Contractual Chains

Joel Watson*

December 2013
Incomplete and under revision – Please do not circulate†

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 fully—where a party is unable to contract directly with others whose actions affect his payoffs. The paper investigates 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. The equilibrium construction features *assurance contracts* and *cancellation penalties*.

*UC San Diego; http://econ.ucsd.edu/∼jwatson/. Some of the work reported here was developed while the author visited Yale in 2010 and when the author was a fellow in residence at the Center for Advanced Study in the Behavioral Sciences at Stanford in 2012-2013. For their helpful comments and discussion, the author thanks Nageeb Ali, Gorm Grønnevet, Keri Hu, Natalia Lazzati, colleagues at Yale and UCSD, and seminar participants at Oxford, Washington University, UC Irvine, University of British Columbia (2012 summer theory conference), the 2012 Decentralization Conference at CalTech, Pittsburgh, Georgetown, the 2012 Econometric Society summer conference, UT Austin McCombs Business School, Columbia University, and Stanford University. The author thanks the NSF for financial support (SES-1227527). THANK Gorm Grønnevet (Senior Research Economist, Telenor Research)

†First version: May 2011. Watson is really dragging his feet on this one!

1 Introduction

In many contractual settings, there is multilateral productive interaction but barriers prevent the parties from contracting as a large group. Instead, parties 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: The developer cares about the subcontractors' work (it influences her benefit of the completed project) but she doesn't contract directly with the subcontractors. I call this an *externality due to lack of direct links*; such an externality is formally defined as a situation in which at least one agent is unable to contract directly with another whose productive action he cares about.

The question investigated here is whether a lack-of-direct-links externality can be internalized via a *chain of bilateral contracts* and, if so, under what conditions. For instance, in the example just described, the developer and general building contractor may form a deal specifying that the developer must pay the building contractor conditional on the subcontractors' satisfactory performance. Furthermore, the general building contractor may have an agreement with each subcontractor that rewards the subcontractor contingent on satisfactory performance. In this example, the contractual chains run from the developer, through the general building contractor, to each subcontractor. I seek to understand whether contractual chains like these would be effective in achieving efficient outcomes and whether they would arise in an equilibrium of the decentralized contracting process.

Many important real settings feature lack-of-direct-links externalities. Examples include the internal organization of firms (where multiple workers have employment contracts with the firm but also care about each others' actions), sales of goods exhibiting network externalities, private provision of public goods, commons problems, risk-sharing arrangements, and the like. Although these types of externalities are common, they have not been methodically studied and are not well understood.

I develop a general model that covers the various applications listed above. In the model, a set of n players interacts in an "underlying game" $\langle A, u \rangle$ with externally enforced monetary transfers and full verifiability. Payoffs are linear in money. Before the underlying game is played, the players can establish contracts. 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 an undirected network of bilateral links, $L \subset N \times N$, where $N = \{1, 2, \ldots, n\}$. Only linked pairs of players $(i, j) \in L$ can communicate and contract. A contract between players i and j specifies transfers between them as a function of the outcome of the underlying game. That is, it is a function $m : A \to \mathbb{R}^n_0$ such that $m_k(a) = 0$ for all $k \notin \{i, j\}$ and every $a \in A$, where $\mathbb{R}^n_0$ denotes the vectors in $\mathbb{R}^n$ whose components sum to zero (balanced transfers). I shall measure efficiency by the players' joint values (sum of payoffs). Since this is meaningful only if it is feasible to transfer money across the network, I assume that L is connected.

A novel aspect of the approach taken in this paper is that it emphasizes the role of the *contracting institution* that facilitates contract formation. Formally, a contracting institution is an extensive game form with payoff-irrelevant messages. The outcomes of this game form are the externally enforced contracts. Critically, contracting institutions are restricted by the network of links and by the following assumptions that represent the notion of voluntary, decentralized contracting:

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

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

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

In this paper, I focus on a narrow "possibility" question: Given that contracting must take place in a decentralized manner, is there a natural contracting institution under which efficient outcomes can be achieved across a wide range of underlying games and networks? If so, what are the minimal requirements of the network and what are the key properties of the contracting institution? In technical terms, the first question asks whether there is

[1]Further discussion on the scope of contracting may be found in Sections 3, 4, and 5. One can also motivate the second 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.

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, among other things, what kind of network structures can be allowed.

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

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

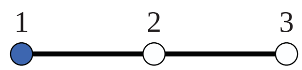

In the figure above, which depicts the network, I indicate with a filled circle that player 1 has an action in the underlying game. In this example, the efficient outcome cannot be obtained with a natural contracting institution that has one-shot, simultaneous contracting. To see this, consider the contracting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

In order for player 1 to select action b in equilibrium, it must be that players 1 and 2 form a contract that pays player 1 at least 1 for choosing this action in the production phase. Because player 2 can guarantee himself a payoff of 0 by naming the null contract, the proposed equilibrium must also have players 2 and 3 form a contract that transfers at least 1 from player 3 to player 2 in the event that player 1 chooses b. Then player 2 is compensated by player 3 for the amount 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 $m(\mathrm{b}) = (\tau, -\tau, 0)$ and $m'(\mathrm{b}) = (0, \tau', -\tau')$, for some values τ and τ' satisfying $1 \leq \tau \leq \tau' \leq 4$. Also, it must be the case that $\tau \geq 1 + m_1(\mathrm{c})$, 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 1 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 1, where player 1's action of b benefits player 3, it would not deliver an efficient outcome if the externality ran the other way.

The next example shows that the problem is worse in cases of complex externalities.

***Example 2*:** Suppose $n = 3$ and $A_i = \{\mathrm{b}, \mathrm{c}\}$ for $i = 1, 2, 3$. Payoffs in the underlying game are given by: $u(\mathrm{b}, \mathrm{b}, \mathrm{b}) = (1, 1, 1)$, $u(\mathrm{c}, \mathrm{c}, \mathrm{c}) = (0, 0, 0)$, 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: $L = \{(1, 2), (2, 1), (2, 3), (3, 2)\}$.

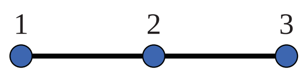

First note that, if the contracting institution specifies a particular order in which players can communicate and establish contracts, then it cannot achieve the efficient action profile $(\mathrm{b}, \mathrm{b}, \mathrm{b})$. This is because the player who is last to contract can decline to do so and 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.

Second, observe that contracting is problematic 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 a finite number of discrete rounds. Suppose that 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 it goes into force and the contracting session ends for these two players. Such an institution would allow all of 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 that, in equilibrium, 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. One can show that there is no PBE that leads to the efficient action profile with high probability.

Despite the conclusion of these examples, my message is positive. This paper's single theorem establishes the existence of a natural contracting institution that virtually implements efficient outcomes. In other words, for every underlying game, every connected network of contractual relationships, and every $\varepsilon > 0$, there is a PBE that achieves an efficient outcome with probability of at least $1 - \varepsilon$. Thus, there is a contracting institution under which, in all productive settings, lack-of-direct-links externalities can be overcome by decentralized contractual chains. Furthermore, a critical property of a well-functioning contracting institution is that it allows for sequential communication and contracting.

The next section provides more details of the general model. Section 3 gives some applications and discusses the related literature. Section 4 contains the theorem and its proof. In Section 5, I offer additional comments on the related literature and I discuss further steps in the research program. [Note: There is a missing section on tangential results and discussion. The next version will also have a better layout of the main analysis.]

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 $u\colon A \to \mathbb{R}^n$ is the payoff function. Payoffs are in monetary units. Let $N = \{1, 2, \ldots, n\}$ denote the set of players. I assume that n is finite and that A is a subset of some finite set $\mathcal{A}$, which is interpreted as the set of all possible action profiles that the external enforcer can recognize. For any subset of players $J \subset N$, write $a_J \equiv (a_i)_{i \in J}$ as the vector of actions for these players. The players commonly know the action space and payoff function.

An external enforcer compels monetary transfers between the players, and he does so as directed by contracts that the players form. The enforcer does not observe the payoff function u or the feasible space of action profiles A. Otherwise, the contracting environment has *full verifiability* in that the enforcer observes the action profile $a \in A$ that the players select. Thus, transfers can be freely conditioned on the outcome of the underlying game. 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 $(j, i) \in L$. The network is generally not transitive; thus, player i may be able to contract with player j, and player j may be able to contract with player k, whereas player i is unable to contract with player k. Assume that L is connected.

Contracting partners can condition transfers between them on actions taken by third parties. For example, a contract between players 1 and 2 could specify that player 1 must pay an amount to player 2 in the event that player 3 selects a particular action in the underlying game. However, 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 $m : \mathcal{A} \to \mathbb{R}^n_0$, where $m_k(a) = 0$ for all $k \notin \{i, j\}$ and every $a \in \mathcal{A}$. 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 such a set, let

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

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

Contracting Institution Formalities

A contracting institution is an extensive game form with payoff-irrelevant messages that map to contracts enforced by the external enforcer. I constrain attention to a particular class of game forms in which the players send messages in discrete rounds $1, 2, \ldots, R$. There is no discounting. In every round, each player i sends a vector of messages, one message for every other player. Let Λ^r_{ij} denote the set of feasible message from player i to player j in round r, and let $\lambda^r_{ij} \in \Lambda^r_{ij}$ 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^r_{ij}, for all i, j, r.

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

$$z^r_{ij} = (\lambda^1_{ij}, \lambda^2_{ij}, \ldots, \lambda^r_{ij}) \in Z^r_{ij} \equiv \Lambda^1_{ij} \times \Lambda^2_{ij} \times \cdots \times \Lambda^r_{ij}.$$

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

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),

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

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 $\mathbb{R}_0^n$.

A contracting institution is called *natural* if the following assumptions hold: First, contracting is private in that (a) a given player can 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. Second, contracting in different relationships takes place independently. Third, players have the option of rejecting contracts. 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.

Notion of Implementation

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 $\mathcal{A}$. The primary objective is to evaluate the performance of a given contracting institution across various networks L and underlying games $\langle A, u \rangle$. Let us say that a contracting institution $\mathcal{C}$ *implement efficient outcomes* if for every product set $A \subset \mathcal{A}$, every payoff function $u : A \to \mathbb{R}^n$, and every connected network network L, there is a PBE of the entire game in which an efficient action profile a^* is played. Let us say that $\mathcal{C}$ *implements ε-efficient outcomes* if for every product set $A \subset \mathcal{A}$, every payoff function $u : A \to \mathbb{R}^n$, and every connected network network L, 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 $\mathcal{C}$ *virtually implements efficient outcomes* if this contracting institution implements ε-efficient outcomes for all $\varepsilon > 0$.

3 Applications and Related Literature

To relate the analysis herein to the previous literature, it is useful to distinguish lack-of-direct-links externalities from the kind of externalities that the literature has focused on. Let us say that an *externality of non-contractibility* is present when a pair of linked players is not able to condition transfers on an action that is payoff relevant to them (perhaps one of their own actions or the action of a third player). Also, let us say that a setting exhibits *sparse contracting* if some pairs of players are not able to contract, irrespective of whether they care about each others' actions in the underlying game. Much of the related literature comprises special cases of the general model proposed here. In these models, there are no externalities due to lack of direct links. Rather, the focus is on externalities of non-contractibility 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, typically 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.

Below are brief descriptions of a few applications and more details on the related literature. Figures 2 and 1 illustrate the network structures in the applications.

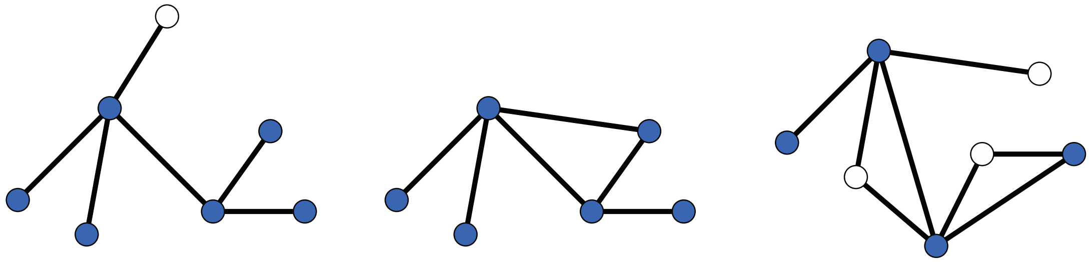

Figure 1: Some applications.

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 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.

Common agency

In the common agency model introduced by Bernhaim 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 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.[3]

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.

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

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

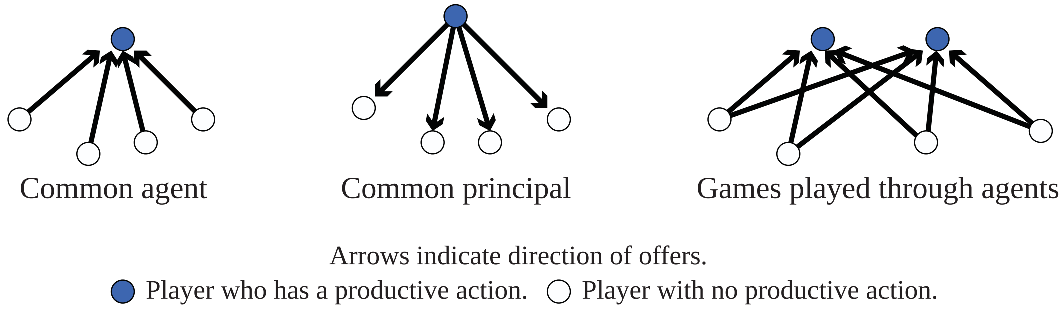

Figure 2: Applications studied in the literature.

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.[4]

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).[5]

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.[6] 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

[4]See 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.

[5]The 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 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.

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

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 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.[7]

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.[8]

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 bargain-

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

[8]There are important examples of contracts on contracts in the real world, such as non-compete covenants, requirements for subcontracting, and the like. This type of contracting is ruled out here, however.

[9]The typical model has a public, dynamic negotiation process. Payoffs are described in terms of a char-

ing 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 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 Efficient Implementation

This section presents the main result.

Theorem: *Fix n and $\mathcal{A}$. There exists a natural contracting institution $\mathcal{C}$ that virtually implements efficient outcomes.*

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

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 and "cancellation penalties." In later rounds the players can unilaterally cancel these contracts. Contracts that are not canceled will then be enforced. If a player is the first to cancel a contract, then this party must pay a cancellation penalty.

Formally, define institution $\mathcal{C}^*$ as follows. Let R be any integer that is at least $n-1$. 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 name (i) arbitrary mappings from $A \times [0,1]$ to transfers between the contracting parties, and (ii) cancellation penalties. That is, for each pair of players (i,j), Λ_{ij}^1 is the set of pairs (g,p). The first element is a mapping $g\colon A \times [0,1] \to \mathbb{R}_0^n$ that satisfies $g_k(a,\sigma) = 0$ for all $k \notin \{i,j\}$, every $a \in A$, and every $\sigma \in [0,1]$. The second element is a vector $p = (p^2, p^3, \ldots, p^R) \in \mathbb{R}^{R-1}$, where p^r is the cancellation penalty for round r. Identify the null message $\underline{\lambda}$ for round 1 as the constant function $g \equiv (0,0,\ldots,0)$ and $p = 0$. If players i and j sent the same message (g,p) in

acteristic 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).

round 1, then call g their "provisional arrangement." In every round $r > 1$, the message space is defined as $\Lambda^r_{ij} \equiv \{\text{cancel}, \underline{\lambda}\}$.

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

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

Fix contracting institution $\mathcal{C}^*$. Consider any underlying game $\langle A, u\rangle$ and any connected network L, where $A \subset \mathcal{A}$. 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 $1 - \varepsilon$. I proceed by first analyzing two special cases. The first contains the key ideas behind the proof; all but the most avid readers will probably want to skip the second and third cases.

Case 1: *The underlying game $\langle A, u\rangle$ has a pure-strategy Nash equilibrium $\underline{a}$ and there is an efficient action profile a^* such that $\underline{a}_i \neq a^*_i$ for each $i \in N$.*

In this case, the efficient action profile has all of the players selecting actions that differ from their underlying-game 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 $t = 2, 3, \ldots, T$. The sequence $\{k^t\}_{t=1}^T$ is called the *path from i to j*. A minimally connected subnetwork K exists by construction and is defined to be symmetric.

For 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:

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

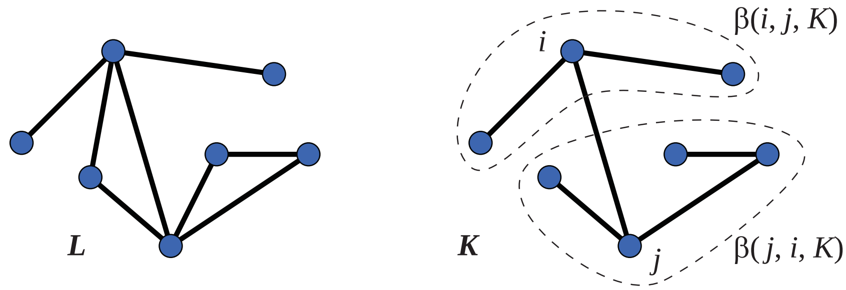

Figure 3: Construction of K and definition of β.

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. Figure 3 illustrates K and β. Note that $\beta(i, j, K) \cap \beta(j, i, K) = \emptyset$ and $\beta(i, j, K) \cup \beta(j, i, K) = N$. Define $\delta(i, j, K)$ to be the length of the largest path between player i and players in the set $\beta(i, j, K)$; this is zero if $\beta(i, j, K) = \{i\}$. 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.

Define $\gamma \equiv \max_{a,a' \in A, i \in N}[u_i(a) - u_i(a')]$, which is the maximum payoff difference for the players in the underlying game. Let $p^* = (p^{2*}, p^{3*}, \ldots, p^{R*})$ be such that $p^{2*} > \gamma$ and $p^{r*} = p^{2*}(r - 1)$ for all $r = 3, 4, \ldots, R$. Let q be a number satisfying $q > p^{R*}$.

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. For players $(i, j) \in K$ and action profile $a \in A$, define

$$Q(a, a^*, i, j, K) = \#\{k \in \beta(i, j, K) \mid a_k \neq a_k^*\}.$$

This is the number of players on i's side of the subnetwork whose actions differ from those in a^*. We can find contracts 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:

  $$m_i^{ij}(a) = -m_j^{ij}(a) = m_i^{ij}(a^*) + [Q(a, a^*, j, i, K) - Q(a, a^*, i, j, K)]q$$

  for all $a \in A$.

The second property simply requires finding transfers to satisfy the inequalities, and these obviously exist given that K is connected. The third property is achieved by construction.

Note that 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—that is, those in $\beta(i, j, K)$, including player i—will select their part of the efficient action profile. If any player in $\beta(i, j, K)$ deviates from $a^*_{\beta(i,j,K)}$ then player i is required to pay q to player j. This penalty is multiplied by the number of players on i's side of the network (including i) who deviate.

Because q exceeds the maximum payoff difference, 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 player i's strategy and beliefs, for any $i \in N$:

1. In round 1 of the contracting phase, and for each $j \in K^i$, player i is supposed to send message (g, p^*), where $g(a, \sigma) = m^{ij}$ for all σ.

2. If the messages sent in round 1 between player i and all of his contracting partners were as prescribed (an event that player i can observe), then in round 2 player i is supposed to send the null message to all of them. Specify the same rule for each subsequent round: If player i detects no deviations with his contracting partners through round $r - 1$, then player i is to send the null message to each of them in round r. Furthermore, if player i and/or some of his contracting partners deviated in an inessential way that could not affect the induced contracts with player i (by, for example, saying "cancel" after forming a provisional arrangement that selects the null contract), then player i should ignore these deviations.

3. Through any round r, if player i's history conforms to specifications 1-2 above, then player i believes that everyone has played as prescribed (even those whom player i is not linked to).

4. If player i's history through round R conforms to specifications 1-2 above, then player i selects a^*_i in the production phase.

5. Let j be a contracting partner of player i for which $m^{ij} \neq \underline{m}$. Suppose player i's history through round $r-1$ conforms to specifications 1-2 above, but that player j's message to player i in round r is a deviation. Suppose player i and his other contracting partners exchange messages in round r as prescribed above. If $R-r+1 \geq \delta(i, j, K)$ then, regardless of his own actions from round $r + 1$, player i believes that at the end of the contracting phase, the contracts between the players in $\beta(j, i, K)$ will all be null and that $\underline{a}_{\beta(j,i,K)}$ will be selected by these players in the production phase.

6. If the event in item 5 occurs, then player i is supposed to send the message "cancel" to all of his contracting partners in the remaining rounds of contracting phase.

7. Suppose that player i deviates in round 1 with a set $J \subset K^i$ of contracting partners. Then, regardless of his own actions from round 2, player i believes that, at the end of the contracting phase, the contracts between the players in $\cup_{j\in J}\beta(j,i,K) \equiv \overline{J}$ will all be null and that $\underline{a}_{\overline{J}}$ will be selected by these players in the production phase.

8. If the event in item 7 occurs, then player i is supposed to send the message "cancel" to all of his contracting partners in the remaining rounds of contracting phase.

9. If the events described in items 5-6 occur, or the events described in items 7-8 occur, then player i is supposed to select action $\underline{a}_i$ in the productive phase.

It is not difficult to confirm that players prefer not to deviate from the behavior prescribed here. That is, this partial description of strategies 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 q to each of his contracting partners in K because he has assurance contracts with them, and this penalty exceeds whatever gain he gets in the underlying game. Likewise, if the events described in items 5-6 occur, or the events described in items 7-8 occur, then at the end of the contracting phase, player i believes that all contracts are null and that everyone else will select $\underline{a}_{-i}$ in the production phase. Player i's best response is to choose $\underline{a}_i$ as prescribed.

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 $u_i(a^*) + M_i(a^*)$. Second, player i might consider deviating in round 1 with one or more of his contracting partners, thereby nullifying the contracts with them. If he does this, the affected contracting partners will perpetuate a wave of cancellations, and the connected players will eventually select their part of $\underline{a}$. After the initial deviation, player i thus has the incentive to cancel all of his contracts in round 2. If he fails to cancel a contract with a player j then he will later have to pay player j a multiple of the assurance penalty q, given that he expects $a_{\beta(i,j,K)} \neq a^*_{\beta(i,j,K)}$. It is better for player i to cancel the contract with player j and forfeit the smaller cancellation penalty p^{2*}. Furthermore, player i strictly prefers to cancel in round 2 rather than later, because the cancellation penalty increases over time. Still, the cancellation penalty outweighs any payoff gain in the underlying game, so player i's best hope for a gain by deviating in round 1 is to deviate with *all* of his contracting partners—that is, by forming null contracts from the beginning. But this sets off a wave of cancellations and ultimately $\underline{a}$ is played in the production phase, which gives player i a payoff of $u_i(\underline{a}) \leq u_i(a^*) + M_i(a^*)$.

The third option for player i is to deviate only after round 1 of the contracting phase, nullifying contracts by canceling them. For instance, player i could wait until round R and catch his contracting partners by surprise with some cancellations. This is also a bad option, however, because player i would then have to pay the cancellation penalty of p^{R*} times the number of contracting partners with whom player i cancels. The penalty outweighs any gain in the underlying game and implies a payoff below $u_i(a^*) + M_i(a^*)$.

To see why the condition $R - r + 1 \geq \delta(i, j, K)$ is important to items 5 and 6 above, observe that this guarantees that there are enough rounds for the contagious sequence of cancellations to reach all of the players in the set $\beta(i, j, K)$. When player i observes j's deviation, player i can conclude that player j is in a sequence of cancellations that will eventually reach all of the players in $\beta(j, i, K)$. This is true because, by round r, already $r - 1$ rounds of cancellations may be thought to have occurred within the set $\beta(j, i, K)$, and further cancellations can occur until the end of the contracting phase. So player i can believe that player j and the others in $\beta(j, i, K)$ believe that all contracts will be nullified, and everyone has the incentive to select their part of $\underline{a}$ in the production phase. If $R - r + 1 < \delta(i, j, K)$ then the cancellation sequence would not reach everyone in $\beta(i, j, K)$ and thus player i could not necessarily believe that everyone expects $\underline{a}$ to be selected in the production phase; in this case, player i might think that $a^*_{\beta(j,i,K)}$ is possible and then not have the incentive to cancel contracts or select $\underline{a}_i$.

The only loose end in the analysis is that I have not described beliefs and behavior for some out-of-equilibrium contingencies, such as when two players deviate at the same time or when player j cancels a contract with player i but there are not enough rounds remaining to ensure that a contagious cancellations would reach all players. A complete the description of the PBE requires stating these beliefs and action choices, and showing that the beliefs are consistent and the actions are sequentially rational. A sketch is provided in the Appendix.

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.

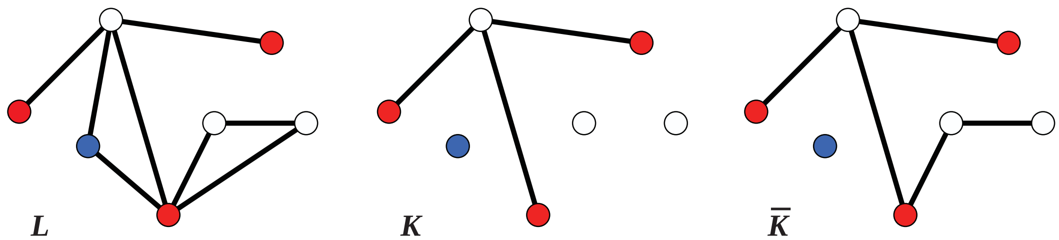

Red denotes players in $\hat{\boldsymbol{N}}$.

Figure 4: Construction of $\overline{K}$ for case 2.

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.[10] The result is a network $\overline{K}$ with the desired properties.

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

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

1. For each $(i, j) \notin \overline{K}$, $m_{ij} = \underline{m}$.

2. For each $(i, j) \in \overline{K}$, m_{ij} is an assurance contract with penalty q.

3. For all $i \in \overline{I}$, $u_i(a^*) + M_i(a^*) \geq u_i(\underline{a})$.

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

4. $[(i,j) \in K, i \in I, j \in \overline{I} \setminus I]$ implies $u_i(\underline{a}) > u_i(a^*) + M_i(a^*) - m_i^{ij}(a^*)$.

5. $[(i,j),(j,k) \in \overline{K}, K \subset \beta(i,j,\overline{K})]$ implies $u_j(\underline{a}) > u_j(a^*) + M_j(a^*) - m_j^{jk}(a^*)$.

If $\overline{K} = K$ then the contracts between these players can be constructed exactly as in Case 1. If $\overline{K} \neq K$ 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 $\overline{K}$. 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 $(i,j) \in \overline{K}$. 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 $\overline{I} \setminus I$. By construction, payments flow in from the periphery of $\overline{I}$ to I conditional on a^* being played.

Consider the strategies as described for Case 1, except now player i's active contracting partners are given by $\overline{K}^i \equiv \{j \mid (i,j) \in \overline{K}\}$ rather than by K^i. Furthermore, make some adjustments to the penalties for the contractual relationships at the periphery of K and beyond. These are players in the set

$$P \equiv \{i \in K \mid \text{there exists } j \in \overline{I} \setminus I \text{ with } (i,k) \in \overline{K}\} \cup (\overline{I} \setminus I).$$

For any relationship $(i,j) \in K$ such that $i \in P$, define $\delta_i \equiv \delta(i,j,\overline{K})$ and prescribe that players i and j name cancellation penalties $\hat{p} = (\hat{p}^2, \hat{p}^3, \ldots, \hat{p}^R)$ defined as follows: $\hat{p}^r = 0$ for $r = 2, 3, \ldots, \delta_i + 1$, $\hat{p}^{\delta_i+2} > \gamma$ and $\hat{p}^r = \hat{p}^{\delta_i+2}(r - \delta_i - 1)$ for all $r = \delta_i + 3, \delta_i + 4, \ldots, R$. Next consider any relationship $(i,j) \in \overline{K}$ such that $i, j \in P$, and without loss of generality order the players so that $\beta(i,j,\overline{K}) \subset \overline{I} \setminus I$. That is, player i is outside of player j with respect to the core group I. Prescribe that players i and j name cancellation penalties $\hat{p}$ defined in the same way. For the rest of the contractual relationships in $\overline{K}$, specify that they name p^* as defined in Case 1. The idea is to give players in the set P a way of canceling without penalty for a number of rounds that equals their distance from the periphery of $\overline{K}$.

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 $j \in \overline{I} \setminus I$. If the contract between player i and player j is nullified in round 1, then player i expects a cancellation sequence to the players in the set $\beta(j,i,\overline{K})$; but since all of these players still play $a^*_{\beta(j,i,\overline{K})} = \underline{a}_{\beta(j,i,\overline{K})}$ 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 in round 1, player i prefers to cancel his other contracts in round 2, when he does not have to pay a cancellation penalty. Furthermore, player i has no interest in canceling the contract with j to begin with, given property 4.

Regarding the other players, note that those at the periphery of $\overline{I}$ are paying in to get a^*. If they nullify in round 1, a cancellation sequence progresses and the outcome is $\underline{a}$, which is worse given property 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 $\overline{I}$ prefers that this not happen. Each player i who is outside of $\overline{I}$ 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 3*:** Suppose $n = 3$, $A_1 = A_2 = \{\mathrm{b}, \mathrm{c}\}$, and player 3 has no action in the underlying game. Payoffs in the underlying game are given by: $u(\mathrm{b}, \mathrm{b}) = (1, 1, 0)$, $u(\mathrm{c}, \mathrm{b}) = (0, 0, 5)$, $u(\mathrm{b}, \mathrm{c}) = (0, 0, 0)$, and $u(\mathrm{c}, \mathrm{c}) = (0, 1, 0)$. The network is $L = \{(1, 3), (2, 3)\}$.

In this example, $\underline{a} = (\mathrm{b}, \mathrm{b})$ is the Nash equilibrium of the underlying game, and the efficient action profile is $a^* = (\mathrm{c}, \mathrm{b})$. Note that $\underline{a}_2 = a^*_2 = \mathrm{b}$ 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 q.

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 q with a probability that is bounded away from zero. If q 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 $\overline{I}$ and networks K and $\overline{K}$ 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 $\overline{K}$ 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 q becomes large as ε gets closer to zero.

Cases 1-3 cover every underlying game $\langle A, u\rangle$. Also, recall that L was an arbitrary connected network. Therefore, $\mathcal{C}^*$ implements ε-efficient outcomes. Because ε was arbitrary, we have that $\mathcal{C}^*$ virtually implements efficient outcomes.

5 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 the model herein, the contracting institution is an object of external design, and it must obey the physical reality represented by the natural contracting assumptions. The contracts are the player design element. These come together to determine the induced game between the players.

The model presented here leaves out some institutional constraints, for instance 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 "complex 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 sketch how to complete the PBE construction for Case 1. The structure is the similar for Cases 2 and 3. The specification of strategies given in the text leaves out what player i should believe and do in some off-path contingencies. The argument below establishes that there is a way to complete the strategy and belief specification to form a PBE.

Fix an underlying game $\langle A, u \rangle$ and consider an artificial game in which the contracts $\mathcal{M}$ and cancellation penalties described in the text for Case 1 are exogenously supplied to the players and then, in rounds 1 through R, each player has the opportunity to cancel any of his contracts as contracting institution $\mathcal{C}$ allows. Thus, interaction occurs just as in contracting institution $\mathcal{C}$ except that the players do not name pairs (g, p) in round 1 of the contracting phase; instead, they decide only whether to cancel the exogenously supplied contracts. Assume that the fee for canceling a contract in round 1 is 0.

Note that there is a finite number of information sets and actions in the game just described. (Each player chooses between "cancel" and $\underline{\lambda}$ in the R rounds of the contracting phase.) We can translate the specification of beliefs and behavior in the text to the artificial game by interpreting the choice of "cancel" in the first round of the artificial game as the selection of the null message in the first round of the actual game. The construction in the text thus specifies beliefs and behavior for some of the information sets in the artificial game, in particular (i) ones in which the players select and observe $\underline{\lambda}$ in each round, and (ii) ones in which players experience nullification of a contract with enough rounds remaining for a cancellation sequence to reach all players.

Fixing these specifications, one can perform the standard PBE existence argument on the set of consistent beliefs and strategies for the remaining information sets. This yields a PBE for the artificial game. Furthermore, applying this specification to the actual game covers all information sets except ones in which players select a message in round 1 that is neither the prescribed message nor the null message. If this happens in a contractual relationship, and if the resulting preliminary arrangement is null, then let us specify that these players continue as though just one of them deviated by selecting the null message.

This leaves contingencies in which players i and j in a contractual relationship both deviated to select the same message (g, p). In this case, assume that player i believes throughout the contracting phase that player j did not deviate in round 1 with any of player j's other contracting partners. Under the assumption that this is true for both players i and j, one can then find an equilibrium specification for the restricted game between players i and j. Use this equilibrium to define the beliefs of player i for all information sets, including ones in which player i deviated with several different contracting partners in round 1, and let i's strategy provide for a sequential best response. By construction, we have sequentially rational strategies and consistent beliefs over all information sets.

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.

Cafeggi, 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.