Smart Contracting in Network Markets*

Darrell Duffie and Chaojun Wang
Stanford GSB Wharton School

February 24, 2026

Abstract

With complete-information bilateral bargaining in network settings, holdup is eliminated when contracts across the network are agreed atomically (all or none) via a smart contract. Applications include over-the-counter trading, syndicated lending, multi-tranche securitizations, third-party financed purchases, and bookbuilding. Under a novel extensive-form bargaining protocol, any firm can give a "greenlight" to the terms of a contract proposed to that firm, which automatically converts those terms into a binding contract if the terms proposed to all other firms also receive greenlights. In any Perfect Bayesian Equilibrium with Markov strategies, firms immediately agree on socially efficient contracts that equalize expected gains across firms.

JEL Classifications: C70, D47, D60, D70, G12, K22, L14

Keywords: Network, bargaining, smart contract, atomic settlement, holdup, over-the-counter

*We are grateful for conversations with Peter DeMarzo, Itamar Drechsler, Vincent Glode, Sergiu Hart, Urban Jermann, David Kreps, Alistair Milne, Julien Prat, Ariel Rubinstein, Nick Roussanov, Andrzej Skrzypacz, Takuo Sugaya, Robert Wilson, Ali Yurukoglu, and Jeff Zwiebel. Comments are welcome. Duffie is also a NBER Research Associate and a Fellow of the Canadian Derivatives Institute.

1 Introduction

In settings with bilateral bargaining between networked agents under complete information, we show that holdup inefficiencies are eliminated if agreements on the contracts being negotiated are reached atomically (all or none) by using smart contracts. Applications include over-the-counter trading networks, syndicated lending, multi-tranche securitizations, purchases requiring third-party financing, "building a book" for a security offering, and other multi-agent markets in which pairs of agents bargain bilaterally over terms that could affect the values of their negotiations with other firms. Under our suggested protocol, the terms of a contract proposed to any firm can be given a "greenlight" by that firm, automatically converting those terms into a binding contract if and only if the terms proposed to all other firms are also assigned greenlights and are thus likewise converted to contracts. We analyze tree-based networks, including hub-spoke and chain networks that are common in many settings. Under mild technical conditions, the unique outcome of any Perfect Bayesian Equilibrium with Markov strategies is immediate agreement on socially efficient contracts that generate equal expected gains for every firm.

The Coase Theorem states that if property rights are well defined and transaction costs are sufficiently low, private parties can bargain to reach an economically efficient outcome without government intervention, regardless of the initial allocation of rights (Coase, 1960). The holdup problem marks a limitation of the Coase Theorem (Klein, Crawford, and Alchian, 1978, Williamson, 1975, 1979). Specifically, the impossibility of writing fully contingent contracts that anticipate all future states is a form of transaction cost. With bilateral contracting among networked agents, holdup is prevalent because ending a negotiation with one counterparty subjects a firm to holdup when negotiating with another counterparty. Even fully contingent contracts cannot always overcome the holdup inefficiencies in this setting.

Consider, as one illustrative example, an over-the-counter asset market in which an intermediary with a low valuation for holding the asset is in a position to buy from a prospective seller and sell to a prospective buyer. The intermediary recognizes that if a trade is first

agreed with the seller, the prospective buyer would then be able to demand a low price, given the intermediary's low valuation. Anticipating this holdup, the intermediary may avoid an initial purchase from the prospective seller. The same problem arises if the intermediary first negotiates with the prospective buyer. In a worked example, we show that our smart-contracting protocol overcomes this holdup problem, achieving efficient trade, whereas excluding smart contracting from an otherwise identical bargaining protocol leads to no trade in any equilibrium (Section 5.1).

When arranging a leveraged loan syndication or issuing the tranches of a collateralized loan obligation, the hub firm (the lead bank or the issuer) bilaterally negotiates terms with each provider of capital. Without the benefit of a greenlight protocol, the hub firm is exposed to holdup because of the need to go back and forth to the various capital providers in an attempt to get deal terms that are collectively consistent with the total required financing or the total collateral pool. Each capital provider has the incentive to shade its demands (make "low-ball" offers) in an attempt to force better terms, knowing that the hub firm needs all capital providers to agree before the deal is done. Holdup can lead to an inefficiently large commitment of capital by the hub firm on "cold" deals or underpricing of "hot" ones, despite mechanisms such as market flex provisions designed to mitigate holdup (Bruche, Malherbe, and Meisenzahl, 2020). Holdup is eliminated by our protocol because no contract is agreed unless all contracts are agreed.

For a third example, consider a buyer's need for third-party financing in order to complete the purchase of an asset. The buyer is exposed to holdup because the all-in profitability of the deal depends on both the price obtained from the seller and the financing terms obtained from the debt capital provider.

Absent externalities, one might conjecture that it is easy to achieve efficiency, at least in hub-and-spoke networks. The hub firm could make one-time take-it-or-leave-it offers to spoke firms, as in Segal (1999), or repeatedly make offers to spoke firms until acceptance is achieved. For example, Manea (2018, p.1275) shows that efficiency is achieved in the limit as the probability p that the hub firm makes offers to the spoke firms approaches 1. Possessing all of the bargaining power, the hub firm would offer socially efficient terms and extract all

surplus. All spoke firms would agree. Implementing a protocol that assigns full bargaining power to certain firms is challenging, because those firms would need a way to commit to avoid considering any counteroffer. Without assigning full bargaining power to certain firms, or without some other approach such as the smart-contract protocol that we model, holdups can be expected to cause inefficiency or even a failure to reach agreement (Gofman, 2011, Manea, 2018).

The rest of the paper is organized as follows. Section 2 outlines prior related work. Section 3 defines the primitive objects of bilateral bargaining in multi-agent games, states the "greenlight" protocol, and provides its equilibrium. Section 4 explains the role of the greenlight protocol in overcoming holdup problems, the limited extent to which prior actions by other agents need to be commonly observed in order to obtain the efficient equilibrium outcomes, and methods used in practice to implement greenlight mechanisms, such as smart contracts or trusted neutral authorities. Section 5 shows how the greenlight protocol resolves potential holdup inefficiencies in one illustrative application to over-the-counter asset trading markets. Section 6 shows that the efficient equilibrium expected outcome of the greenlight bargaining protocol is equivalent to the outcome of an axiomatic solution that we call Coordination-Proof bargaining, which refines the set of "Nash-in-Nash" outcomes of Horn and Wolinsky (1988) to a unique bargaining outcome. Section 7 concludes.

2 Relationship to prior work

Our marginal contribution is the design of a new protocol for bilateral bargaining among networked agents that, under natural conditions, leads to efficient equilibrium contracts. We are not aware of other bargaining protocols that avoid the holdup problem in this general setting. Some practical potential examples of our greenlight protocol approach are discussed in Section 4.3. Our results build on models of efficient bargaining in a network of only two agents (Binmore, 1980, Rubinstein, 1982, Binmore, Rubinstein, and Wolinsky, 1986).

There has been a significant amount of prior work on bilateral bargaining among net-

worked agents.[1] Notably, Stole and Zwiebel (1996) and Brügemann, Gautier, and Menzio (2019) analyzed protocols that permit costless renegotiation of any previously reached agreement. Under these protocols, failure to reach an agreement (a breakdown) cannot be renegotiated and thus leads to holdup inefficiencies. The contingent contracting protocols modeled by de Fontenay and Gans (2014) are also unable to resolve the holdup problem because, in the event of a breakdown between two players, A and B, a third player C can hold up B. For example, in a hub-spoke network, intra-firm wage bargaining with the protocol of de Fontenay and Gans (2014) allows for fully contingent contracts yet can lead to over-hiring relative to the efficient benchmark because this protocol implements the Shapley value, as in Brügemann et al. (2019).

In settings involving matching with transfers, holdup can cause inefficient matching or costly delays.[2] As a result, Elliott and Nava (2019) "find inefficiencies to be a common feature" despite the selection of "a protocol that favors the existence of efficient equilibria" in the absence of a centralized mechanism that guarantees efficient matching, as in Polanski (2007), or a folk-theorem type of mechanism, as in Abreu and Manea (2012a).[3] In network

[1]This literature includes Gul (1989), Corominas-Bosch (2004), Polanski (2007), Abreu and Manea (2012a,b), Condorelli, Galeotti, and Renou (2017), Manea (2018), Collard-Wexler, Gowrisankaran, and Lee (2019), and Elliott and Nava (2019). Lee, Whinston, and Yurukoglu (2021, Section 2.1) provides an overview of some of the key results. There is a separate literature addressing strategic multilateral bargaining (which includes coalitional and legislative bargaining), in which one agent proposes actions for all members (or for a selected coalition of members), who then jointly vote on the proposal (Baron and Ferejohn, 1989, Merlo and Wilson, 1995, 1998, Jackson and Moselle, 2002, Snyder Jr., Ting, and Ansolabehere, 2005, Battaglini and Coate, 2007, Battaglini, Nunnari, and Palfrey, 2012, Hart and Mas-Colell, 1996, Chatterjee, Dutta, Ray, and Sengupta, 1993). Our theory applies instead to settings in which pairs of firms *bilaterally* negotiate the terms of their contracts without observing ongoing negotiations between other pairs of firms. Another literature uses the mechanism design approach to address bilateral contracting in network settings (Crémer and Riordan, 1987). Under complete information, holdup is absent because any given mechanism limits each player's choice of cash and non-cash terms to those prescribed by the mechanism and therefore assumes away her ability to appropriate surplus from others.

[2]See Corominas-Bosch (2004), Gale and Sabourian (2006), Abreu and Manea (2012b), and Elliott and Nava (2019).

[3]Another literature considers decentralized bargaining in *large* matching markets, in which either the number of players is infinite or players who reach agreement are replaced by exact replicas (Rubinstein and Wolinsky, 1985, Gale, 1987, Binmore and Herrero, 1988, Manea, 2011, Lauermann, 2013). Holdup is precluded by two assumptions in this setting: a player is allowed to reach an agreement with at most one connected player, and any player's exit does not affect other players' matching opportunities (Elliott and Talamàs, 2025). In the large literature on matching without transfers (Haeringer and Wooders, 2011, Ferdowsian, Niederle, and Yariv, 2025), holdup is absent because players have no means of appropriating surplus from one another.

markets, holdup can lead to inefficient allocation and intermediation when coupled with competition, or can cause a complete failure to trade (Manea, 2018, Gofman, 2011).[4] In the setting of vertical contracting, Collard-Wexler et al. (2019) assumes a "feasibility" condition on payoffs that implies that holdup is not sufficiently severe to impede equilibrium efficiency.[5] As we show, this assumption fails to hold in our example of an over-the-counter asset trading market in which the intermediary has a relatively low valuation for holding the asset. In this example, without the greenlight mechanism, holdup is severe enough to prevent efficient trade, whereas the greenlight protocol generates efficient trade (Section 5.1). Our network structure, contracting spaces, and payoff dependencies are general enough to encompass matching with transfer and intermediation in networks under complete information.

There is also a substantial amount of prior work on axiomatic solutions to bilateral and multilateral bargaining in network markets.[6] By refining the "Nash-in-Nash" outcomes of Horn and Wolinsky (1988) to a unique bargaining outcome, we offer a new axiomatic solution concept, Coordination-Proof bargaining, that generates agreement on the same terms achieved in our noncooperative extensive-form bargaining setting by the greenlight protocol.

Lee, Martin, and Townsend (2024a,b) explain that smart contracts effectively eliminate the risk that an agreed trade will not be settled with an exchange of cash and assets, but settlement on a public blockchain reveals the trade to potential future counterparties, introducing a holdup friction when negotiating the terms of future trades. Holden and Malani (2017) and Gans (2019) describe how smart contracts can be used to enforce contracts.[7] Related to this, Danos, Krivine, and Prat (2021) design a smart contract that improves the enforcement of payment obligations of exogenously agreed trades. In our model, rather than

[4] Condorelli et al. (2017) sidestep holdup by endowing sellers with full bargaining power and focus on the effect of asymmetric information instead.

[5] This assumption "is necessary for there to exist a Nash-in-Nash limit equilibrium in which all agreements form immediately…" (Collard-Wexler et al., 2019, p.168).

[6] Classics include Shapley (1953), Myerson (1977a,b), Horn and Wolinsky (1988), Jackson and Wolinsky (1996). More recent work includes Hatfield, Kominers, Nichifor, Ostrovsky, and Westkamp (2013).

[7] Holden and Malani (2017) write: "A basic problem in contracting is holdup: after one party has made relationship-specific investments, the other party refuses to perform unless the first one offers better terms than the original contract." Similarly, Gans (2019) writes: "Finally, the buyer could hold up payment to the seller if the product is of low quality, however, the buyer could also do that if the product is as expected."

treating the enforcement of agreed trades, agents endogenously bargain over the terms of trade. We show that the greenlight protocol overcomes the holdup problem associated with negotiating the terms of trade, leading in equilibrium to agreement on efficient terms. In the setting of asset trading, the greenlight mechanism also avoids holdup related to the enforcement of trade settlements if settlement is implemented with smart contracts. We provide some examples.

Holden and Malani (2017) provide a qualitative discussion of how smart contracts can address holdup problems in two-player bargaining settings. Townsend (2020) and Gans (2019) survey various types of informational frictions and incentive problems that are mitigated with smart contracts.

3 Efficiency of strategic bargaining with greenlights

This section defines the primitive objects of bilateral bargaining in multi-agent games and provides a protocol based on a holdup-proof mechanism that we call "greenlighting." We define the extensive form of the resulting bargaining game and provide its equilibrium solution, showing that it is efficient and generates equal expected gains for all participants.

3.1 Primitives

A finite set N of n players, called "firms," is identified with the set of vertices of a graph $G = (N, E)$ whose set E of edges defines the pairs of firms that bilaterally bargain over their contract terms. For each edge (i, j), some non-cash contract term s_{ij} is to be chosen from a given finite set S_{ij}. This choice, along with some cash payment $y_{ij} \in \mathbb{R}$ from i to j, defines the terms of a potential bargain between firms i and j. The entire profile of non-cash terms of firm i is some $s_i = (s_{ij})_{(i,j)\in E} \in S_i = \prod_j S_{ij}$.

Firm i assigns a value to its profile of non-cash terms given by some $f_i : S_i \to \mathbb{R} \cup \{-\infty\}$, allowing for the possibility that some profiles are not feasible and are therefore given the value $-\infty$. For instance, in a market with three firms, A, B, and C, suppose that B is endowed with 5 units of a physical asset, while its counterparties A and C are endowed with

none. Clearly, it is not feasible for B to sell 4 units to A and 3 units to C. The total value to firm i for some profile s_i of non-cash terms and associated payments $y_i = (y_{ij})_{(i,j)\in E}$ is

$$v_i(s_i, y_i) = f_i(s_i) - \sum_{(i,j)\in E} y_{ij}. \tag{1}$$

In summary, the game has primitive objects (G, S, f), where $S = (S_i)_{i\in N}$ and $f = (f_i)_{i\in N}$.

We maintain complete information throughout. The network structure G and all firms' feasible actions S and preferences f are common knowledge.[8]

Because the payments (y_{ij}) are transfers that net to zero, the total social value of the profiles $s = (s_i)_{i\in N}$ is $F(s) = \sum_{i\in N} f_i(s_i)$. We assume that there is a unique socially efficient profile $s^* = (s_i^*)_{i\in N}$ of non-cash terms, that is, a unique profile that maximizes $F(s)$. The uniqueness property is generically guaranteed.[9]

3.2 The greenlight bargaining protocol

We now define and characterize the equilibrium of a *greenlight* extensive-form-game protocol for any network bilateral bargaining primitives (G, f, S), as defined in the previous subsection. From this point, we assume that G is a tree, meaning that there is a unique sequence of edges connecting any two distinct nodes.

We need some preliminary definitions. For any given firm i, the graph G can be viewed as a directed tree with root i. The set $c(i)$ of firms that are immediately connected to i are the "child" firms of i. A child firm j of i has its own unique set $c(j)$ of child firms, and so on. Fixing such a directed tree, firm j can consider a collection $\{(s_{jk}, y_{jk}) : k \in c(j)\}$ of terms to offer its respective child firms. We consider the following "greenlight" mechanism for converting such offers to contracts (agreed bargains). When firm k receives an offer from its parent firm j, firm k can pick any probability g_k for the event that the offer is given a greenlight. All offers automatically become contracts only in the event that all offers

[8] While asymmetric information is relevant in many market settings, we are not aware of tractable general models of extensive-form network bargaining games that allow for asymmetric information.

[9] That is, parameterizing the model by the payoff values $\{f_i(s_i) : s_i \in S_i, i \in N\}$, treated as an element of an open set $\mathcal{P}$ of possible parameters in a Euclidean space (ignoring the $-\infty$ values of f_i), there is a closed set of Lebesgue measure zero of parameters in $\mathcal{P}$ for which there is more than one maximand of F.

have been given a greenlight by their receivers. The next section discusses how greenlight mechanisms arise in practice.

Formally, the extensive form of the game is as follows. We fix a probability space $(\Omega, \mathcal{F}, P)$. The bargaining process has a sequence of stages, numbered $0, 1, 2, \ldots$. The information available to all firms at the beginning of stage t is defined by a sub-σ-algebra $\mathcal{F}_t$ of $\mathcal{F}$ that increases in t as information accumulates. There is no initial information, in that $\mathcal{F}_0$ is trivial.[10] The *status-quo* value $f_i(s_i^0)$ to firm i is without loss of generality normalized to zero. The value to firm i for finalizing agreement on terms (s_i, y_i) in stage t is $\delta^t v_i(s_i, y_i)$, for some common discount factor $\delta \in (0, 1)$. Time discounting makes delay costly, which is needed for equilibrium uniqueness in two-person strategic bargaining (Rubinstein, 1982, Binmore et al., 1986). The size of the discount factor δ plays no role in the expected equilibrium payoffs because we will show that there is always immediate equilibrium agreement at stage 0 at terms that are efficient and have expected cash terms that do not depend on δ. Appendix C replaces time discounting with a probabilistic breakdown of bargaining and allows for *endogenous* continuation values upon breakdown. We show that efficiency is robust to this alternative specification.

Each stage t of the bargaining process proceeds as follows. One of the firms, denoted $I(t)$, is selected at random, uniformly among the firms, to make the first offers of terms, as illustrated in Figure 1. This selection of the initial proposer is observable only to firm $I(t)$ until the beginning of the next stage. Treating the graph G as a directed tree with root $I(t)$, firm $I(t)$ makes offers to each of its child firms. Once a child firm j receives an offer (s_{ij}, y_{ij}) from its parent i, it chooses a probability g_j for greenlighting its parent's offer and chooses the offers $(s_{jk}, y_{jk})_{k \in c(j)}$ to its own child firms, and so on, until each firm other than the root firm $I(t)$ has received an offer from its parent firm and has chosen a probability with which to greenlight this offer. These choices, g_j and $(s_{jk}, y_{jk})_{k \in c(j)}$, must be based only on (measurable with respect to) the information contemporaneously available to firm j, which is the join $\mathcal{F}_{jt}$ of three sources of information: *(i)* the information $\mathcal{F}_t$ commonly available at the beginning of stage t, *(ii)* the event of whether firm j is selected to make the first offers in

[10] That is, the events in $\mathcal{F}_0$ have probability zero or 1.

stage t, $\mathbb{1}_{j=I(t)}$, and *(iii)* for a non-root node, the information contained in its parent's offer to firm j. We say that firm j follows a Markov strategy if its choices, g_j and $(s_{jk}, y_{jk})_{k\in c(j)}$, are measurable with respect to the firm's current-stage information, namely *(ii)* whether firm j is selected to make the first offers in stage t, $\mathbb{1}_{j=I(t)}$, and *(iii)* (for a non-root node) the information contained in its parent's offer to firm j.

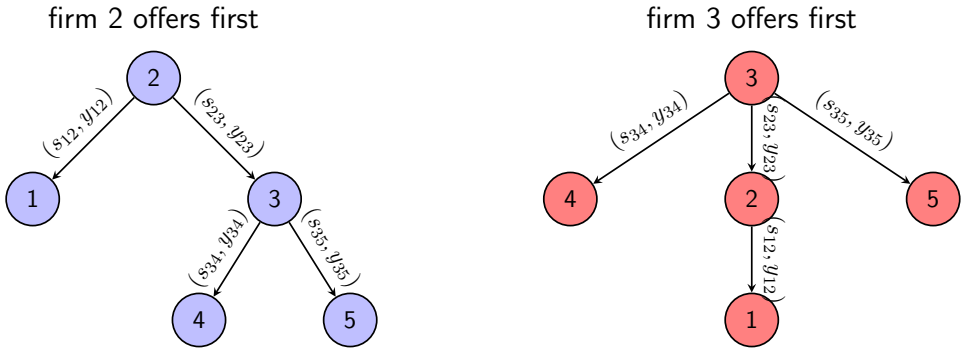


Figure 1: An illustration of the order in which offers are made at stages t and $t+1$ of a bargaining game for a fixed network of 5 firms in which the initial proposer of offers in stage t is $I(t) = 2$ and in stage $t+1$ is $I(t+1) = 3$.

The stage-t greenlight events are then realized, independently. If all offers receive a greenlight, then the terms of all offers automatically become binding agreements and the game ends. If the outcome is that one or more offers do not receive a greenlight, the game goes to the next stage, $t+1$. This "all-or-nothing" aspect of the protocol is the essence of "atomic" contracting.

As a tie-breaking rule, if some firm j is indifferent between any two choices for its offers and greenlight probability, $((s_{jk}, y_{jk})_{k\in c(j)}, g_j)$ and $((s'_{jk}, y'_{jk})_{k\in c(j)}, g'_j)$, then firm j chooses between them based on which would receive the higher conditional expected number of greenlights from its child firms.[11] If the two alternative sets of offers would receive the same expected number of greenlights from the child firms, then firm j chooses between

[11] This expectation is $\sum_{k\in c(j)} g_k$. If more than two such choices are tied, tie-breakers are applied pairwise in any order. This tie-breaking rule ensures that there is no coordination failure that could arise if a firm does not greenlight an offer or offers terms for which it does not expect to receive greenlights simply based on a self-fulfilling belief that some other offers would not receive greenlights.

$((s_{jk}, y_{jk})_{k\in c(j)}, g_j)$ and $((s'_{jk}, y'_{jk})_{k\in c(j)}, g'_j)$ based on which assigns the higher probability of a greenlight to the offer that j received from its parent firm, that is, based on the higher of g_j and g'_j. In the event that there is still a tie between $((s_{jk}, y_{jk})_{k\in c(j)}, g_j)$ and $((s'_{jk}, y'_{jk})_{k\in c(j)}, g'_j)$, then firm j chooses one of them arbitrarily.

The full history of play through stage t becomes observable to all firms at the beginning of stage $t+1$. The common information $\mathcal{F}_{t+1}$ at the beginning of stage $t+1$ is then the join of $\mathcal{F}_t$, the information revealed by $I(t)$ and all of the stage-t offers and greenlight events. Appendix B shows that little of this information is actually needed to support the equilibrium described in the following main result of the paper, whose proof is found in Appendix A. The game induced by (G, S, f, δ) with the above-defined greenlight bargaining protocol is denoted $\Gamma(G, S, f, \delta)$.

We consider Perfect Bayesian Equilibrium (PBE) with Markov strategies, but allow for non-Markovian deviations.

Theorem 1. *$\Gamma(G, S, f, \delta)$ has a Perfect Bayesian Equilibrium (PBE) with Markov strategies. All PBE with Markov strategies implement the unique socially efficient profile s^* of non-cash terms in stage 0 with probability 1. All PBE with Markov strategies have the same expected value $F(s^*)/n$ for each firm.*

The result states that the expected gain $E(v_i(s_i, y_i))$ is equalized across firms. In fact, gains are equalized not only in expectation; gains are actually equalized almost surely in the limit as the time-discount effect gets small. For example, with fixed rates of time preference, as the time between bargaining rounds shrinks to zero, all firms gain equally in the limit.

Proposition 1. *Fix a sequence $\{\delta_k\}$ of discount factors converging to 1. Let $(s^*, y(k))$ be the realized terms of any PBE with Markov strategies for $\Gamma(G, S, f, \delta_k)$. For each firm i, $v_i(s_i^*, y(k)_i) \to F(s^*)/n$ almost surely.*

4 Discussion

This section includes a discussion of the role of the greenlight bargaining protocol in overcoming holdup problems, the extent to which common observability of all prior offers and greenlight decisions is actually required, and methods used in practice to implement the greenlight protocol.

4.1 Greenlights overcome holdup inefficiencies

Without the smart-contract mechanism that we introduced, holdups can block efficiency. For example, suppose there are three firms in a network, in which firm 1 is connected to firm 2, which is connected to firm 3. Suppose that firm 1, when negotiating with firm 2, believes that firm 2 may have completed a bargain with firm 3. This affects the *status-quo* value of firm 2 with respect to its negotiation with firm 1, possibly putting firm 2 at a relative disadvantage. Anticipating this holdup effect, firm 2 may initially negotiate a bargain with firm 3 that reduces its exposure to holdup by firm 1, but potentially also reduces the total gain of the three firms. The next section includes a worked-out example of this holdup effect for an over-the-counter asset trading market. The greenlight mechanism eliminates this holdup inefficiency because it prevents firm 1 from exploiting the potential existence of a committed agreement between firms 2 and 3. The same intuition applies to the general game setting described in the previous section.

There are non-trivial bargaining settings (G, S, f) with no such holdup problem. For example, suppose the payoff functions f_i are additive in that, for each edge (i, j), there is some $f_{ij} : S_{ij} \to \mathbb{R}$ with the property[12] that

$$f_i(s_i) = \sum_{(i,j)\in E} f_{ij}(s_{ij}).$$

In this case, there is no holdup problem and efficiency can be achieved without the greenlight mechanism.

[12]This additivity property includes an assumption of finiteness of f_{ij} and rules out feasibility constraints.

Holdup is also absent in a protocol that assigns all the bargaining power to "upstream" firms in a network. In the limiting case of Manea (2018, p.1275) in which the probability that an offer is made by the hub firm approaches one, the sole asset owner in the network repeatedly makes sale offers to one of its connected potential intermediaries until acceptance, then the next asset owner does the same, and so on. Endowed with all the bargaining power, each asset owner i chooses the intermediary j that maximizes the total surplus while extracting all that surplus, anticipating that j will do the same. Our protocol, however, does not assign full bargaining power to any firm, because every firm has the opportunity to make counteroffers in subsequent stages. In Section 5.1, we define an alternative bargaining game $\Gamma^{\text{no greenlight}}(G, S, f, \delta)$ that is identical to $\Gamma(G, S, f, \delta)$ with the exception that there is no greenlight protocol. Only for the special case in which a fixed firm is selected to make the first offer in every stage does the bargaining game $\Gamma^{\text{no greenlight}}(G, S, f, \delta)$ reduce to the limiting case of Manea (2018). Otherwise, $\Gamma^{\text{no greenlight}}(G, S, f, \delta)$ involves holdup and can result in inefficient terms, as we show in Proposition 2. Moreover, we have verified in unreported work that the conclusion of Proposition 2 holds, under some adjusted valuation parameters, even if each firm is randomly selected to make the first offer with an arbitrary non-zero probability.

4.2 Common observability of past offers is not crucial

Within a given stage of our model, a firm is unable to observe ongoing negotiations between other pairs of firms. However, we assume that, at the beginning of any bargaining stage, all actions in prior bargaining stages become observable. This is unrealistic in many practical settings. One may be concerned that the equilibrium efficiency result in Theorem 1 is due not only to the greenlight approach but also to the large amount of common information. For example, one might worry that our efficiency result relies somehow on potential punishments and rewards, as in folk-theorem sorts of mechanisms like that of Abreu and Manea (2012a). Restricting to PBE with Markov strategies rules out the achievement of efficiency through such folk-theorem types of mechanisms. More generally, Appendix B shows that the

equilibrium outcome of Theorem 1 remains intact even if, for some arbitrarily large number T of stages, at the beginning of any stage t before T, no firm has received any information about prior offers to other firms or the prior greenlight outcomes of other firms.

Relative to the perfect observability of other players' actions that is commonly assumed in the literature,[13] our informational requirements are mild.

We also do not rely on equilibrium refinements that restrict off-path beliefs—such as passive beliefs[14]—thereby broadening the scope of applications. For example, consider the contracting of terms with investors in the various tranches of a family of collateralized loan obligations (CLOs) backed by the same asset pool. If the CLO issuer (the hub firm) offers a larger quantity to Buyer 1 (a spoke firm) than anticipated, it is reasonable for Buyer 1 to think that other buyers have demanded less than anticipated.

4.3 How greenlight mechanisms arise in practical settings

Greenlight mechanisms can arise in practice through the existence of a trusted neutral agent that can inform all firms whether or not all offers have been given a greenlight, a contractible signal that converts offers into binding contracts.

For example, consider a commercial real estate transaction between a buyer, a seller, and a bank providing financing to the buyer. A title firm can inform all three contracting parties whether the purchase agreement is final and the financing contract is final because *(i)* the buyer and the seller contractually agreed to complete the purchase contingent on financing from the bank to the buyer, and *(ii)* the bank and the buyer agreed on mortgage financing, secured by the real estate, contingent on transfer of title from the seller to the buyer.

A practical alternative to a trusted neutral agent for converting greenlit agreements into binding contracts is a smart contract, conceived by Szabo (1996) long before its popular implementation in blockchain applications. Szabo describes a smart contract as "a set of promises, specified in digital form, including protocols within which the parties perform

[13] Examples include Stole and Zwiebel (1996), Brügemann et al. (2019), Manea (2011, 2018), Elliott and Nava (2019).

[14] See Hart, Tirole, Carlton, and Williamson (1990), McAfee and Schwartz (1994), Segal (1999), de Fontenay and Gans (2013), Collard-Wexler et al. (2019).

their promises." Smart contracts offer a single, general-purpose implementation of the "all-or-none" greenlight bargaining protocol, eliminating the need for bespoke trusted neutral agents.

For example, in the real estate transaction just described,[15] the closing funds and the title of the real estate can be tokenized and recorded on a ledger, such as a blockchain. Initially, the real estate token is recorded on the ledger as the property of the seller. The tokenized funds are recorded as properties of the bank and the buyer. The buyer's offer to the seller and the buyer's request to the bank for a mortgage are then recorded on the ledger. These are irrevocable if both the seller and the bank agree within a stipulated time, but have no effect if either the seller or the bank fails to agree by the stipulated time. The seller can then place a record on the ledger indicating its greenlight for an agreement to the buyer's offer, contingent on financing from the bank and receipt of funds. The bank can place a record on the ledger indicating its greenlight for an agreement to provide the needed secured financing to the buyer, contingent on a greenlight for the purchase agreement. If and when these two greenlight messages are recorded on the ledger, the ledger software automatically and irrevocably executes the purchase and financing agreements, transferring the mortgage principal from the bank to the buyer, the real estate title from the seller to the buyer, the closing funds from the buyer to the seller, and recording the mortgage obligation of the buyer to the bank, secured by the real estate. (This mechanism does not, on its own, guarantee future mortgage payments by the buyer.)

Smart contracting in tokenized asset markets has become a standard method for the *settlement* of agreed trades (World Economic Forum, 2025). Our results show how smart contracts enable a protocol supporting the *bilateral negotiation* of terms of trade that are collectively efficient in network settings.

[15] For background, see Freyermuth and Tosato (2023).

5 An application to asset trading networks

As one illustrative application, we show how the smart contracting approach resolves potential holdup inefficiencies when applied to an over-the-counter asset trading market.

After the global financial crisis, stricter bank capital regulations increased the cost to the shareholders of large dealer banks of absorbing trades onto their balance sheets (Andersen, Duffie, and Song, 2019). As a result, large dealers have sometimes reduced their provision of market liquidity and have increasingly shifted from direct principal-based intermediation of trades to riskless-principal trades or, similarly, "agency trades," in which the dealer acts effectively as an agent that enables trades between buying and selling customers.[16] Viewed through the lens of our model, dealers may conduct riskless-principal and agency trades not only to lower their inventory costs but also to mitigate holdup problems that arise when the dealer has significant inventory costs. As an alternative to riskless-principal or agency intermediation by dealers, trade could be mediated with smart contracts, as explained in this section.

Let $G = (N, E)$ be a tree network composed of n trading firms. An asset can be bilaterally traded between any two connected firms. Initially, firm i owns an integer number $q_i^0 \geq 0$ units of the asset. The total amount of the asset available in the market is $Q = \sum_{i \in N} q_i^0$. The non-cash term $s_{ij} = q$ means that firm i transfers q units of the asset to firm j.

Each firm i has an exogenous fixed valuation $u_i(q) \in \mathbb{R} \cup \{-\infty\}$ for q units of the asset, with short sales precluded by our assumption that $u_i(q) = -\infty$ for $q < 0$. The trading game is the network bargaining game with primitives (G, S, f), where

$$S_{ij} = \{-Q, -Q+1, \ldots, Q-1, Q\}, \qquad (i,j) \in E, \tag{2}$$

$$f_i(s_i) = u_i\left(q_i^0 - \sum_{(i,j) \in E} s_{ij}\right) - u_i\left(q_i^0\right). \tag{3}$$

[16] Bessembinder, Jacobsen, Maxwell, and Venkataraman (2018), Bao, O'Hara, and (Alex) Zhou (2018) and Choi, Huh, and Seunghun Shin (2024) provide such empirical evidence.

Thus, we assume integer trade sizes. The *status-quo* (no-trade) value of firm i is

$$f_i\left(s_i^0\right) = u_i(q_i^0) - u_i(q_i^0) = 0.$$

The socially efficient allocation is $\left\{q_i^0 - \sum_{(i,j)\in E} s_{ij}\right\}_{i\in N}$, where s^* is the unique solution of $\max_s F(s)$ with $F(s) = \sum_{i\in N} f_i(s_i)$.

5.1 Inefficient allocations without smart contracts

We first consider the asset trading game without the greenlight bargaining protocol. Every stage proceeds exactly as in the original game setting of Section 3.2, but for two differences:

1. The probability to greenlight an offer (s_{ij}, y_{ij}) is replaced by the probability to agree to the offer. An agreement converts the offered terms (s_{ij}, y_{ij}) into a binding contract immediately.

2. Bargaining is restricted to connected firms (i, j) that have not yet reached an agreement in prior stages.

The game is otherwise identical to the original game defined in Section 3.2. In stage t, one of the firms, denoted $I(t)$, is selected at random, uniformly among the firms who have not yet reached agreement with one or more edge counterparties, to make the first offers of terms. This selection of the initial proposer is observable only to firm $I(t)$ until the beginning of the next stage. Treating the graph G as a directed tree with root $I(t)$, firm $I(t)$ makes offers to each of its child firms. Once each of these child firms j receives an offer (s_{ij}, y_{ij}) from its parent i, firm j chooses a probability a_j for agreeing to its parent's offer and chooses the offers $(s_{jk}, y_{jk})_{k\in c(j)}$ to its own child firms, and so on, until each firm other than the root firm $I(t)$ has received an offer from its parent firm and has chosen the probability with which it agrees to this offer. These choices, a_j and $(s_{jk}, y_{jk})_{k\in c(j)}$, must be based only on (be measurable with respect to) the information contemporaneously available to firm j, which is the join $\tilde{\mathcal{F}}_{jt}$ of the information $\tilde{\mathcal{F}}_t$ commonly available at the beginning of stage t and (for a non-root node) the information contained in its parent's offer to firm j.

The outcomes of all agreement events are then realized. The offers that receive an agreement are immediately incorporated into each firm's terms, which update from $\left(s_i^{t-1}, y_i^{t-1}\right)$ to (s_i^t, y_i^t) for firm $i \in N$. Firm $i \in N$ receives the value $v_i\left(s_i^t, y_i^t\right) - v_i\left(s_i^{t-1}, y_i^{t-1}\right)$ in stage t, which is discounted by δ^t to account for its time value. The *status-quo* (no-trade) value of firm i is $v_i\left(s_i^0, y_i^0\right) = u_i(q_i^0) - u_i(q_i^0) = 0$. The full history of play through stage t becomes observable to all firms at the beginning of stage $t+1$. The common information $\tilde{\mathcal{F}}_{t+1}$ available to all firms at the beginning of stage $t+1$ is then the join of $\tilde{\mathcal{F}}_t$, the information revealed by $I(t)$ and all of the stage-t offers and agreement events. The same tie-breaking rule that was defined in the protocol for the original game of Section 3.2 is also applied to this game with no greenlight protocol. This game, denoted $\Gamma^{\text{no greenlight}}(G, S, f, \delta)$, solely replaces "greenlighting an offer" by "agreeing to an offer" but otherwise preserves all elements of $\Gamma(G, S, f, \delta)$, including the sequence of events and the information structure.

We first show an example of how, without the greenlight protocol, holdup impedes efficient trades.

The trading network consists of $n = 3$ firms that are connected in a chain, as shown in Figure 2. Initially, firm 1 owns 1 unit of the asset, while firms 2 and 3 own none of the asset. Firm 1 is a prospective seller because it is the unique asset owner and does not have the highest valuation, in that $u_1(q) = \$99.7q$, $u_2(q) = \$99.2q$, and $u_3(q) = \$100q$. The three firms have the same one-stage discount factor $\delta < 1$. The above valuations imply that the unique socially efficient allocation is obtained by having firm 1 sell its asset to firm 2 at stage zero, and for firm 2 to sell the asset to firm 3 at stage zero.

Proposition 2. *There exists some $\underline{\delta} \in (0, 1)$ such that for every discount factor $\delta \in (\underline{\delta}, 1)$, in any PBE of the example three-firm trading game $\Gamma^{\text{no greenlight}}(G, S, f, \delta)$, there is no trade.*

The result is proved as follows. In a given PBE (not necessarily with Markov strategies), if a trade is first completed between firms 1 and 2, the *expected* price at which firm 2 sells the asset to firm 3, conditional on them reaching a trade agreement, is the average of the two firms' valuations, $(\$99.2 + \$100)/2 = \$99.6$, which is lower than the valuation of firm 1, \$99.7. As a result, no gains from trade exist between firms 1 and 2. Therefore, no trade

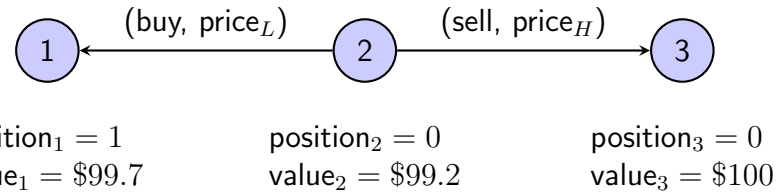


Figure 2: An over-the-counter asset trading market with firm 2 as intermediary. Initially, firm 1 has one unit of the asset, while firms 2 and 3 have none. There is no short selling. The firms' respective asset valuations are $u_1(q) = \$99.7q$, $u_2(q) = \$99.2q$, and $u_3(q) = \$100q$. The greenlight mechanism achieves the unique socially efficient allocation, in which firm 1 sells its asset to firm 2 at stage zero, and firm 2 sells the asset to firm 3 at stage zero. Without the greenlight mechanism, there is no trade in equilibrium. The schematic shows a stage in which the intermediary offers terms to firms 1 and 3.

occurs between either pair of firms.

The efficient allocation arises only if firm 2 intermediates the asset from firm 1 to firm 3. However, because firm 2 has the lowest of the three firms' valuations, firm 2 has a low *status-quo* value when trying to sell its purchased asset to firm 3. Firm 2 is thus held up by firm 3 whenever the two bargain, leading to a low resale price of \$99.6. Anticipating the low resale price, firm 2 finds it unprofitable to buy the asset from firm 1. The same form of holdup occurs if, instead, a trade is first completed between firms 2 and 3.

If trades are simultaneously completed between both pairs of firms, in that firm 2 is selected to make offers to firms 1 and 3, who both agree to their respective offers, then firm 3 could deviate to reject the offer of firm 2. The deviation is profitable because (a) time discounting can be arbitrarily small, and (b) firm 3 can then hold up firm 2 in subsequent bargaining, knowing that firm 2 has completed a trade with firm 1.

That firm 3 can demand a price of \$99.6 that is even below the valuation of firm 1 precisely corresponds to the failure of the "feasibility" assumption on firm payoffs in Collard-Wexler et al. (2019), which requires that "each firm would prefer maintaining all of its agreements . . . to dropping any set of its agreements (holding all other agreements fixed)."

5.2 Efficient trade using greenlights

The greenlight protocol eliminates holdup, and in particular implements the socially efficient allocation q^* in all PBE with Markov strategies, as stated by the following immediate corollary of Theorem 1.

Corollary 1. *For any primitives (G, S, f, δ) of a network asset trading game, all PBE with Markov strategies of $\Gamma(G, S, f, \delta)$ implement the unique socially efficient allocation q^* in stage 0 with probability 1. In all PBE with Markov strategies, all firms have the same expected gain in value $F(q^*)/n$.*

When applied to the illustrative three-firm chain network depicted in Figure 2, this result implies that in stage zero firm 2 buys the asset from firm 1 at an expected price of \$99.8 and resells it to firm 3 at an expected price of \$99.9. Both trades occur with probability 1 in all PBE with Markov strategies. The three firms share equally in the trade surplus, each earning an expected profit of $(\$100 - \$99.7)/3 = \$0.1$. As δ approaches 1, the realized prices are arbitrarily close to these expected prices (Proposition 1).

In general, it is possible that some firms in a trading network play no role as potential buyers or sellers. In the stated equilibrium, these "irrelevant" firms nevertheless earn the same relative payoff as active trading firms because the greenlight mechanism gives any firm in the network a veto power over the trades. Naturally, buyers, sellers, and intermediaries can achieve a Pareto improvement by excluding irrelevant firms in a pre-trade stage of the network formation, or by changing the trade protocol so that the greenlight decisions of these irrelevant firms are ignored.

It is also possible that some firms contribute little to social efficiency, yet earn the same share of the total social surplus as firms whose actions may contribute much more. Consider the three firms depicted in Figure 2, with asset valuations $u_1(q) = \$99.7q$, $u_2(q) = \$98.8q$, and $u_3(q) = \$100q$, but with the alternative initial asset holdings $q_1^0 = 1$, $q_2^0 = 1$, and $q_3^0 = 0$. Under the greenlight protocol, Corollary 1 implies that in stage zero firm 2 buys one unit of the asset from firm 1 at an expected price of \$100.2 and sells two units of the asset to firm 3 at an expected price of \$99.75. Both trades occur with probability 1 in all PBE

with Markov strategies. The three firms share equally in the trade surplus in any PBE with Markov strategies, each earning an expected profit of $(\$100 \times 2 - \$99.7 - \$98.8)/3 = \$0.5$. However, firms 2 and 3 would be strictly better off if they were to trade only one unit with each other, excluding firm 1, earning $(\$100 - \$98.8)/2 = \$0.6$ each. That firms 2 and 3 may want to exclude firm 1 does not dissuade them from using the greenlight protocol. After firms 2 and 3 have traded one unit with each other, they would then want to include firm 1 and bilaterally bargain in the chain network using the greenlight protocol. Following these two steps, firms 2 and 3 would each earn an expected profit of \$0.6 from trading one unit with each other (excluding firm 1), and then the three firms would each earn an expected profit of \$0.1 by trading the second unit with atomic settlement.

In summary, the greenlight protocol's equal-gain property does not erode its attractiveness or ability to reach efficiency. On the contrary, a network of firms would always want to use the greenlight protocol to split any "money left on the table" until they reach efficiency. This argument applies not only to asset trading markets, but to contracting with general primitives (G, S, f).

6 Axiomatic solution: Coordination-Proofness

We now propose a new axiomatic solution concept, Coordination-Proof bargaining, which refines the set of "Nash-in-Nash" outcomes of Horn and Wolinsky (1988) to a unique bargaining outcome when firms' no-agreement values are given by their respective *status-quo* values.[17] The Coordination-Proof bargaining is outcome-identical to the contractual terms implemented in any PBE with Markov strategies of the extensive-form bargaining game $\Gamma(G, S, f, \delta)$, in the limit as the discount factor δ goes to 1.

A network bargaining problem involves multiple bilateral bargaining problems. A firm can leverage information learned from one negotiation when negotiating another, and can coordinate actions across different concurrent negotiations. The "Nash-in-Nash" bargaining

[17] Horn and Wolinsky (1988, p.412) do not insist on a special specification for no-agreement values, and note that different specifications are plausible depending on context.

set of Horn and Wolinsky (1988) implicitly avoids consideration of such coordination. A bargaining outcome is said to be Nash-in-Nash if each pair of connected firms reaches the Nash bargaining solution, defined in Nash (1950a) for two agents, by taking as given the bargaining outcomes of all other directly connected firms. Here, the first "Nash" refers to Nash Bargaining (Nash, 1950a), while the second "Nash" refers to Nash equilibrium in strategic games (Nash, 1950b, 1951). Therefore, "Nash-in-Nash" applies the idea of Nash equilibrium to network bilateral bargaining problems, with every bilateral bargaining process being conducted separately and simultaneously. As a result, a Nash-in-Nash outcome can be implemented with "delegated-agent" extensive-form games in which a firm delegates each of its bilateral bargaining processes to a separate agent. Each bilateral bargaining process is conducted between two delegated agents who use alternating offers, as in Binmore et al. (1986), without being able to coordinate with, or even observe, the bargaining outcomes of other bargaining processes for the same firm.

To visualize the potential importance of this form of coordination, we consider a variant of the three-firm asset trading network depicted in Figure 2. The setting is modified by supposing that firm 1 initially owns $q_1^0 = 2$ units of the asset rather than 1 unit. The initial asset holdings of firms 2 and 3 remain 0. As with the greenlight protocol, we assume atomic settlement, so that a failure to reach an agreement in either of the bilateral bargaining processes results in each of the three firms receiving its *status-quo* (no-trade) value, $u_i(q_i^0)$.

This bargaining problem has two Nash-in-Nash outcomes. One of these is socially efficient, having firm 1 sell 2 units of the asset at the price \$99.8 to firm 2, who sells 2 units at the price \$99.9 to firm 3. The three firms each earn a profit of $2 \times (\$100 - \$99.7)/3 = \$0.2$. The other Nash-in-Nash outcome is not socially efficient, having firm 1 sell only 1 unit of the asset at price \$99.8 to firm 2, who sells 1 unit at the price \$99.9 to firm 3. In this outcome, each of the three firms earns a profit of $(\$100 - \$99.7)/3 = \$0.1$. The efficient outcome can be achieved by having the middle firm coordinate both bargaining processes. The inefficient outcome is not proof to coordination, because the efficient outcome is both achievable through coordination and Pareto superior to the inefficient one.

Motivated by this example, we define a new axiomatic solution concept, Coordination-

Proof bargaining, for a general network bilateral bargaining problem (G, S, f) of Subsection 3.1, where firms' *status-quo* values are normalized to zero.

Definition 1. Coordination-Proof bargaining

(i) Suppose (G, S, f) has two firms, connected. That is, $|N| = 2$, $|E| = 1$. We say that (s, y) is a Coordination-Proof bargaining solution if and only if it is prescribed by the Nash bargaining solution, in that

$$s = \underset{s' \in S}{\operatorname{argmax}} \left[f_1(s') + f_2(s') \right]$$
$$y_{12} = \frac{1}{2} \left[f_1(s) - f_2(s) \right].$$

(ii) Working inductively, we fix an integer $e > 1$ and assume that a Coordination-Proof bargaining solution has been defined for any (G, S, f) with fewer than e edges. Then, for any (G, S, f) with $|E| = e$ edges:

(a) (s, y) is defined to be *Coordination-Achievable* if, for any sub-tree $G' = (N', E')$ of G, $(s_{|E'}, y_{|E'})$ is a Coordination-Proof bargaining solution in the bilateral bargaining problem (G', S', f'), where

$$S'_{ij} = S_{ij}, \quad (i,j) \in E',$$
$$f'_i(s'_i) = f_i\left(s'_i, (s_{ij})_{(i,j) \in E \setminus E'}\right) - \sum_{(i,j) \in E \setminus E'} y_{ij}, \qquad s'_i \in S'_i.$$

(b) $(s^{\mathrm{CP}}, y^{\mathrm{CP}})$ is defined to be a *Coordination-Proof bargaining solution* if it is Coordination-Achievable and if there does not exist another Coordination-Achievable (s, y) such that $v_i(s_i, y_i) > v_i(s_i^{\mathrm{CP}}, y_i^{\mathrm{CP}})$ for all i in N.

Thus, a bargaining outcome is Coordination-Proof if it is Coordination-Achievable and not strictly Pareto dominated by any other Coordination-Achievable bargaining outcome, where coordination achievability requires that no sub-tree of firms can benefit by deviating to a Coordination-Achievable bargaining outcome.

The notion of Coordination-Proof bargaining is inspired by and adapted from the equilibrium concept of Coalition-Proof Nash equilibrium of Bernheim, Peleg, and Whinston (1987) for noncooperative games. Underlying the noncooperative games considered in Bernheim et al. (1987) is the assumption that "players can freely discuss their strategies, but cannot make binding commitments." Non-binding communication among a coalition of players allows the coalition members to coordinate their choices in a self-enforceable manner. In our setting, Coordination-Achievability mirrors this notion of self-enforceability, while Coordination-Proofness mirrors Coalition-Proofness.

Theorem 2. *For any (G, S, f), there is a unique Coordination-Proof bargaining solution. The associated outcome $(s^{\mathrm{CP}}, y^{\mathrm{CP}})$ consists of the socially efficient profile $s^{\mathrm{CP}} = s^*$ of non-cash terms and payments y^{CP} achieving equal sharing of gains, in that,*

$$v_i(s_i^*, y_i^{\mathrm{CP}}) = \frac{F(s^*)}{n}, \qquad i \in N.$$

A proof is found in Appendix A.

7 Conclusion

For bilateral bargaining in network markets, we have shown that holdup problems can be eliminated by a novel "greenlight" bargaining protocol. This protocol can be implemented by smart contracts or by assuming the existence of a trusted neutral authority, such as a notary, that can provide contractible information about agreements in the network of firms.

Applications include over-the-counter trading networks, syndicated lending, multi-tranche securitizations, purchases requiring third-party financing, "building a book" for a security offering, and other multi-agent market settings in which pairs of agents bargain bilaterally over terms that could potentially affect their no-agreement values in negotiations with other firms. With the greenlight protocol, the terms of a contract proposed to any firm can be given a "greenlight" by that firm, which automatically converts those terms to a binding contract if and only if the contract terms proposed to all other directly or indirectly con-

nected firms are also given greenlights and are thus likewise converted to contracts. This all-or-none automatic conversion from greenlit contingent commitments to final contracts is known as "atomic." Atomic settlement is already a common smart-contract approach used in decentralized finance for the execution of agreed trades and is also emerging in transactions involving conventional financial assets and investors (World Economic Forum, 2025). Atomic settlement has been applied, for example, to foreign exchange trades and Treasury repos.

Our main result is that for any tree network based on strategic bilateral bargaining with the greenlight protocol, the unique outcome of Perfect Bayesian Equilibrium with Markov strategies is immediate socially efficient agreements, with equal expected gains for every firm. Equilibrium is also shown to exist. As time discounting becomes negligible, the outcomes for the cash terms converge to their ex-ante expectations.

As an illustration, we analyze the efficiency benefit of using the greenlight bargaining protocol in over-the-counter asset trading networks. Without the greenlight protocol, the holdup problem can impede efficient trade even in a simple example of three firms in a chain. Under mild technical conditions, the unique outcome of any Perfect Bayesian Equilibrium with Markov strategies is immediate agreement on socially efficient contracts that generate equal gains from trade for every firm.

In our general setting, we also show that the expected equilibrium outcome of the greenlight bargaining protocol is equivalent to the outcome of an axiomatic solution, Coordination-Proof bargaining. This is also the limit equilibrium outcome of the bargaining protocol in each state of the world (almost surely), as the cost of time delay goes to zero.

Appendices

A Proofs

To prove Theorem 1 and Proposition 1, we first establish the properties of a closely related game. We augment the game primitives (G, S, f) to $(G, S, f, \hat{v})$, where $\hat{v} = (\hat{v}_i)_{i \in N}$ are exogenous constants that determine the continuation values for each firm in the case of no agreement.

We consider a bargaining game with the primitives $(G, S, f, \hat{v})$, but having only one stage. The stage proceeds in the exact same way as in the original game setting of Section 3.2, but for one difference: If one or more offers do not receive a greenlight, then each firm j obtains the payoff $\delta \hat{v}_j$, which can be thought of as the discounted value of receiving $\hat{v}_j$ in one period. The game then ends. The game is otherwise identical to the original game defined in Section 3.2. That is, one of the firms, denoted i, is selected at random, uniformly among the firms, to make the first offers of terms. This selection of the initial proposer is observable only to firm i. Treating graph G as a directed tree with root i, firm i makes offers to each of its child firms. Each of these child firms then chooses a probability for greenlighting its parent's offer and makes offers to its own child firms, and so on, until each firm other than the root firm i has received an offer from its parent firm and has chosen a probability with which to greenlight this offer. The offers and greenlight probability by any non-root firm j must be based only on the information contained in the offer that j receives from its parent firm. If all offers receive a greenlight, then the terms of all offers automatically become binding agreements and the game ends, with any firm $j \in G$ receiving the value $v_j(s_j, y_j)$ defined by (1). The same tie-breaking rule that was defined in the protocol for the original game of Section 3.2 is also applied to this auxiliary game, denoted $\Gamma(G, S, f, \hat{v})$.

Lemma 1. *All PBE of game $\Gamma(G, S, f, \hat{v})$ are payoff equivalent. In the event that firm i is*

selected to make the first offers of terms, the payoff of firm i is

$$\max\left\{\max_{s} F(s) - \delta \sum_{j\in N\setminus\{i\}} \hat{v}_j \ , \ \delta\hat{v}_i \right\}.$$

Otherwise, the payoff of firm i is $\delta\hat{v}_i$. The expected payoff of firm i is

$$\frac{1}{n}\left[\max\left\{\max_{s} F(s) - \delta \sum_{j\in N\setminus\{i\}} \hat{v}_j \ , \ \delta\hat{v}_i \right\} + \delta(n-1)\hat{v}_i\right]. \tag{A.1}$$

Proof. We fix a PBE (σ, μ) of $\Gamma(G, S, f, \hat{v})$ throughout this proof, where σ is the equilibrium strategy profile and μ is the belief system. For any pair $(i, j) \in E$ of connected firms, treating graph G as a directed tree with root i, we let $G_{ij} = (N_{ij}, E_{ij})$ be the branch of the directed tree G that has root j, and $c(k)$ be the set of "child" firms of firm $k \in N$. We show that

Claim 1. For any pair $(i, j) \in E$ of connected firms, if firm j receives from firm i an offer (s_{ij}, y_{ij}) that satisfies

$$\max_{(s_{k\ell})_{k\ell\in E_{ij}}} F_{ij}\left(s_{ij}, (s_{k\ell})_{k\ell\in E_{ij}}\right) + y_{ij} - \delta \sum_{k\in N_{ij}\setminus\{j\}} \hat{v}_k \geq \delta\hat{v}_j, \tag{A.2}$$

where

$$F_{ij}\left(s_{ij}, (s_{k\ell})_{k\ell\in E_{ij}}\right) = f_j(s_{ij}, (s_{jk})_{k\in c(j)}) + \sum_{k\in N_{ij}\setminus\{j\}} f_k(s_k),$$

then firm j will greenlight the offer with probability 1.

Our proof of this claim is based on induction over the number h of generations of the tree G_{ij} (that is, the length of the longest downstream path starting from root node j). When $h = 0$, firm j has no child firms in G_{ij}. Upon receiving an offer (s_{ij}, y_{ij}) that satisfies (A.2), regardless of the belief of j, it is weakly optimal for j to greenlight the offer. The tie-breaking rule implies that firm j greenlights the offer with probability 1. Suppose that the claim holds with h generations of firms. If G_{ij} has $h + 1$ generations, G_{jk} has h generations for every

child firm $k \in c(j)$ of j. We let $\left(s^*_{jk}(s_{ij}), y^*_{jk}(s_{ij})\right)_{k\in c(j)}$ be such that

$$\begin{aligned}\left(s^*_{jk}(s_{ij})\right)_{k\in c(j)} &\in \operatorname*{argmax}_{(s_{jk})_{k\in c(j)}} f_j(s_j) + \sum_{k\in c(j)} \max_{(s_{\ell m})_{\ell m\in E_{jk}}} F_{jk}\left(s_{jk}, (s_{\ell m})_{\ell m\in E_{jk}}\right), \\ y^*_{jk}(s_{ij}) &= \delta \sum_{\ell\in N_{jk}} \hat{v}_\ell - \max_{(s_{\ell m})_{\ell m\in E_{jk}}} F_{jk}\left(s^*_{jk}(s_{ij}), (s_{\ell m})_{\ell m\in E_{jk}}\right).\end{aligned} \tag{A.3}$$

Then for every child firm $k \in c(j)$ of j, the offer $\left(s^*_{jk}(s_{ij}), y^*_{jk}(s_{ij})\right)$ from its parent firm j satisfies

$$\max_{(s_{\ell m})_{\ell m\in E_{jk}}} F_{jk}\left(s^*_{jk}(s_{ij}), (s_{\ell m})_{\ell m\in E_{jk}}\right) + y^*_{jk}(s_{ij}) - \delta \sum_{\ell\in N_{jk}\setminus\{k\}} \hat{v}_\ell = \delta\hat{v}_k.$$

The induction hypothesis implies that firm k would greenlight the offer $\left(s^*_{jk}(s_{ij}), y^*_{jk}(s_{ij})\right)$ from its parent firm j. Then for firm j, greenlighting an offer (s_{ij}, y_{ij}) that satisfies (A.2) while offering $\left(s^*_{jk}(s_{ij}), y^*_{jk}(s_{ij})\right)_{k\in c(j)}$ to its child firms $k \in c(j)$ would yield a payoff of at least

$$\begin{aligned}& f_j\left(s_{ij}, \left(s^*_{jk}(s_{ij})\right)_{k\in c(j)}\right) + y_{ij} - \sum_{k\in c(j)} y^*_{jk}(s_{ij}) \\ =\ & \max_{(s_{k\ell})_{k\ell\in E_{ij}}} F_{ij}(s_{ij}, (s_{k\ell})_{k\ell\in E_{ij}}) - \delta \sum_{k\in N_{ij}\setminus\{j\}} \hat{v}_k + y_{ij} \\ \geq\ & \delta\hat{v}_j.\end{aligned}$$

It is thus weakly optimal for firm j to greenlight the offer regardless of the belief of j. The tie-breaking rule implies that firm j greenlights the offer with probability 1. Claim 1 therefore follows by induction.

We now return to the proof of Lemma 1. If firm i is selected to propose the first offers, the highest possible payoff that it could receive is

$$\max\left\{\max_s F(s) - \delta \sum_{j\in N\setminus\{i\}} \hat{v}_j\ ,\ \delta\hat{v}_i\right\}. \tag{A.4}$$

If

$$\max_s F(s) - \delta \sum_{j\in N\setminus\{i\}} \hat{v}_j < \delta\hat{v}_i,$$

then the highest possible payoff (A.4) is equal to $\delta\hat{v}_i$, which is also the lowest possible payoff

of firm i. Thus, the payoff of firm i must be equal to (A.4).

If

$$\max_s F(s) - \delta \sum_{j \in N \setminus \{i\}} \hat{v}_j \geq \delta \hat{v}_i,$$

then Claim 1 and the tie-breaking rule imply that firm i can actually attain its highest possible payoff (A.4) by offering (s_i^*, y_i^*), where

$$y_{ij}^* = \delta \sum_{k \in N_{ij}} \hat{v}_k - \max_{(s_{k\ell})_{k\ell \in E_{ij}}} F_{ij}\left(s_{ij}^*, (s_{k\ell})_{k\ell \in E_{ij}}\right).$$

Precisely, for every other firm $j \in N \setminus \{i\}$, there exists some offer $(s_{jk}, y_{jk})_{k \in c(j)}$ that satisfies condition (A.2) and would receive greenlights with probability 1 (Claim 1). Then, the tie-breaking rule implies that all subsequent offers from every other firm $j \in N \setminus \{i\}$ will receive greenlights with probability 1. Offer (s_{ij}^*, y_{ij}^*) also satisfies condition (A.2) and will receive a greenlight with probability 1 (Claim 1). Therefore, (s_i^*, y_i^*) and all subsequent offers will receive greenlights with probability 1, allowing firm i to attain its highest possible payoff (A.4).

Taken together, the payoff of firm i must be equal to (A.4) if $i = I(t)$.

If another firm $j \in N \setminus \{i\}$ is selected to propose the first offers, then the payoff of firm j is

$$\max \left\{ \max_s F(s) - \delta \sum_{k \in N \setminus \{j\}} \hat{v}_k \ , \ \delta \hat{v}_j \right\}.$$

Since the lowest possible payoff of any other firm $k \in N \setminus \{i, j\}$ is $\delta \hat{v}_k$, the highest possible payoff of firm i is at most $\delta \hat{v}_i$, which is also the lowest possible payoff of i. Hence, the payoff of firm i must be equal to $\delta \hat{v}_i$ if $i \neq I(t)$.

Therefore, the expected payoff of firm i must be given by (A.1). Lemma 1 follows. □

Proof of Theorem 1. We extend the proof of Binmore et al. (1986).

Step 1 (equilibrium payoff): A subgame that starts from the beginning of a given stage t of the game is identical to the original game $\Gamma(G, S, f, \delta)$. Each such subgame starts with the random selection of $I(t)$, as in the original game $\Gamma(G, S, f, \delta)$. Under a profile of Markov

strategies, each firm's expected payoff $\hat{v}_{i,t}$ in every subgame that starts from the beginning of stage t is identical, regardless of the history of play leading up to that subgame. Thus, if agreement is not reached in stage $t-1$, the continuation value of firm i is common knowledge and equal to $\hat{v}_{i,t}$, regardless of firms' (possibly differing) beliefs. The subgame that starts from the beginning of stage $t-1$ then reduces to the one-stage game $\Gamma(G, S, f, \hat{v}_{i,t})$.

We let $\bar{v}_i$ and $\underline{v}_i$ be the highest and the lowest expected payoff of firm i in any PBE with Markov strategies ($i \in N$). Lemma 1 implies that the highest attainable expected payoff of firm i is

$$\frac{1}{n}\left[\max\left\{\max_s F(s) - \delta \sum_{j\in N\setminus\{i\}} \underline{v}_j \ , \ \delta\bar{v}_i \right\} + \delta(n-1)\bar{v}_i\right].$$

Therefore, for every i,

$$\bar{v}_i = \frac{1}{n}\left[\max\left\{ F(s^*) - \delta \sum_{j\in N\setminus\{i\}} \underline{v}_j \ , \ \delta\bar{v}_i \right\} + (n-1)\delta\bar{v}_i\right].$$

Likewise, for every i,

$$\underline{v}_i = \frac{1}{n}\left[\max\left\{ F(s^*) - \delta \sum_{j\in N\setminus\{i\}} \bar{v}_j \ , \ \delta\underline{v}_i \right\} + (n-1)\delta\underline{v}_i\right].$$

This constitutes a system of equations for $(\bar{v}_i, \underline{v}_i)_{i\in N}$ that can be expressed in the form $(\bar{v}_i, \underline{v}_i)_{i\in N} = \mathcal{M}\left((\bar{v}_i, \underline{v}_i)_{i\in N}\right)$. One solution to this system of equations is given by

$$\bar{v}_i = \underline{v}_i = \frac{1}{n} F(s^*).$$

This is the unique solution because $\mathcal{M}$ is a contraction map using the L_1 metric.[18]

Step 2 (equilibrium non-cash terms). Since the unique equilibrium payoffs of all firms add up to $F(s^*)$, it must be that the socially efficient profile s^* of non-cash terms is immediately implemented in stage 0 with probability 1.

Step 3 (constructing a PBE with Markov strategies): We construct the following PBE

[18] For any $x = (x_1, \ldots, x_m)$ and $y = (y_1, \ldots, y_m)$, the L_1 distance between x and y is $\sum_k |x_k - y_k|$.

with Markov strategies.

In each stage t, if firm j is selected to propose the first offer, $j = I(t)$, then j offers (s_j^*, y_j^*) to its child firms, where

$$y_{jk}^* = \delta \frac{|N_{jk}|}{n} F(s^*) - \max_{(s_{\ell m})_{\ell m \in E_{jk}}} F_{jk}\left(s_{jk}^*, (s_{\ell m})_{\ell m \in E_{jk}}\right).$$

If $j \neq I(t)$, then upon receiving an offer (s_{ij}, y_{ij}) from its parent firm i, j offers to its child firms $\left(s_{jk}^*(s_{ij}), y_{jk}^*(s_{ij})\right)_{k \in c(j)}$, as defined by (A.3).

Any non-root firm j greenlights any offer (s_{ij}, y_{ij}) from its parent firm i that satisfies

$$\max_{(s_{k\ell})_{k\ell \in E_{ij}}} F_{ij}\left(s_{ij}, (s_{k\ell})_{k\ell \in E_{ij}}\right) + y_{ij} \geq \delta \frac{|N_{ij}|}{n} F(s^*).$$

This strategy profile is Markovian and satisfies the tie-breaking rule. One can also verify that it is sequential with respect to any consistent belief system. Therefore, coupling the above strategy profile with any consistent belief system constitutes a PBE with Markov strategies. □

Proof of Proposition 1. Lemma 1 and Theorem 1 imply that if i is selected to make the first offers of terms, the payoff of firm i is

$$v_i(s_i^*, y(k)_i) = F(s^*)\left(1 - \frac{\delta_k(n-1)}{n}\right).$$

Otherwise,

$$v_i(s_i^*, y(k)_i) = \delta_k \frac{F(s^*)}{n}.$$

Therefore, $v_i(s_i^*, y(k)_i) \to F(s^*)/n$ almost surely. □

Proof of Proposition 2. The proof immediately follows the statement of Proposition 2. □

Proof of Theorem 2. We fix the primitives (G, S, f).

Step 1: Any Coordination-Achievable bargaining outcome (s, y) has equal sharing of gains, in that

$$v_i(s_i, y_i) = \frac{F(s)}{n}, \qquad i \in N.$$

To show this, we fix a Coordination-Achievable bargaining outcome (s, y) throughout Step 1. Given any pair $(i, j) \in E$ of connected firms, we let $N' = \{i, j\}$, $E' = \{(i, j)\}$, and $G' = (N', E')$. We consider the two-firm bilateral bargaining problem (G', S', f') between i and j, with

$$\begin{aligned} S'_{ij} &= S_{ij}, \\ f'_i(s'_{ij}) &= f_i\left(s'_{ij}, (s_{ik})_{(i,k)\in E\setminus E'}\right) - \sum_{(i,k)\in E\setminus E'} y_{ik}, \qquad s'_{ij} \in S'_{ij}, \\ f'_j(s'_{ij}) &= f_j\left(s'_{ij}, (s_{jk})_{(j,k)\in E\setminus E'}\right) - \sum_{(j,k)\in E\setminus E'} y_{jk}, \qquad s'_{ij} \in S'_{ij}. \end{aligned}$$

The definition of Coordination-Achievability implies that the terms (s_{ij}, y_{ij}) are prescribed by the Nash Bargaining solution, which has equal sharing of gains for i and j. Thus,

$$v_i(s_i, y_i) = v_j(s_j, y_j).$$

Since the above identity holds for any pair (i, j) of connected firms, it must be that the Coordination-Achievable outcome (s, y) has the same value $F(s)/n$ for every firm.

Step 2: The bargaining outcome (s^*, y^{CP}), defined in Theorem 2, is Coordination-Achievable and strictly Pareto dominates all other Coordination-Achievable bargaining outcomes.

One can verify by definition that $(s^{\mathrm{CP}}, y^{\mathrm{CP}})$ is Coordination-Achievable. Given any Coordination-Achievable bargaining outcome (s, y) that is different from $(s^{\mathrm{CP}}, y^{\mathrm{CP}})$, it must be that $s \neq s^{\mathrm{CP}}$, because otherwise equal sharing of gains uniquely pins down the payment y^{CP} and then $(s, y) = (s^{\mathrm{CP}}, y^{\mathrm{CP}})$. Thus, for each firm i in N,

$$v_i(s_i, y_i) = \frac{F(s)}{n} < \frac{F(s^*)}{n} = \frac{F(s^{\mathrm{CP}})}{n} = v_i(s^*_i, y^{\mathrm{CP}}_i).$$

It follows that $(s^{\mathrm{CP}}, y^{\mathrm{CP}})$ is the unique Coordination-Proof solution of (G, S, f).

□

B Less information does not undermine efficiency

We discussed the potential concern that the assumed information structure of our bargaining protocol allows each firm at each stage to observe all past actions of all other firms, including their offers and greenlight decisions, as well as which other firms were designated to make the first offers in past stages. This large amount of information may be viewed as unrealistic. This appendix shows that the efficient equilibrium outcome of Theorem 1 remains intact even if, for some arbitrarily large fixed integer $T > 0$, at the beginning of every stage $t \leq T$, firm i does not have any of this information. As in the original version of the game, equilibrium agreement to the efficient terms is achieved immediately at stage 0.

For each firm i, we define a more parsimonious information structure given by sub-σ-algebras $\{\mathcal{F}^*_{it} : t \geq 0\}$ that are defined inductively as follows. We begin by letting $\mathcal{F}^*_{i0}$ be generated by (i) the indicator $\mathbb{1}_{i=I(0)}$ of the event that firm i is selected to be the stage-0 root firm, that is, the initial proposer in stage 0 and (ii) the offer made to firm i in stage 0 if firm i is not the initial proposer. In stage t, the information $\mathcal{F}^*_{it}$ available to firm i when it makes its offer and greenlight decision is the join of the information $\mathcal{F}^*_{i,t-1}$ available to firm i in the previous stage, the information generated by the indicator of the event that agreement was not reached at stage $t-1$, the indicator of whether firm i was selected to make the first offers in stage t, $\mathbb{1}_{i=I(t)}$, and (for a non-root node) the information contained in its parent's offer in stage t.

This information $\mathcal{F}^*_{it}$ seems more natural, in most settings, than the larger information set $\mathcal{F}_{it}$ assumed in Theorem 1. For the following result, we assume that, for an arbitrary fixed integer $T > 0$, firm i makes an offer and greenlight decision at stage t that are based only on (are measurable with respect to) the information in $\mathcal{F}^*_{it}$ for $t \leq T$, and based on $\mathcal{F}_{it}$ for $t > T$. We denote the corresponding bargaining game by $\Gamma(T)$.

Proposition 3. *For any fixed integer $T > 0$, all PBE with Markov strategies of $\Gamma(T)$ implement the unique socially efficient profile s^* of non-cash terms in stage 0 with probability 1. All PBE with Markov strategies have the same value $F(s^*)/n$ for each firm.*

Proof of Proposition 3. We fix a PBE with Markov strategies of $\Gamma(T)$. The subgame of $\Gamma(T)$

that starts from the beginning of stage $T+1$ is identical to the original game $\Gamma(G, S, f, \delta)$, whose equilibria are characterized by Theorem 1. In the remainder of the proof, when computing current present values as of a given stage t, we normalize the equilibrium payoffs by the discount factor δ^t, so as to get the discounted value from the perspective of stage t. Theorem 1 implies that every firm has an expected current value of $F(s^*)/n$ at stage $T+1$, conditional on arriving at stage $T+1$, regardless of the history of play leading up to stage $T+1$. Thus, if agreement is not reached in stage T, every firm's continuation value is common knowledge and equal to $F(s^*)/n$, regardless of firms' (possibly differing) beliefs. The game that starts from the beginning of stage T then reduces to the one-stage game $\Gamma(G, S, f, F(s^*)/n)$. Lemma 1 implies that in stage T, every firm has an expected current value of $F(s^*)/n$ conditional on arriving at stage T. Recursively, every firm has an expected current value at stage t of $F(s^*)/n$ conditional on arriving at stage t for each t in $\{T, T-1, \ldots, 0\}$. Therefore, every firm has an initial equilibrium value of $F(s^*)/n$. Since the unique equilibrium payoffs of all firms add up to $F(s^*)$, it must be that the socially efficient s^* is immediately implemented in stage 0 with probability 1.

□

C Efficiency is robust to endogenous breakdown values

Section 3 uses time discounting to capture the cost of delayed agreement. As a result, firms' *status-quo* values determine their bargaining positions (Rubinstein, 1982, Binmore et al., 1986). An alternative specification of the delay cost is a probabilistic breakdown of bargaining, in which case the firms' breakdown values determine their bargaining positions (Binmore et al., 1986). Upon breakdown, it is possible that a firm would reach bargains with a subset of its edge counterparties, giving the firm an endogenous continuation value. Is our efficiency result robust to allowing for endogenous breakdown values? This appendix shows that the answer is "Yes."

Given some model primitives (G, S, f), we augment the network bilateral bargaining game $\Gamma(G, S, f, \delta)$ with the greenlight protocol of Section 3 as follows: There is no time

discounting. If the bargaining protocol has not ended by stage $t \geq 0$, it proceeds to stage $t+1$ with probability $\delta \in (0,1)$, and breaks down with probability $1-\delta$. For example, a smart contract could randomly expire. The occurrence of a breakdown is common knowledge, so the common information $\mathcal{F}^b_{t+1}$ available to all firms at the beginning of stage $t+1$ is now the join of $\mathcal{F}^b_t$, the information revealed by $I(t)$, all of the stage-t offers and greenlight events, and the occurrence of a breakdown.

Instead of terminating the bargaining game and assigning firms with exogenously fixed values upon a breakdown, we assume that a continuation game $\Gamma(\mathscr{B})$ is entered upon a breakdown. Here, $\mathscr{B}$ denotes the full history of play leading up to the breakdown. Defining a breakdown continuation game is possible because the history $\mathscr{B}$ is common knowledge at the onset of a breakdown. In general, the continuation game $\Gamma(\mathscr{B})$ could have distinct possible structures, depending on which history $\mathscr{B}$ of play resulted in the breakdown, however we assume otherwise.

Assumption 1. The continuation game $\Gamma(\mathscr{B})$ does not depend on the history $\mathscr{B}$ of play leading up to the breakdown.

Given this assumption, we let Γ^b denote the unique continuation game that follows any breakdown, regardless of the prior history $\mathscr{B}$ of play. Further, we make:

Assumption 2. Markov strategies are well defined in the continuation game Γ^b. A PBE with Markov strategies of the continuation game Γ^b exists. All PBE with Markov strategies of Γ^b have the same expected value, denoted as v^b_i, for each firm i.

We let $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$ denote the augmented game. We make Assumption 2 to avoid creating equilibrium multiplicity or causing equilibrium non-existence in the greenlight bargaining game $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$ simply because its breakdown continuation game Γ^b has no equilibria or has multiple equilibria with different values for the firms. For generality, we do not insist on any special specification of the breakdown continuation game Γ^b, nor do we impose any assumption on its equilibrium profile $v^b = \left\{v^b_i\right\}_{i \in N}$ of expected firm values.

All of our results for $\Gamma(G, S, f, \delta)$ with exogenous *status-quo* values continue to hold for $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$ with endogenous breakdown values v^b, after substituting v^b for the *status-*

quo values. The results' statements are otherwise unchanged. The following corollary states that the results of Theorem 1 and Proposition 1 carry over to this setting.

Corollary 2. *A PBE with Markov strategies of* $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$ *exists. All PBE with Markov strategies implement the unique socially efficient profile* s^* *of non-cash terms in stage 0 with probability 1. All PBE with Markov strategies have the same expected value* $\left[F(s^*) - \sum_j v_j^b\right]/n + v_i^b$ *for each firm* i*. Fix a sequence* $\{\delta_k\}$ *converging to 1. Let* $(s^*, y(k))$ *be the realized terms of any PBE with Markov strategies for* $\Gamma\left(G, S, f, \delta_k, \Gamma^b\right)$*. For each firm* i*,* $v_i(s_i^*, y(k)_i) \to \left[F(s^*) - \sum_j v_j^b\right]/n + v_i^b$ *almost surely.*

Proof. Given a subgame Γ, a strategy profile σ, and a belief system μ of $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$, we let $\sigma|_\Gamma$ and $\mu|_\Gamma$ be the restriction of σ and μ to the subgame Γ. We define $\sigma|_{\Gamma(G,S,f,\delta)}$ as the restriction of σ to the game $\Gamma(G, S, f, \delta)$, in that at any information set of $\Gamma(G, S, f, \delta)$, $\sigma|_{\Gamma(G,S,f,\delta)}$ prescribes the same action for the firm acting there as σ. Likewise, we let $\mu|_{\Gamma(G,S,f,\delta)}$ be the restriction of μ to the game $\Gamma(G, S, f, \delta)$, in that $\mu|_{\Gamma(G,S,f,\delta)}$ prescribes the same distribution at any information set of $\Gamma(G, S, f, \delta)$ as μ.

By Assumptions 1 and 2, and based on the definition of PBE, a strategy profile and a belief system (σ, μ) constitute a PBE with Markov strategies of $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$ if and only if $\left(\sigma|_{\Gamma(\mathscr{B})}, \mu|_{\Gamma(\mathscr{B})}\right)$ is a PBE with Markov strategies of the breakdown continuation game Γ^b for any breakdown event $\mathscr{B}$ and $\left(\sigma|_{\Gamma(G,S,f,\delta)}, \mu|_{\Gamma(G,S,f,\delta)}\right)$ is a PBE with Markov strategies of $\Gamma\left(G, S, f, \delta, v^b\right)$, where $\Gamma\left(G, S, f, \delta, v^b\right)$ is the game $\Gamma(G, S, f, \delta)$ of Section 3 with exogenous *status-quo* values v^b (instead of normalizing v^b to zero).

Therefore, Corollary 2 immediately follows from Theorem 1 and Proposition 1.

□

References

Abreu, D. and M. Manea (2012a): "Bargaining and Efficiency in Networks," *Journal of Economic Theory*, 147, 43–70.

——— (2012b): "Markov Equilibria in a Model of Bargaining in Networks," *Games and Economic Behavior*, 75, 1–16.

Andersen, L., D. Duffie, and Y. Song (2019): "Funding Value Adjustments," *The Journal of Finance*, 74, 145–192.

Bao, J., M. O'Hara, and X. (Alex) Zhou (2018): "The Volcker Rule and Corporate Bond Market Making in Times of Stress," *Journal of Financial Economics*, 130, 95–113.

Baron, D. P. and J. A. Ferejohn (1989): "Bargaining in Legislatures," *American Political Science Review*, 83, 1181–1206.

Battaglini, M. and S. Coate (2007): "Inefficiency in Legislative Policymaking: A Dynamic Analysis," *American Economic Review*, 97, 118–149.

Battaglini, M., S. Nunnari, and T. R. Palfrey (2012): "Legislative Bargaining and the Dynamics of Public Investment," *The American Political Science Review*, 106, 407–429.

Bernheim, B., B. Peleg, and M. D. Whinston (1987): "Coalition-Proof Nash Equilibria I. Concepts," *Journal of Economic Theory*, 42, 1–12.

Bessembinder, H., S. Jacobsen, W. Maxwell, and K. Venkataraman (2018): "Capital Commitment and Illiquidity in Corporate Bonds," *The Journal of Finance*, 73, 1615–1661.

Binmore, K. (1980): "Nash Bargaining Theory II (Now Published in The Economics of Bargaining, Edited by K. Binmore and P. Dasgupta, Basil Blackwell, Oxford, 1987.)," *STICERD - Theoretical Economics Paper Series*.

Binmore, K., A. Rubinstein, and A. Wolinsky (1986): "The Nash Bargaining Solution in Economic Modelling," *The RAND Journal of Economics*, 17, 176–188.

Binmore, K. G. and M. J. Herrero (1988): "Matching and Bargaining in Dynamic Markets," *The Review of Economic Studies*, 55, 17–31.

Bruche, M., F. Malherbe, and R. R. Meisenzahl (2020): "Pipeline Risk in Leveraged Loan Syndication," *The Review of Financial Studies*, 33, 5660–5705.

Brügemann, B., P. Gautier, and G. Menzio (2019): "Intra Firm Bargaining and Shapley Values," *The Review of Economic Studies*, 86, 564–592.

Chatterjee, K., B. Dutta, D. Ray, and K. Sengupta (1993): "A Noncooperative Theory of Coalitional Bargaining," *The Review of Economic Studies*, 60, 463–477.

Choi, J., Y. Huh, and S. Seunghun Shin (2024): "Customer Liquidity Provision: Implications for Corporate Bond Transaction Costs," *Management Science*, 70, 187–206.

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

Collard-Wexler, A., G. Gowrisankaran, and R. S. Lee (2019): ""Nash-in-Nash" Bargaining: A Microfoundation for Applied Work," *Journal of Political Economy*, 127, 163–195.

Condorelli, D., A. Galeotti, and L. Renou (2017): "Bilateral Trading in Networks," *The Review of Economic Studies*, 84, 82–105.

Corominas-Bosch, M. (2004): "Bargaining in a Network of Buyers and Sellers," *Journal of Economic Theory*, 115, 35–77.

Crémer, J. and M. H. Riordan (1987): "On Governing Multilateral Transactions with Bilateral Contracts," *The RAND Journal of Economics*, 18, 436–451.

Danos, V., J. Krivine, and J. Prat (2021): "Reversible and Composable Financial Contracts," Extended abstract.

de Fontenay, C. and J. S. Gans (2013): "Bilateral Bargaining with Externalities," Tech. rep., Melbourne Business School Discussion Paper No. 2004-32.

de Fontenay, C. C. and J. S. Gans (2014): "Bilateral Bargaining with Externalities," *The Journal of Industrial Economics*, 62, 756–788.

Elliott, M. and F. Nava (2019): "Decentralized Bargaining in Matching Markets: Efficient Stationary Equilibria and the Core," *Theoretical Economics*, 14, 211–251.

Elliott, M. and E. Talamàs (2025): "Bargaining Foundations for Price Taking in Matching Markets," *International Economic Review*, 66, 1561–1575.

Ferdowsian, A., M. Niederle, and L. Yariv (2025): "Strategic Decentralized Matching: The Effects of Information Frictions," Working Paper.

Freyermuth, R.W., C. K. O. and A. Tosato (2023): "Crypto in Real Estate Finance," *Alabama Law Review*, 75, 93–156.

Gale, D. (1987): "Limit Theorems for Markets with Sequential Bargaining," *Journal of Economic Theory*, 43, 20–54.

Gale, D. and H. Sabourian (2006): "Markov Equilibria in Dynamic Matching and Bargaining Games," *Games and Economic Behavior*, 54, 336–352.

Gans, J. S. (2019): "The Fine Print in Smart Contracts," Tech. rep., NBER Working Paper No. 2544, National Bureau of Economic Research.

Gofman, M. (2011): "A Network-Based Analysis of Over-the-Counter Markets," Dissertation, University of Chicago, Chicago, IL.

Gul, F. (1989): "Bargaining Foundations of Shapley Value," *Econometrica*, 57, 81–95.

Haeringer, G. and M. Wooders (2011): "Decentralized Job Matching," *International Journal of Game Theory*, 40, 1–28.

Hart, O., J. Tirole, D. W. Carlton, and O. E. Williamson (1990): "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity. Microeconomics*, 1990, pp. 205–286.

Hart, S. and A. Mas-Colell (1996): "Bargaining and Value," *Econometrica*, 64, 357–380.

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

Holden, R. and A. Malani (2017): "Can Blockchain Solve the Holdup Problem in Contracts?" Tech. rep., University of Chicago.

Horn, H. and A. Wolinsky (1988): "Bilateral Monopolies and Incentives for Merger," *The RAND Journal of Economics*, 19, 408–419.

Jackson, M. O. and B. Moselle (2002): "Coalition and Party Formation in a Legislative Voting Game," *Journal of Economic Theory*, 103, 49–87.

Jackson, M. O. and A. Wolinsky (1996): "A Strategic Model of Social and Economic Networks," *Journal of Economic Theory*, 71, 44–74.

Klein, B., R. G. Crawford, and A. A. Alchian (1978): "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process," *The Journal of Law & Economics*, 21, 297–326.

Lauermann, S. (2013): "Dynamic Matching and Bargaining Games: A General Approach," *American Economic Review*, 103, 663–689.

Lee, M. J., A. Martin, and R. M. Townsend (2024a): "Optimal Design of Tokenized Markets," Staff Reports (Federal Reserve Bank of New York), Federal Reserve Bank of New York.

——— (2024b): "Zero Settlement Risk Token Systems," Staff Reports (Federal Reserve Bank of New York), Federal Reserve Bank of New York.

Lee, R. S., M. D. Whinston, and A. Yurukoglu (2021): "Structural Empirical Analysis of Contracting in Vertical Markets," in *Handbook of Industrial Organization*, ed. by K. Ho, A. Hortaçsu, and A. Lizzeri, Elsevier, vol. 4 of *Handbook of Industrial Organization, Volume 4*, 673–742.

Manea, M. (2011): "Bargaining in Stationary Networks," *American Economic Review*, 101, 2042–2080.

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

McAfee, R. P. and M. Schwartz (1994): "Opportunism in Multilateral Vertical Contracting: Nondiscrimination, Exclusivity, and Uniformity," *The American Economic Review*, 84, 210–230.

Merlo, A. and C. Wilson (1995): "A Stochastic Model of Sequential Bargaining with Complete Information," *Econometrica*, 63, 371–399.

——— (1998): "Efficient Delays in a Stochastic Model of Bargaining," *Economic Theory*, 11, 39–55.

Myerson, R. B. (1977a): "Graphs and Cooperation in Games," *Mathematics of Operations Research*, 2, pp. 225–229.

——— (1977b): "Values of Games in Partition Function Form," *International Journal of Game Theory*, 6, 23–31.

Nash, J. (1950a): "The Bargaining Problem," *Econometrica*, 18, 155–162.

——— (1951): "Non-Cooperative Games," *Annals of Mathematics*, 54, 286–295.

Nash, J. F. (1950b): "Equilibrium Points in N-Person Games," *Proceedings of the National Academy of Sciences*, 36, 48–49.

Polanski, A. (2007): "Bilateral Bargaining in Networks," *Journal of Economic Theory*, 134, 557–565.

Rubinstein, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, pp. 97–109.

Rubinstein, A. and A. Wolinsky (1985): "Equilibrium in a Market with Sequential Bargaining," *Econometrica*, 53, 1133–1150.

Segal, I. (1999): "Contracting with Externalities," *The Quarterly Journal of Economics*, 114, 337–388.

Shapley, L. S. (1953): "A Value for N-Person Games," *Contributions to the theory of games*, 2, 307–317.

Snyder Jr., J. M., M. M. Ting, and S. Ansolabehere (2005): "Legislative Bargaining under Weighted Voting," *American Economic Review*, 95, 981–1004.

Stole, L. and J. Zwiebel (1996): "Intra-Firm Bargaining under Non-Binding Contracts," *Review of Economic Studies*, 63, 375–410.

Szabo, N. (1996): "Smart Contracts: Building Blocks for Digital Markets," Tech. rep., Phonetic Sciences Amsterdam.

Townsend, R. M. (2020): *Distributed Ledgers: Design and Regulation of Financial Infrastructure and Payment Systems*, The MIT Press.

Williamson, O. E. (1975): *Markets and Hierarchies, Analysis and Antitrust Implications: A Study in the Economics of Internal Organization*, Free Press.

——— (1979): "Transaction-Cost Economics: The Governance of Contractual Relations," *The Journal of Law & Economics*, 22, 233–261.

World Economic Forum (2025): "Asset Tokenization in Financial Markets: The Next Generation of Value Exchange," World Economic Forum Insight Report, May.