Contractual Chains

Joel Watson, April 2024

B Supplemental Appendix (on-line only)

B.1 Statements and Proofs of Additional Results

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

Result 1: *For any given $n \geq 3$, there exists an underlying game with partial verifiability and a connected network such that (i) feasible contracts exist to support efficient production, and yet (ii) regardless of the contracting institution, there is no sequential equilibrium of the grand game in which an efficient action profile of the underlying game is played with positive probability.*

Proof: Consider the following example of team production with three players. Player 1, the manager, has no action in the underlying game, so $A_1 = \{1\}$. Players 2 and 3, the workers, have action spaces $A_2 = A_3 = \{0, 1\}$, where 1 stands for high effort and 0 represents low effort. Payoffs are given by $u_1(a) = 5(a_2 + a_3)$, $u_2(a) = -3a_2$, and $u_3(a) = -3a_3$. The network consists of only the pairs $(1, 2)$ and $(1, 3)$, meaning the manager can contract with each worker individually.

Assume the external enforcer can verify only whether or not $a = (1, 1, 1)$. That is, there is partial verifiability represented by the partition of A with these two elements: $\{(1, 1, 1)\}$ and $\{(1, 0, 1), (1, 1, 0), (1, 0, 0)\}$. The verifiability constraint requires every contract m to specify $m(1, 0, 1) = m(1, 1, 0) = m(1, 0, 0)$. Efficiency requires $a = (1, 1, 1)$ to be chosen, maximizing the players' joint value.

There are contracts m^{12} and m^{13} that, if formed in equilibrium, would give players 2 and 3 both the incentive to choose high effort, such as contracts specifying bonus payments $m_2^{12}(1, 1, 1) = 4$ and $m_3^{13}(1, 1, 1) = 4$, and $m_2^{12}(a) = m_3^{13}(a) = 0$ for every $a \neq (1, 1, 1)$. With these contracts and high effort from both workers, player 1's payoff would be 2. There cannot be such an equilibrium outcome, however, for then player 1 would strictly gain by refusing to contract with player 3 while behaving with player 2 as directed. This deviation would not be observed by player 2 and therefore would not affect player 2's choice of high effort, and it would lead player 3 to choose low effort, resulting in a payoff of 3 for player 1.

In fact, in this example, regardless of the contracting institution, every equilibrium of the grand game has $a = (1, 1, 1)$ chosen with probability zero. Here is the formal analysis:

For any given equilibrium, let f be the joint distribution of (a_2, a_3) on the equilibrium path. Consider that in some equilibrium contingency at the end of the contracting phase, player 2's contract with player 1 is m^{12} and player 2 is supposed to select high effort with positive probability. Let ζ be the probability that, in this contingency, player 2 thinks player 3 will select high effort. Note that player 2 receives $m_2^{12}(1, 1, 1)$ if and only if both workers choose high effort, and otherwise player 2 receives $m_2^{12}(1, 0, 0) = m_2^{12}(1, 1, 0) = m_2^{12}(1, 0, 1)$. Player 2's incentive condition requires $\zeta m_2^{12}(1, 1, 1) + (1-\zeta) m_2^{12}(1, 0, 0) - 3 \geq m_2^{12}(1, 0, 0)$, which simplifies to

$$m_2^{12}(1, 1, 1) - m_2^{12}(1, 0, 0) \geq 3/\zeta.$$

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

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

$$f(1, 1) \cdot 3 \cdot \frac{f(1, 1) + f(1, 0)}{f(1, 1)} = 3f(1, 1) + 3f(1, 0).$$

If player 1 were to deviate by refusing to contract with player 3 while still contracting with player 2 as specified by the equilibrium, then player 1 would save this expected bonus without changing player 2's action in the underlying game. There would be an associated loss in player 1's relationship with player 3 of no more than $(5 - 3)[f(1, 1) + f(0, 1)]$, which is the expected surplus generated by player 3, because player 3's equilibrium payoff must be weakly greater than zero. In equilibrium, player 1 must be dissuaded from deviating and so we must have $3f(1, 1) + 3f(1, 0) \leq 2f(1, 1) + 2f(0, 1)$, which simplifies to $f(1, 1) \leq 2f(0, 1) - 3f(1, 0)$. The same steps apply to player 1 considering whether to refuse to contract with player 3, which implies $f(1, 1) \leq 2f(1, 0) - 3f(0, 1)$.

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

This example proves the result for the case of $n = 3$. It is easy to see that for any larger number of players, we can specify an underlying game in which players 1-3 interact as described above and the other players have no action choices and always zero payoffs. Let the network comprise exactly every pair $(1, j)$ for $j = 2, 3, \ldots, n$. The analysis above extends to this setting, where the deviations we check for player 1 include player 1 refusing to contract with every player $j \leq 4$ in addition to either player 2 or player 3, and we obtain the same conclusion regarding inefficient outcomes. □

For more on team production with partial verifiability and private contracting, see Goldmanis and Ray (2021).

Result 2: *For any given $n \geq 4$, there exists an underlying game $\langle A, u \rangle$ and a network L such that every player has a link in L (although L is disconnected) and, regardless of the contracting institution, there is no sequential equilibrium of the grand game in which an efficient action profile of the underlying game is played with positive probability.*

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

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

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

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

(i) $\Lambda^{\ell}_{ij}(h^{\ell-1}_{ij}, h^{\ell-1}_{ji}) = \{\underline{\lambda}\}$ for all $\ell > \hat{r}^{ij}$, $h^{\ell-1}_{ij}$, and $h^{\ell-1}_{ji}$; and

(ii) $\mu^{ij}(h_{ij}, h_{ji}, \cdot) \equiv \underline{m}$ for every $h_{ij} = (\lambda^{\underline{r}}_{ij}, \lambda^{\underline{r}+1}_{ij}, \ldots, \lambda^{\overline{r}}_{ij})$ and $h_{ji} = (\lambda^{\underline{r}}_{ji}, \lambda^{\underline{r}+1}_{ji}, \ldots, \lambda^{\overline{r}}_{ji})$ for which either $\lambda^{\hat{r}^{ij}}_{ij} = \underline{\lambda}$ or $\lambda^{\hat{r}^{ij}}_{ji} = \underline{\lambda}$ or both.

Result 3: *For any given $n \geq 4$, there exists an underlying game $\langle A, u \rangle$, a connected network L, and a number $\delta < 1$ such that, for every contracting institution exhibiting dated commitment, there is no sequential equilibrium of the grand game in which an efficient action profile of the underlying game is played with a probability greater than δ.*

Proof: A version of the example shown in Figure 3 suffices to prove this result. Let L be the network comprising exactly the pairs $(1, 2), (2, 3), \dots, (n-1, n)$. This is a linear, connected network. Let the underlying game be one in which only players 1 and n have choices to make, each chooses between actions 0, 1, and 2, and the payoffs of players 1, 2, $n-1$, and n are, in this order, shown in the table of Figure 3 with a_4 replaced by a_n. The other players get a payoff of 0 regardless of a_1 and a_n. Note that the efficient action profile entails $a_1 = a_n = 1$.

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

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

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

We would therefore have a contradiction if $\delta > 6/7$. The claim holds for any δ strictly between $6/7$ and 1. $\square$

Result 4: *For any given $n \geq 3$, there exists an underlying game $\langle A, u \rangle$, a connected network L, and a number $\delta < 1$ such that, for every contracting institution with strictly fewer than $n-1$ contracting rounds, no sequential equilibrium of the grand game has an efficient action profile in the underlying game played with a probability greater than δ.*

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

Suppose player 1 deviates from her equilibrium strategy by sending message $\underline{\lambda}$ to player 2 in every contracting round, ensuring that she has the null contract with player 2, and by choosing $a_1 = 2$ in the production phase. This deviation can have no effect on contracting between player $n-1$ and n, and also on player n's action in the production phase, because there aren't enough rounds through which this deviation can alter play in such a fashion as to disrupt contracting between players $n-1$ and n, and player n would not detect any deviation from the equilibrium path. Thus, with player 1's deviation, player n still must choose $a_n = 1$ with probability at least δ. Player 1's deviation gives this player an expected payoff of at least $\delta 9 + (1-\delta)2$, which must be a lower bound on player 1's equilibrium expected payoff. The other players' equilibrium payoffs are bounded as described in the proof of Result 3, and we reach a contradiction if $\delta > 6/7$ as before. $\square$

Result 5: *Take as given any integer $n \geq 3$, any finite set of underlying games G, and any integer $\kappa \in [2, n]$. Let $\mathcal{L}$ be the set of all connected networks of diameter weakly less than κ. There exists a contracting institution (representing private, independent, and voluntary contracting) satisfying $\overline{r} - \underline{r} \leq 2\kappa - 2$ that implements efficient outcomes.*

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

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

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

Notice that none of the analysis used to identify the fundamental elements for a given tuple $(\langle A, u\rangle, a^*, \underline{\alpha})$ requires a^* to be efficient. All that was required is that a^* is more efficient than $\underline{\alpha}$. Thus, we can repeat the construction of the fundamental elements by substituting for a^* any action profile $\tilde{a}$, provided that $\underline{\alpha}$ is a Nash equilibrium of $\langle A, u\rangle$ and $\sum_{i\in N} u_i(\tilde{a}) > \sum_{i\in N} u_i(\underline{\alpha})$. Then for any connected network L, all of the fundamental elements are well defined (not necessarily uniquely) and satisfy the conditions stated in Section 4.2. Further, we need not have limited the set of initial tuples to just one pair a^* and $\underline{\alpha}$ for each underlying game; that is, we could allow multiple combinations.

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

in which $\tilde{a}$ is played in the production phase and transfers are $(y^{ij})_{i\neq j}$ on the equilibrium path, so that the payoff vector is $u(\tilde{a}) + Y$.

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

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

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

Result 6: *Take as given any integer $n \geq 2$, any finite set of action profiles $\mathcal{A}$, and any permissible set $\mathcal{S}$ of scenarios. There exists a contracting institution (representing private, independent, and voluntary contracting) such that the following is true for every scenario $(\langle A, u\rangle, \tilde{a}, \underline{\alpha}, L, (y^{ij})_{i\neq j}) \in \mathcal{S}$: In the case in which $\langle A, u\rangle$ is the underlying game and L is the network, there is a sequential equilibrium of the grand game that yields the payoff vector $u(\tilde{a}) + Y$.*

Proof: Take as given any integer $n \geq 2$, any finite set of action profiles $\mathcal{A}$, and any finite set $\mathcal{S}$ of permissible scenarios. Fix the fundamental elements for these scenarios to satisfy the conditions of permissibility. For every $(\langle A, u\rangle, \tilde{a}, \underline{\alpha}, (y^{ij})_{i\neq j}) \in \mathcal{S}$, all of the steps described in Sections 4.3 and 4.4 to define feasible contracts and target conditional arrangements go through without alternation except for replacing a^* with $\tilde{a}$, and instead of applying Lemma 2 we can directly construct $(\breve{m}^{ij})_{(i,j)\in K}$ to have the required properties, by using the permissibility conditions. Specifically, we set $\breve{m}^{ij}$ to be the $\tilde{a}$-assurance contract with baseline transfer y^{ij}. The equilibrium construction then goes through as described in Appendix A.3, with no modifications. □

B.2 Additional Discussion

This paper's novel approach of constrained contracting-institution design allows us to analyze settings without limiting ourselves to a single model of contract negotiation, while also requiring that assumptions about contracting, such as its voluntary nature, are expressed separately from other assumptions on the contracting process and technology. If we had adopted one of the prior literature's simple models of contract formation from the start, we would have a limited view of contractual linkages and would not have found the main result.

Advantages of the noncooperative modeling approach

On the comparison between fully noncooperative models and cooperative matching of matching and coalitional bargaining, consider the example of a collaboration agreement discussed in Section 3.2, where, among other things, player 4's productive action directly affects player 1's payoff in the underlying game. Compare this to a supply-chain setting in which player 4 may provide an intermediate good to player 3, who in turn may provide an intermediate good to player 1. For the latter setting, suppose player 1's payoff is a function of only the type and quantity of the intermediate good delivered by player 3, and player 3's cost of producing the good for player 1 depends on the intermediate good supplied by player 4. Thus, player 1 cares about player 4's productive action only to the extent that it affects the negotiated terms of her contract with player 3.

Because these two settings are distinguished by different production technologies, they are differentiated unambiguously by a model that explicitly accounts for productive interaction, as accomplished herein by specifying the noncooperative underlying game. A modeling approach that abstracts from the underlying game by specifying payoffs as a function of an abstract set of contracts is not well suited to make the distinctions that these two examples illustrate. For instance, Fleiner et al. (2018), Fleiner et al. (2019), and others in the cooperative matching literature assume that a player's payoff depends on only the contracts this player signs, which would not allow for the externality in the collaboration-agreement example. Matching models that allow payoffs to be a function of the entire set of contracts formed, such as in Rostek and Yoder (2020, 2022) and Pycia and Yenmez (2019) for two-sided markets, can capture LDL externalities to some extent, but it is not clear how they could distinguish between, say, the collaboration-agreement and supply-chain examples without an explicit account of the production technology.[23]

[23]The supply chain example features what some call a pecuniary externality, though it may be more in-

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

In summary, the noncooperative approach, taken herein and in line with Jackson and Wilkie (2005) and Ellingsen and Paltseva (2016), has advantages that complement other approaches to the study of contractual networks.[24] The noncooperative approach provides a good foundation for precisely defining the technologies of production and enforcement, including the extent of verifiability. This structure allows one to sort out alternative methods of linking contractual relationships, such as conditioning transfers on third-party productive actions as opposed to contracts on contracts. Importantly, it also gives contracts their natural meaning, enabling predictions on the actual form that contracts take in applications.[25] Recall that contracts on contracts were ruled out in the present modeling exercise by defining contracts as mappings from the outcome of the underlying game and by the independence requirement on contracting institutions. Contracts on contracts would be enabled by altering either or both of these assumptions.

Barriers to efficiency and design in applications

Despite the emphasis of this modeling exercise on attaining efficiency, the Theorem should not be regarded as a claim that efficient outcomes will always be reached in reality, but rather as a reference for applications and a benchmark for further theoretical analysis. Applications vary technologically and may not fit with the assumptions made here, with respect to production and enforcement technologies as well as the contracting institution. Moreover, even under favorable conditions, efficiency relies on the players coordinating to achieve not just

structive to avoid the term and instead say the downstream result of a possible contracting or market distortion. Other entries in the matching literature include Ostrovsky (2008), Hatfield and Kominers (2012, 2015), Hatfield et al. (2013), Manea (2018), and Bando and Hirai (2021).

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

[25] Regarding contracts on contracts, Peters and Szentes (2008) tackle one of the key modeling components in their analysis of interactive promises. They examine settings in which players can make unilateral commitments about how they will play in the underlying game, and each player's promises can be conditioned on the promises of others. The authors develop a mathematical apparatus to handle the infinite regress issue, and they prove a folk theorem implying the existence of efficient equilibria. Their work suggests that interactive contracts require the external enforcement system to develop a sophisticated language for the cross-referencing of promises.

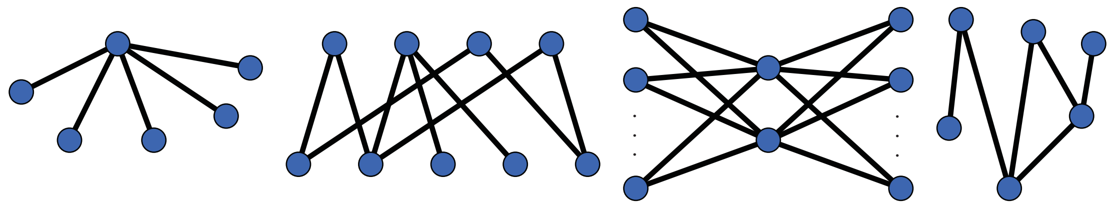

Figure 9: A few common networks.

an equilibrium but the right equilibrium from a potentially large set.

When evaluating barriers to efficiency, it may be helpful to categorize examples in terms of prominent aspects of their networks and the structure of their underlying games. Figure 9 illustrates four classes of networks. The networks shown would be suitable to model, from left to right, (i) vertical contracting with a single supplier, as well as common-agent or common-principal settings;[26] (ii) vertical contracting in a bipartite supply network and two-sided markets;[27] (iii) platforms and general intermediation networks;[28] and (iv) community interaction with an arbitrary contractual network.[29]

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

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

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

[28]Rochet and Tirole (2003, 2006) and Armstrong (2006) look at various pricing alternatives for intermediaries; Rysman (2009) provides an overview of this area of literature. Other entries include Weyl (2010) on pricing strategies and monopoly power, Lee (2014) on tipping points for platform adoption, Reisinger (2014) on two-part tariffs and equilibrium selection, Edelman and Wright (2015) on the microstructure of interaction between agents on the two sides of the market, and Hagiu and Wright (2015) on vertical integration.

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

On the efficacy of simpler contracting institutions

Following the discussion in Section 5.5, I comment further on whether efficient contracting requires options to adjust externally enforced elements, as the SCO institution facilitates in rounds 1 through $\overline{r}$. For instance, consider a "Simultaneous Contracting and Sequential Communication" (SCSC) institution that has one round of simultaneous contract creation, determining the induced game $\langle A, u + M\rangle$ for the production phase, followed by multiple rounds of messages that do not affect the contracts but are used by the players to coordinate on actions to take in the production phase. Whether an SCSC institution can implement efficient outcomes is not addressed by Results 3 and 4, which show that sequential contracting is essential but do not distinguish between external and self-enforced options.

I conjecture that SCSC institutions would not be effective generally but perhaps could implement efficient outcomes for some special classes of underlying games. The best hope might be for underlying games in which $\underline{\alpha}_i(a_i^*) = 0$ for every player i; but even if an efficient equilibrium exists (which is not clear), the construction remains challenging and still requires players to coordinate on some sort of assurance contracts.[30] There is greater doubt for dealing with other underlying games. Consider the case in which $\underline{\alpha}_i(a_i^*) > 0$ for at least one player i. A disruption in contracting with such a player may not motivate her partner to pass along word of the disruption. For example, if $\underline{\alpha}_i(a_i^*) = 1$, then this player would still choose her part of the efficient action profile in response to a belief that the other players have flipped to $\underline{\alpha}$, and so off-equilibrium-path beliefs and behavior could not be as simple as coordination on $\underline{\alpha}$. More broadly, there is a direct conflict between assurance penalties and making $\underline{\alpha}$ a Nash equilibrium of the "induced game" in the production phase following a deviation. The proof of the Theorem (and SCO institution) deals with these problems by use of conditional arrangements that switch to different contracts upon cancellation and, by varying the forcing arrangements as a function of ϕ, induce players to choose actions in the production phase that, under their on-equilibrium-path contracts, would expose them to high assurance penalties.[31] Characterizing what SCSC and alternative institutions can achieve is

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

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

a worthwhile topic for future research; I have not been able to make much progress in the context of the current exercise.

Other Comments

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

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

There is one sense in which allowing for money burning could be marginally useful. Recall that in the SCO contracting institution there are rounds in which the feasible messages for a pair of players depend on their past messages to each other. For instance, if a pair made a conditional arrangement at some round $r < 0$, then they are restricted to silence in rounds $r + 1, \dots, 0$. The reason for introducing this restriction was technical: It reduced the number of information sets in the grand game, simplifying the process of constructing a sequential equilibrium with the desired properties (which is complicated still). With money-burning, we could assume that the message spaces in each round are history independent, but obtain effectively the same restriction as before by having players burn money if players send non-null messages when not allowed.