 , and continuing as on the equilibrium path), his payoff will be strictly less than \underline{w}_9 . This follows from Lemma 2(c) and that all players other than himself and player 10 will not detect any deviation from the prescribed path. Player 9 instead sends message $\underline{\lambda}$ to player 8 in round -1 , effectively declining to form the target conditional arrangement with her, and player 9 plans to cancel his contract with player 11 in round 1 when doing so is free. In turn, Player 8 is put in the same position and in round 0 sends message $\underline{\lambda}$ to player 6, declining to establish their target conditional arrangement.

Players 1-6 played as though on the equilibrium path through round 0. Player 6, upon receiving the default message from player 8 in round 0, believes that this message was the

first and only deviation from the equilibrium path. Player 6 believes further that player 8 will cancel her contract with player 9 and will choose $a_8^8 \neq a_8^*$ in the production phase when $\phi = 8$. This makes player 6 liable for an assurance penalty in her contracts with players 3, 5, and 7. As in the first scenario, in round 1 player 6 then cancels her contracts with these players, leading to a wave of cancellations that flows across the network.

At the end of the contracting phase, every contract is cancelled or null, the players in N^K all correctly think as much, and $\underline{\alpha}$ is played in the high-probability event of $\phi = 0$. Thus, by declining to contract with player 9 in round -2 , player 10's expected payoff becomes \underline{w}_{10} rather than the equilibrium value $u_{10}(a^*) + \dot{M}_{10}(a^*)$. From Lemma 2(b), she prefers not to deviate. Likewise, the other choices described above are rational, such as player 6 cancelling contracts in round 1.

Additional notes

The logic given in the two scenarios above is incomplete. The formal constructive proof provides the precise beliefs and behavior, and verifies sequential rationality, for all of the personal histories that would be encountered in the two scenarios above and all others that compose Ξ . Also included are information sets in which players have observed unilateral deviations that I will classify as *insignificant variations*. These are departures from the prescribed equilibrium-path actions that would have no material effect if the players ignore them and continue as on the equilibrium path.

Note also that the examples discussed in Section 3 are essentially special cases of the construction sketched above. In the collaboration-agreement setting shown in Figure 3, players 1 and 4 are active, $K = L$, and all players are in the core group \hat{N}^K . In equilibrium, each of the three contractual relationships establishes its target conditional arrangement in round 0.

In the setting with a peripheral beneficiary shown in Figure 4, players 1 and 3 are active and therefore the core group is $\hat{N}^K = \{1, 2, 3\}$. A monetary contribution from player 4 is needed to provide incentives to the others, so the adequate network K is the same as L . In equilibrium, players 3 and 4 form their conditional arrangement in round -1 , whereas the pairs (1, 2) and (2, 3) form theirs in round 0.

5 Elaboration and Discussion

5.1 Two extensions

We observed the importance of sequential contracting, with Result 3 showing that at least $n - 1$ rounds of contracting is needed for general efficient implementation. If we restrict attention to networks with bounded diameter (defined as the greatest distance between any two players), then a shorter contracting phase will suffice if the bound is small enough. Result 4 in Appendix B.1 states that *given $n \geq 3$, any finite set of underlying games, and any integer $\kappa \in [2, n]$, and letting \mathcal{L} be the set of all connected networks of diameter weakly less than κ , there is a contracting institution satisfying $\bar{r} - \underline{r} \leq 2\kappa - 2$ that implements efficient outcomes.*

Although an efficient equilibrium is featured in the Theorem, one might expect other equilibria to exist, achieving a variety of different outcomes, and this is indeed the case. Result 5 in Appendix B.1 is analogous to folk theorems in repeated games. Stating the result requires additional terminology and definitions that are not repeated here. Essentially the result shows that the SCO contracting institution with suitably large contract spaces supports multiple equilibria with a range of payoff vectors above underlying-game Nash-equilibrium values for core players. In the special case of no peripheral players, as illustrated in Figure 3, all feasible payoff vectors above $u(\underline{\alpha})$ can be approximately achieved by equilibria of the grand game. Thus, the present modeling exercise shares a theme of prior models of interactive contracts that produce folk-style results (for instance, Peters and Szentes 2012).

5.2 Notes about option contracts and penalties

I next comment on the interpretation of the SCO contracting institution. As defined, the institution has $2n - 2$ rounds of messages in the contracting phase and, following the exogenous random draw ϕ , the output of the institution for a pair of players (i, j) is their “contract” m^{ij} . A different, perhaps more realistic, interpretation is that the contracting phase comprises just the first n rounds of messages (rounds \underline{r} through 0) and the conditional arrangements are interpreted as contracts. The later rounds are then dates at which the players can exercise options in their individual contracts, through their continued communication along edges of the network. That is, the contract for a pair of players specifies a transfer as a function of communication in rounds 0 to \bar{r} , the random draw ϕ , and the verifiable action profile a .

On a related note, in the proof of the Theorem, penalties γ and ψ were chosen for convenience to suffice for all contracting pairs and underlying games, and therefore are 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 practical issue goes beyond the present modeling exercise.

5.3 Summary of technical and conceptual contributions

Presented here is the first general analysis of technological and institutional requirements for internalizing LDL externalities. The model’s noncooperative game-theoretic structure allows for a precise account of the production and enforcement technologies. Treating the contracting institution as a payoff-irrelevant design component allows one to identify properties of the underlying game, the degree of verifiability, and the set of feasible contracting partners that are sufficient for efficiency in a best-case scenario regarding equilibrium selection. The Theorem shows that, in the setting of private bilateral contracting, global verifiability of productive actions and a connected network of feasible contracting partners is enough.

The Theorem gives a distinctly different message than do prior analyses of specific applications featuring bilateral contracting in settings with multilateral productive interaction. For instance, inefficient outcomes are in some cases predicted by McAfee and Schwartz’s (1994) analysis of contracting in a star network, Segal’s (1999) study of similar settings,

Prat and Rustichini’s (2003) analysis of games played through agents, and De Fontenay and Gans’ (2014) model of contracting on a network with LDL externalities. The analysis here helps to show why efficiency is not reached in these other models. The first three assume a contracting institution with only one or two rounds, not allowing for sequential contract formation. Also, all but Prat and Rustichini (2003) effectively disallow contracting parties to condition transfers on others’ productive actions.¹⁷

The modeling exercise features novel steps to deal with significant analytical challenges. The design problem and equilibrium constructions are complex because both contracting and productive actions are modeled noncooperatively, there are many information sets and asymmetric information throughout the grand game, and the stringent requirements of sequential equilibrium is imposed. The general modeling framework is new to the literature, requiring fresh analysis including how elements of the equilibrium construction are organizing generically. Further, this paper is the first to employ the partial-equilibrium construction method for sequential equilibrium utilizing the theorem of Watson (2023).

It is worth expounding on the strengths of the fully noncooperative modeling approach taken herein, in comparison to the approach of cooperative matching theory and coalitional bargaining theory.¹⁸ By specifying payoffs as a function of an abstract set of contracts that the players form, these two other lines of research account for productive actions as though taken by an external enforcer. Further, contracting is analyzed using a cooperative stability concept. Without an explicit account of the player’s inalienable productive actions, one cannot 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. The noncooperative approach provides a foundation for distinguishing types of externalities and understanding what is required to internalize them. These points are elaborated with an example in Section B.2.

5.4 Implications for applications

The analysis herein does not exactly pin down either the manner in which contracting must take place to achieve efficient outcomes or the precise form of equilibrium contracts. However, the minor results presented in Appendix B.1 identify some of the necessary ingredients, namely (a) a connected network of contractual relationships, (b) contractual linkages in the form of transfers conditioned on the productive actions of third parties, (c) endogenous sequential contracting, and (d) commitment with opportunities for parties to adjust contracts based on their personal experience.

¹⁷McAfee and Schwartz (1994) look at both private contracting and public contracting. De Fontenay and Gans (2014) assume that disagreement between two parties induces their link to break, rendering them unable to contract, and that this is publicly observed (thus contracting is not entirely private). The other two technically don’t even have LDL externalities.

¹⁸In coalitional bargaining models, centralized contracting is possible because the grand coalition can form a contract. Subgroups 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).

The Theorem identifies additional elements that can be successfully employed, such as (a) endogenous sequencing of contractual commitments starting with passive, peripheral parties and ending with the core group of active parties; (b) limited options to cancel contracts; (c) assurance penalties that motivate play of efficient productive actions and engender waves of cancellation following disruptions; and (d) cancellation penalties that encourage parties to cancel contracts when vulnerable and discourage them from canceling late.¹⁹

The modeling exercise may help us recognize elements that support or deter efficient contracting in real settings. For example, in many collaboration agreements (as illustrated in Section 3.2), performance guarantees and cross-firm management arrangements establish linkages across contractual relationships (Bernstein and Peterson 2020). Sequential contract formation and option contracts are ubiquitous, notably in procurement and supply chains. For instance, in design-build competitions, bidders are typically teams of companies that will provide complementary products and services (such as architectural and construction firms), and a preliminary agreement is formed within each team before the eventual winning team negotiates a contract with the buyer.

Despite the emphasis here on attaining efficiency, the Theorem should not be regarded as a claim that efficient outcomes will always be reached in reality, but rather as a reference for applications and a benchmark for further theoretical analysis. Applications vary technologically and may not fit with the assumptions made here, with respect to production and enforcement technologies as well as the contracting institution. Moreover, even under favorable conditions, efficiency relies on the players coordinating to achieve not just an equilibrium but the right equilibrium from a potentially large set. When evaluating barriers to efficiency, it may be helpful to categorize examples in terms of prominent aspects of their networks and the structure of their underlying games. Classes of applications are briefly listed in Appendix B.2.

The theory may eventually provide input to the design of markets and enforcement systems, as they facilitate contract creation. An example is the set of legal and procedural rules for eminent domain, where cases typically involve a number of property owners and land-use externalities. Platforms that facilitate contracting in related markets such as in the health-care sector essentially set aspects of the contracting institution. Legal infrastructure and regulation determine verifiability and other aspects of the enforcement technology. Organizations may play a role in designing the contracting institution, such as when a procuring party (for instance, a municipality) sets the rules for a design-build competition.

5.5 Variations for Further Study

The general modeling platform may provide a good foundation for exploring theoretical variations. One category is to characterize the performance of alternative contracting institutions, such as ones that appear in real settings but may not implement efficient outcomes.

¹⁹Cancellation waves may remind one of contagious punishments in socially repeated games (Kandori 1992), but they have a different structure. In the latter, a player participates in a contagious punishment due to a shift in intertemporal trade-offs and expectations about play in future matches. A player is motivated to participate in an out-of-equilibrium cancellation wave because her contracts contain a “poison pill” that makes her vulnerable when others on her side of the network will depart from a^* in the production phase.

We could also ask whether there is a contracting institution that performs better than the one described here, by more strongly implementing efficient outcomes or by achieving distributional goals. Another question is whether bargaining power would interfere with attainment of efficient outcomes. The SCO institution gives individual players no appreciable bargaining power because negotiation takes place through simultaneous demands. It remains to be seen what can be achieved with, for instance, offer-response protocols. A related practical issue to explore further is whether efficient contracting requires options to adjust externally enforced transfers, as the SCO institution facilitates in rounds 1 through \bar{r} , or could be accomplished with a simpler institution. This is taken up in Appendix B.2, where I conclude that dynamic adjustment of externally enforced contracts appears to be needed in general.

A second category of conceptual variation relates to the technologies of production and external enforcement. One could consider variations in the extent of verifiability. For instance, in some applications, productive actions are only partially verifiable, and we could use general results on what can be implemented. Appendix B.2 provides an example with partial verifiability in which there exist contracts that would induce the players to choose the efficient productive action profile, and yet there is no efficient equilibrium of the grand game under any contracting institution. Another realistic form of limited verifiability is *local* rather than *global* verifiability—for example, where contracting parties can provide evidence of the productive actions that they and their contracting partners take, but not the actions that others choose. One can also explore settings in which aspects of contracts can be verified, allowing some form of contracts on contracts, which the parties may find useful when dealing with limited verifiability in the dimensions described above. Further, one could look at variations regarding observability, such as where contracting is publicly observed rather than private.²⁰

Additional directions for further research on the technical front include modeling alienable (contractually assigned) productive actions; dynamic production; multilateral contracting, as in Ellingsen and Paltseva (2016) in the noncooperative arena and Rostek and Yoder (2022) on the cooperative side; and endogenous contracting networks, where players invest to establish links.²¹

Further, note that the model presented here leaves out some institutional constraints, such 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.²² Finally, it would be helpful to explore ways of extending the Nash program into the present setting.

²⁰McAfee and Schwartz (1994) look at both private contracting and public contracting. De Fontenay and Gans (2014) assume that disagreement between two parties induces their link to break, rendering them unable to contract, and that this is publicly observed (thus contracting is not entirely private).

²¹The last topic overlaps with the literature on network-based production and games played on endogenous networks (surveys include Jackson and Zenou 2015, and Bramoullé and Kranton 2016). One could examine whether incentives to isolate and free-ride can be overcome by pressure to join a chain of relationships.

²²A 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.

6 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. The model points to four distinct barriers to inefficiency in practice: (1) suboptimal contracting institutions that, for instance, do not provide players with the opportunity to solidify or adjust contracts in sequence; (2) limited verifiability of productive actions, as demonstrated in Appendix B.2; (3) institutional rules or enforcement technologies that artificially limit the feasible space of contracts; and (4) coordination problems in equilibrium selection. The first item may also include problems related to exertion of bargaining power.

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 with 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).²³ By precisely accounting for the productive technology, enforcement technology, and contracting institution in a general way, the modeling platform developed here lends itself to further exploration in both abstract and applied directions.

A Appendix

A.1 Proofs of lemmas

The lemmas are restated and proved here.

Lemma 1: *Take as given any underlying game in G and any connected network L . There exists a network $K \subset L$ that is essential.*

Proof: Because L is connected and $N^L = N$, there exists a minimally connected network L' satisfying $N^{L'} = N$. By definition, $\gamma > \underline{w}_i(A, u, L')$ for all $i, k \in N$, and ε has been set to satisfy Inequality 1. Therefore L' is adequate. Because the space of networks is finite and the proper subset relation is transitive and irreflexive, there must exist a subset of L' (possibly L' itself) that is essential. \square

Lemma 2: *Take as given a nontrivial underlying game $\langle A, u \rangle \in G$ and its essential network K . There exists a collection of contracts $\{\check{m}^{ij}\}_{(i,j) \in K}$ such that the following conditions hold, where $\check{M} \equiv \sum_{(i,j) \in K, i < j} \check{m}^{ij}$.*

²³As 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.

- a) For each pair $(i, j) \in K$, $\check{m}^{ij} = \check{m}^{ji}$ is an a^* -assurance contract.
- b) For each player $i \in N^K$, $u_i(a^*) + \check{M}_i(a^*) > \underline{w}_i$.
- c) For each pair $(i, j) \in K$ satisfying $\rho(j) = \rho(i) + 1$, $u_i(a^*) + \check{M}_i(a^*) - \check{m}_i^{ij}(a^*) < \underline{w}_i$.

Proof: In this proof, I write $\underline{w}_i(A, u, \cdot)$ to show its dependence on the chosen subnetwork, because K will be compared to a further subnetwork K' .

Because K is essential, it is adequate and therefore satisfies Inequality 2. That K is connected implies the existence of baseline transfers $(\tau^{ij})_{(i,j) \in K}$ such that $u_i(a^*) + \sum_{\{j, k \mid (j, k) \in K, j < k\}} \tau_i^{jk} > \underline{w}_i(A, u, K)$ for every $i \in N^K$. To see this, note that summing the left side over all $i \in N^K$ yields the left side of Inequality 2 because the transfers are balanced in the set of contracting partners. For each pair $(i, j) \in K$, let \check{m}^{ij} be the a^* -assurance contract with baseline transfer τ^{ij} . Then conditions (a)-(b) hold.

Condition (c) also must hold. To see why, take any pair $(i, j) \in K$ such that $\rho(j) = \rho(i) + 1$. Because player j and all other players in $\beta(j, i)$ are peripheral, we can remove them from network K to form subnetwork K' that is minimally connected and contains every active player. Because K' is not adequate, we know that

$$\sum_{k \in N^{K'}} u_k(a^*) \leq \sum_{k \in N^{K'}} \underline{w}_k(A, u, K'). \quad (3)$$

Because the transfers are balanced and the only contracted transfer between players in $N^{K'}$ and players in $N^K \setminus N^{K'}$ is the transfer for pair (i, j) , we have that $\sum_{k \in N^{K'}} \check{M}_k(a^*) = \check{m}_i^{ij}(a^*)$. Adding $\sum_{k \in N^{K'}} \check{M}_k(a^*) - \check{m}_i^{ij}(a^*) = 0$ to the left side of Inequality 3 and rearranging terms yields

$$u_i(a^*) + \check{M}_i(a^*) - \check{m}_i^{ij}(a^*) + \sum_{k \in N^{K'} \setminus \{i\}} \left[u_k(a^*) + \check{M}_k(a^*) - \underline{w}_k(A, u, K') \right] \leq \underline{w}_i(A, u, K'). \quad (4)$$

As noted above, $\underline{w}_k(A, u, K') = \underline{w}_k(A, u, K)$ for every $k \in N^{K'}$. Using this with condition (b) implies $u_k(a^*) + \check{M}_k(a^*) > \underline{w}_k(A, u, K')$ for every $k \in N^{K'}$. Therefore the bracketed terms on the left side of Inequality 4 are strictly positive, implying condition (c). \square

A.2 Proof of the theorem: partial construction and existence

Consider any number of players n and finite set G of underlying games. Let the contracting institution be the SCO contracting institution defined in Section 4.1 with the distribution of ϕ as specified in Subsection 4.2 and the where $\overline{\mathcal{M}}^{ij}$ is a finite subset of \mathcal{M}^{ij} containing the contracts identified in Subsection 4.3 and the null contract.

To prove the Theorem, we must show that, for each underlying game $\langle A, u \rangle \in G$ and connected network L , there is an efficient sequential equilibrium of the grand game. This subsection describes the equilibrium constructions, which will utilize all of the elements

developed in subsections 4.2–4.4 including the essential network K , the featured contracts, and the target conditional arrangements.

For now, let us leave out the case in which the underlying game is trivial and also leave out the case in which $|\hat{N}^K| \leq 2$ (where $\bar{N} = \hat{N}^K$ is implied). The latter case requires a variation in the equilibrium construction that will be described at the end of this section. The former case is easy to handle and is also discussed at the end of this section.

Take as given a finite set of underlying games G and let the contracting institution be the SCO institution, with ε and γ defined in Subsection 4.2 and $\bar{\mathcal{M}}$ defined to be any finite set of contracts that includes those described in Subsection 4.3 as well as the null contract. Various other elements defined in subsections 4.2–4.4 will be referenced below.

Let I denote the set of information sets (personal histories) in the grand game. This is quite a large set, with a lot of overlapping private information. Constructing a sequential equilibrium requires us to specify the belief and action choice at every information set. The system of beliefs must be fully consistent and the strategies sequentially rational.²⁴

Rather than describe the complete equilibrium strategies, I will specify the actions to be taken at a number of key information sets denoted by Ξ , including all that will be on the equilibrium path and some that will be off the equilibrium path. I also will specify the beliefs at these information sets about the actions taken at the other information sets in Ξ . I will show that the actions specified for Ξ are sequentially rational regardless of choices made at the other information sets. Then I will find a specification of fully mixed strategies for Ξ that support the specified beliefs and satisfy the conditions needed to apply the theorem of Watson (2023), which guarantees the existence of a sequential equilibrium of the entire grand game that coincides on Ξ with the construction here.

For every pair $(i, j) \in K$, define $r^{ij} \equiv -\max\{\rho(i), \rho(j)\}$. This will be the round in which the pair (i, j) is supposed to form their conditional arrangement. Recall that $h_{ij} = (\lambda_{ij}^r, \dots, \lambda_{ij}^{\bar{r}})$ denotes the sequence of messages that player i sends to player j in the contracting phase. The equilibrium prescribed path of play is described next.

Definition 8: For each pair of players $(i, j) \in K$, the **prescribed message sequence** is defined by $\lambda_{ij}^{r^{ij}} = \lambda_{ji}^{r^{ij}} = \check{c}^{ij}$, and $\lambda_{ij}^r = \lambda_{ji}^r = \underline{\lambda}$ for each round $r \neq r^{ij}$. For each pair of players $(i, j) \notin K$, the prescribed message sequence is $\lambda_{ij}^r = \lambda_{ji}^r = \underline{\lambda}$ for every round $r \in \{\underline{r}, \dots, \bar{r}\}$.

That is, players i and j who are supposed to contract are prescribed to send each other the null message until round r^{ij} , send each other message \check{c}^{ij} in round r^{ij} to form this conditional arrangement, and send the null message to each other thereafter. Players who are supposed to not contract, or who are not linked, are prescribed to send each other the null message in every round.

Definition 9: The **prescribed path** of play is for the players to send their prescribed message sequences to each other in the contracting phase and then select a^* in the production phase regardless of ϕ .

²⁴Full consistency rules out a variety of beliefs such as the following. Player i , upon seeing a surprise message from player j , concludes that player k has deviated, in a setting in which information about k 's supposed deviation could not have reached player j .

Note that, in the prescribed path, players linked in K make conditional arrangements that without cancellation will lead to the set of contracts $\{\check{m}^{ij}\}_{(i,j) \in K}$ identified by Lemma 2. It is clear that a^* is a Nash equilibrium of the induced game $\langle A, u + \check{M} \rangle$, so if the players reach the production phase on the equilibrium path then it is rational for each of them to choose her part of a^* . The difficulty from here is in formulating beliefs and behavior for off-equilibrium-path contingencies, demonstrating that players do not have the incentive to deviate, and showing that the beliefs are fully consistent.

Terminology for key information sets

To describe the key information sets, some additional terminology will be helpful. We start with classifications of the sequence of messages sent between a pair of players $(i, j) \in K$. The first definition below describes message sequences that conform to the prescribed path except for “insignificant” variations, where the target conditional arrangement was made in round r^{ij} , but in each earlier round one (at most) of the players sent a non-null message.

Definition 10: For $(i, j) \in K$, say that (h_{ij}, h_{ji}) is the *prescribed message sequence except for insignificant variations* if $\lambda_{ij}^{r^{ij}} = \lambda_{ji}^{r^{ij}} = \check{c}^{ij}$, for each $r < r^{ij}$ either $\lambda_{ij}^r = \underline{\lambda}$ or $\lambda_{ji}^r = \underline{\lambda}$ or both, and $\lambda_{ij}^r = \lambda_{ji}^r = \underline{\lambda}$ for each round $r > 0$. For $(i, j) \notin K$, say that (h_{ij}, h_{ji}) is the *prescribed message sequence except for insignificant variations* if for every $r < 1$ either $\lambda_{ij}^r = \underline{\lambda}$ or $\lambda_{ji}^r = \underline{\lambda}$ or both.

Note that, in the definition above, it is not necessary to state conditions for $r \in \{r^{ij} + 1, \dots, 0\}$ in the case of $(i, j) \in K$ or conditions for $r \geq 1$ in the case of $(i, j) \notin K$ because, given the other conditions, the players would be restricted to silence in these rounds. The next definition refers to bilateral message sequences in which a player has unilaterally blocked formation of a target conditional arrangement.

Definition 11: For any ordered pair $(i, j) \in K$, the *ij-decline sequence* is defined by: $\lambda_{ji}^{r^{ji}} = \check{c}^{ij}$, $\lambda_{ij}^{r^{ij}} = \underline{\lambda}$, and $\lambda_{ij}^r = \lambda_{ji}^r = \underline{\lambda}$ for every $r \neq r^{ij}$. Say that (h_{ij}, h_{ji}) is an *ij-decline sequence except for insignificant variations* if $\lambda_{ij}^{r^{ij}} \neq \check{c}^{ij} = \lambda_{ji}^{r^{ji}}$ and, for every $r \in \{r, \dots, 0\} \setminus \{r^{ij}\}$, either $\lambda_{ij}^r = \underline{\lambda}$ or $\lambda_{ji}^r = \underline{\lambda}$ or both.

The *ij-decline sequence* has the players behaving as on the prescribed path through round r^{ij} , when player j offers the target conditional arrangement but player i sends the null message, so a conditional arrangement is not formed; the players send the null message to each other thereafter. Insignificant variations involve player i sending any message other than \check{c}^{ij} in round r^{ij} while player j sends \check{c}^{ij} , and at least one of the players silent in the other rounds. The next definition refers to sequences in which a pair of players send each other the prescribed-path messages (forming their target conditional arrangement) until one of them cancels.

Definition 12: For any ordered pair $(i, j) \in K$ and any round $r \geq 1$, the *ijr-cancel sequence* is defined by: $\lambda_{ij}^{r^{ij}} = \lambda_{ji}^{r^{ji}} = \check{c}^{ij}$, $\lambda_{ij}^r = \text{“cancel”}$, $\lambda_{ji}^r = \underline{\lambda}$, and $\lambda_{ij}^\ell = \lambda_{ji}^\ell = \underline{\lambda}$ for $\ell < r^{ij}$. Say that (h_{ij}, h_{ji}) is an *ijr-cancel sequence except for insignificant variations* if

$\lambda_{ij}^{r^{ij}} = \lambda_{ji}^{r^{ji}} = \check{c}^{ij}$, $\lambda_{ij}^r = \text{“cancel”}$, $\lambda_{ji}^r = \underline{\lambda}$, and for every $\ell < r^{ij}$ either $\lambda_{ij}^\ell = \underline{\lambda}$ or $\lambda_{ji}^\ell = \underline{\lambda}$ or both.

Next I describe particular full sequences of messages between all players in the contracting phase. These sequences proceed as in the prescribed path until a round in which one relationship experiences a disruption in the formation of a conditional arrangement, and this disruption triggers a particular contagion to other relationships.

Definition 13: For any ordered pair $(i, j) \in K$ satisfying $\rho(i) \geq \rho(j)$, the *ij-initiated transit sets*, denoted by $P_{ij}^{r^{ij}}, P_{ij}^{r^{ij}+1}, \dots, P_{ij}^{\bar{r}}$, are defined inductively as follows:

- $P_{ij}^{r^{ij}} = \{(i, j)\}$.
- If $r^{ij} < 0$ then for $r \in \{r^{ij}, \dots, -1\}$ and given P_{ij}^r , let $P_{ij}^{r+1} = \{(k, k') \in K \mid \rho(k) \geq \rho(k'), \exists k'' \neq k' \text{ s.t. } (k'', k) \in P_{ij}^r\}$.
- Then, letting $\bar{P}_{ij} \equiv \cup_{r \in \{r^{ij}, \dots, 0\}} P_{ij}^r$, let $P_{ij}^1 = \{(k, k') \in K \mid (k, k') \notin \bar{P}_{ij}, (k', k) \notin \bar{P}_{ij}, \text{ and } \exists k'' \text{ s.t. } (k'', k) \in \bar{P}_{ij} \text{ or } (k, k'') \in \bar{P}_{ij}\}$.
- Finally, for $r \in \{1, \dots, \bar{r} - 1\}$ and given P_{ij}^r , let $P_{ij}^{r+1} = \{(k, k') \in K \mid \exists k'' \neq k' \text{ s.t. } (k'', k) \in P_{ij}^r\}$.

By construction, the *ij-initiated transit sets* are disjoint. In the definition of P_{ij}^1 , the condition of $(k, k'') \in \bar{P}_{ij}$ applies to (i, j) .

Let the *full sequence of messages* in the contracting phase be denoted by $h = (h_{ij})_{i,j \in N; i \neq j}$, and note that this accounts for the sequence of messages between every pair of players.

Definition 14: For any ordered pair $(i, j) \in N \times N$ satisfying $(i, j) \in K$ and $\rho(i) \geq \rho(j)$, the *ij-trigger sequence* is the full sequence of messages uniquely defined by:

- For every $(k, k') \in \bar{P}_{ij}$, $(h_{kk'}, h_{k'k})$ is the *kk'-decline sequence*.
- For every $r \in \{1, \dots, \bar{r}\}$ and $(k, k') \in P_{ij}^r$, $(h_{kk'}, h_{k'k})$ is a *kk'r-cancel sequence*.
- For every $(k, k') \notin \cup_{r \in \{\underline{r}, \dots, \bar{r}\}} P_{ij}^r$, $(h_{kk'}, h_{k'k})$ is the *prescribed message sequence*.

Say that ***h*** is an *ij-trigger sequence except for insignificant variations* if the conditions above hold in the weaker sense of “except for insignificant variations.”

Recall that h_{ij}^r is the sequence of messages from player *i* to player *j* through round *r* of the contracting phase. Note that for a given player *i*, $(h_{ij}^r, h_{ji}^r)_{j \neq i}$ is the sequence of messages between player *i* and all other players through round *r*. Also, for a given sequence of messages \tilde{h}_{ij} through the entire contracting phase, let \tilde{h}_{ij}^r refer to the truncation to round *r*.

Definition 15: For any player $i \in N$, say that $(\mathbf{h}_{ij}^r, \mathbf{h}_{ji}^r)_{j \neq i}$ is *consistent with the prescribed path* if, for every $j \neq i$, (h_{ij}, h_{ji}) is the *prescribed message sequence except for insignificant variations*.

Definition 16: For any players $i, j, k \in N$ and $r \in \{\underline{r} - 1, \dots, \bar{r}\}$, say that $(\mathbf{h}_{ik'}^r, \mathbf{h}_{k'i}^r)_{k' \neq i}$ is *consistent with a jk-trigger sequence* if it is not consistent with the prescribed path and

there exists $(\tilde{h}_{k'k''})_{k',k'' \in N, k' \neq k''}$ that is a jk -trigger history except for insignificant variations, such that for every $k' \neq i$, $h_{ik'}^r = \tilde{h}_{ik'}^r$ and $h_{k'i}^r = \tilde{h}_{k'i}^r$. Say that $(\mathbf{h}_{ik'}^r, \mathbf{h}_{k'i}^r)_{k' \neq i}$ is **consistent with a trigger sequence** if there exist $j, k \in N$ such that $(h_{ik'}^r, h_{k'i}^r)_{k' \neq i}$ is consistent with a jk -trigger sequence.

A sequence of messages between a given player and the other players can be consistent with multiple trigger sequences. For example, in the example shown in Figure 8, the 86-trigger sequence and 76-trigger sequence would present the same way to player 2 (in round 2 when player 2 receives the cancellation message from player 5).

An information set for player i is a personal history through some round r of the contracting phase. In the case of $r < \bar{r}$, player i 's personal history is exactly $(h_{ij}^r, h_{ji}^r)_{j \neq i}$. In the case of $r = \bar{r}$, player i 's personal history is given by $(h_{ij}, h_{ji})_{j \neq i}$ and the realization of the random draw ϕ . In both cases, let us say that the *personal history is consistent with a jk -trigger sequence* if $(h_{ij}^r, h_{ji}^r)_{j \neq i}$ satisfies this condition, and likewise say that it is *consistent with the prescribed path* if $(h_{ij}^r, h_{ji}^r)_{j \neq i}$ satisfies this condition. The null history at the beginning of the grand game is trivially consistent with the prescribed path. The terminology just developed allows the key information sets to be easily defined.

Definition 17: *The set of key information sets Ξ is defined to comprise, for each $i \in N$, every personal history for player i that is consistent with the prescribed path, and every personal history for player i that is consistent with a trigger sequence.*

Note that many information sets are not in Ξ . Examples include a personal history in which player i established a conditional arrangement with some player j for which $(i, j) \in L$ and yet $(i, j) \notin K$, or where $(i, j) \in K$ but these players formed a conditional arrangement that is not their target one and/or formed their conditional arrangement in a round other than r^{ij} . In these personal histories, player i detects simultaneous deviations by players i and j . Also absent from Ξ are some personal histories consistent with unilateral deviations, such as when player i deviates from the equilibrium path to cancel a contract in a round that implies payment of a cancellation penalty.

Prescribed actions at key information sets

Recall that $d_i^r = (\lambda_{ij}^r)_{j \neq i}$ denotes player i 's action in round r of the contracting phase (the vector comprising the messages that player i sends to each other player). Denote by \check{d}_i^r player i 's action in round r on the prescribed path. Further, for any ordered pair (j, k) such that there exists a jk -trigger sequence, and for any player i , denote by $\tilde{d}_i^r(j, k)$ player i 's action in round r of the jk -trigger sequence.

Consider the information sets in Ξ belonging to a given player i . The prescribed actions are specified as follows. Listed first are the personal histories consistent with the prescribed path, followed by those consistent with a trigger sequence.

Strategy-PP: *For each personal history through round $r < \bar{r}$ that is consistent with the prescribed path, in round $r + 1$ player i chooses action \check{d}_i^{r+1} . For each personal history*

through round \bar{r} that is consistent with the prescribed path, in the production phase player i chooses action a_i^* .

Strategy-TS: Consider any personal history of player i that is consistent with a trigger sequence. If this personal history is through any round $r < \bar{r}$, then in round $r + 1$ player i chooses action $\tilde{d}_i^r(j, k)$, for any j and k such that the personal history is consistent with a jk -trigger sequence.²⁵ If this personal history is through round \bar{r} , then player i 's action in the production phase is determined as follows:

- (A) If $i \in \underline{N}$ then player i chooses $\underline{\alpha}_i$, which is a_i^* , regardless of ϕ .
- (B) If $i \in \overline{N}$ and there is another core player j such that player i 's personal history is consistent with a ji -trigger sequence or an ij -trigger sequence, and if $|\hat{\beta}(i, j)| = 1$, then player i chooses $\underline{\alpha}_i$ in the case of $\phi = 0$ and \hat{a}_i^ϕ in the case of $\phi \in N$.
- (C) If $i \in \overline{N}$ and there is a peripheral player j such that player i 's personal history is consistent with a ji -trigger sequence, then player i chooses $\underline{\alpha}_i$ in the case of $\phi = 0$ and \underline{a}_i^ϕ in the case of $\phi \in N$.
- (D) Otherwise, player i chooses $\underline{\alpha}_i$ in the case of $\phi = 0$ and a_i^ϕ in the case of $\phi \in N$.

In case B, player i is a core player and has just one partner in the core group with whom she is supposed to contract, but the target conditional arrangement for this pair was declined by her partner or herself. In this event, and since in the trigger sequence she also cancels all conditional arrangements with peripheral players, player i enters the production phase with only the null contract. Player i 's personal history in this case is consistent with a ji -trigger sequence or an ij -trigger sequence. In the former subcase, it may also be consistent with other trigger sequences, such as one initiated by a decline in some round $r < 0$ that led player j to decline with player i . In Case C, a peripheral player declined with player i in round -1 , and then player i declined with all core partners in round 0 and cancelled with other peripheral partners in round 1. Case D covers all instances in which player i enters the production phase with a conditional arrangement cancelled with at least one other core player; in this event, player i is supposed to choose her part of a^ϕ for every $\phi \in N$.

Beliefs at key information sets

I next describe partial beliefs of the players at the information sets in Ξ , specifically the marginal over the actions taken (or to be taken) by the players at all information sets in Ξ . This leaves out the belief of a player at an information set in Ξ about actions taken at information sets in $I \setminus \Xi$, and it leaves unaddressed the beliefs of the players at these other information sets.

Consider the information sets in Ξ belonging to a given player i . The partial beliefs, described as appraisals (probability distribution over the space of strategy profiles), are specified as follows.

²⁵There are cases in which there is more than one pair (j, k) with which the jk -trigger sequence player i 's personal history is consistent, but it is not difficult to confirm that $\tilde{d}_i^r(j, k)$ is the same for them.

Belief-PP: For each personal history that is consistent with the prescribed path, player i believes that the actions taken at the information sets in Ξ are exactly as prescribed by Strategy-PP, except for any inconsistencies observed by player i . That is, (1) player i believes that in every prior round r , the other players sent exactly the messages described by \hat{d}_j^r except for those messages that, in player i 's observation, constitute insignificant variations; and (2) at unreached information sets, players would behave as prescribed.²⁶

Belief-TS: For each personal history that is consistent with a trigger sequence, let T be the (j, k) pairs such that player i 's personal history is consistent with a jk -trigger sequence and the network- K distance between $\{j, k\}$ and i is minimized among such pairs. Then player i believes that actions taken at information sets in Ξ are as prescribed by Strategy-TS for one or more jk -trigger sequences where $(j, k) \in T$, except for any inconsistencies observed by player i . Inconsistencies are resolved as described in the previous case.

In other words, for personal histories consistent with the prescribed path, player i believes that play has and will proceed according to the prescribed path, except for any observed discrepancies (which player i believes are insignificant variations). Likewise, for personal histories consistent with a trigger sequence, player i believes that play has and will proceed according to a trigger sequence, except for any observed discrepancies (which player i believes are insignificant variations). In the latter case, as noted above, it is possible that player i 's personal history is consistent with multiple trigger sequences. In this case, player i believes that the actual trigger sequence playing out is among those having the trigger-decline action occurring in a relationship closest to player i .

Rationality at key information sets

Remember that, because the key information sets are a proper subset of all information sets, the specified behavior at these information sets only partially defines the player's strategy profile. Likewise, we have only partially defined the player's beliefs at these information sets. Nonetheless, we can verify that the prescribed behavior is sequentially rational at the information sets in Ξ given the partial beliefs. The following list provides the details for every information set in Ξ belonging to any player $i \in N$. All items on the list pertain to a player $i \in N^K$, whereas only the first two are relevant for a player $i \in N \setminus N^K$.

1. Personal histories through round \bar{r} that are consistent with the prescribed path: Player i is in the production phase. From Belief-PP, player i believes that the other players will select a_{-i}^* . Because player i has assurance contracts with those she was supposed to contract with, she prefers to choose a_i^* , for any deviation would cost her ψ (per assurance contract), which exceeds the maximal payoff gain in the underlying game.

²⁶For example, in a personal history that is consistent with the prescribed path, player i may have received a message $\lambda_{ji}^r \neq \underline{\lambda}$ from player j in round r , where $r^{ij} > \underline{r}$. This errant message did not disrupt the contracting with player j and so it was an insignificant variation in player i 's experience. Player i would then believe that player j 's action in round r was the vector formed from \hat{d}_j^r by replacing $\underline{\lambda}$ with λ_{ji}^0 as the message sent to player i (leaving all other messages unchanged).

2. Personal histories through any round $r \in \{0, \dots, \bar{r} - 1\}$ that are consistent with the prescribed path: By adhering to the prescribed path, as specified, player i expects to eventually obtain the payoff $u_i(a^*) + \check{M}_i(a^*)$. If she deviates from \check{d}_i^{r+1} by cancelling her conditional arrangement with another core player (and regardless of whether she cancels with multiple other players) then, regardless of how she behaves later, her payoff must fall strictly below $u_i(a^*) + \check{M}_i(a^*)$.

This is because the maximum that player i could gain by altering play in the underlying game is strictly less than γ , and by cancelling she is forced to pay at least one cancellation penalty of at least γ . Further, since there are no loops in network K , and given what is feasible in the continuation (in particular, that pairs of players who did not make conditional arrangements must remain silent with each other), player i 's cancellation cannot lead another player to eventually cancel with her, 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 satisfying $(i, k) \in K$ with whom player i retains the conditional arrangement, all of the players in $\beta(k, i)$ are expected to play their part of a^* in the production phase, unaware of player i 's deviation.

Finally, deviating by cancelling conditional arrangements only with peripheral players will result in a lost transfer from these players, in addition to the cancellation penalty if $r > 0$, lowering player i 's payoff.

3. Personal histories through round \bar{r} that are consistent with a trigger sequence: Let (j, k) be a pair such that player i 's personal history is consistent with a jk -trigger sequence and the distance between j and i is minimized among such pairs. Note that player i in such a contingency has reached the production phase with all of her target conditional arrangements either cancelled or declined. She believes every other player will choose the relevant action specified by Strategy-TS, including that every passive player k' will choose $a_{k'}^*$. Consider cases as delineated in Strategy-TS(A)-(D):

- (A) If $i \in \underline{N}$ then the prescribed action a_i^* is clearly best given that all other passive players do the same. This is true regardless of ϕ and whether player i has experienced a cancellation, because $a_i^\phi = a_i^*$ for $\phi > 0$ in every forcing contract that results from a target conditional arrangement having been cancelled.
- (B) Next take the case of $i \in \overline{N}$, there is another core player j such that player i 's personal history is consistent with a ji -trigger sequence or an ij -trigger sequence, and $|\hat{\beta}(i, j)| = 1$. Here, player i 's contract was declined with the only core player she was supposed to contract with, and she cancelled in round 1 with any peripheral contracting partners. Player i believes that target conditional arrangements between all other pairs of core players were established and then cancelled, so that a_{-i}^ϕ will be played in the event of $\phi > 0$ and $\underline{\alpha}_{-i}$ will be played in the event of $\phi = 0$. Player i 's payoff is exactly as in the underlying game, and so \hat{a}_i^ϕ is optimal in the event of $\phi > 0$ and $\underline{\alpha}_i$ is optimal in the event of $\phi = 0$.
- (C) In the case of $i \in \overline{N}$ and there is a peripheral player j such that player i 's personal history is consistent with a ji -trigger sequence, player i has declined in round 0 with the core players she was supposed to contract with, and she cancelled in round 1 with

any other peripheral contracting partners. She believes that these actions perpetuated the ji -trigger sequence (the other players abide by Strategy-TS), leading the other players to select $\underline{\alpha}_{-i}$ in the case of $\phi = 0$ and $\underline{a}_{-i}^{i\phi}$ in the case of $\phi \in N$. By construction of $\underline{a}^{i\phi}$, it is optimal for player i to choose $\underline{\alpha}_i$ in the case of $\phi = 0$ and $\underline{a}_i^{i\phi}$ in the case of $\phi \in N$, as specified.

- (D) For every remaining trigger sequence for $i \in \overline{N}$, player i made the target conditional arrangement with at least one other core player and all of her conditional arrangements were cancelled. In the case of $\phi = 0$, all of her contracts are null except for constant transfers and she believes the other players will select $\underline{\alpha}_{-i}$, to which $\underline{\alpha}_i$ is a best response. In the case of $\phi > 0$, she has only contracts that force a_i^ϕ (the penalty ψ outweighs any deviation gain in the underlying game) or are null except for constant transfers, and she has at least one of the former. Thus, regardless of what she believes the other players will choose, player i 's optimal action is $\underline{\alpha}_i$ in the case of $\phi = 0$ and a_i^ϕ in the case of $\phi \in N$.

4. Personal histories through any round $r \in \{1, \dots, \bar{r} - 1\}$ that are consistent with a trigger sequence: Given the definition of trigger sequence, player i has no choice to make (restricted to silence with everyone else) except for the subcase in which another player k cancelled a conditional arrangement with player i in round r and player i earlier established a conditional arrangement with at least one other player (not yet cancelled). It must be that $\rho(k) \leq \rho(i)$ and $|\hat{\beta}(k, i)| > 1$ (that is, there are multiple active players on k 's side of network K). Player i is supposed to cancel all remaining conditional arrangements.

Because player i believes that the players in $\beta(k, i)$ will continue to play according to Strategy-TS, player i believes that in the production phase, $a_{\beta(k, i)} \neq a_{\beta(k, i)}^*$ for at least one value of ϕ . This follows from the fact that $\hat{a}_j^j = a_j^j \neq a_j^*$ for every $j \in \overline{N}$. Therefore, player i expects to eventually pay the assurance penalty of ψ with probability of at least ε , for every outstanding conditional arrangement that she does not cancel (because from Lemma 2(a) the resulting contracts are assurance contracts).

From the definition of ψ , the expected penalty exceeds $(n - 1)\gamma$ and thus exceeds the maximal gain in the underlying game. The arranged cancellation penalties are strictly below $(n - 1)\gamma$, and so player i prefers to cancel all outstanding conditional arrangements. In fact, player i prefers to do so immediately (in round $r + 1$), since the cancellation penalty increases with r and no other player in $\beta(i, k)$ would otherwise cancel with her (implied by K having no loops).

5. Personal histories through round $r = 0$ that are consistent with an ik -trigger sequence for some $k \in N$: It is the case that $\rho(i) \geq \rho(k)$, which implies that $|\hat{\beta}(k, i)| \geq 1$ because only pairs of core players are scheduled to form conditional arrangements in round 0. Player i is supposed to cancel all remaining conditional arrangements.

The logic for class 4 of personal histories, based on player i believing that $a_{\beta(k, i)} \neq a_{\beta(k, i)}^*$ for at least one value of ϕ , holds here as well, implying that player i prefers to cancel all outstanding conditional arrangements in the current round 1. Note that this includes the subcase of $i \in \overline{N}$ and player i having initiated the trigger sequence by declining with exactly one other core player k in round 0. It also includes the case in which player i is peripheral,

where player i believes that $\underline{a}_{-i}^{\bar{k}\phi}$ will be chosen by the other players in the event of $\phi > 0$, where \bar{k} is the closest active player to player i . Recall that, by construction, $\underline{a}_{-i}^{\bar{k}\phi} \neq a_{-i}^*$ for some value of ϕ .

6. Personal histories through round $r = 0$ that are consistent with a ki -trigger sequence for some $k \in N$ satisfying $\rho(k) = \rho(i) = 0$: Player i is supposed to cancel all remaining conditional arrangements. As in class 4 and class 5 of personal histories, $|\hat{\beta}(k, i)| \geq 1$, player i believes that $a_{\beta(k, i)} \neq a_{\beta(k, i)}^*$ for at least one value of ϕ , and so player i prefers to cancel all outstanding conditional arrangements in the current round 1.

7. Personal histories through round $r = 0$ that are consistent with a ji -trigger sequence for some $j \in N$ satisfying $\rho(j) = \rho(i) + 1$: In this class, from Belief-TS, player i believes that a ji -trigger sequence is in process. In round $r^{ij} + 1$, player i declined with every player k satisfying $(i, k) \in K$ and $\rho(k) \leq \rho(i)$, which is a single player if $\rho(i) > 0$ and possibly multiple players if $\rho(i) = 0$. Player i is supposed to cancel all remaining conditional arrangements. As in the previous class of personal histories, $|\hat{\beta}(k, i)| \geq 1$ and player i believes that $a_{\beta(k, i)} \neq a_{\beta(k, i)}^*$ for at least one value of ϕ , and so player i prefers to cancel all outstanding conditional arrangements in the current round 1.

8. Personal histories through round $r < 0$ that are consistent with a ji -trigger sequence for some $j \in N$: As in the previous class, from Belief-TS, player i believes that a ji -trigger sequence is in process, and here it must be that $\rho(j) = \rho(i) + 1$. Player i is supposed to send the null message to everyone else. If $r \neq -\rho(j)$ then she has no incentive to deviate because non-null messages would be interpreted as insignificant variations and ignored by the others.

If $r = -\rho(j)$ then there is at least one player k for which $(k, i) \in K$ and $r^{ik} = r + 1$ (exactly one such player in the case of $r < -1$), and this player would expect to receive message \check{c}^{ik} . By sending something other than this expected message with every such player k , player i continues the ji -trigger sequence and expects to receive the payoff \underline{w}_i . By sending message \check{c}^{ik} to each such player, player i makes them believe that they are on the prescribed path except for insignificant variations, and by continuing as though on the prescribed path player i expects to get the payoff $u_i(a^*) + \check{M}_i(a^*) - \check{m}_i^{ij}(a^*)$. From Lemma 2, this value is strictly less than \underline{w}_i , so player i prefers not to deviate in this way.

Other deviations in the continuation cannot improve player i 's payoff either. For instance, if $\rho(i) = 0$ and there are multiple other core players with whom player i is supposed to contract, declining with some but not all of them will put player i in the position addressed in the previous classes, where player i expects to pay an assurance penalty for some values of ϕ .

Further, if player i sends message \check{c}^{ik} to each player k described above, pretending with them to be on the prescribed path, and plans to cancel conditional arrangements with any of them later, then player i expects to pay a cancellation penalty. Planning to cancel with only a player j' for which $\rho(j') = \rho(i) + 1$ would not entail a cancellation penalty if it is done in round 1, but then player i loses a positive transfer from this player given that she has induced each player k described above to play as though on the prescribed path.

9. Personal histories through round $r < 0$ that are consistent with an ij -trigger sequence for some $j \in N$: In this class, player i earlier initiated a trigger sequence by declining with

player j . Player i is supposed to send the null message to everyone else. Given that player i believes the ij -trigger sequence is in progress and expects to receive the null message from all other player, deviating would not affect player i 's expected payoff because any non-null message would be viewed as an insignificant variation and ignored by the recipient.

10. Personal histories through any round $r < 0$ that are consistent with the prescribed path: By adhering to the prescribed path, as specified, player i expects to eventually obtain the payoff $u_i(a^*) + \check{M}_i(a^*)$. In the case of $r = -1$ and $\rho(i) = 0$, player i could deviate by declining with every player k with whom she is supposed to establish a conditional arrangement in round 0 (by sending a message other than \check{c}^{ik}). It would then be optimal for her to cancel all conditional arrangements with peripheral players in round 1 (by the same logic described in class 4). As in class 3(C), she would believe that the other players will select $\underline{\alpha}_{-i}$ in the case of $\phi = 0$ and $\underline{a}_{-i}^{i\phi}$ in the case of $\phi \in N$. The best that player i could then do in the production phase is to choose $\underline{\alpha}_i$ in the case of $\phi = 0$ and $\underline{a}_i^{i\phi}$ in the case of $\phi \in N$, leading to the expected payoff \underline{w}_i . From Lemma 2(b), her expected payoff would be strictly below $u_i(a^*) + \check{M}_i(a^*)$, and therefore player i does not want to deviate in this manner.

She can do no better by declining with only some of the players with whom she is supposed to contract, for then she would be on the hook for an assurance penalty or a cancellation penalty. Deviating by sending non-null messages to players expecting to receive the null message would not further affect player i 's expected payoff, since these messages would be ignored as insignificant variations.

The analysis is much the same in the case of $r < -1$ and $\rho(i) = -r - 1$. Here there is one player k for whom $(i, k) \in K$ and $r^{ik} = r + 1$, and it is the case that $\rho(k) = \rho(i) - 1$. If in round $r + 1$ player i declines with player k by not sending message \check{c}^{ik} , then player i would believe that play will proceed according to the ik -trigger sequence, resulting in expected payoff \underline{w}_i .

Next take the case of $r < -1$, $\rho(i) = -r - 2$, and the existence of a player j for whom $(i, j) \in K$ and $\rho(j) = \rho(i) + 1$. In the current round $r + 1$, player i is supposed to send message \check{c}^{ij} to such a player. Deviating to decline the target conditional arrangement nullifies the contract with this other player and, if player i were to otherwise behave as though on the prescribed path, then she would expect to obtain the payoff $u_i(a^*) + \check{M}_i(a^*) - \check{m}_i^{ij}(a^*)$, which is below $u_i(a^*) + \check{M}_i(a^*)$ by Lemma 2(c). Declining in this manner and planning to decline again with all others in the following period would give player i an expected payoff of at most \underline{w}_i . No other deviation is worthwhile, using the logic laid out above.

Translating the partial construction to a fully described sequential equilibrium

The penultimate step of the proof is to specify a sequence of fully mixed behavior strategies for the information sets in Ξ that converges to the partial strategy defined by Strategy-PP and Strategy-TS, and that induces the partial beliefs defined by Belief-PP and Belief-TS. I shall use the term *situation* in place of information set, to be consistent with Watson (2023). Beliefs are expressed as appraisals (probability distributions over the space of strategy profiles).

Denote by s a strategy profile in the grand game and note that it can be expressed as a mapping from I to the action space in the grand game, such that for each $\xi \in I$, $s(\xi)$ is a

feasible action at situation ξ . Let S denote the space of strategy profiles in the grand game, let S_{Ξ} denote the set of strategy profiles restricted to Ξ , and for any $s \in S$ let s_{Ξ} be the restriction to Ξ . For each $\xi \in I$, let $S(\xi)$ denote the set of strategy profiles that reach ξ (the path of play passes through situation ξ) and let $S(\xi)_{\Xi} \equiv \{s_{\Xi} \mid s \in S(\xi)\}$.

A probability distribution $\pi \in \Delta S_{\Xi}$ is called *fully mixed* if it has full support, and it is called a *behavior strategy on Ξ* if it exhibits independence across these situations. Let π^* denote the behavior strategy on Ξ defined by Strategy-PP and Strategy-TS. Note that, as constructed, π^* is a degenerate distribution (a pure strategy profile).

Define a sequence $\{\pi^{\kappa}\}_{\kappa=1}^{\infty}$ of fully mixed behavior strategies on Ξ as follows. For each personal history of player i through any round $r < \bar{r}$ that is consistent with the prescribed path, player i randomizes independently across the components of the message vector. In the case of $r \leq 0$, for each other player j and each feasible message λ_{ij} that differs from what \check{d}_i^r specifies to be sent to player j , player i sends message λ_{ij} to player j with probability $(1/\kappa)^{n-r}$. The remaining probability (which converges to 1 as $\kappa \rightarrow \infty$) is put on the message prescribed by \check{d}_i^r . In the case of $r \in \{1, \dots, \bar{r} - 1\}$, for each other player j and each feasible message λ_{ij} that differs from what \check{d}_i^r specifies to be sent to player j , player i sends message λ_{ij} to player j with probability $(1/\kappa)^{2n}$. The remaining probability (which converges to 1 as $\kappa \rightarrow \infty$) is put on the message prescribed by \check{d}_i^r .

Similarly, for each personal history of player i through any round $r < \bar{r}$ that is consistent with a jk -trigger sequence, player i randomizes independently across the components of the message vector. For each other player j' and each feasible message $\lambda_{ij'}$ that differs from what $\check{d}_i^r(j, k)$ specifies to be sent to player j' , player i sends message $\lambda_{ij'}$ to player j' with probability $(1/\kappa)^{2n}$. The remaining probability is put on the message prescribed by $\check{d}_i^r(j, k)$.

Finally, for each personal history of player i through round \bar{r} that is consistent with the prescribed path, player i chooses each action $a_i \neq a_i^*$ with probability $(1/\kappa)$ and puts the remaining probability on a_i^* (which is what Strategy-PP prescribes). Likewise, for each personal history of player i through round \bar{r} that is consistent with a jk -trigger sequence, player i puts probability $(1/\kappa)$ on each action other than that prescribed by Strategy-TS, and puts the remaining probability on the action prescribed by Strategy-TS.

The sequence $\{\pi^{\kappa}\}_{\kappa=1}^{\infty}$ clearly converges to π^* , because at each situation in Ξ , the probability put on the action prescribed by π^* converges to 1 as κ approaches ∞ . It is also clear that for each $\xi \in \Xi$, the conditional distribution $\pi^{\kappa}(\cdot \mid S(\xi)_{\Xi})$ converges. Let $q^{\xi} \equiv \lim_{\kappa \rightarrow \infty} \pi^{\kappa}(\cdot \mid S(\xi)_{\Xi})$ and note that this is the appraisal of the player on the move at ξ about the behavior at the situations in Ξ . It is easy to see that q^{ξ} is as described by Belief-PP and Belief-TS, whichever is the relevant case.

For example, consider a personal history of player i through round $r < 1$ that is consistent with the prescribed path but where the message from player j in round r differed from what \check{d}_j^r specifies. The unexpected message from player j is an insignificant variation in the communication with player i . The probability that this occurred due to a tremble of player j 's hand in round r (a tremble that affected only the message to i) is on the order of $(1/\kappa)^{n-r}$, whereas the probability that it followed from a deviation by any player in a previous round that caused player j to send the unexpected message in round r is at most on the order of $(1/\kappa)^{n-r+1}$. Thus, in the limit as $\kappa \rightarrow \infty$, player i believes play has been exactly on the

prescribed path except for the insignificant variation observed in round r . The same logic works for the case of multiple insignificant variations.

For another example, consider a personal history of player i through round $r \geq 1$ that is consistent with a trigger sequence, where player j cancelled with player i in round r . The probability that the cancellation occurred due to a tremble of player j 's hand in round r is on the order of $(1/\kappa)^{2n}$, whereas the probability that it followed from a decline choice at round 0 by some player is at least on the order of $(1/\kappa)^n$, because there exists such a decline sequence that would reach player i in round r . Further, any decline sequence initiated prior to round 0 occurs on the order of at most $(1/\kappa)^{n+1}$. Thus, in the limit as $\kappa \rightarrow \infty$, player i believes play has been exactly on a $j'k$ -trigger sequence, where $j', k \in \hat{N}^K$ (they are core players, so the decline action that initiated the trigger sequence occurred in round 0). The same logic works when insignificant variations are included.

Note that in the case just described, player i believes that the trigger sequence in process was initiated by a single decline choice in round 0, rather than being initiated earlier. Likewise, for a personal history through round $r \leq 0$ that is consistent with a trigger sequence and where some player j declined with player i , player i believes that a ji -trigger sequence is in process, rather than a trigger sequence that was initiated earlier.

The appraisal system for the partial equilibrium construction is given by $Q \equiv (q^\xi)_{\xi \in \Xi}$, where Ξ is the union of Ξ and artificial situations that represent the beginning of the game (see Watson 2023 for an explanation). The appraisals include the specification of π^* . We can extend the appraisals to include nature's choices by taking the product of each q^ξ and nature's behavior strategy, since nature moves after the contracting phase and the players observe nature's choice.

To summarize, Q is fully consistent (Kreps and Wilson 1982) in the partial game because it was constructed from a sequence of fully mixed behavior strategies. Further, we have verified that Q is sequentially rational, regardless of the players' behavior at situations $I \setminus \Xi$ (what Watson 2023 calls Ξ -sequentially rational). This means that Q is a Ξ -partial sequential equilibrium, as defined by Watson (2023).

The last step is to use the theorem of Watson (2023) to establish the existence of a sequential equilibrium in the entire grand game that coincides with Q on Ξ . To do this, we must verify that the *rectangular margin-support condition* holds, which is that $\{s \in S(\xi) \mid s_\Xi \in \text{supp } q^\xi\}$ is a Ξ -product set, for every $\xi \in \Xi$ that is a situation for a strategic player (not nature). This is straightforward given all of the work we have done to construct Q .

Observe that, for each $\xi \in \Xi$, the appraisal q^ξ puts zero probability on strategy profiles that pass through any situation in $I \setminus \Xi$ before reaching ξ . This is clear from the fact that every situation reached through a history that is the prescribed message sequence except for insignificant variations is itself a personal history consistent with the prescribed path. Likewise, every situation reached through a history that is an ij -trigger sequence except for insignificant variations is itself a personal history consistent with either the ij -trigger sequence or the prescribed path.

Thus, for each $\xi \in \Xi$ and $s \in S$ such that $s_\Xi \in \text{supp } q^\xi$, it must be that $s \in S(\xi)$ regardless of the behavior specified for $I \setminus \Xi$. That is, $\{s \in S(\xi) \mid s_\Xi \in \text{supp } q^\xi\} = S(\xi)_\Xi \times S_{I \setminus \Xi}$, and so this is a product set.

The theorem of Watson (2023) then establishes the existence of a sequential equilibrium in the entire grand game, given by an appraisal system P , such that for each $\xi \in \Xi$, $p_{\Xi}^{\xi} = q^{\xi}$. In particular, the players' equilibrium behavior at the situations in Ξ is exactly as described by Q ; it is given by π^* . The equilibrium path is the prescribed path, and so the grand game ends with play of action profile a^* in the underlying game, which is efficient.

Specifications for the cases held aside

Recall that, for the complicated construction completed above, we left aside the case in which the underlying game is trivial and also the case in which $|\hat{N}^K| \leq 2$ (implying $\bar{N} = \hat{N}^K$). I next describe how to deal with these cases.

In the case of a trivial underlying game, where the underlying game has an efficient Nash equilibrium, we can construct an efficient sequential equilibrium of the grand game as follows. The prescribed path entails all players sending the null message to each other in every round, and then choosing \underline{a} . Define Ξ as the situations that are consistent with the prescribed path, where insignificant variations may have occurred. At the situations in Ξ , the players are prescribed to behave as on the prescribed path, ignoring insignificant variations, which is clearly sequentially rational. The partial-sequential-equilibrium construction and full equilibrium existence work as before.

In the case in which $|\hat{N}^K| \leq 2$, the core group comprises exactly two active players or exactly one active player. In the subcase in which $\underline{a}_i \neq a_i^*$ for $i \in \hat{N}^K$ (which is implied by $|\hat{N}^K| = 1$), the equilibrium construction is exactly as described in the previous section except that profiles a^{ϕ} , \hat{a}_i^j , and \underline{a}^{ik} do not come into play because trigger sequences result in contracts not being formed by core players (rather than being cancelled). Strategy-TS is modified to specify simply that player i chooses \underline{a}_i in the underlying game. It is not difficult to see that the rest of the construction holds together.

In the subcase in which $|\hat{N}^K| = 2$ and yet $\underline{a}_j = a_j^*$ for one player $j \in \hat{N}^K$ (it cannot be both), we can treat this player as passive and perform the equilibrium construction as though there is exactly one active player (the subcase covered in the previous paragraph). This is because on the equilibrium path and in any trigger sequence, the lone active player i plays only a_i^* or \underline{a}_i , and thus player j optimally responds with a_j^* with any specified assurance contract or the null contract.

B Proposed Supplementary Appendix (on-line only)

B.1 Statements and Proofs of Minor Results

In this section are formal statements of the minor results mentioned in Sections 3 and 5 of the paper, along with technical details and proofs. The first three were noted in Section 3 and the last two in Section 5. These results are proved by extending what has already been presented in the main body of this article and in Appendix A.2.

Result 1: *For any given $n \geq 4$, there exists an underlying game $\langle A, u \rangle$ and a network L such that every player has a link in L (although L is disconnected) and, regardless of the contracting institution, there is no sequential equilibrium of the grand game in which an efficient action profile of the underlying game is played with positive probability.*

Proof: For any given $n \geq 4$, a version of the example shown in Figure 2 suffices to prove this result. Let κ be the largest integer less than $n/2$. Let L be the network comprising exactly the pairs

$$(1, 2), (2, 3), \dots, (\kappa - 1, \kappa), (\kappa + 1, \kappa + 2), (\kappa + 2, \kappa + 3), \dots, (n - 1, n).$$

Note that this network is disconnected due to the missing link between players κ and $\kappa + 1$. Let the underlying game be one in which only players 1 and n have choices to make, each chooses between action 0 and action 1, and the payoffs of players 1, 2, $n - 1$, and n are, in this order, shown in the table of Figure 2 with a_4 replaced by a_n . The other players get a payoff of 0 regardless of a_1 and a_n . The logic presented in section 3.1 concerning the incentives of players 1 and 2 applies here without alteration. \square

For the second result, here is a precise definition: Let us say that the *contracting institution exhibits dated commitment* if for every pair of players (i, j) , there is a round \hat{r}^{ij} such that

- (i) $\Lambda_{ij}^\ell(h_{ij}^{\ell-1}, h_{ji}^{\ell-1}) = \{\underline{\lambda}\}$ for all $\ell > \hat{r}^{ij}$, $h_{ij}^{\ell-1}$, and $h_{ji}^{\ell-1}$; and
- (ii) $\mu^{ij}(h_{ij}, h_{ji}, \cdot) \equiv \underline{m}$ for every $h_{ij} = (\lambda_{ij}^r, \lambda_{ij}^{r+1}, \dots, \lambda_{ij}^{\bar{r}})$ and $h_{ji} = (\lambda_{ji}^r, \lambda_{ji}^{r+1}, \dots, \lambda_{ji}^{\bar{r}})$ for which either $\lambda_{ij}^{\hat{r}^{ij}} = \underline{\lambda}$ or $\lambda_{ji}^{\hat{r}^{ij}} = \underline{\lambda}$ or both.

Result 2: *For any given $n \geq 4$, there exists an underlying game $\langle A, u \rangle$, a connected network L , and a number $\delta < 1$ such that, for every contracting institution exhibiting dated commitment, there is no sequential equilibrium of the grand game in which an efficient action profile of the underlying game is played with a probability greater than δ .*

Proof: A version of the example shown in Figure 3 suffices to prove this result. Let L be the network comprising exactly the pairs $(1, 2), (2, 3), \dots, (n - 1, n)$. This is a linear, connected network. Let the underlying game be one in which only players 1 and n have choices to make, each chooses between actions 0, 1, and 2, and the payoffs of players 1, 2, $n - 1$, and n are, in this order, shown in the table of Figure 3 with a_4 replaced by a_n . The other

players get a payoff of 0 regardless of a_1 and a_n . Note that the efficient action profile entails $a_1 = a_n = 1$.

To obtain a proof by contradiction, consider any contracting institution that exhibits dated commitment and suppose there is an equilibrium of the grand game in which the efficient action profile is played with a probability of at least δ . Without loss of generality, we can assume that $\hat{r}^{12} \geq \hat{r}^{n-1n}$ holds for the given contracting institution; if this inequality is reversed, substitute player n for player 1 and player $n - 1$ for player 2, and the logic is the same.

Consider the equilibrium paths of play in which efficient actions $a_1 = 1$ and $a_n = 1$ are chosen. By presumption, the set of pure strategy profiles that induce these various paths of play are assigned probability of at least δ by the equilibrium (generally mixed) strategy. Suppose player 1 deviates from her equilibrium strategy by sending message $\underline{\lambda}$ to player 2 in every round \hat{r}^{12} situation/information set and by also choosing $a_1 = 2$ in the production phase, and otherwise follows her equilibrium strategy. This deviation can have no effect on contracting between player $n - 1$ and n , and also on player n 's action in the production phase, because their contract would have to be set by round \hat{r}^{12} and they do not communicate after. Thus, with player 1's deviation, player n still must choose $a_n = 1$ with probability at least δ .

Player 1's deviation therefore gives this player an expected payoff of at least $\delta 9 + (1 - \delta)2$, which must be a lower bound on player 1's equilibrium expected payoff. The equilibrium expected payoffs of players 2 and $n - 1$ are bounded below by 2 (because, for instance, if player 2 refused to contract then player 1 must then choose $a_1 = 0$ or $a_1 = 2$), player n 's equilibrium expected payoff is bounded below by 4 (by refusing to contract and then choosing $a_n = 0$, this is the lowest payoff possible), and the equilibrium expected payoffs of all other players are bounded below by 0. The sum of lower bounds is $\delta 9 + (1 - \delta)2 + 8$, which can be no greater than the maximal joint value of 16. This inequality simplifies to $\delta \leq 6/7$. We would therefore have a contradiction if $\delta > 6/7$. The claim holds for any δ strictly between $6/7$ and 1. \square

Result 3: *For any given $n \geq 3$, there exists an underlying game $\langle A, u \rangle$, a connected network L , and a number $\delta < 1$ such that, for every contracting institution with strictly fewer than $n - 1$ contracting rounds, no sequential equilibrium of the grand game has an efficient action profile in the underlying game played with a probability greater than δ .*

Proof: The steps to prove this result are nearly identical to the steps for Result 2. In the case of $n \geq 4$, use the same example employed for Result 2. (For the case of $n = 3$, it suffices to consider an underlying game in which players 1 and 3 are playing a prisoners' dilemma and player 2 has no choice and gets a constant payoff of 0.) For any given contracting institution with strictly fewer than $n - 1$ contracting rounds, suppose there is an equilibrium of the grand game in which the efficient action profile is played with a probability of at least δ .

Suppose player 1 deviates from her equilibrium strategy by sending message $\underline{\lambda}$ to player 2 in every contracting round, ensuring that she has the null contract with player 2, and by choosing $a_1 = 2$ in the production phase. This deviation can have no effect on contracting between player $n - 1$ and n , and also on player n 's action in the production phase, because there aren't enough rounds through which this deviation can alter play in such a fashion as to disrupt contracting between players $n - 1$ and n , and player n would not detect any deviation

from the equilibrium path. Thus, with player 1's deviation, player n still must choose $a_n = 1$ with probability at least δ . Player 1's deviation gives this player an expected payoff of at least $\delta 9 + (1 - \delta)2$, which must be a lower bound on player 1's equilibrium expected payoff. The other players' equilibrium payoffs are bounded as described in the proof of Result 2, and we reach a contradiction if $\delta > 6/7$ as before. \square

Result 4: *Take as given any integer $n \geq 3$, any finite set of underlying games G , and any integer $\kappa \in [2, n]$. Let \mathcal{L} be the set of all connected networks of diameter weakly less than κ . There exists a contracting institution (representing private, independent, and voluntary contracting) satisfying $\bar{r} - \underline{r} \leq 2\kappa - 2$ that implements efficient outcomes.*

Proof: This result is proved by noticing that, in the proof of the Theorem, for all of the steps the necessary number of contracting rounds is bounded by parameters of the network L . Specifically, we need $\lfloor \underline{r} \rfloor$ to be weakly greater than the largest periphery index (to allow peripheral players to establish conditional arrangements in order of periphery index), and we need \bar{r} to be weakly greater than one less than the maximal distance between core players (to allow a sequence of cancellations to progress across the core group following a decline between a pair of core player at round 0). For a network of diameter κ , the maximal periphery index and the maximal distance between core players are both $\kappa - 1$, and therefore the total number of rounds needed for the proof is $(\kappa - 1) + (\kappa - 2) + 1 = 2\kappa - 2$. The addition of 1 here is to account for round 0. Under these conditions, the proof of the Theorem goes through without alteration. \square

On the topic of multiple equilibria and the range of equilibrium values, the next result is analogous to folk theorems in repeated games. For intuition and to develop terminology, let us review the analysis underlying the Theorem. Recall that, in the proof of the Theorem, for each underlying game in G , we took a^* to be an arbitrarily chosen efficient action profile and $\underline{\alpha}$ to be an arbitrarily chosen Nash equilibrium. Therefore, we started with a set of tuples $(\langle A, u \rangle, a^*, \underline{\alpha})$, one for each $\langle A, u \rangle \in G$. From this set, with further arbitrary selection, we derived elements \underline{N} , \bar{N} , a^i , and \hat{a}_i^j for $i, j \in N$. Global parameters ε and γ were selected in relation to the set of underlying games, to satisfy the conditions described in Section 4.2 such as Inequality 1. Likewise, upon fixing a connected network L , we derived a profile \underline{a}^{ik} for every $i \in \hat{N}^K$ and $k \in N$, a value \underline{w}_i for every $i \in N^K$, and a special subnetwork K called essential. And for every $i \in N^K$, we defined the periphery index $\rho(i)$.

Let us call all of these derived elements, collectively, the *fundamental elements* in relation to the given set of tuples $(\langle A, u \rangle, a^*, \underline{\alpha})$. Determination of fundamental elements is generally not unique. Recall that all of this structure led to the identification of contracts $\{\check{m}^{ij}\}_{(i,j) \in K}$ and target conditional arrangements, for each underlying game and network, and ultimately to the construction of an efficient equilibrium in the grand game.

Notice that none of the analysis used to identify the fundamental elements for a given tuple $(\langle A, u \rangle, a^*, \underline{\alpha})$ requires a^* to be efficient. All that was required is that a^* is more efficient than $\underline{\alpha}$. Thus, we can repeat the construction of the fundamental elements by substituting for a^* any action profile \tilde{a} , provided that $\underline{\alpha}$ is a Nash equilibrium of $\langle A, u \rangle$ and $\sum_{i \in N} u_i(\tilde{a}) > \sum_{i \in N} u_i(\underline{\alpha})$. Then for any connected network L , all of the fundamental elements are well defined (not necessarily uniquely) and satisfy the conditions stated in Sec-

tion 4.2. Further, we need not have limited the set of initial tuples to just one pair a^* and $\underline{\alpha}$ for each underlying game; that is, we could allow multiple combinations.

Take as given n and \mathcal{A} . Use the term *scenario* for any tuple $(\langle A, u \rangle, \tilde{a}, \underline{\alpha}, L, \{y^{ij}\}_{i \neq j})$ with the properties that $\langle A, u \rangle$ is an n -player game with $A \subset \mathcal{A}$, $\tilde{a} \in A$, $\underline{\alpha}$ is a Nash equilibrium of $\langle A, u \rangle$, $\sum_{i \in N} u_i(\tilde{a}) > \sum_{i \in N} u_i(\underline{\alpha})$, $L \in N \times N$ is a connected network, and $y^{ij} = y^{ji} \in \mathbb{R}_0^n(i, j)$ for $i \neq j$. Let $Y \equiv \sum_{i < j} y^{ij}$. We will want to know whether, for underlying game $\langle A, u \rangle$ and network L , there is a sequential equilibrium of the grand game in which \tilde{a} is played in the production phase and transfers are $\{y^{ij}\}_{i \neq j}$ on the equilibrium path, so that the payoff vector is $u(\tilde{a}) + Y$.

The proof of the Theorem focuses on, for each underlying game and network, a scenario in which \tilde{a} is efficient (called a^*). We can explore the prospect of multiple equilibria by looking at a set of scenarios that share the same underlying game and network. Call a set \mathcal{S} of scenarios *permissible* if the following conditions hold. First, \mathcal{S} is finite. Second, global parameters γ and ε suffice for all underlying games. That is, for every $(\langle A, u \rangle, \tilde{a}, \underline{\alpha}, \{y^{ij}\}_{i \neq j}) \in \mathcal{S}$, $i \in N$, and $a \in A$, we have $\gamma > 2|u_i(a)|$ and $\sum_{i \in N} u_i(\tilde{a}) > \sum_{i \in N} [(1 - n\varepsilon)u_i(\underline{\alpha}) + n\varepsilon\gamma]$ and also $\varepsilon < 1 - \underline{\alpha}_i(\tilde{a}_i)$ for each player i for whom $\underline{\alpha}_i(\tilde{a}_i) < 1$ (corresponding to the inequalities in Section 4.2). Third, letting $Y \equiv \sum_{i < j} y^{ij}$, it is the case that:

- for every $(i, j) \notin K$, y^{ij} is the 0 vector;
- $u_i(\tilde{a}) + Y_i > \underline{w}_i$ for each player $i \in N^K$; and
- for each pair $(i, j) \in K$ satisfying $\rho(j) = \rho(i) + 1$, $u_i(\tilde{a}) + Y_i - y_i^{ij} < \underline{w}_i$.

Note that the second and third conditions correspond to conditions b and c in Lemma 2.

Result 5: *Take as given any integer $n \geq 2$, any finite set of action profiles \mathcal{A} , and any permissible set \mathcal{S} of scenarios. There exists a contracting institution (representing private, independent, and voluntary contracting) such that the following is true for every scenario $(\langle A, u \rangle, \tilde{a}, \underline{\alpha}, L, \{y^{ij}\}_{i \neq j}) \in \mathcal{S}$: In the case in which $\langle A, u \rangle$ is the underlying game and L is the network, there is a sequential equilibrium of the grand game that yields the payoff vector $u(\tilde{a}) + Y$.*

Proof: Take as given any integer $n \geq 2$, any finite set of action profiles \mathcal{A} , and any finite set \mathcal{S} of permissible scenarios. Fix the fundamental elements for these scenarios to satisfy the conditions of permissibility. For every $(\langle A, u \rangle, \tilde{a}, \underline{\alpha}, \{y^{ij}\}_{i \neq j}) \in \mathcal{S}$, all of the steps described in Sections 4.3 and 4.4 to define feasible contracts and target conditional arrangements go through without alternation except for replacing a^* with \tilde{a} , and instead of applying Lemma 2 we can directly construct $\{\check{m}^{ij}\}_{(i,j) \in K}$ to have the required properties, by using the permissibility conditions. Specifically, we set \check{m}^{ij} to be the \tilde{a} -assurance contract with baseline transfer y^{ij} . The equilibrium construction then goes through as described in Appendix A.2, with no modifications. \square

B.2 Additional Discussion

This paper's novel approach of constrained contracting-institution design allows us to analyze settings without limiting ourselves to a single model of contract negotiation, while

also requiring that assumptions about contracting, such as its voluntary nature, are expressed separately from other assumptions on the contracting process and technology. If we had adopted one of the prior literature's simple models of contract formation from the start, we would have a limited view of contractual linkages and would not have found the main result.

Advantages of the noncooperative modeling approach

On the comparison between fully noncooperative models and cooperative matching of matching and coalitional bargaining, consider the example of a collaboration agreement discussed in Section 3.2, 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 differentiated unambiguously by a model that explicitly accounts for productive interaction, as accomplished herein by specifying the noncooperative underlying game. A modeling approach 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 would not allow for the externality in the collaboration-agreement example. Matching models that allow payoffs to be a function of the entire set of contracts formed, such as in Rostek and Yoder (2020, 2022) and Pycia and Yenmez (2019) for two-sided markets, can capture LDL externalities 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.²⁷

With respect to modeling contract formation, the goal of the Nash program is to establish a mathematical equivalence between stability concepts for cooperative models and equilibrium play of noncooperative protocols. 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.

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.²⁸ The noncooperative approach provides

²⁷The supply chain example features what some call a pecuniary externality, though it may be more instructive to avoid the term and instead say the downstream result of a possible contracting or market distortion. Other entries in the matching literature include Ostrovsky (2008), Hatfield and Kominers (2012, 2015), Hatfield et al. (2013), Manea (2018), and Bando and Hirai (2021).

²⁸Additional related papers from the prior literature include Guttman (1978), Danziger and Schnytzer




Figure 9: A few common networks.

a good foundation for precisely defining the technologies of production and enforcement, including the extent of verifiability. This structure allows one to sort out alternative methods of linking contractual relationships, such as conditioning transfers on third-party productive actions as opposed to contracts on contracts. Importantly, it also gives contracts their natural meaning, enabling predictions on the actual form that contracts take in applications.²⁹

Barriers to efficiency in applications

Considering barriers to efficiency, it may be helpful to categorize examples in terms of prominent aspects of their networks and the structure of their underlying games. Figure 9 illustrates 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;³⁰ (ii) vertical contracting in a bipartite supply network and two-sided markets;³¹ (iii) platforms and general intermediation networks;³² and (iv) community interaction with an arbitrary contractual network.³³

(1991), Guttman and Schnytzer (1992), Varian (1994), and Yamada (2003).

²⁹Regarding 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 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 work suggests that interactive contracts require the external enforcement system to develop a sophisticated language for the cross-referencing of promises.

³⁰Representative entries in the literature include Bernheim and Whinston (1986a,b), Segal (1999), and Galasso (2008). Martimort (2007) surveys the related literature.

³¹Kranton 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).

³²Rochet 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.

³³Regarding more general networks, an important but relatively under-developed 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.

On the efficacy of simpler contracting institutions

Following the discussion in Section 5.5, I comment further on whether efficient contracting requires options to adjust externally enforced elements, as the SCO institution facilitates in rounds 1 through \bar{r} . For instance, consider a “Simultaneous Contracting and Sequential Communication” (SCSC) institution that has one round of simultaneous contract creation, determining the induced game $\langle A, u + M \rangle$ for the production phase, followed by multiple rounds of messages that do not affect the contracts but are used by the players to coordinate on actions to take in the production phase. Whether an SCSC institution can implement efficient outcomes is not addressed by Results 2 and 3, which show that sequential contracting is essential but do not distinguish between external and self-enforced options.

I conjecture that SCSC institutions would not be effective generally but perhaps could implement efficient outcomes for some special classes of underlying games. The best hope might be for underlying games in which $\underline{\alpha}_i(a_i^*) = 0$ for every player i ; but even if an efficient equilibrium exists (which is not clear), the construction remains challenging and still requires players to coordinate on some sort of assurance contracts.³⁴ There is greater doubt for dealing with other underlying games. Consider the case in which $\underline{\alpha}_i(a_i^*) > 0$ for at least one player i . A disruption in contracting with such a player may not motivate her partner to pass along word of the disruption. For example, if $\underline{\alpha}_i(a_i^*) = 1$, then this player would still choose her part of the efficient action profile in response to a belief that the other players have flipped to $\underline{\alpha}$, and so off-equilibrium-path beliefs and behavior could not be as simple as coordination on $\underline{\alpha}$. More broadly, there is a direct conflict between assurance penalties and making $\underline{\alpha}$ a Nash equilibrium of the “induced game” in the production phase following a deviation. The proof of the Theorem (and SCO institution) deals with these problems by use of conditional arrangements that switch to different contracts upon cancellation and, by varying the forcing arrangements as a function of ϕ , induce players to choose actions in the production phase that, under their on-equilibrium-path contracts, would expose them to high assurance penalties.³⁵ Characterizing what SCSC and alternative institutions can achieve is a worthwhile topic for future research; I have not been able to make much progress in the context of the current exercise.

An example of partial verifiability

As a prompt for one direction of future research, here is an example showing that the Theorem does not extend to settings with partial verifiability of productive actions. Consider a

³⁴We can look for a variant of assurance contracts such that (i) a^* is a Nash equilibrium of $\langle A, u + M \rangle$; (ii) $u_i(a^*) > u_i(\underline{\alpha})$ for every player i ; and (iii) $\underline{\alpha}$ is a Nash equilibrium of $\langle A, u + M' \rangle$ and $M'_i(\underline{\alpha}) \leq 0$, for every total transfer function M' formed from M by removing any of player i 's contracts, for each player i . We would aim to construct an equilibrium in which any disruption in contracting leads to a wave of messages that coordinates the players on $\underline{\alpha}$ rather than a^* . It is not clear whether condition (iii) can be satisfied always.

³⁵Incidentally, an SCSC institution would require knife-edge indifference conditions to deal with peripheral players. One could ask whether we could have disperse forcing arrangements as part of the on-path contracts in an SCSC institution and use a weaker form of virtual implementation, but there is still the problem of working out what behavior would look like in the production phase following contract declines, where penalties would be paid for some values of ϕ .




Figure 10: A setting with partial verifiability.

simple case of team production with partial output verification, where $n = 3$. The network and underlying game are shown in Figure 10.³⁶

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: $\bar{\omega} = \{(1, 1)\}$ and $\underline{\omega} = \{(0, 1), (1, 0), (0, 0)\}$; 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, in every equilibrium of the grand game, action profile $(1, 0, 0)$ is played with probability 1 in the production phase. The basic logic is easy to describe. 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 $\bar{\omega}$ and the transfer in the event of $\underline{\omega}$ must be at least 2. Such a margin gives player 2 the incentive to select high effort because, with player 3 choosing high effort, player 2's effort choice determines whether $\bar{\omega}$ or $\underline{\omega}$ 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 $\bar{\omega}$ or $\underline{\omega}$ is obtained. Yet $\underline{\omega}$ 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.

Here is the formal analysis. 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 $m_2^{12}(\bar{\omega})$ if and only if both workers choose high effort, and otherwise player 2 receives $m_2^{12}(\underline{\omega})$, player 2's incentive condition requires $\zeta m_2^{12}(\bar{\omega}) + (1 - \zeta)m_2^{12}(\underline{\omega}) - 2 \geq m_2^{12}(\underline{\omega})$, which simplifies to

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

³⁶See also Goldmanis and Ray's (2021) model of team production with partial verifiability and private bilateral contracting, which report general results.

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

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

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 contradicts rationality. Thus, workers select low effort for sure in equilibrium.

Other Comments

The contracts defined here specify balanced transfers (no money thrown away). One might ask whether the analysis would change if money burning were allowed or, in a related vein, if contracting partners could commit to make transfers to, but not from, third parties. In fact, expanding the range of contracts in this way would not affect the current analysis and results.

Further, it would not simplify the proof of the main result or make it a trivial finding. To see this, recall that weak implementation is easy to obtain in social-choice settings with common knowledge of the state, more than two agents, and the existence of outcomes that severely punish the players. The planner asks the agents to report the state. If they all send the same report then the planner compels the desired outcome for this state. If all but one agent gives the same report, so that a single agent reports differently, then the planner compels an outcome that severely punishes this agent (and possibly others). Money burning can be used to severely punish players, but such a mechanism cannot be replicated by a contracting institution as modelled here, because of the assumption that contracting is independent and voluntary.

There is one sense in which allowing for money burning could be marginally useful. Recall that in the SCO contracting institution there are rounds in which the feasible messages for a pair of players depend on their past messages to each other. For instance, if a pair made a conditional arrangement at some round $r < 0$, then they are restricted to silence in

rounds $r + 1, \dots, 0$. The reason for introducing this restriction was technical: It reduced the number of information sets in the grand game, simplifying the process of constructing a sequential equilibrium with the desired properties (which is complicated still). With money-burning, we could assume that the message spaces in each round are history independent, but obtain effectively the same restriction as before by having players burn money if players send non-null messages when not allowed.

References

- Armstrong, M. (2006), “Competition in Two-Sided Markets.” *The RAND Journal of Economics* 37 (3): 668–691.
- Bando, K. and T. Hirai (2021), “Stability and Venture Structures in Multilateral Matching,” *Journal of Economic Theory* 196: 105292.
- 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.
- Bramoullé, Y. and R. Kranton (2016), “Games Played on Networks,” in *Oxford Handbook of the Economics of Networks*, edited by Y. Bramoullé, A. Galeotti, and B. Rogers.
- 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. Jankó, A. Tamura, and A. Teytelboym (2018), “Trading Networks with Bilateral Contracts,” unpublished manuscript.
- Fleiner, T., R. Jagadeesan, Z. Jankó, 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.
- Goldmanis, M. and K. Ray (2021), “Team Incentives Under Private Contracting,” *RAND Journal of Economics*, 52 (2, Summer): 334–35.
- 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.
- Jackson, M.O. and Y. Zenou (2014), “Games on Networks,” in the *Handbook of Game Theory*, Vol. 4, edited by P. Young and S. Zamir.

- 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.
- 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 (2020), “Matching with Complementary Contracts,” *Econometrica* 88 (5): 1793–1827.
- Rostek, M. and N. Yoder (2022), “Matching with Multilateral Contracts,” manuscript in circulation.

- 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.
- Watson, J. (2023), “Partially Constructed Sequential Equilibrium,” unpublished draft.
- Weyl, G. (2010), “A Price Theory of Multi-Sided Platforms,” *American Economic Review* 100 (September): 1642