

Contractual Chains

Joel Watson*

2011 Preliminary draft, under revision
Please do not circulate or display publicly.

Abstract

This paper examines a model in which individuals contract with one another along the branches of a fixed network and then select verifiable productive actions with externally enforced transfers. Special cases of the model include most settings of contracting with externalities that have been studied in the literature. The model allows for a type of externality that the previous literature has not explored — where a party is unable to contract directly with others whose actions affect his payoffs. The paper explores the prospects for efficient outcomes under various contract-formation protocols (“contracting institutions”) and network structures. There are contracting institutions that always yield efficient equilibria for any connected network. A critical property is that the institutions allow for sequential contract formation or cancellation.

*UC San Diego; <http://econ.ucsd.edu/~jwatson/>. Aspects of the work reported here were developed while the author visited Yale in 2010 and the Center for Advanced Study in the Behavioral Sciences (Stanford) in 2012. The author thanks colleagues at these institutions and at UCSD for comments and support.

1 Introduction

In many contractual settings, there is multilateral productive interaction but the parties are not able to contract as a group. Instead, they must form contracts in a decentralized way, whereby pairs of agents establish contracts independently of one another. As an example, consider a subcontracting arrangement that is common in construction projects. A developer may contract with a general building contractor, who separately establishes agreements with subcontractors for the provision of some of the constituent components. This example features an externality, since the developer cares about the subcontractors' work (it influences his benefit of the completed project) but he doesn't contract directly with the subcontractors. The question addressed here is whether such an externality can be internalized and, if so, under what conditions.

The subcontracting example is one of many real settings in which some economic agents are unable to contract directly with those whose productive actions they care about. This feature is present in the internal organization of firms (where multiple workers have employment contracts with the firm but care about each others' actions), sales of goods exhibiting network externalities, private provision of public goods, commons problems, and the like. There are also related applications in which parties contract directly with those whose actions affect their welfare, but still they contract in a decentralized manner and cannot contract with everyone.

I develop a general model that covers these various examples and applications. In the model, a set of players interacts in an "underlying game" with externally enforced monetary transfers. Before the underlying game is played, the players can establish contracts. A contract specifies transfers as a function of the outcome of the underlying game. Due to physical limitations and the scope of the enforcement institution, contracting must take place in a decentralized and private way. Specifically, only bilateral contracts are feasible and, furthermore, there are limits on which pairs of players can communicate and establish contracts. These limitations are formally represented by a network structure and some "natural" contracting assumptions. Only linked pairs of players can communicate and contract.

To sort out the possible barriers to achieving efficient outcomes, and to relate my modeling exercise to the previous literature, it is useful to distinguish between two types of externalities. An *externality of non-contractibility* is present when a pair of linked players is not able to condition transfers on an action that is payoff relevant to them (either one of their own actions or the action of a third player). An *externality due to lack of direct links* is present when a player is unable to contract directly with others whose actions affect his payoffs. In addition to these externalities, a setting may exhibit *sparse contracting* in that some pairs of players may not be able to contract, irrespective of whether they care about each others' actions in the underlying game. Furthermore, it is important to account for the *contracting institution*, which specifies the protocol by which the players can communicate and establish contracts.

Much of the related literature comprises special cases of the general model presented here. This includes models of common agency (Bernheim and Whinston 1986a,b), com-

mon principal (Segal 1999), games played through agents (Prat and Rustichini 2003), and games with side payments (Jackson and Wilkie 2005). In these models, there are no externalities due to lack of direct links. Rather, the focus is on externalities of non-contractibility (as in Segal 1999), or problems of sparse and/or decentralized contracting as opposed to contracting by all players as a group. Each of these models specifies a single contracting institution; it is standard to focus on a protocol with simultaneous commitments or offers. Holding aside common agency, which produces efficient equilibria, the literature finds that inefficient outcomes are generally inescapable even without externalities of non-contractibility. The related literature is discussed in more detail in the next section.

As noted above, I am mainly interested in externalities due to lack of direct links, but my modeling exercise also provides insights regarding decentralized and sparse contracting in general. This means that the external enforcer observes the action profile played in the underlying game. In line with much of the related literature, I also assume that every contractual relationship can condition transfers on the action profile, so there are no externalities of non-contractibility.

I take a new approach to the analysis of decentralized contracting — one that emphasizes the role of the contracting institution. The issue of efficiency is posed as a possibility question: Given that contracting must take place in a decentralized manner, is there a natural contracting institution under which efficient outcomes can be achieved across a wide range of underlying games and networks?¹ If so, what are the minimal requirements of the network? In technical terms, the first question asks whether there is a contracting institution that *implements efficient outcomes*, meaning that for every underlying game and every allowed network structure, there is a perfect Bayesian equilibrium in which an efficient action profile is played in the underlying game. The second question asks what kind of network structures can be allowed.

I determine that the answer to the first question is affirmative, and the answer to the second question is that the network must be connected. These are provided in the paper's single theorem, with some supporting examples. Thus, there exists a contracting institution with the following property: For any underlying game and network of contractual relationships, an efficient outcome can be obtained in equilibrium as long as all players are indirectly connected to each other on the network. Furthermore, an example and a lemma reveal a critical property of such a well-functioning contracting institution: that it allow for sequential communication and contracting.

The result has several implications. First, with the right kind of contracting institution, contractual chains can form to internalize externalities due to lack of direct links. Second, these chains must be developed in the context of a connected network and sequential contracting process. Third, sparse contracting is not a barrier to efficiency. Thus, some of the previous literature's negative findings about efficiency (such as in Prat and Rustichini 2003) depend on the specific contracting institution that was assumed for the analysis.

¹Jackson and Wilkie (2005) allude to this tack in their comment that the "structure" of contracting is critical to whether efficient outcomes will be reached, although they argue that different protocols would not change the supported outcomes in their model.

The next section defines the main ingredients of the general model. Section 3 gives some applications and discusses the related literature. In Section 4, I present some simple examples that demonstrate the challenges of implementing efficient outcomes. Section 5 contains the formal definition of contracting institutions, the theorem, and the proof of the theorem. In Section 6, I offer additional comments on the related literature and I discuss further steps in the research program.

2 The Setting

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

An external enforcer compels monetary transfers between the players, and he does so as directed by contracts that the players form. The enforcer does not observe the payoff function u . Otherwise, the contracting environment has *full verifiability* in that the enforcer observes the action profile $a \in A$. Thus, transfers can be freely conditioned on the outcome of the underlying game. However, contracts can be formed only in a restricted set of bilateral relationships. This set of contractual relationships is given by a fixed network $L \subset N \times N$, with the interpretation that players i and j can form a contract if and only if $(i, j) \in L$. Contracting by larger groups of agents is not possible. I sometimes call a pair $(i, j) \in L$ *contracting partners*.

Network L is undirected and thus symmetric, so $(i, j) \in L$ implies that $(j, i) \in L$. Also, L is generally not transitive; thus, player i may be able to contract with player j , and player j may be able to contract with player k , whereas player i is unable to contract with player k .

Contracting partners can condition transfers between them on actions taken by third parties. For example, a contract between players 1 and 2 could specify that player 1 must pay an amount to player 2 in the event that player 3 selects a particular action in the underlying game. However, assume that a contract may not impose a transfer on any third party. For instance, the enforcer will not enforce a contract between players 1 and 2 that specifies a transfer to or from player 3.² Thus, a contract between players i and j is a function $m: A \rightarrow \mathbb{R}_0^n$, where $m_k(a) = 0$ for all $k \notin \{i, j\}$ and every $a \in A$. Here \mathbb{R}_0^n is the subset of vectors in \mathbb{R}^n whose components sum to zero (balanced transfers). Denote by \underline{m} the null contract that always specifies a transfer of zero.

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

²Due to full verifiability, the analysis here would not be affected if one allowed contracting partners to commit to make transfers to, but not from, third parties.

such a set, let

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

A set \mathcal{M} is *feasible* if, for every $m \in \mathcal{M}$, there exist $(i, j) \in L$ such that m is a contract between i and j .

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

The protocol by which the players form contracts will be called the *contracting institution*. It is a mechanism \mathcal{C} that specifies how the players can send messages to one another and how these messages translate into contracts that the enforcer enforces. I impose the following *natural contracting assumptions* on \mathcal{C} :

- Players have the option of rejecting contracts.
- Contracting is private, so (a) the contracting institution allows a player to receive messages from only those to whom he or she is linked in the network, and (b) a player cannot observe messages exchanged between other players.
- Contracting in different relationships takes place independently, so the contract formed between players i and j does not depend on the messages sent by, or received by, any other player k .

Call a contracting institution *natural* if it satisfies these assumptions. Section 5 provides the formal development, along with a more detailed account of the contracting phase.

The natural contracting assumptions represent the physical constraints and the limitations of the external enforcement system that necessitate decentralized contracting rather than centralized planning. The first assumption ensures that the definition of “contract” is conventional, in that it requires the consent of both parties. This is a standard requirement for external enforcement in modern legal systems. The second assumption imposes the network limitation — that communication and contract formation take place only between linked players — and thus it describes real physical constraints on how the players can interact. The third assumption embodies the principle that any two contracting parties are free to form whatever contract they desire, uninhibited by others in the society.

One can also motivate the third assumption on the basis of limitations of the enforcement system. Consider a legal system in which courts resolve disputes by hearing evidence and then issuing judgments that compel transfers between the parties. In reality, contracting partners rarely go to court. Instead, they voluntarily make the transfers that their contracts require, with the understanding that failure to comply would trigger a lawsuit, at which point the court would compel the required transfers. When a pair of contracting partners appears in court, the judge can observe their contract and the verifiable outcome of the underlying game (for evidence of these can be provided by the contracting parties). However,

the judge will not readily observe the contracts written in other contractual relationships, or at least it may be prohibitively costly for the parties to gather and provide such evidence to the court. Thus, it is feasible to enforce transfers as a function of only the outcome of the underlying game and not the messages sent in other contractual relationships.³

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

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

Note that, because of private contracting, the players have asymmetric information at the time of productive interaction. For example, player i does not observe the contract formed between two other players j and k . I analyze behavior using the concept of perfect Bayesian equilibrium (PBE) in pure strategies.

In the analysis that follows, I hold fixed the number of players n and the space of action profiles A . The primary objective is to evaluate the performance of a given contracting institution across various networks L and payoff functions u . Let us say that a contracting institution \mathcal{C} and network L *implement efficient outcomes* if for every payoff function $u : A \rightarrow \mathbb{R}^n$, there is a PBE of the entire game in which an efficient action profile a^* is played. Let us say that (\mathcal{C}, L) *implements ε -efficient outcomes* if for every payoff function $u : A \rightarrow \mathbb{R}^n$, there is a PBE of the entire game in which an efficient action profile a^* is played with probability at least $1 - \varepsilon$. Finally, let us say that (\mathcal{C}, L) *virtually implements efficient outcomes* if this contracting institution and network implement ε -efficient outcomes for all $\varepsilon > 0$.

3 Applications and Related Literature

Below are brief descriptions of a few applications, along with some notes on the related literature. Figures 1 and 2 illustrate the network structures in the applications.

Common agency

In the common agency model introduced by Bernheim and Whinston (1986a,b), there is a single “agent” and multiple “principals.” In the underlying game, only the agent has an action; the principals all care about the agent’s action. The network of links is a star in

³In reality, it may be possible to provide evidence of communication in “nearby” contractual relationships, such as ones involving one of the current disputants. For instance, in a contractual dispute between players 1 and 2, player 2 might be able to introduce some evidence about her communication with player 3. Such a possibility is not formally considered here, but it may be interesting to consider in future work.

which the agent is directly linked to every principal and there are no other links. The contracting institution involves the principals simultaneously making commitments regarding how much money to transfer to the agent as a function of the agent's action. The agent does not have the option of declining the principals' offers, so in a sense the setting has unilateral commitments rather than contracting; however, the transfers are constrained to be nonnegative, so the agent would never refuse.

The common agency setting does not have externalities of the types described in the introduction. The principals can all contract on the action that they care about (the agent's action). Furthermore, the principals can all contract with the party who takes this action (the agent). There is sparse contracting since the principals cannot contract with each other (they are not connected in the network). Efficient equilibria always exist in settings of complete information, so the large literature on common agency focuses mainly on interesting questions of equilibrium characterization. Recent entries, such as Prat and Rustichini (1998) and Bergemann and Valimaki (2003), examine settings in which the principals make their commitments sequentially; as with the static case, there are efficient equilibria.⁴

Common principal

In Segal's (1999) model of contracting with externalities (see also Segal and Whinston 2003), a principal and multiple agents are connected by a star network with the principal in the center. The principal is the only player with an action in the underlying game. The contracting institution requires the principal to simultaneously offer contracts to the agents, and then the agents decide whether to accept or reject the contracts. Cases of private and public contracting are studied. The principal's action is multidimensional, with one component for each agent (representing, for instance, the amount to trade with this agent). Agents care about not only their own level of trade with the principal, but they also care about the principal's level of trade with other agents.

Segal (1999) assumes that a contract with one agent cannot specify transfers as a function of the action components for other agents, so this setting exhibits externalities due to non-contractibility. Unless the principal can make the trade with one agent conditioned on communication with another (a departure from natural contracting), equilibria are typically inefficient. The common principal framework has no externalities due to lack of direct links, although contracting is sparse.⁵

A variation on Segal's model involves a setting in which the agents have actions in the underlying game and they care about each others' actions. This case, which to my knowledge the previous literature has not examined, presents an interesting version of the lack-of-direct-links externality whereby contracts are clustered on a central party (a star network).⁶

⁴The literature also looks at settings with incomplete information, where efficient outcomes are generally not attained. For a survey of the literature, see Martimort (2007).

⁵See Galasso (2008) for a recent study that looks at various bargaining protocols and provides additional references. Moller (2007) performs an analysis of the common-principal problem with sequential contracting.

⁶The modeling exercise herein does not have much to add to what the literature has already learned about the common-principal setting in the case where only the principal has productive actions. Under the

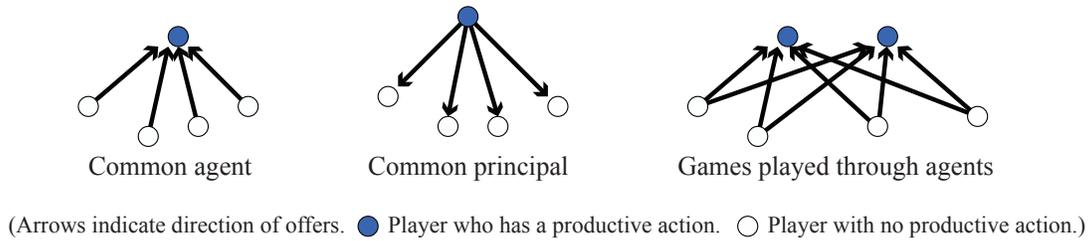


Figure 1: Applications studied in the literature.

Games played through agents

Prat and Rustichini (2003) examine a setting that extends the common agency framework to have multiple agents. Only the agents have actions in the underlying game. The other players (the principals) simultaneously make unilateral commitments regarding how much money to transfer to the agents as a function of the agents’ actions. Prat and Rustichini’s model covers some important applications, including settings of lobbying, multiple auctions, vertical restraints, and two-sided matching. The authors examine the case in which a transfer to one agent can be conditioned on the actions of all of the agents, so there are no externalities of non-contractibility. There are also no externalities due to a lack of direct links, because it is assumed that each agent does not care about the actions of the other agents. Contracting is sparse since agents cannot contract with one another and, likewise, principals cannot contract with one another. Still, efficient equilibria are not guaranteed, although they arise in some cases.

Subcontracting and supply chains

In the simplest version of this application, a buyer contracts with one or more “first-level” suppliers to provide related goods or services. These suppliers do not perform all of the required tasks, however. The first-level suppliers contract with a second level of suppliers who can perform the additional tasks. The second-level folks may, in turn, contract with agents at a third level, and so on. The buyer’s payoff depends on all of the suppliers’ actions, but the buyer can contract only with the first-level suppliers. One example is a typical construction project, whereby a property owner contracts with a commercial building contractor, who separately contracts with “subcontractors” to provide individual components.

In the setting just described, the various suppliers take actions in the underlying game

assumption that a principal-agent pair can specify transfers as a function the principal’s entire vector of trades (as I assume here), the common-principal model converges to the common-agency model except with a different contracting institution. In the common-agency case, the player in the center (called the agent) has the productive action and this player receives offers from the others. In the common-principal case, the player in the center (called the principal) has the productive action but this player makes offers to the others. Since the common-agency models have efficient equilibria, it is clear that in any setting with a star network and where the center player is the only one with an action in the underlying game, there exists a natural contracting institution that implements efficient outcomes.

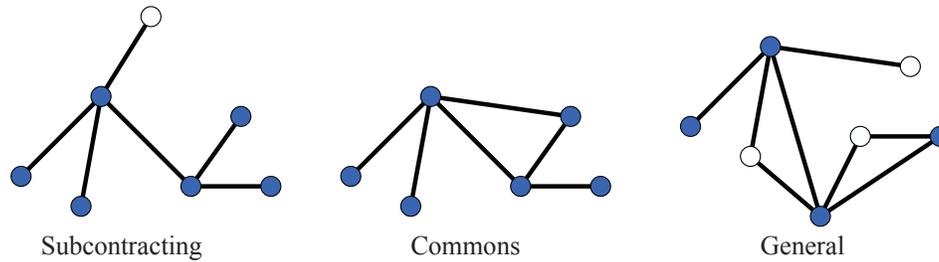


Figure 2: Other applications.

and the buyer obtains a benefit as a function of the action profile. The network limits the buyer from contracting directly with all of the suppliers, but there is a path from the buyer to each supplier through the subcontracting relationships.

Commons

Commons problems (most prominently described by Hardin 1968, less prominently by Lloyd 1833) feature a dense system of interdependence, whereby each player's action affects the welfare of everyone else in the society. Standard examples include extraction of an exhaustible common-property resource, resource use in an open-access environment, population growth, and other individual activity that creates a public good or bad. In the underlying game of a commons problem, every player has an action and everyone's payoff depends on everyone else's actions (sometimes through an aggregate). Although in some cases it may be possible for the entire society to establish a contract at once (the grand coalition), in other cases agreements can be made only in smaller groups. The framework herein addresses the case in which all contracting is bilateral.

Organizations and the firm

An important, but relatively undeveloped, theory of the firm views an organizational structure as a nexus of contracts (Jensen and Meckling 1976). See Laffont and Martimort (1997) for a survey of the literature.⁷ Because of various transaction costs, contracting takes place bilaterally or in small groups rather than by the entire set of players collectively. In this application, the underlying game represents the productive interaction between the players, and the network represents the limitations on contracting that arise due to transaction costs. A key question is whether the network leads to a nexus of contracts that supports efficient outcomes. The special case of a star network captures the example of a team of workers who all have bilateral employment contracts with a manager, and this is an example of a common-principal model where agents have productive actions.

Other Models in the Literature

I conclude this section with notes on some other related papers in the literature. Jackson and Wilkie (2005) study a model in which players can make unilateral commitments before

⁷Other organizational structures comprise networks of contracts but may not be regarded as traditional firms. Cafeggi (2008) provides one legal perspective.

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

Ellingsen and Paltseva (2011) expand on Jackson and Wilkie's (2005) model by allowing the players to form agreements rather than just make binding promises. A contract is formed only if both players accept the deal. Ellingsen and Paltseva allow unilateral promises as well. They show that, regardless of the underlying game, there exist efficient equilibria.⁸

Peters and Szentes (2008) examine a setting in which players can make unilateral commitments about how they will play in the underlying game. The authors call these "contracts," but let me describe them as "promises" to relate them to the objects in Jackson and Wilkie's model. The key feature in Peters and Szentes' model is that each player's promise 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 modeling exercise reveals that interactive contracts require the external enforcement system to develop a sophisticated language for the cross-referencing of promises.

The model herein shares a theme of Jackson and Wilkie (2005) in terms of looking at a wide class of contracts with full verifiability. I carry the theme further by exploring various contracting institutions and decentralized contracting along network links. Also, in comparison to the papers just discussed, I examine contracts (rather than unilateral promises) and private contracting. I assume that contracts can condition transfers on the verifiable productive actions of others but not on the contracts written elsewhere in society.

Further afield is the growing literature on coalitional bargaining. In these models, the grand coalition can form a contract (so centralized contracting is possible) but subgroups of players can shape the final agreement by first making agreements in their smaller coalitions. The incentives of coalitions to manipulate in this way sometimes precludes the attainment of an efficient outcome.⁹ De Fontenay and Gans (2007) depart from the coalitional bargaining literature by presenting a model with only bilateral bargaining by players linked on a network. There are externalities due to both the lack of direct links and non-contractibility (links pairs can contract on only their own "trade" and cannot condition transfers on actions taken in other relationships). Disagreement permanently severs the link and leads the other

⁸Related papers from the prior literature include Guttman (1978), Danziger and Schnytzer (1991), Guttman and Schnytzer (1992), Varian (1994), and Yamada (2003).

⁹The typical model has a public, dynamic negotiation process. Payoffs are described in terms of a characteristic function over coalition structures or a state variable that can be changed over time by the active coalition. A representative sample of contributions is: Chatterjee et al. (1993), Seidmann and Winter (1998), Gomes (2005), Gomes and Jehiel (2005), Bloch and Gomes (2006), Hafalir (2007), and Hyndman and Ray (2007).

pairs to contract from scratch. Contracting takes place sequentially and privately, except that everyone knows when a link is broken. The authors show that, under passive beliefs, equilibria produce the Myerson value of a characteristic function defined by bilaterally efficient actions, but efficiency generally cannot be attained.

4 Examples and Intuition

In this section, I present a few simple examples and some intuition to show the difficulties of implementing efficient action profiles. I begin with a minimal requirement on the network. Note that if every player with a nontrivial action set is a member of at least one bilateral contractual relationship, then there exists a set of feasible contracts \mathcal{M} that support a^* . One can specify, for instance, that player i 's contract with some other player specify a large punishment if player i fails to choose a_i^* . But if connections are lacking to internalize externalities, then players may not have the incentive to form contracts that support a^* . Here is an example:

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

Imagine that players 1 and 2 have a contract that transfers 5 from player 1 to player 2 in the event that $a_1 = c$. This contract would lead player 1 to select b in the underlying game, which is the efficient outcome. But under no natural contracting institution would players 1 and 2 form such a contract. To see this, notice that player 1 can guarantee himself a payoff of 1 by refusing to contract and then choosing action c . Likewise, player 2 can guarantee himself 0 by refusing to contract. Under any contract between players 1 and 2 that induces player 1 to select b in the underlying game, the joint payoff of these players would be 0, so it is irrational for one or both players to form such a contract. Clearly the problem is that, without a contractual chain from player 1 to player 3, there is no way for player 3 to indirectly compensate player 1 for action b . Thus, a connected network is a minimal criterion for the attainment of efficient outcomes.

Assume, therefore, that the network L is connected and let us consider the issue of timing in the contracting phase. In many of the models in the existing literature, players form contracts (or make unilateral commitments) simultaneously. As the next example shows, with simultaneous contracting it is generally not possible to achieve efficient outcomes.

Example 2: Consider the same underlying game as in Example 1, so that $n = 3$, $A_1 = \{b, c\}$, players 2 and 3 have no productive actions, $u(b) = (0, 0, 4)$, and $u(c) = (1, 0, 0)$. Suppose the network is $L = \{(1, 2), (2, 3)\}$, so the network is connected.

In this example, the efficient outcome cannot be obtained with a natural contracting institution that has one-shot, simultaneous contracting. To see this, consider the contracting institution in which linked players form contracts simultaneously via the "Nash demand game." That is, players 1 and 2 each name contracts that map A_1 to transfers between

them, and simultaneously players 2 and 3 do the same. Note that player 2 declares two contract demands, one for player 1 and one for player 3.

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

Unfortunately, this specification of behavior in the contracting phase is not consistent with individual incentives. If Player 3 deviates by naming the null contract in his relationship with player 2, then he will get a payoff of 4 rather than the payoff of $4 - \tau'$ that he would have received under the specified behavior. He gets 4 because his deviation does not disrupt the formation of contract m between players 1 and 2. At the productive phase, player 2 would not know that any deviation occurred, and he would select b.

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

On the other hand, if the order of contracting is reversed, so that players 2 and 3 first negotiate, followed by players 1 and 2, then there is an efficient equilibrium. In this equilibrium, player 2 will not make a deal with player 1 unless he was previously successful in contracting with player 3. But although this contracting institution functions well for the externality of Example 2, where player 1's action of b benefits player 3, it would not deliver an efficient outcome if the externality ran the other way.

There are two points here. First, to obtain efficient outcomes a natural contracting institution must allow contracts to be established or adjusted in sequence, so a player can condition his decision to contract with another player on the success or failure of his negotiations with others. Second, the order matters critically. If a contracting institution specifies a fixed order in which the various linked pairs negotiate and form contracts, there will always be an underlying game in which this order fails to yield an efficient equilibrium. In fact, underlying games with complex externalities are particularly problematic. The next example — a commons setting — is such a case.

Example 3: Suppose $n = 3$ and $A_i = \{b, c\}$ for $i = 1, 2, 3$. Payoffs in the underlying game are given by: $u(b, b, b) = (1, 1, 1)$, $u(c, c, c) = (0, 0, 0)$, and if at least one player selects c and another player selects b then those selecting c get 4 and those selecting b get

–12. The network is as before: $L = \{(1, 2), (2, 3)\}$.

If the contracting institution is rigid as to the order in which players can communicate and establish contracts, then it cannot achieve the efficient action profile (b, b, b) . This is because the player who is last to contract can decline to do so and choose action c . Any player who is not linked with this last player will not know of the deviation and will still select b in the productive phase, so the deviator would get a payoff of 4 and thus strictly prefer to deviate.

This example's complex externalities pose problems even for institutions that give relationships flexibility over the time at which they establish contracts. Consider, for instance, an institution that allows for contracting in discrete rounds. Suppose that there is flexibility as to when any given relationship can establish a contract. For example, in each round a linked pair of players interacts as in the Nash demand game over contracts. If the players name the same contract then this contract goes into force.¹⁰ Such an institution would allow the players to coordinate on forming contracts in sequence (with, say, players 1 and 2 establishing their contract first, followed by players 2 and 3) or at the same time. But suppose players i and j are supposed to form their contract at the end of the sequence. If player j deviates by naming the null contract, then he is poised to get the cheater's payoff of 4 if player k (the third player) will still choose b in the productive phase.¹¹

The logic just presented for Example 3 is incomplete and merely meant to convey the difficulties of implementing efficient outcomes. At this point, it is clear that a well-functioning contracting institution must facilitate sequential formation or adjustment of contracts. The question from here is whether there is a natural contracting institution that performs well.

5 Implementation with Natural Contracting Institutions

This section presents the main result of this paper: There exist a contracting institution that, along with any connected network, virtually implements efficient outcomes. I begin the section with a formal description of contracting institutions and the natural contracting assumptions. I then present the theorem. At the end of this section is a lemma on the necessity of sequential contract formation or adjustment.

Contracting Institution Formalities

A contracting institution is an extensive game form with payoff-irrelevant messages that map to contracts enforced by the external enforcer. I constrain attention to a particular

¹⁰One might assume that the contract that the players agreed to supersedes any previous agreement, but this is not necessary to what comes next.

¹¹If the contracting institution allows further time, so that player i can let player k know of i 's failure to contract with j , then perhaps player k could be induced to select c in productive phase, undermining player j 's change of gaining from the deviation. But player i may not have the incentive to give player k this news, for i and j would both gain in the underlying game if they select c while player k chooses b .

class of game forms in which the players send messages in discrete rounds $1, 2, \dots, R$, where $R = \infty$ is allowed. There is no discounting. In every round, each player i sends a vector of messages, one message for every other player. Let Λ_{ij}^r denote the set of feasible message from player i to player j in round r , and let $\lambda_{ij}^r \in \Lambda_{ij}^r$ denote the message that player i sends. The message space is arbitrary and can include descriptions of contracts. Assume that there is a “null message” $\underline{\lambda}$ that is an element of Λ_{ij}^r , for all i, j, r .

Let z_{ij}^r denote the sequence of messages sent from player i to player j from round 1 to round r . That is,

$$z_{ij}^r = (\lambda_{ij}^1, \lambda_{ij}^2, \dots, \lambda_{ij}^r) \in Z_{ij}^r \equiv \Lambda_{ij}^1 \times \Lambda_{ij}^2 \times \dots \times \Lambda_{ij}^r.$$

Let $\underline{z}^r = (\underline{\lambda}, \underline{\lambda}, \dots, \underline{\lambda})$ be the sequence of r null messages. A full history of messages through round r is given by $h^r = \{z_{ij}^r\}_{i \neq j}$. A corresponding history of messages between players i and j is given by (z_{ij}^r, z_{ji}^r) .

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

The contracting institution specifies a mapping from the outcome of the contracting phase (the full history of communication through R , as well as the random draw σ) to the contracts formed between the various pairs of players. For any such outcome (h^R, σ) , let $\mu(i, j, h^R, \sigma)$ denote the contract formed between players i and j . Because $\mu(i, j, h^R, \sigma)$ and $\mu(j, i, h^R, \sigma)$ would be the same contract, we can avoid the redundancy by defining $\mu(i, j, \cdot)$ only for $i < j$. Then, for a given sequence of messages and random draw (h^R, σ) , the function describing the sum of contracted transfers is

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

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

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

Part b of the private contracting assumption requires that, at the end of round r , player i 's personal history is exactly the sequence of messages to and from this player:

$$h_i^r = \{z_{ij}^r, z_{ji}^r\}_{j \neq i} \in H_i^r \equiv \times_{j \neq i} (Z_{ij}^r \times Z_{ji}^r).$$

It follows that player i 's strategy in the contracting phase specifies, for each round r , a mapping from H_i^r to $\times_{j \neq i} \Lambda_{ij}^r$.

The assumption of independent contracting requires that the contract between any two players i and j depends only on the sequence of messages between these two players. That

is, $\mu(i, j, h^R, \sigma)$ depends on h^R only through the component (z_{ij}^R, z_{ji}^R) . We can therefore write $\mu(i, j, (z_{ij}^R, z_{ji}^R), \sigma)$.

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

$$\mu(i, j, (\underline{z}^R, z_{ji}^R), \sigma) = \mu(i, j, (z_{ij}^R, \underline{z}^R), \sigma) = \underline{m},$$

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

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

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

Efficient Implementation

With the definition of natural contracting institutions in hand, I proceed to the main result.

Theorem: *Fix n and A . There exists a natural contracting institution \mathcal{C} such that, for every connected network L , (\mathcal{C}, L) virtually implements efficient outcomes.*

Note that the natural contracting institution identified here performs well for *all* connected networks and all underlying games. It is not defined contingent on the network.

The rest of this subsection contains the proof, with one minor part contained in the Appendix. The proof features a contracting institution with a two-part messaging structure, where tentative contracts are formed and then players have the option of canceling them. In the first round of messaging, the contracting pairs engage in a Nash demand protocol that determines their tentative contracts. In later rounds the players can unilaterally cancel these contracts. Contracts that are not canceled will then be enforced.

Formally, define institution \mathcal{C}^* as follows. There is open-ended messaging, so that $R = \infty$. The public randomization device is a draw from the uniform distribution on the unit interval. Message spaces in the first round are such that the players can name arbitrary mappings from $A \times [0, 1]$ to transfers between the contracting parties. That is, for each pair of players (i, j) , Λ_{ij}^1 is the set of mappings of the form $g : A \times [0, 1] \rightarrow \mathbb{R}_0^n$, where $g_k(a, \sigma) = 0$ for all $k \notin \{i, j\}$, every $a \in A$, and every $\sigma \in [0, 1]$. Identify the constant function $g \equiv (0, 0, \dots, 0)$ as the null message $\underline{\lambda}$ for round 1. In every round $r > 1$, the message space is defined as $\Lambda_{ij}^r \equiv \{\text{cancel}, \underline{\lambda}\}$.

For each pair (i, j) , function $\mu(i, j, \cdot)$ is defined in a straightforward way. Suppose players i and j named the same function g in round 1 and they sent message $\underline{\lambda}$ in all later rounds. Then for random draw σ , the contract between them is $m \equiv g(\cdot, \sigma)$. If the players named different functions in round 1 and/or one or both of them sent the message “cancel” in a later round, then their contract is the null contract \underline{m} .

Contracting institution \mathcal{C}^* is natural. To see this, note that players have the option of rejecting contracts by sending message $\underline{\lambda}$ in the first round. Contracting is private because we look at the effective game form that restricts messages to $\underline{\lambda}$ between pairs of players that are not in the network. Finally, contracting in different relationships takes place independently, since the contract between players i and j is a function of only the messages that players i and j exchange.

Fix contracting institution \mathcal{C}^* . Take any payoff function u and any connected network L . We must show that, for any number $\varepsilon > 0$, there is a PBE of the entire game in which an efficient action profile a^* is played with probability of at least $1 - \varepsilon$. Let a^* be any efficient action profile. I proceed by first analyzing two special cases.

Case 1: *The underlying game $\langle A, u \rangle$ has a pure-strategy Nash equilibrium \underline{a} , and $\underline{a}_i \neq a_i^*$ for each $i \in N$.*

In this case, the efficient action profile has all of the players selecting actions that differ from their Nash equilibrium actions. I will show that there is a PBE in which a^* is played with probability 1. The first step is to find a minimally connected sub-network $K \subset L$, where each pair of players is connected (indirectly or directly) by exactly one path. That is, for each pair of players i and j , there is exactly one sequence $\{k^t\}_{t=1}^T \subset N$ with the following properties: $k^1 = i$, $k^T = j$, and $(k^{t-1}, k^t) \in K$ for all $t = 2, 3, \dots, T$. The sequence $\{k^t\}_{t=1}^T$ is called the *path from i to j* . A minimally connected subnetwork K exists by construction and is defined to be symmetric.

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

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

In words, $\beta(i, j, K)$ is the set of players that, relative to player i , are on “the other side” of network K from player j . Note that $\beta(i, j, K) \cap \beta(j, i, K) = \emptyset$ and $\beta(i, j, K) \cup \beta(j, i, K) = N$. Also, for each player i , define $K^i \equiv \{j \mid (i, j) \in K\}$ as player i ’s *active contracting partners*, which is the set of players that player i is supposed to establish non-null contracts with.

I next specify a set of contracts $\mathcal{M} = \{m^{ij}\}_{i < j}$, where m^{ij} denotes the contract for the pair (i, j) . In some places below I refer to the contract m^{ij} without specifying $i < j$; in the case of $i > j$, it is understood that $m^{ij} = m^{ji}$ but we do not include m^{ij} in \mathcal{M} to avoid double counting in the sum M . Let γ be a large number, exceeding the maximum difference between payoffs in the underlying game. We can find a set of contracts $\mathcal{M} = \{m^{ij}\}_{i < j}$ with these properties:

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

$$m_i^{ij}(a_{\beta(i,j,K)}, a_{\beta(j,i,K)}^*) = -\gamma \quad \text{for all } a_{\beta(i,j,K)} \neq a_{\beta(i,j,K)}^*, \quad \text{and}$$

$$m_i^{ij}(a_{\beta(i,j,K)}^*, a_{\beta(j,i,K)}) = \gamma \quad \text{for all } a_{\beta(j,i,K)} \neq a_{\beta(j,i,K)}^*.$$

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

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

Assuming γ is large enough, clearly a^* is a Nash equilibrium of $\langle A, u + M \rangle$, so the set of contracts supports a^* . I next show that there is a PBE of the entire game in which these contracts are formed and a^* is played. Here is a partial description of the strategies:

1. Players are supposed to name the contracts \mathcal{M} in round 1. That is, for all $i < j$, players i and j send the same message to each other in round 1: the function g specifying $g(a, \sigma) = m^{ij}$ for all σ . In each subsequent round r , each player i sends message $\underline{\lambda}$ to all of his active contracting partners (those in K^i), as long as all of them played as specified in round 1 and the messages they exchanged between rounds 1 and $r-1$ were all $\underline{\lambda}$. Otherwise, player i sends message “cancel” to all of his active contracting partners. As for any other player $j \notin K^i$ but with $(i, j) \in L$, player i sends message “cancel” in each round $r > 1$.

2. In the production phase, each player i selects a_i^* as long as there were no deviations in messages sent to and from player i in the contracting phase. On the other hand, if all of player i 's contracts are null (for instance, due to cancellations), then player i selects \underline{a}_i . Finally, take an intermediate case in which player i 's contracts with some players are as specified in \mathcal{M} but his contracts are null with all other players. That is, consider any set J that is a proper subset of

$$\{j \mid m^{ij} \neq \underline{m}\} \setminus \{\emptyset\}$$

and define $J' \equiv N \setminus J$. Then let $\bar{J} = \cup_{j \in J} \beta(j, i, K)$ be the set of players on the other side of network K from i , relative to the players in J . Let $\bar{J}' = \cup_{j \in J'} \beta(j, i, K)$. Further, define $M' \equiv \sum_{j \in J'} m^{ij}$ to be the sum of the existing contracts. In this case, let player i select a best response to $(\underline{a}_{\bar{J}}, a_{\bar{J}}^*)$, where player i 's payoffs is defined by $u_i + M'_i$.

Part 2 here does not explain what player i should do in the production phase in a situation in which he has formed a contract \hat{m}^{ij} with some player j such that $\hat{m}^{ij} \neq m^{ij}$ and $\hat{m}^{ij} \neq \underline{m}$. Note that this contingency could be reached only if player i and j simultaneously both deviated in the first round of the contracting phase, so the specification of subsequent behavior in the production phase will not be critical to the equilibrium construction. But it must be done in a way that is sequentially rational, and this is described in the Appendix.

Note that if some player i has contractual relationships with both players j and k , and if k deviates in a message to i , then player i is supposed to cancel the contract with both players j and k , as well as with all of his other active contracting partners. In this way, a deviation by one player will trigger a contagious sequence of contract cancellations that flows across the network.

The next step in the equilibrium construction is to describe the players' beliefs. Suppose that, throughout the game, each player believes that there are no deviations other than any that were observed by this player. For instance, player i believes that all past messages exchanged between other players j and k were as specified above. Also suppose that players ignore deviations by linked players who are not active contracting partners; this is not critical but it keeps the equilibrium simple in that signaling occurs only through cancellations by contracting partners. These beliefs are consistent with Bayes' rule by construction.

The behavior described above is sequentially rational. To see this, first note that the behavior specified for the production phase has each player best responding to his belief about the actions of the other players. For instance, player i observes no deviations in the contracting phase then he believes that contracts \mathcal{M} were formed and the other players will select a_{-i}^* , to which a_i^* is a best response given player i 's induced payoff function $u_i + M_i$. If player i were to deviate from a_i^* , he would have to pay γ to each of his contracting partners in K because he has assurance contracts with them.

Likewise, if all of player i 's contracts are null then this player believes that all other contracts in the society are also null, due to the contagious sequence of cancellations. Thus, player i believes that the other players will select \underline{a}_{-i} in the production phase, to which \underline{a}_i is a best response since that his payoff is given by u_i . In the intermediate case in which some of player i 's contracts were nullified out of equilibrium, player i believes that the other players will choose $(\underline{a}_{-j}, a_{-j}^*)$, and player i is to select a best response.¹² Contingencies involving multilateral deviations in the first round of the contracting phase are discussed in the Appendix.

As for rationality in the contracting phase, first examine player i 's specified behavior with any player $j \notin K^i$ but with $(i, j) \in L$. These pairs are supposed to name the null contract and cancel in every round, which is clearly rational for i under the expectation that j will do so. Regarding player i 's more substantive relation to other players, consider three options for player i . First, player i could behave as prescribed. Given that the other players do the same, player i can anticipate getting the payoff $u_i(a^*) + M_i(a^*)$.

¹²In this case, calculating player i 's best response for payoff function $u_i + M_i'$ is the same as for u_i , because player i expects that he will have to pay the penalty γ to all those with which he has non-null contracts, regardless of his action.

Second, player i might consider deviating in the contracting phase with a strict subset of contracting partners. That is, player i could either no match the contract named in round 1 by i 's contracting partner, or he could later cancel a contract. Such a move would serve only to create a null contract where there was supposed to be a non-null contract. Suppose that player i nullifies the contract with player j but forms the specified contract m^{ij} with some other player k . Then a cancellation sequence will ensue for all of the players in the set $\beta(j, i, K)$ and these players will select $\underline{a}_{\beta(j, i, K)}$ in the underlying game. Further, the players in the set $\beta(k, i, K)$ observe no deviations and will select $a_{\beta(k, i, K)}^*$ in the underlying game. Because m^{ij} is an assurance contract, player i will then have to pay the penalty γ to player k , regardless of what happens in player i 's other contractual relationships. Furthermore, no players will be paying such a penalty to player i . Since γ is large, player i strictly prefers not to deviate in this way.

Third, player i could, in any round of the contracting phase, nullify all of his contracts. Given that the other players behave as specified, if player i does this then it leads to a contagious sequence of contract cancellations across the entire network. The players then select \underline{a} in the underlying game and player i obtains the payoff $u_i(\underline{a})$. Because $u_i(a^*) + M_i(a^*) \geq u_i(\underline{a})$, player i has no incentive to deviate in this way.

Case 2: *The underlying game $\langle A, u \rangle$ has a pure-strategy Nash equilibrium \underline{a} and, for each $i \in N$, either $\underline{a}_i \neq a_i^*$ or a_i^* is a best response to a_{-i} .*

I will show that in this case, as in Case 1, there is a PBE in which a^* is played with probability 1. For any $J \subset N$, let $U_J(a) \equiv \sum_{i \in J} u_i(a)$. Let \hat{N} be the set of players for which $\underline{a}_i \neq a_i^*$. These are the players whose actions will trigger penalties under assurance contracts.

We can find a set $K \subset L$ that minimally connects \hat{N} . Some players outside of \hat{N} may be included, but they are not *peripheral* to K ; that is, each has two or more links. Let I be the set of players connected by K .

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

A simple construction algorithm suffices to deliver \overline{K} . Starting with K , we will add a link (i, j) for some player $i \in I$ and some player $j \in N \setminus I$ for which $(i, j) \in L$. Since L is connected, there is such a player. If $U_{I \cup \{j\}}(a^*) \geq U_{I \cup \{j\}}(\underline{a})$ then the algorithm terminates. Otherwise, we continue by adding another player k who is linked via network

L with a player in $I \cup \{j\}$ but who is not yet included, again checking the joint value inequality. We continue in this way until the joint value inequality holds. The algorithm must reach its goal, because $U_N(a^*) \geq U_N(\underline{a})$. Once this algorithm terminates, we conduct a paring routine in which any peripheral player whose removal would not flip the joint value comparison is removed from the network.¹³ The result is a network \bar{K} with the desired properties.

Note that, unlike in Case 1, the network of active contracting pairs \bar{K} is generally not complete. Thus, there may be players who are supposed to have null contracts with everyone else.

Let γ be large as before. We can find a set of contracts $\mathcal{M} = \{m^{ij}\}_{i < j}$ with the following properties:

1. For each $(i, j) \notin \bar{K}$, $m_{ij} = \underline{m}$.
2. For each $(i, j) \in \bar{K}$, m_{ij} is an assurance contract with penalty γ .
3. For all $i \in \bar{I}$, $u_i(a^*) + M_i(a^*) \geq u_i(\underline{a})$.
4. $[(i, j) \in K, i \in I, j \in \bar{I} \setminus I]$ implies $u_i(\underline{a}) > u_i(a^*) + M_i(a^*) - m_i^{ij}(a^*)$.
5. $[(i, j), (j, k) \in \bar{K}, K \subset \beta(i, j, \bar{K})]$ implies $u_j(\underline{a}) > u_j(a^*) + M_j(a^*) - m_j^{jk}(a^*)$.

If $\bar{K} = K$ then the contracts between these players can be constructed exactly as in Case 1. If $\bar{K} \neq K$ then one can find a set of contracts that meets the requirements by using the following algorithm. First, create the contracts between the players in I and those linked to I in network \bar{K} . Clearly, one can do this so that $u_i(a^*) + M_i(a^*) = u_i(\underline{a})$ for all $i \in I$, and so that $m_i^{ij}(a^*) > 0$ for all $i \in I, j \notin I$ with $(i, j) \in \bar{K}$. For any such player j , we look to see whether this player is part of a longer chain away from I . If so, we arrange his other contracts so that $u_j(a^*) + M_j(a^*) = u_j(\underline{a})$. Continue this construction for other players in $\bar{I} \setminus I$. By construction, payments flow in from the periphery of \bar{I} to I conditional on a^* being played.

Consider the strategies exactly as described for Case 1, except now player i 's active contracting partners are given by $\bar{K}^i \equiv \{j \mid (i, j) \in \bar{K}\}$ rather than by K^i . No player in the set I can profitably deviate by nullifying any contracts with other players in I , for the same reasons as explained in Case 1.

Next consider a players $i \in I$ who has a non-null contract with a player $j \in \bar{I} \setminus I$. If the contract between player i and player j is nullified, then player i expects a cancellation sequence to the players in the set $\beta(j, i, \bar{K})$; but since all of these players still play $a_{\beta(j, i, \bar{K})}^* = \underline{a}_{\beta(j, i, \bar{K})}$ in the productive phase, player i would not be subject to any penalty if he kept the prescribed contracts with his links in I . However, player i no longer gets the positive transfer from player j and he would then expect a payoff below $u_i(\underline{a})$, given property 4 above. Thus, if his contract with j is nullified, player i strictly prefers to cancel

¹³The paring routine is important because when adding a player, the joint value difference may rise more than it did in a previous round.

his other contracts. Furthermore, player i has no interest in canceling the contract with j to begin with, given property 3.

Regarding the other players, note that those at the periphery of \bar{I} are paying in to get a^* . If they nullify, a cancellation sequence progresses and the outcome is \underline{a} , which is worse given property 3 above. A player i along a path outside of I will cancel inward if an outward link cancels, due to the anticipated payoff otherwise falling below $u_i(\underline{a})$ as property 5 above implies. Also, this player will cancel outward if an inward link cancels, due to the assurance contract with the outward link. Overall, if any non-null contract is nullified then it leads to a full cancellation sequence and \underline{a} . Every player in \bar{I} prefers that this not happen. Each player i who is outside of \bar{I} expects that all other players will end up selecting either a_{-i}^* or \underline{a}_{-i} , to which the prescribed $a_i^* = \underline{a}_i$ is a best response at the production phase.

Case 3: *The underlying game $\langle A, u \rangle$ has a mixed-strategy Nash equilibrium $\underline{\alpha}$ and/or there is a player i for whom $\underline{\alpha}_i$ is a point mass on a_i^* and a_i^* is not a best response to a_{-i}^* .*

A separate construction in this case is required to deal with the following kind of example.

Example 4: Suppose $n = 3$, $A_1 = A_2 = \{b, c\}$, and player 3 has no action in the underlying game. Payoffs in the underlying game are given by: $u(b, b) = (1, 1, 0)$, $u(c, b) = (0, 0, 5)$, $u(b, c) = (0, 0, 0)$, and $u(c, c) = (0, 1, 0)$. The network is $L = \{(1, 3), (2, 3)\}$.

In this example, $\underline{a} = (b, b)$ is the Nash equilibrium of the underlying game, and the efficient action profile is $a^* = (c, b)$. Note that $\underline{a}_2 = a_2^* = b$ and also b is not a best response for player 2 to c being played by player 1; thus, this example does not fit into Cases 1 or 2 but is in Case 3.

Here is the problem with trying the equilibrium construction used in Case 2. The contract between players 2 and 3 would need to specify a transfer of at least 1 to player 2 when a^* is played; otherwise, player 2 would prefer to nullify the contract and induce \underline{a} . However, suppose that player 3 nullifies with player 2 and keeps the prescribed contract with player 1. According to the construction used in Case 2, player 2 must believe that player 3 actually does cancel with player 1, so that player 1 will then select b in the production phase and player 2 selects the best response b . However, in this case, player 3 is better off because (c, b) is played, he does not have to transfer any money to player 2, and he pays no penalty to player 1 since player 2 selects b .

With a slight adjustment in strategies, the construction used for Case 2 can be applied to this case. Define \hat{N} to be the set of players with more than one feasible action in the underlying game, and let \hat{a} be an action profile in which $\hat{a}_i \neq a_i^*$ for all $i \in \hat{N}$. The basic idea is for the players to use the public randomization device to contract on randomizing between a^* and \hat{a} . Specifically, players form contracts that coordinate on a^* with probability $1 - \varepsilon$ (that is, if $\sigma > \varepsilon$) and \hat{a} with probability ε (if $\sigma \leq \varepsilon$). Contracts are of the

assurance form so, for example, if $\sigma > \varepsilon$ and a^* is not played then the relevant players have to pay the penalty γ .

If one of player i 's contracts is nullified and as a result player i expects some others to play their part of the Nash strategy profile $\underline{\alpha}$, then with positive probability the productive outcome will differ from what player i is assuring his other contracting partners. Thus, unless player i cancels his other contracts, he anticipates paying the penalty γ with a probability that is bounded away from zero. If γ is set high enough, player i strictly prefers to cancel his other contracts and the logic from Cases 1 and 2 goes through.

Here are a few more details. If $U_N(a^*) = U_N(\underline{\alpha})$ then finding a PBE that achieves an efficient outcome is trivial, so let us assume that $U_N(a^*) > U_N(\underline{\alpha})$. Fix $\varepsilon > 0$. Sets I and \bar{I} and networks K and \bar{K} are defined exactly as in Case 2, using $\underline{\alpha}$ in place of \underline{a} and using the correlated mixed strategy $(a^*, 1 - \varepsilon; \hat{a}, \varepsilon)$ in place of a^* . As long as ε is small enough, the joint payoff of $(a^*, 1 - \varepsilon; \hat{a}, \varepsilon)$ strictly exceeds the joint payoff of $\underline{\alpha}$ for the entire set of players, and \bar{K} is well defined. The incentives for the players works out as described for Case 2 and we have a PBE in which efficient action profile a^* is played with probability $1 - \varepsilon$. Note that the required γ becomes large as ε gets closer to zero.

Cases 1-3 cover every payoff function u . Also, recall that L was an arbitrary connected network. Therefore, for every connected network L , (\mathcal{C}^*, L) implements ε -efficient outcomes. Because ε was arbitrary, we have that (\mathcal{C}^*, L) virtually implements efficient outcomes.

On the Need for Sequential Contracting

Example 3 illustrated that, to obtain efficient outcomes, the contracting institution must allow for sequential commitment to, or cancellation of, contracts. This subsection provide a more general version of this idea. To state the result, I begin with a definition. Fix a contracting institution in which R is finite. For a sequence of messages z_{ij}^{r-1} and a message $\lambda_{ij}^r \in \Lambda_{ij}^r$, denote by $z_{ij}^{r-1} \lambda_{ij}^r$ the resulting sequence through round r . Let us say that the *contracting institution allows player i to nullify contracts at R* if for every other player j , there is a message $\lambda^0 \in \Lambda_{ij}^R$ such that

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

for all σ and all sequences $z_{ij}^{R-1} \in Z_{ij}^{R-1}$ and $z_{ji}^R \in Z_{ji}^R$. That is, if player i sends message λ^0 to player j at the end of the contracting phase, then their contract is null regardless of the messages they sent earlier.

Lemma: Fix n and A , and suppose that for every player i , A_i contains at least two elements. Consider any contracting institution \mathcal{C} with finite R , and suppose that it allows some player i to nullify contracts at R . Let $k \neq i$ be another player and let L be any network in which $(i, k) \notin L$. Then (\mathcal{C}, L) fails to implement 1-efficient outcomes.

To proof of this lemma, one can simply use a version of Example 3 for any given network. In other words, this result shows that for every $\varepsilon < 1$, there is a payoff function

u such that in no equilibrium is an efficient action profile played with probability at least $1 - \varepsilon$. Clearly then, if a contracting institution is to perform well, it must give players the opportunity to adjust the terms of some other their contractual relationships in response to the contracting in their other relationships.

6 Conclusion

The modeling exercise presented here demonstrates that there are natural contracting institutions that support efficient outcomes through decentralized contracting, assuming the network of contracting partners is connected and there are no externalities of non-contractibility. In particular, externalities due to lack of direct links can be internalized through contractual chains.

In one sense, this modeling exercise follows Hurwicz's (1994) prescription of incorporating "natural" constraints into problems of institutional design. This is in contrast to an extreme perspective that posits a centralized policymaker who has complete control over the design of the game form in which economic agents (players) will be engaged. In reality, this design problem is constrained by an exogenous physical reality. Constraints include actions and communication routes that are naturally available to the players and from which they cannot be excluded, as well as barriers to the inclusion of some other communication channels.

As Jackson and Wilkie argue (2005), 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, there will be both "external planner" and "player" design elements. In my model, 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.

My model surely leaves out some institutional constraints, for instance having to do with limits on the information or sophistication of the external enforcer. It would be useful in future research 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. A simple illustration along these lines is given by comparing the results of Jackson and Wilkie (2005) and Ellingsen and Paltseva (2011). 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.

On future research, it would also be useful to examine special classes of underlying games, for instance ones with "linear externalities" (as in Example 2) as opposed to "com-

plex externalities” (as in Example 3). Many further issues arise. For instance, is the contracting institution developed in the proof a realistic one? Are there physical constraints (for instance, limiting to finite rounds of communication) that would preclude this institution from operating? If so, are there feasible institutions that can perform well, at least in specific cases? It may also be useful to develop a hybrid model in which the contracting phase is modeled in coalitional/cooperative form, which may produce a simpler technical structure suitable for applied analysis.

Appendix

Here I complete the description of equilibrium strategies and the sequential rationality check for Case 1 of the proof. The structure is the similar for Cases 2 and 3.

The specification of strategies given in the text leaves out what player i should do if he enters the production phase with a set of contracts that could exist only if he and at least one other player simultaneously deviated in the first round of the contracting phase. That is, player i has formed a contract \hat{m}^{ij} with some player j such that $\hat{m}^{ij} \neq m^{ij}$ and $\hat{m}^{ij} \neq \underline{m}$.

I proceed by first looking at the special case in which player i has a single non-null contract, it is with player j , and the contract m^{ij} is not consistent with \mathcal{M} . That is, player i has an out-of-equilibrium, non-null contract with player j and all other contracts with player i are null. In this case, player i is to believe that every other pair of players has a null contract. This belief is consistent with the prescribed strategies and with Bayes’ rule because players are supposed to cancel contracts following any deviation; such cancellation by players i and j , other than with each other, would propagate through the network since K is connected. So, player i believes he is in the position of (i) having an out-of-equilibrium contract m^{ij} with player j , and (ii) having mutual knowledge with player j that this is the only non-null contract in existence and everyone else will select $\underline{a}_{-\{i,j\}}$ in the production phase.

For this special case, we can find a Nash equilibrium of the two-player game

$$\langle A_i \times A_j, (u_i(\cdot, \underline{a}_{-\{i,j\}}), u_j(\cdot, \underline{a}_{-\{i,j\}})) + m^{ij}(\cdot, \underline{a}_{-\{i,j\}}) \rangle,$$

which is just the game between player i and j , with contract m^{ij} , under the assumption that the other players select their part of the Nash equilibrium \underline{a} of the underlying game. Let $\alpha_i(m^{ij})$ and $\alpha_j(m^{ij})$ be the equilibrium strategies (possibly mixed). Let us prescribe that player i select $\alpha_i(m^{ij})$ in the productive phase in this contingency.

Next consider the general case in which player i has any number of out-of equilibrium contracts and some may be null. Let I be the subset of K^i with whom player i has a contract that is not consistent with \mathcal{M} , and let I' be the subset of K^i with whom i has a contract that is consistent with \mathcal{M} . Further, let J be the subset of $N \setminus (K^i \cup \{i\})$ with whom player i has a contract that is not consistent with \mathcal{M} , and let J' be the subset of $N \setminus (K^i \cup \{i\})$ with whom player i has a contract that is consistent with \mathcal{M} . Also, let \tilde{M} be the sum of player i ’s contracts.

The prescribed strategies say that the players in I would have canceled all contracts with their active partners, and so the cancellation sequence should reach all players in $\bar{I} \equiv \cup_{j \in I} \beta(j, i, K)$. However, this set could intersect J and the players in the intersection have not canceled all of their contracts (since they have non-null contracts with player i). Let $\tilde{I} \equiv \bar{I} \setminus (I \cup J)$. Let us assume player i believes that the players in \tilde{I} have nullified all of their contracts, they believe all contracts are null, and they select $a_{\tilde{I}}$.

Since contracts between player i and the players in set I' are in accordance with \mathcal{M} , posit that player i believes all of the players in $\bar{I}' \equiv \cup_{j \in I'} \beta(j, i, K)$ detect no deviations and thus will select $a_{\bar{I}'}$.

Finally, note that the players in $I \cup J$ have out-of-equilibrium contracts with player i . Assume player i believes that each of these players thinks all contracts other than his contract with i are null. Thus, player i believes that such a player j will select $\alpha_j(m^{ij})$.

By construction, $\tilde{I} \cup \bar{I}' \cup I \cup J = N \setminus \{i\}$ and the component sets are disjoint. We can prescribe that player i select a best response to $(a_{\tilde{I}}, a_{\bar{I}'}, \alpha_j(m^{ij}))$. By construction, player i 's beliefs are consistent and he best responds in the productive phase.

References

- Bergemann, D. and J. Valimaki (2003): "Dynamic Common Agency," *Journal of Economic Theory* 111, 23-48.
- Bernheim, D. and M. Whinston (1986a): "Common Agency," *Econometrica* 54, 923-942.
- Bernheim, D. and M. Whinston (1986b): "Menu Auctions, Resource Allocations and Economic Influence," *Quarterly Journal of Economics* 101, 1-31.
- Bloch, F. and A. Gomes (2006): "Contracting with Externalities and Outside Options," *Journal of Economic Theory* 127 (1), 172-201.
- Cafaggi, Fabrizio (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.
- 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 (2007): "Bilateral Bargaining with Externalities," unpublished paper, University of Melbourne.
- Ellingsen, T. and E. Paltseva (2011): "Non-cooperative Contracting," unpublished paper, Department of Economics, Stockholm School of Economics.
- Galasso, A. (2008): "Coordination and Bargaining Power in Contracting with Externalities," *Journal of Economic Theory* 143 (1), 558-570.

- Gomes, A. (2005): "Multilateral Contracting with Externalities," *Econometrica* 73 (4), 1329-1350.
- Gomes, A. and Jehiel, P. (2005): "Dynamic Processes of Social and Economic Interactions: On the Persistence of Inefficiencies," *Journal of Political Economy* 113, 626-667.
- Guttman, J.M. (1978): "Understanding Collective Action: Matching Behavior," *American Economic Review* 68, 251-255.
- Guttman, J.M. and A. Schnytzer (1992): "A Solution to the Externality Problem Using Strategic Matching," *Social Choice and Welfare* 9, 73-88.
- Hafalir, I.E. (2007): "Efficiency in Coalition Games with Externalities," *Games and Economic Behavior* 61, 242-258.
- Hardin, G. (1968): "The Tragedy of the Commons," *Science* 162, 1243-1248.
- Hurwicz, L. (1994): "Economic Design, Adjustment Processes, Mechanisms, and Institutions," *Economic Design* 1, 1-14.
- Hyndman, K. and D. Ray (2007): "Coalition Formation with Binding Agreements," *Review of Economic Studies* 74, 1125-1147.
- Jackson, M.O. and S. Wilkie (2005): "Endogenous Games and Mechanisms: Side Payments Among Players," *Review of Economic Studies* 72 (2), 543-566.
- Jensen, M.C. and W.H. Meckling (1976): "Theory of the firm: Managerial behavior, agency costs and ownership structure," *Journal of Financial Economics* 3 (4), 305-360.
- Laffont, J.-J. and D. Martimort (1997): "The Firm as a Multicontract Organization," *Journal of Economics and Management Strategy* 6 (2), 201-234.
- Lloyd, W.F. (1833): *Two Lectures on the Checks to Population*, Oxford: Oxford University Press.
- Martimort, D. (2007): "Multi-Contracting Mechanism Design," in *Advances in Economic Theory: Proceedings of the World Congress of the Econometric Society*, edited by Blundell, Newey and Persson, Cambridge University Press.
- Peters, M. and B. Szentes (2008): "Definable and Contractible Contracts," unpublished paper, University of British Columbia.
- Prat, A. and A. Rustichini (2003): "Games Played Through Agents," *Econometrica* 71 (4), 989-1026.
- Segal, I. (1999): "Contracting with Externalities," *Quarterly Journal of Economics* 114 (2), 337-88.
- Segal, I. and M. Whinston (2003): "Robust Predictions for Bilateral Contracting with Externalities," *Econometrica* 71 (3), 757-791.

- Seidmann, D.J. and E. Winter (1998): "Exploring Gains from Trade in Multilateral Bargaining: A Theory of Gradual Coalition Formation," *Review of Economic Studies* 65, 793-815.
- Varian, H. (1994b): "A Solution to the Problem of Externalities when Agents are Well-Informed," *American Economic Review* 84, 1278-1293.
- Yamada, A. (2003): "Efficient Equilibrium Side Contracts," *Economics Bulletin* 3 (6), 1-7.