Smart Contracting in Network Markets*

Darrell Duffie and Chaojun Wang
Stanford GSB Wharton School

May 15, 2026

Latest Version

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, Piotr Dworczak, Vincent Glode, Brett Green, Sergiu Hart, Mathew Jackson, Urban Jermann, David Kreps, John Kuong, Alistair Milne, Julien Prat, Ariel Rubinstein, Nick Roussanov, Jakob Sagi, Andrzej Skrzypacz, Takuo Sugaya, Tomasz Sadzik, Brian Waters, Robert Wilson, Xiaobo Yu, 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 many economic and financial settings, firms connected by a network of contractual relationships bargain bilaterally. A firm's negotiations are interdependent: The marginal value to the firm of its contractual term with a given counterparty depends on its terms with other counterparties. This interdependency exposes the firm to holdup, because ending a negotiation with one counterparty possibly enables another counterparty to extract a larger rent from the firm in exchange for agreeing to a given term that now has a higher marginal value to the firm. For example, after an over-the-counter (OTC) dealer buys an asset from a seller, the prospective buyer can exploit the dealer's costly inventory to demand a lower resale price; after a buyer signs an asset-purchase agreement with a seller, its lender can demand a higher spread for needed financing because failure to pay triggers penalties; and in land assembly, after a developer buys several parcels, a remaining landowner can exploit the developer's need for the last parcel to demand a higher price. To reduce exposure to holdup, firms may settle on a socially inefficient set of contracts in equilibrium.

The large literature on bilateral bargaining between networked agents "has emphasized interdependencies and externalities across firms and agreements because they are often fundamental to bilateral oligopoly environments." (Collard-Wexler, Gowrisankaran, and Lee, 2019). Watson (2024) shows how externalities can be internalized through monetary transfers contingent on agents' externality-generating actions, even among indirectly connected agents. We show that holdup inefficiencies arising from payoff interdependencies are eliminated if agreements on the contracts being negotiated are reached atomically (all or none) by using smart contracts. Section 2 discusses prior work in greater detail.

Section 3.1 describes the primitives of bilateral bargaining among networked firms. In an arbitrary network of firms, each pair of connected firms negotiates a bilateral contract consisting of a non-cash term and a cash transfer; non-cash terms determine real value, while cash redistributes surplus. A firm's payoff is an arbitrary joint function of all its non-cash terms with its connected firms, allowing flexible payoff interdependencies, while cash terms enter linearly. The primitive objects consisting of the network, the set of possible non-cash terms, and firms' payoff functions are common knowledge.

An extensive-form bargaining protocol must keep each contractual negotiation bilateral, in that third parties can neither propose nor vote on the terms of a contract.

Section 3.2 proposes a bargaining protocol based on atomic agreement applicable to all tree-based networks, including hub-spoke and chain networks that are common in many settings. Under our suggested bargaining 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. In each stage, a firm is randomly selected to make the first offers of terms to its child firms, each of which chooses whether to greenlight the offer from its parent firm and then makes offers to its own child firms, and so on. If all offers receive a greenlight, the offers become binding contracts and the game ends. If one or more offers do not receive a greenlight, the game proceeds to the next stage with no agreement. Within a stage, the ongoing negotiation of a contract—including its offered terms, the firm making the offer, and whether the offer receives a greenlight—remains unobservable to third parties. Under a tie-breaking rule, the unique outcome of any Perfect Bayesian Equilibrium with Markov strategies is immediate agreement on socially efficient contracts. The no-cycle restriction, the restriction to Markov equilibrium strategies, and the tie-breaking rule, along with propagating offers, become unnecessary if firms' moves are sequential and publicly observable (Appendix C).

Atomic agreement is distinct from the concept of unanimous agreement typically considered in multilateral and legislative bargaining (e.g., Baron and Ferejohn (1989)). Atomicity requires agents to agree *at the same time*, whereas unanimity requires joint consent on the *same terms*. Our greenlight protocol implements atomicity but not unanimity, because firms cannot observe—let alone propose or vote on—the bilateral terms between other firms. Even in the fully sequential variant with publicly observable actions of Appendix C, an early mover cannot veto subsequent offers made by other firms. More fundamentally, atomicity and unanimity target distinct economic frictions. While unanimity resolves the inefficiency of "external costs imposed by collective action" (Buchanan and Tullock, 1965), Appendix

B shows that unanimity without atomicity has no bite against the holdup problem. In contrast, Appendix C shows that atomicity alone suffices to eliminate the holdup inefficiency if actions are publicly observable. In practice, unanimous agreement is neither feasible nor desirable in most economic and finance settings, where contractual terms between two firms are bilaterally negotiated and remain unobservable to third parties.

One concern with atomic agreement is the rigidity of its all-or-none requirement. In practice, a sub-network of firms would form side agreements—excluding the other firms whose greenlights are not desired—if doing so benefits themselves. Subsection 5.2 gives an example in which allowing for such flexibility seemingly gives rise to a blocking coalition. However, such threat of being excluded will be fully anticipated, and therefore the firms exposed to the exclusion risk would be willing to give up more profit to the other firms in the grand greenlight protocol in order not to be excluded. Appendix E formalizes this intuition, showing that the viability and the efficiency of our protocol is robust to accommodating the threat of exclusion. The accommodation gives every firm an endogenous no-agreement value in the grand greenlight bargaining protocol involving all firms. The threat of exclusions is priced in those no-agreement values. In that extension, firms achieve the socially efficient terms in the grand greenlight protocol, each receiving a share of the *residual* surplus relative to the total no-agreement value as an *addition* to its own no-agreement value.

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, as in a special case of Manea (2018, p.1275). 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.

Consider, as one illustrative example, an OTC asset market in which a dealer 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 dealer 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 dealer's low valuation. Anticipating this holdup, the dealer may avoid an initial purchase from the prospective seller. The same problem arises if the dealer first negotiates with the prospective buyer. In a worked example, we show that our greenlight bargaining protocol overcomes this holdup problem, achieving efficient trade, whereas excluding the greenlight protocol from an otherwise identical bargaining game leads to no trade in any equilibrium (Section 5.1). Without assigning full bargaining power to certain firms, or without some other approach such as atomic agreement that we model, holdups can be expected to cause inefficiency even with competition (Manea, 2018).

When arranging a leveraged loan syndication, the lead bank may commit to borrower terms before completing the syndication to investors. Once the borrower commitment is fixed, the lead bank's disagreement payoff in bilateral investor negotiations is lower because retaining an unplaced loan share is costly. Investors can therefore bargain against the lead bank's underwriting exposure. The same logic can arise in a multi-tranche securitization when the issuer warehouses collateral or fixes a collateral pool before all tranche terms are finalized. 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). Our protocol eliminates holdup because no contract is agreed unless all contracts are agreed.

Non-smart-contracting institutions exist to offer partial greenlight-like solutions to holdup inefficiencies. An OTC dealer conducts riskless-principal trades in an attempt to synchronize a purchase and a sale. A closing company certifies and executes both a sales agreement and a mortgage agreement. However, the existing institutions are restricted to three-firm chains, because synchronizing multiple bilateral negotiations without an obvious central coordinating party is practically difficult. Smart contracting adds two benefits: *(i)* It provides a scalable implementation of atomic agreement for general tree-based networks; *(ii)* it offers a single, general-purpose implementation of atomic agreement, eliminating the need for bespoke trusted neutral agents.

The rest of the paper is organized as follows. 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 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.[1] 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).

Several potential remedies have been attempted to target holdup inefficiencies on bilateral bargaining among networked agents: *(i)* Competition cannot resolve inefficiencies arising from holdups. Sellers' incentives to pursue intermediation chains that avoid holdups can lead to inefficient intermediation and allocation in tree-based network markets (Manea, 2018).[2] *(ii)* Stole and Zwiebel (1996) and Brügemann, Gautier, and Menzio (2019) analyzed proto-

[1] Lee, Whinston, and Yurukoglu (2021, Section 2.1) provides an overview of some of the key results among a significant amount of prior work on bilateral bargaining among networked agents. A separate literature addresses 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] Condorelli, Galeotti, and Renou (2017) sidestep holdup by not allowing buyers to ever make counter offers and focus on the effect of asymmetric information instead.

cols that permit costless renegotiation of any previously reached agreement. Under these protocols, however, failure to reach an agreement (a breakdown) cannot be renegotiated and thus holdup inefficiencies remain. *(iii)* Contingent contracts modeled by de Fontenay and Gans (2014) are practically challenging to negotiate because the number of contingencies contained in each contract is exponential in the number of other possible bilateral contracts in the network. Perhaps also due to this complexity, no bargaining protocol with contingent contracts has been shown to ensure efficiency of its equilibrium outcomes in the presence of holdup, to our knowledge.

In settings involving matching with transfers, holdup can cause inefficient matching or costly delays.[3] 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 central coordinating mechanism as in Polanski (2007). [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 dealer 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.

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

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

[4]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.

[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).

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.[6] 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 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 an arbitrary graph $G = (N, E)$ whose set E of edges defines the pairs of firms that bilaterally bargain

[6] 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."

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 flexible *interdependencies* across the negotiations of firm i, in that the marginal value to i of its non-cash term s_{ij} with firm j depends on its negotiated non-cash terms $(s_{ij'})_{j'\neq j}$ with other connected firms. Including $-\infty$ in the range of f_i allows 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

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

The *status-quo* value $f_i(s_i^0)$ to firm i is without loss of generality normalized to zero.

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.[7]

Because firms have transferable utilities, a profile of contract terms $(s_{ij}, y_{ij})_{(i,j)\in E}$ is Pareto efficient if and only if it maximizes the total social value $F(s) = \sum_{i\in N} f_i(s_i)$ of the non-cash terms s. We assume that there is a unique socially efficient profile s^* of non-cash terms, that is, a unique profile that maximizes $F(s)$. The uniqueness property is generically guaranteed.[8]

[7] 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.

[8] 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.

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. For this subsection, 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 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.[9]

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

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

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 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)* and *(iii)*.

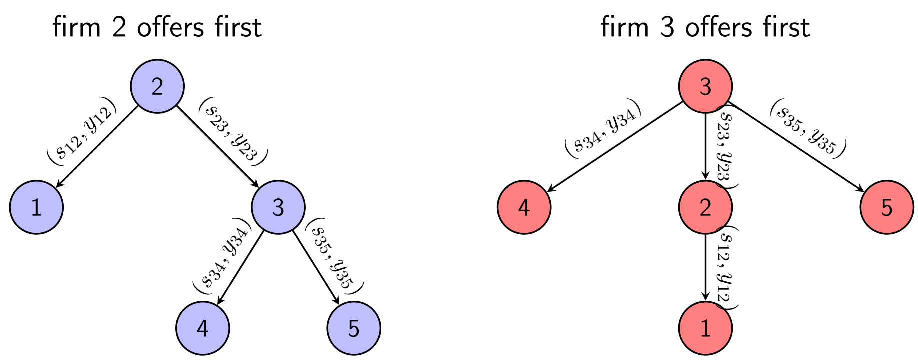

Figure 1: An illustration of the order in which offers are made at stages t and $t+1$ of a bargaining game for a fixed network of 5 firms in which the initial proposer of offers in stage t is $I(t) = 2$ and in stage $t+1$ is $I(t+1) = 3$.

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.

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 E replaces time discounting with a probabilistic breakdown of bargaining and allows firms to play an arbitrary continuation game that gives each of them an *endogenous* continuation value upon breakdown. We show that efficiency is robust to accommodating endogenous breakdown continuation values.

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.[10] If the two alternative sets of offers would receive the same expected number of greenlights from the child firms, then firm j chooses between $((s_{jk}, y_{jk})_{k\in c(j)}, g_j)$ 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 D 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 no-cycle restriction, the restriction to Markov equilibrium strategies, the tie-breaking rule, and propagating offers are required *solely* to eliminate equilibrium multiplicity arising from the game's imperfect information. Appendix C shows that if firms' moves are sequential and public, these four assumptions become unnecessary; greenlighting alone is sufficient to obtain the same efficient outcome in every subgame-perfect equilibrium.

The game induced by (G, S, f, δ) with the above-defined greenlight bargaining protocol

[10] 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.

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. *All PBE with Markov strategies have the same realized value $v_i^*(\delta)$ for each firm i. Fix a sequence $\{\delta_k\}$ of discount factors converging to 1. For each firm i, $v_i^*(\delta_k) \to F(s^*)/n$ almost surely.*

That every firm—first proposer or not—has nearly the same *realized* value in equilibrium when discounting vanishes shows that our greenlight protocol gives every firm equal bargaining power, rather than assigning all bargaining power to certain firms. In the special case of $n = 2$ firms, the greenlight protocol reduces to the canonical random-proposer bargaining model of Binmore (1987), where each firm is selected to propose with probability 1/2 in every bargaining round.[11]

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

[11] Our greenlight bargaining protocol can also be generalized to accommodate any distribution of bargaining power by randomly selecting each firm to make the first offers with an arbitrary non-zero probability in every stage. Then the distribution of surplus will be proportional to the selection probabilities. Since our main result is efficiency, not the division of surplus, we focus our analysis on the canonical case of equal bargaining power.

greenlight decisions is actually required, and methods used in practice to implement the greenlight protocol.

4.1 Atomic agreement overcomes holdup inefficiencies

Without the greenlight bargaining protocol based on atomic agreement, 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 marginal value to firm 2 of its negotiation with firm 1 precisely due to the interdependency between the two negotiations, possibly allowing firm 1 to extract more rent from firm 2 in exchange for a given non-cash term. 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 protocol 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.

Assuming no payoff interdependency would preclude holdup. Suppose the payoff functions f_i are separable 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.

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,

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

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 offers 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 D 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 information about any prior offer between other firms or whether the offer received a greenlight.

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 an unexpectedly large quantity to Buyer 1 (a spoke firm), 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 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.

[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).

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. The next section develops an application to asset trading networks, for which smart contracting provides the benefit of scaling relative to existing institutions by implementing atomic agreement beyond three-firm chains.

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.

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

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 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, riskless-principal intermediation is a partial greenlight-like solution: it synchronizes the dealer's purchase and sale, avoiding the dealer's exposure to bargaining after taking costly inventory onto its balance sheet or promising to deliver an asset that the dealer does not yet have in its inventory. However, empirical evidence on riskless-principal trades is concentrated in three-firm chains: one dealer matches and offsets two sides. Longer multi-dealer chains are documented, but we know of no evidence that they may be organized as riskless-principal trades.[17] Riskless-principal intermediation is not scalable to longer chains, and more generally to larger tree-based networks, because synchronizing multiple trades between different trading pairs without an obvious central coordinating party is practically difficult.] Smart contracting provides a solution and scales the greenlight protocol to general tree-based networks: An algorithm monitors the greenlight progress of multiple simultaneous bilateral negotiations between different pairs of counterparties in real time, automatically and immediately converting them into binding agreements once a greenlight shows up on every bilateral negotiation.

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.

[16] Zitzewitz (2010), Harris (2015), Ederington, Guan, and Yadav (2015), Schultz (2017), Bessembinder, Jacobsen, Maxwell, and Venkataraman (2018), Bao, O'Hara, and (Alex) Zhou (2018), Goldstein and Hotchkiss (2020), O'Hara and Zhou (2021), Choi, Huh, and Seunghun Shin (2024), Craig, Harris, and Shohfi (2026), and Jacobsen and Venkataraman (2025) provide empirical evidence from corporate bond markets. Sirri (2014) and Li and Schürhoff (2019) provide related evidence from municipal bond markets, and Kamate and Kumar (2024) from OTC currency derivatives.

[17] Li and Schürhoff (2019) study both riskless-principal trades and longer intermediation chains, and identify riskless-principal trades only in three-firm chains.

Each firm i has an exogenous fixed valuation $u_i(q) \in \mathbb{R} \cup \{-\infty\}$ for q units of the asset, with naked short sales precluded by our assumption that $u_i(q) = -\infty$ for $q < 0$.[18] We assume integer trade sizes. The trading game is the network bargaining game with primitives (G, S, f), where

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

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

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 $\Gamma(G, S, f, \delta)$ 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 $\Gamma(G, S, f, \delta)$. In stage t, one of the firms, denoted $I(t)$, is selected at random, uniformly among the firms that have not yet

[18] This restriction is equivalent to infinite practical cost of a failure to deliver: If firm 2 promises the asset to firm 3 before securing it from firm 1 and then cannot deliver, firm 2 faces a direct penalty, damage to its client relationship with firm 3, and reputational costs.

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 with whom no agreement has yet been reached. 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 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 for the original game of Section 3.2 with the greenlight protocol is also applied to this game with no greenlight protocol. This game, denoted as $\Gamma^{\text{no greenlight}}(G, S, f, \delta)$, 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 moves 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 $u_1(q) = \$99.7q$, $u_2(q) = \$98.8q$, 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.

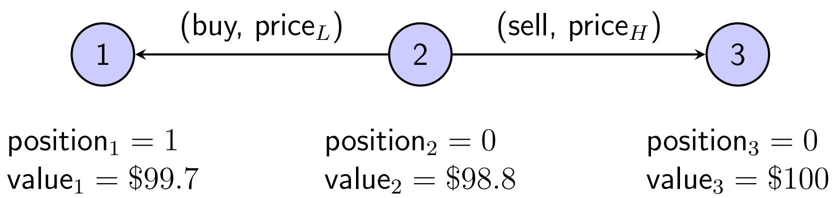

Figure 2: An over-the-counter asset trading market with firm 2 as a dealer. Initially, firm 1 has one unit of the asset, while firms 2 and 3 have none. There is no short selling. The firms' respective asset valuations are $u_1(q) = \$99.7q$, $u_2(q) = \$98.8q$, and $u_3(q) = \$100q$. The greenlight mechanism 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 dealer offers terms to firms 1 and 3.

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

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, $(\$98.8 + \$100)/2 = \$99.4$, 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 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 resell 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 \$99.4. 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 with a strictly positive probability, because short selling is precluded, $f_2(0,1) = -\infty$.

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 \$99.4 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. All PBE with Markov strategies have the same realized value $v_i^*(\delta)$ for each firm i, and the same expected value $F(q^*)/n$ for all firms. Fix a sequence $\{\delta_k\}$ of discount factors converging to 1. For each firm i, $v_i^*(\delta_k) \to F(q^*)/n$ almost surely.*

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).

5.3 A blocking coalition?

Consider three firms depicted in Figure 3, with asset valuations $u_1(q) = \$99.7q$, $u_2(q) = \$98.8q$, and $u_3(q) = \$100q$, but with alternative initial asset holdings $q_1^0 = 1$, $q_2^0 = 1$, and $q_3^0 = 0$. The efficient allocation transfers both units to firm 3 and creates total surplus \$1.5. With exogenous no-agreement values, Corollary 1 implies that all three firms earn expected profit \$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. One may be concerned that the greenlight protocol is not coalition-proof, and can be "blocked" by firms 2 and 3 in this example.

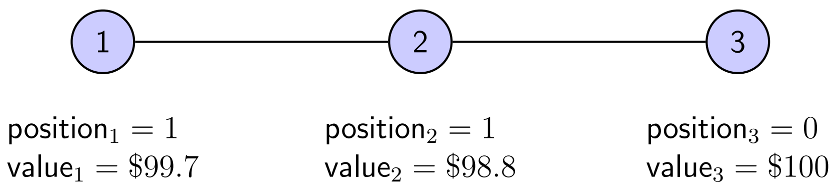

Figure 3: Initially, firms 1 and 2 have one unit of the asset each, while firm 3 has none.

However, such threat of being excluded will be fully anticipated by firm 1, who would therefore be willing to give up more profit in the grand greenlight protocol in order not to be excluded. Appendix E formalizes this intuition by allowing firms to play an arbitrary continuation game if the grand greenlight protocol breaks down with no agreement. Firms' ability to exclude each other by reaching side agreements in a smaller network is priced in their endogenous continuation values v^b in that breakdown game. In that augmented game, firms achieve the socially efficient terms s^* in the grand greenlight bargaininig protocol, each receiving an equal share of the residual surplus $F(s^*) - \sum_j v_j^b$ on top its own v_i^b both in expectation and in the limit as the breakdown probability vanishes.

Appendix E.1 describes one natural specification of the breakdown continuation game Γ^{nf}. In Γ^{nf}, a firm can make offers to any subset of its child firms that it selects. A non-root firm that does not receive an offer neither gives a greenlight nor makes offers in the current stage. Otherwise, Γ^{nf} is identical to the grand greenlight bargaining protocol $\Gamma(G, S, f, \delta)$.

Figure 4 illustrates how Γ^{nf} allows firms to form a sub-tree through bilateral side agreements and exclude firms whose greenlights are not desired.

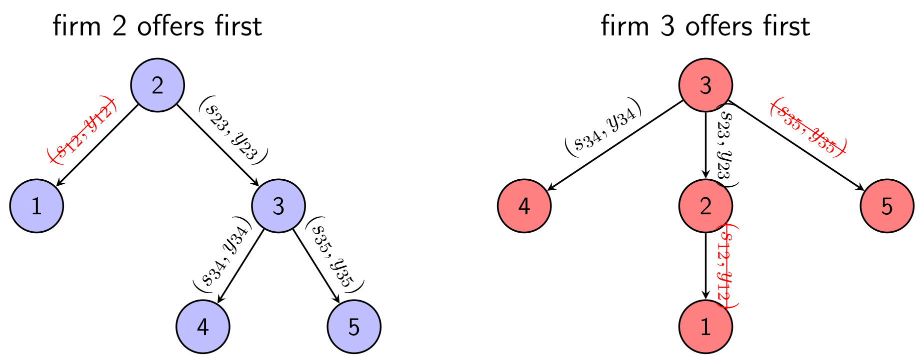

Figure 4: An example of a breakdown continuation game in which a firm can make offers to any subset of its child firms, forming a smaller sub-tree and excluding firms whose greenlight is not desired.

When the three firms in Figure 3 play the breakdown continuation game Γ^{nf}, their expected values are $v^b = (0.3, 0.6, 0.6)$ as the breakdown probability of Γ^{nf} vanishes. Appendix E.2 derives the continuation values and formalizes the argument for general primitives (G, S, f).

Therefore, the grand green light bargaining protocol continues to implement the efficient terms s^* immediately. The exclusion threat changes the division of surplus to $(0.3, 0.6, 0.6)$ in the limit.

6 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 an trusted neutral authority, such as a notary, that can provide contractible information about agreements in the network of firms. Smart contracting provides relative to existing non-smart-contracting institutions: It is scalable to general tree networks and general purpose.

Applications include over-the-counter trading networks, purchases requiring third-party financing, land assembly, syndicated lending, multi-tranche securitizations, "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 connected 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.

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 $\Gamma(G, S, f, \delta)$ 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 $\Gamma(G, S, f, \delta)$. 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 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, we let $G_{ij} = (N_{ij}, E_{ij})$ be the branch of 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.

We prove this claim 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 make 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 make 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 $(\bar{v}_i, \underline{v}_i)_{i\in N}$ expressed in the fixed-point 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.[19]

Step 2 (constructing a PBE with Markov strategies): We construct the following PBE with Markov strategies.

In each stage t, if firm j is selected to make the first offers, $j = I(t)$, then j offers (s_j^*, y_j^*)

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

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.

Step 3 (equilibrium non-cash terms). Since the unique expected equilibrium payoffs of all firms add up to $F(s^*)$ and the maximizer s^* of F is unique, if $F(s^*) > 0$, it must be that the socially efficient profile s^* of non-cash terms is immediately implemented in stage 0 with probability 1. If $F(s^*) = 0$, we fix a PBE with Markov strategies (σ, μ) of $\Gamma(G, S, f, \delta)$. The offers constructed in Step 2 satisfy the condition in Claim 1 and thus would receive greenlights with probability 1. The tie-breaking rule thus implies that all offers from every firm $j \neq I(0)$ other than the initial proposer $I(0)$ receive greenlights with probability 1 in stage 0 under (σ, μ). Lemma 1 implies that the payoff of firm $I(0)$ selected to make the first offers in stage 0 is zero under (σ, μ). Firm $I(0)$ can obtain its equilibrium payoff of zero while inducing the maximum possible expected number of greenlights from its child firms with its offers constructed in Step 2. The tie-breaking rule implies that all offers from $I(0)$ also receive greenlights with probability 1 in stage 0 under (σ, μ). The non-cash terms implemented by (σ, μ) with probability 1 in stage 0 must be s^*, as the unique expected equilibrium payoffs of all firms add up to $F(s^*)$ and the maximizer s^* of F is unique. □

Proof of Proposition 1. We fix a PBE with Markov strategies. Lemma 1 and Theorem 1 imply that if i is selected to make the first offers of terms, the payoff of firm i is

$$F(s^*)\left(1-\frac{\delta(n-1)}{n}\right).$$

Otherwise, the payoff of firm i is $\delta\,\frac{F(s^*)}{n}$. Proposition 1 follows. $\square$

Proof of Proposition 2. The proof immediately follows the statement of Proposition 2. $\square$

B Unanimous but sequential voting cannot fix holdup

This appendix shows that unanimous voting on each bilateral contract does not resolve the holdup inefficiency of Section 5.1 when contracts become binding sequentially rather than atomically.

We consider a unanimous-but-sequential protocol that proceeds exactly as the greenlight bargaining protocol of Section 3.2, but for three differences:

1. A firm makes offers to its child firms one at a time.

2. Each offer (s_{ij}, y_{ij}), along with the firm making the offer, is observed and subject to a public voting by all other firms sequentially, with a fixed order and observable ballots, as in Baron and Ferejohn (1989). Only unanimous approval converts the offered terms (s_{ij}, y_{ij}) into a binding contract and does so immediately.

3. Offer-making is restricted to connected firms (i, j) that have not yet reached a binding contract in prior stages.

The game has perfect information because of 1 and 2. The game is otherwise identical to the original game $\Gamma(G, S, f, \delta)$ defined in Section 3.2. In stage t, one of the firms, denoted $I(t)$, is selected at random, uniformly among the firms that have not yet reached binding contracts with one or more edge counterparties, to make the first offers of terms. Treat the graph G as a directed tree with root $I(t)$. When a firm makes offers, it does so to each of its child firms with whom no binding contract has yet been reached, in an exogenously fixed order. After firm i makes an offer (s_{ij}, y_{ij}), each firm k other than i chooses a probability

α_k^{ij} for voting in favor of the offer sequentially in an exogenously fixed voting order. The outcome of the vote is then realized. If every firm votes in favor, then (s_{ij}, y_{ij}) becomes a binding contract immediately. If any firm votes against the offer, then no contract is formed between firms i and j. After all offers from firm i have been voted on, each child firm of i, in an exogenously fixed order, makes offers 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 the offer has been voted on. Each choice, α_k^{ij} or (s_{ij}, y_{ij}), must be based only on (be measurable with respect to) information contained in the full history of play leading up to that choice, which consists of all earlier offers, firms making the offers, and ballots on the offers.

The binding offers in stage t 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 same tie-breaking rule that was defined for the original game of Section 3.2 is adapted to this game with unanimous voting. If a firm is indifferent between two offers that it could make, it chooses the offer that is more likely to receive unanimous approval. If a firm is indifferent between voting in favor of or against an offer, it votes in favor. We denote this game by $\Gamma^{\text{unanimous}}(G, S, f, \delta)$.

Since $\Gamma^{\text{unanimous}}(G, S, f, \delta)$ has perfect information, PBE is equivalent to subgame-perfect equilibrium (SPE). We apply the unanimous-but-sequential protocol to the three-firm chain trading example of Section 5.1.

Proposition 3. *In the three-firm chain trading example of Section 5.1, no trade occurs in any SPE of* $\Gamma^{\text{unanimous}}(G, S, f, \delta)$.

Proof. Fix a history where a trade between firms 1 and 2 has been unanimously approved and thus completed first. From that point onward, firm 1 will vote in favor of any offered terms (s_{23}, y_{23}) between firms 2 and 3, because (s_{23}, y_{23}) does not affect the payoff of firm 1. Thus, the *expected* price at which firm 2 sells the asset to firm 3 is the average of the two firms' valuations, ($98.8 + $100)/2 = $99.4, 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 occurs between either pair of firms. A history where a trade is first completed between firms 2 and 3 is also off-path in any SPE because short selling is precluded, $f_2(0,1) = -\infty$. Therefore, no trade occurs in any subgame-perfect equilibrium of $\Gamma^{\text{unanimous}}(G, S, f, \delta)$. $\square$

C Atomicity is sufficient if information is perfect

Four assumptions imposed on the original greenlight bargaining game $\Gamma(G, S, f, \delta)$ in Section 3.2—the no-cycle restriction, the restriction to Markov equilibrium strategies, the tie-breaking rule, and propagating offers—are required *solely* to eliminate equilibrium multiplicity arising from the game's imperfect information. We show that these four assumptions are not necessary for a fully sequentialized variant of $\Gamma(G, S, f, \delta)$ that has perfect information.

We define the fully sequentialized variant for any arbitrary graph G, which may contain cycles and be disconnected. The fully sequentialized variant proceeds exactly as the original greenlight bargaining game $\Gamma(G, S, f, \delta)$, but for two differences:

1. Firms move one at a time. Conditional on firm $I(t) = i$ moving first in stage t, the remaining firms move sequentially according to an arbitrary exogenous order $\prec_i$ on $N \setminus \{i\}$, which need not be based on shortest-path distance to i.

2. Each offer (s_{ij}, y_{ij}), along with the firm making the offer and whether it receives a greenlight, is observed by all other firms.

Therefore, the game has perfect information. The game is otherwise identical to the original game $\Gamma(G, S, f, \delta)$. In stage t, one of the firms, denoted $I(t)$, is selected at random, uniformly among the firms, to make the first offers of terms. The identity of firm $I(t)$ is publicly observed. Firm $I(t)$ moves first, and the remaining firms then move sequentially according to $\prec_{I(t)}$. When firm j moves, it has observed all previous moves in the stage. Firm j first chooses a probability g_j for greenlighting the offers (s_{ij}, y_{ij}) it receives from its connected firms i that moved earlier in stage t. These greenlight events are then realized, independently. If one or more offers do not receive a greenlight from j, the stage ends immediately with no agreement and the game proceeds to stage $t + 1$. If firm j greenlights all offers received, it then chooses a publicly observable offer (s_{jk}, y_{jk}) to each connected firm k that moves later in stage t. These choices, g_j and (s_{jk}, y_{jk}), must be based only on (measurable with respect to) information contained in the full history of play leading up to those choices, which consists of all earlier offers and firms making the offers. After the last firm moves, if every offer receives a greenlight, then the terms of all stage-t offers become binding agreements simultaneously and the game ends. No tie-breaking rule is imposed. We denote this game by $\Gamma^{\mathrm{sp}}(G, S, f, \delta)$. In the special case where G is a tree and $\prec_i$ is based on shortest-path distance to i, $\Gamma^{\mathrm{sp}}(G, S, f, \delta)$ fully sequentializes the simultaneous moves of the original game $\Gamma(G, S, f, \delta)$ while preserving their order, and makes offers and greenlight outcomes publicly observable.

Since $\Gamma^{\mathrm{sp}}(G, S, f, \delta)$ has perfect information, PBE is equivalent to subgame-perfect equilibrium (SPE). We do not restrict equilibrium strategies to be Markov.

Proposition 4. *An SPE of $\Gamma^{\mathrm{sp}}(G, S, f, \delta)$ exists. If G is connected, then all SPE have the same realized value $v_i^{\mathrm{sp}}(\delta)$ for each firm i, and the same expected value $F(s^*)/n$ for all firms. Fix a sequence $\{\delta_k\}$ of discount factors converging to 1. For each firm i, $v_i^{\mathrm{sp}}(\delta_k) \to F(s^*)/n$ almost surely.*

Like Theorem 1, we need a lemma to prove Proposition 4. For any vector $\hat{v} = (\hat{v}_i)_{i \in N}$, let $\Gamma^{\mathrm{sp}}(G, S, f, \hat{v})$ denote the one-stage version of $\Gamma^{\mathrm{sp}}(G, S, f, \delta)$ in which, if some offer does not receive a greenlight, the game ends immediately and firm i receives $\delta \hat{v}_i$.

Lemma 2. *If G is connected, then all SPE of $\Gamma^{\mathrm{sp}}(G, S, f, \hat{v})$ are payoff equivalent. If 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 therefore

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

Proof. We let G be an arbitrary graph and $N_1, \ldots, N_M$ be its connected components, We define $F_m(s) = \sum_{j \in N_m} f_j(s_j)$ for every $m \leq M$. We give every firm i the choice of a greenlight probability g_i, even if firm i receives no offers, and let $\tilde{\Gamma}^{\mathrm{sp}}(G, S, f, \hat{v})$ denote this augmented game. We show that

Claim 2. If $F_m(s^*) > \delta \sum_{j \in N_m} \hat{v}_j$ for every $m \leq M$, then all SPE of $\tilde{\Gamma}^{\mathrm{sp}}(G, S, f, \hat{v})$ are payoff equivalent. If firm i is the first to make offers of terms in some connected component G_m, the payoff of firm i is

$$F_m(s^*) - \delta \sum_{j \in N_m \setminus \{i\}} \hat{v}_j. \tag{C.5}$$

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

We prove this claim based on induction over the number n of firms in G. The claim is immediate with $n = 1$ firm. Assume the claim holds with fewer than n firms. Let i be the first firm to move in G, and let G_m be the component containing i.

Given arbitrary offers (s_i, y_i) by i, we define the primitives (G', S', f') for the residual network with the remaining $n-1$ firms: We let $G' = G - i := (N', E')$ be the residual network, where $N' = N \setminus \{i\}$ and $E' = \{(j,k) \in E : j, k \in N'\}$, and let $S' = S_{|E'}$ be the restriction of the non-cash-term space S to the residual network, and define the residual payoff functions by

$$f'_j(s'_j) = \begin{cases} f_j(s'_j, s_{ij}) + y_{ij}, & \text{if } (i,j) \in E, \\ f_j(s'_j), & \text{otherwise.} \end{cases}$$

The subgame following the offers (s_i, y_i) of i is identical to $\tilde{\Gamma}^{\mathrm{sp}}\left(G', S', f', (\hat{v}_j)_{j \neq i}\right)$.

Every $j \in N_m \setminus \{i\}$ can secure $\delta \hat{v}_j$ by withholding a greenlight, while total surplus in G_m is at most $\max_s F_m(s)$. Hence, the payoff of i in any SPE is at most given by (C.5).

To show that (C.5) is also a lower bound of the equilibrium payoff of i, fix some $\varepsilon > 0$. Let $\mathcal{C}$ be the connected components of $G_m - i$. In each component $C \in \mathcal{C}$, choose a firm $j_C \in C$ connected to i and a number $\varepsilon_C > 0$, with $\sum_C \varepsilon_C < \varepsilon$. We let firm i offer the non-cash terms s_i^* and cash terms

$$y_{ijC} = \delta \sum_{k\in C} \hat{v}_k + \varepsilon_C - \sum_{\ell\in C} f_\ell(s_\ell^*),$$

All other cash terms from i into C are zero. Then the maximal total payoff of firms in C is

$$\delta \sum_{\ell\in C} \hat{v}_\ell + \varepsilon_C,$$

which is strictly greater than $\delta \sum_{\ell\in C} \hat{v}_\ell$. The maximal total payoff in any other component $G_{m'}$ is unchanged ($m' \neq m$). The induction hypothesis therefore applies to $\tilde{\Gamma}^{\text{sp}}(G', S', f', (\hat{v}_j)_{j\neq i})$. All residual components reach agreement, and i obtains

$$f_i(s_i^*) - \sum_{C\in\mathcal{C}} y_{ijC} = F_m(s^*) - \delta \sum_{j\in N_m\setminus\{i\}} \hat{v}_j - \sum_{C\in\mathcal{C}} \varepsilon_C.$$

Letting $\varepsilon \downarrow 0$ gives the lower bound. Therefore, the payoff of i is equal to (C.5) in any SPE. Then any other firm j in G_m must receive $\delta\hat{v}_j$ in any SPE, and the induction hypothesis gives the claimed payoffs in all other components. This proves Claim 2.

We now return to proof of Lemma 2. Any SPE of $\Gamma^{\text{sp}}(G, S, f, \hat{v})$ can be augmented to become an SPE of $\tilde{\Gamma}^{\text{sp}}(G, S, f, \hat{v})$ by letting $g_j = 1$ for any firm j that does not receive offers. If $F(s^*) > \delta \sum_j \hat{v}_j$, Lemma 2 is a special case of Claim 2. If $F(s^*) \leq \delta \sum_j \hat{v}_j$, every firm can secure $\delta\hat{v}_j$, and total payoff is no larger than $\delta \sum_j \hat{v}_j$; hence every firm receives exactly $\delta\hat{v}_j$. The stated payoff formula follows from these two cases, and the expected-payoff formula follows by averaging over the uniformly selected first mover. □

We now prove Proposition 4.

Proof of Proposition 4. Step 1 (equilibrium payoff): We let $\bar{v}_i^{\text{sp}}$ and $\underline{v}_i^{\text{sp}}$ be the highest and the lowest expected payoff of firm i in any SPE of $\Gamma^{\text{sp}}(G, S, f, \delta)$ ($i \in N$).

Because agreement is atomic, every stage that ends with no agreement leads to another subgame that is identical to $\Gamma^{\text{sp}}(G, S, f, \delta)$. Lemma 2 implies that the highest attainable expected payoff of firm i is

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

Therefore, for every i,

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

Likewise, for every i,

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

This constitutes the same system of equations used in the proof of Theorem 1. The same contraction-map argument therefore implies

$$\bar{v}_i^{\text{sp}}=\underline{v}_i^{\text{sp}}=\frac{F(s^*)}{n},\qquad i\in N.$$

Hence, all SPE of $\Gamma^{\text{sp}}(G,S,f,\delta)$ have the same expected value $F(s^*)/n$ for all firms. Lemma 2 also implies that in any SPE, if firm i is selected to make the first offers of stage 0, then

$$v_i^{\text{sp}}(\delta)=F(s^*)\left(1-\frac{\delta(n-1)}{n}\right),$$

whereas if another firm is selected first, then

$$v_i^{\text{sp}}(\delta)=\delta\,\frac{F(s^*)}{n}.$$

Therefore, all SPE have the same realized value $v_i^{\text{sp}}(\delta)$ for each firm i, and $v_i^{\text{sp}}(\delta_k)\to F(s^*)/n$ almost surely if $\delta_k\to 1$.

Step 2 (constructing an SPE): We construct the following SPE.

At any public history h in a stage, let N_m be the connected component containing the current mover i after earlier movers are removed, and define

$$F_m(s;h) = \sum_{k\in N_m} \left[f_k\left((s_{k\ell})_{\ell\in N_m,\,(k,\ell)\in E},\, \left(s^h_{jk}\right)_{j\notin N_m,\,(j,k)\in E} \right) + \sum_{j\notin N_m,\,(j,k)\in E} y^h_{jk} \right].$$

Firm i greenlights all received offers if and only if

$$\max_s F_m(s;h) \geq \delta\, \frac{|N_m| F(s^*)}{n}.$$

If this inequality holds, we fix some $s^m \in \operatorname{argmax}_s F_m(s;h)$. Firm i offers non-cash terms (s^m_{ij}) to its connected firms. For each connected component C of $G_m - i$, i chooses a connected firm $j_C \in C$ and offers cash terms

$$y_{ij_C} = \delta\, \frac{|C|}{n} F(s^*) - \sum_{k\in C} \left[f_k\left((s^m_{k\ell})_{\ell\in N_m,\,(k,\ell)\in E},\, \left(s^h_{jk}\right)_{j\notin N_m,\,(j,k)\in E} \right) + \sum_{j\notin N_m,\,(j,k)\in E} y^h_{jk} \right],$$

$$y_{ik} = 0 \quad \text{for every other connected firm } k \text{ in } C.$$

Under the above strategy profile, one can verify that the current mover in any given subgame has no profitable deviation. The above strategy profile thus constitutes an SPE. □

D 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 structure $\mathcal{F}^*_{it}$ is 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 5. *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 5. 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 $\{T, T-1, \dots, 0\}$. 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.

□

E Efficiency is robust to any game upon breakdown

In practice, a sub-network of firms could form side agreements—excluding the other firms whose greenlights are not desired—if doing so benefits themselves. Is our efficiency result robust to allowing for such threat of exclusion? This appendix shows that the answer is "Yes."

Given some model primitives (G, S, f), we augment the original greenlight bargaining game $\Gamma(G, S, f, \delta)$ of Section 3.2 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 breakdown continuation game Γ^b. The game Γ^b has a PBE with Markov strategies.

We let $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$ denote the augmented game. Assumption 2 rules out only non-existence of continuation play after breakdown. It allows Γ^b to have multiple Markov PBE with different values for the firms. Given any Markov PBE of the breakdown continuation game, let $v^b = \{v_i^b\}_{i \in N}$ denote the equilibrium profile of expected firm values. For generality, we do not insist on any special specification of the breakdown continuation game Γ^b. The proof of our efficiency result for the augmented game $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$ uses only the property that $\sum_i v_i^b \leq F(s^*)$, which is a direct consequence of s^* maximizing the total social value $F(s)$.

The following corollary states that (i) our efficiency result continues to hold for $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$, and (ii) the division of surplus for the original greenlight bargaining game $\Gamma(G, S, f, \delta)$ with exogenous *status-quo* values carries over to $\Gamma\left(G, S, f, \delta, \Gamma^b\right)$ after substituting the expected firm values v^b induced by its breakdown continuation play for the *status-quo* values. Different Markov PBE of Γ^b may change the division of surplus through v^b.

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. In any such PBE, letting* v^b *be the expected firm values induced by its Markov continuation play in* Γ^b*, firm* i *has expected value* $\left[F(s^*) - \sum_j v_j^b\right]/n + v_i^b$*. Fix a sequence* $\{\delta_k\}$ *of discount factors converging to 1. For each* k*, fix a PBE with Markov strategies of* $\Gamma\left(G, S, f, \delta_k, \Gamma^b\right)$*, let* $(s^*, y(k))$ *be the realized terms, and let* $v^b(k)$ *be the expected firm values induced by its Markov continuation play in* Γ^b*. For each firm* i*,*

$$v_i(s_i^*, y(k)_i) - \left(\frac{F(s^*) - \sum_j v_j^b(k)}{n} + v_i^b(k)\right) \to 0 \quad \textit{almost surely.}$$

Proof. Given a Markov strategy profile σ, a belief system μ, and a subgame Γ 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. Since σ is Markovian and any breakdown continuation subgame $\Gamma(\mathscr{B})$ is identical to the game Γ^b (Assumption 1), $\sigma|_{\Gamma(\mathscr{B})}$ is a Markov strategy profile of Γ^b and does not depend on the breakdown history $\mathscr{B}$. Let v^b denote the expected firm values in the game Γ^b under $\sigma|_{\Gamma(\mathscr{B})}$. These values do not depend on the breakdown history $\mathscr{B}$. 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.

We have $\sum_i v_i^b \leq F(s^*)$, since s^* maximizes the total social value $F(s)$. By Assumptions 1 and based on the definition of PBE, (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 after any breakdown history $\mathscr{B}$ 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 3.2 with exogenous *status-quo* values v^b (instead of normalizing v^b to zero).

Therefore, Corollary 2 immediately follows from Assumption 2, Theorem 1, and Proposition 1. $\square$

E.1 Bilateral network formation after breakdown

We describe one natural specification of the breakdown continuation game Γ^b. It allows firms to form a sub-tree through bilateral side agreements and exclude firms whose greenlights are not desired.

The breakdown continuation game proceeds in stages $\tau = 0, 1, 2, \ldots$. Every stage proceeds exactly as this section's greenlight bargaining game $\Gamma(G, S, f, \delta, \Gamma^b)$ with probabilistic breakdown, but for two differences:

1. A firm can make offers to any subset of its child firms that it selects. A non-root firm

that does not receive an offer neither gives a greenlight nor makes offers in the current stage.

2. In the event of a breakdown, all contractual terms remain at *status-quo*.

The breakdown continuation game is otherwise identical to the original game $\Gamma(G, S, f, \delta)$. In stage τ, one firm, denoted $I(\tau)$, 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 $I(\tau)$ until the beginning of the next stage. Treating the graph G as a directed tree with root $I(\tau)$, firm $I(\tau)$ makes offers to any selected subset 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 offers to any selected subset of its own child firms, and so on, until each firm that has received an offer from its parent firm 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 $\mathcal{F}_{j\tau}$ of three sources of information: *(i)* the information $\mathcal{F}_\tau$ commonly available at the beginning of stage τ, *(ii)* the event of whether firm j is selected to make the first offers in stage τ, $\mathbb{1}_{j=I(\tau)}$, and *(iii)* for a non-root node, the information contained in its parent's offer to firm j, if j receives such an offer. 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)* and *(iii)*. The root firm $I(\tau)$ and the firms that receive offers form a sub-tree G' of G.

The stage-τ greenlight events are then realized, independently. If every offer in the sub-tree G' receives a greenlight, then the terms of all these offers automatically become binding agreements while the other contractual terms in G remain at *status quo*, and the breakdown continuation game ends. If one or more offers fail to receive a greenlight, no offer becomes binding. The game then proceeds to the next stage, $t + 1$, with probability $\lambda \in (0, 1)$. With probability $1 - \lambda$, the breakdown continuation game ends and all firms receive their *status-quo* values.

The full history of play through stage τ becomes observable to all firms at the beginning of stage $\tau+1$. The common information $\mathcal{F}_{t+1}$ at the beginning of stage $t+1$ is then the join of $\mathcal{F}_\tau$, the information revealed by $I(\tau)$ and all of the stage-τ offers and greenlight events. We denote the resulting bilateral network-formation game by $\Gamma^{\mathrm{nf}}(G, S, f, \lambda)$.

Proposition 6. *The bilateral network-formation game $\Gamma^{\mathrm{nf}}(G, S, f, \lambda)$ has a PBE with Markov strategies, and thus satisfies Assumption 2.*

"Attaching" this same game to follow every breakdown history satisfies Assumption 1. Any firm can withhold its greenlight, force breakdown of the original greenlight bargaining game, and move play to the more flexible breakdown continuation game $\Gamma^{\mathrm{nf}}(G, S, f, \lambda)$. The breakdown continuation game relaxes the rigidity the original greenlight bargaining game $\Gamma(G, S, f, \delta)$ by allowing firms to form side agreements through bilateral offers and greenlights in a smaller sub-tree. It enables blocking coalition by allowing firms to exclude firms whose greenlights are not desired. Their effects are absorbed through the endogenous breakdown continuation values v^b induced by a Markov PBE of $\Gamma^{\mathrm{nf}}(G, S, f, \lambda)$. Corollary 2 then shows that these endogenous values change the division of surplus, while the grand greenlight protocol continues to implement the socially efficient non-cash terms immediately.

Proof of Proposition 6. Step 1 (equilibrium payoff): For any sub-tree $G' = (N', E')$ of G, let $S' = \prod_{(k,\ell)\in E'} S_{k\ell}$. For any $s' \in S'$, let $F'(s')$ denote the total social value when the contracts in E' have terms s' and all other contracts have their *status-quo* terms. For any vector $v \in [0, F(s^*)]^N$, let $\mathcal{A}_i(v)$ be the set of pairs (G', s') that solve

$$\max_{\substack{G'=(N',E') \text{ sub-tree of } G,\, i\in N' \\ s'\in S'}} \left\{ F'(s') - \lambda \sum_{j\in N'\setminus\{i\}} v_j \right\}.$$

The set of sub-trees and non-cash profiles is finite. For each $v \in [0, F(s^*)]^N$, the sets $\mathcal{A}_i(v)$ are therefore nonempty.

A vector $v \in [0, F(s^*)]^N$ solves the breakdown-continuation-payoff equations if there are distributions π_i over $\mathcal{A}_i(v)$ such that, for every firm i,

$$\begin{aligned} v_i =& \frac{1}{n} \sum_{(G',s')\in\mathcal{A}_i(v)} \pi_i(G',s') \left[F'(s') - \lambda \sum_{j\in N'\setminus\{i\}} v_j \right] \\ &+ \frac{\lambda v_i}{n} \sum_{k\in N\setminus\{i\}} \sum_{\substack{(G',s')\in\mathcal{A}_k(v)\\ i\in N'}} \pi_k(G',s'). \end{aligned} \tag{E.6}$$

Let a correspondence map each v to the set of payoff vectors generated by the right side of (E.6) as the distributions $\{\pi_i\}_{i\in N}$ range over distributions on $\mathcal{A}_i(v)$. This correspondence has nonempty, compact, convex values and a closed graph. It maps $[0, F(s^*)]^N$ into itself. Kakutani's fixed-point theorem gives a vector v^{nf} and distributions $\{\pi_i^{\text{nf}}\}_{i\in N}$ over $\mathcal{A}_i(v^{\text{nf}})$ that solve (E.6).

Step 2 (constructing a PBE with Markov strategies): We construct the following PBE with Markov strategies.

For any pair $(i,j) \in E$ of connected firms, let $G_{ij} = (N_{ij}, E_{ij})$ be the branch of G that has root j, and $c(k)$ be the set of child firms of firm k. Suppose firm j receives from firm i an arbitrary offer (s_{ij}, y_{ij}). For any sub-tree $\tilde{G} = (\tilde{N}, \tilde{E})$ of G_{ij} with $j \in \tilde{N}$ and any $\tilde{s} \in \tilde{S} = \prod_{(k,\ell)\in\tilde{E}} S_{k\ell}$, define

$$F_{ij}^{\tilde{G}}(s_{ij}, \tilde{s}) = \sum_{k\in\tilde{N}} f_k(\bar{s}_k),$$

where $\bar{s}$ is the non-cash profile that assigns s_{ij} to contract (i,j), assigns $\tilde{s}_{k\ell}$ to every contract $(k,\ell) \in \tilde{E}$, and assigns *status-quo* terms to all other contracts. Firm j greenlights the offer (s_{ij}, y_{ij}) if and only if

$$\max_{\substack{\tilde{G}=(\tilde{N},\tilde{E})\text{ sub-tree of }G_{ij},\ j\in\tilde{N}\\ \tilde{s}\in\tilde{S}}} \left\{ F_{ij}^{\tilde{G}}(s_{ij}, \tilde{s}) - \lambda \sum_{\ell\in\tilde{N}\setminus\{j\}} v_\ell^{\text{nf}} \right\} + y_{ij} \geq \lambda v_j^{\text{nf}}. \tag{E.7}$$

Letting $(\tilde{G}^*, \tilde{s}^*)$ be a selected maximizer in (E.7) and treating $\tilde{G}^*$ as a directed tree with root j, firm j makes offers to each of its child firms in $\tilde{G}^*$. For every such child firm k, let $\tilde{G}_{jk}^* = (\tilde{N}_{jk}^*, \tilde{E}_{jk}^*)$ be the branch of $\tilde{G}^*$ that has root k. Firm j offers child firm k the non-cash term $\tilde{s}^*_{jk}$ and the cash term

$$y^*_{jk}(s_{ij}, y_{ij}) = \lambda \sum_{\ell \in \tilde{N}^*_{jk}} v^{\mathrm{nf}}_\ell - F^{\tilde{G}^*_{jk}}_{jk}\left(\tilde{s}^*_{jk}, \tilde{s}^*|_{\tilde{E}^*_{jk}}\right). \tag{E.8}$$

Firm j makes no offers to child firms outside $\tilde{G}^*$.

If firm i is selected to make the first offers, it draws (G', s') according to π^{nf}_i. Treat G' as a directed tree with root i. For each child firm j of i in G', let $G'_{ij} = (N'_{ij}, E'_{ij})$ be the branch of G' that has root j. Firm i offers child j the non-cash term s'_{ij} and the cash term

$$y_{ij} = \lambda \sum_{\ell \in N'_{ij}} v^{\mathrm{nf}}_\ell - F^{G'_{ij}}_{ij}\left(s'_{ij}, s'|_{E'_{ij}}\right).$$

It makes no offers to child firms outside G'.

Let μ^{nf} be any belief system consistent with the constructed strategy profile. We now verify sequential rationality. Consider any non-root information set at which firm j has received an arbitrary parent offer (s_{ij}, y_{ij}). The maximization in (E.7) and the cash terms in (E.8) maximize firm j's payoff. Greenlighting is optimal when condition (E.7) holds. Thus, non-root firms are sequentially rational after every on-path or off-path offer. If selected firm i deviates to any offer profile that receives all greenlights, every reached non-root firm must receive at least $\lambda v^{\mathrm{nf}}_j$. The selected firm's payoff is therefore no larger than the objective value attained by any $(G', s') \in \mathcal{A}_i(v^{\mathrm{nf}})$. Thus the selected firm is sequentially rational.

The strategy profile is Markovian. It is sequentially rational for μ^{nf}. The fixed-point equations (E.6) give the payoff vector v^{nf}. Therefore, coupling the above strategy profile with μ^{nf} constitutes a PBE with Markov strategies.

□

E.2 The three-firm chain example of Subsection 5.2

We derive firms' expected values v^b in the breakdown continuation game $\Gamma^b = \Gamma^{\mathrm{nf}}(G, S, f, \lambda)$ for the three-firm chain example of Subsection 5.2 in which the firms' initial asset holdings are $(1, 1, 0)$. We show that v^b is unique across all Markov PBE.

Lemma 3. *For the three-firm chain trading example of Subsection 5.2 in which the firms' initial asset holdings are* $(1, 1, 0)$*, all Markov PBE of the breakdown continuation game* $\Gamma^{\text{nf}}(G, S, f, \lambda)$ *have the same expected value* v_i^b *for each firm* i*, where*

$$v^b = (0.5, 0.5, 0.5) \qquad \textit{when } \lambda < \frac{3}{5},$$
$$v_1^b(\lambda) = \frac{0.3}{\lambda}, \qquad v_2^b(\lambda) = v_3^b(\lambda) = \frac{1.2}{3-\lambda}, \qquad \textit{when } \lambda \geq \frac{3}{5}.$$

Proof. Let $v^b = (v_1^b, v_2^b, v_3^b)$ be the expected firm values in any given Markov PBE of $\Gamma^{\text{nf}}(G, S, f, \lambda)$. With the notation of (E.6), let $\pi_i(G)$ be the probability with which firm i selects the full chain G when selected to make the first offers. Firm 1 selects G with probability 1, because the only positive-surplus sub-tree available to firm 1 is G.

The payoff gain from selecting G rather than the sub-tree $\{2, 3\}$ for selected firm $i \in \{2, 3\}$ is

$$\left[1.5 - \lambda\left(v_1^b + v_j^b\right)\right] - \left[1.2 - \lambda v_j^b\right] = 0.3 - \lambda v_1^b,$$

where $j \in \{2, 3\} \setminus \{i\}$ is the other firm. Thus firms 2 and 3 select G if $\lambda v_1^b < 0.3$, select the sub-tree $\{2, 3\}$ if $\lambda v_1^b > 0.3$, and can mix between the two only if $\lambda v_1^b = 0.3$.

If $\lambda v_1^b < 0.3$, then $\pi_2(G) = \pi_3(G) = 1$. The fixed-point equations (E.6) are

$$v_i^b = \frac{1.5 - \lambda \sum_{j \neq i} v_j^b}{3} + \frac{2\lambda v_i^b}{3}, \qquad i = 1, 2, 3.$$

The unique solution is $v^b = (0.5, 0.5, 0.5)$. This case requires $\lambda < 3/5$.

If $\lambda v_1^b > 0.3$, then $\pi_2(G) = \pi_3(G) = 0$. The fixed-point equations for firms 2 and 3 are

$$v_2^b = \frac{1.2 - \lambda v_3^b}{3} + \frac{2\lambda v_2^b}{3}, \qquad v_3^b = \frac{1.2 - \lambda v_2^b}{3} + \frac{2\lambda v_3^b}{3}.$$

These equations imply $v_2^b = v_3^b = 1.2/(3 - \lambda)$. Firm 1's fixed-point equation gives

$$v_1^b = \frac{1.5 - 2.4\lambda/(3 - \lambda)}{3}.$$

This implies

$$\lambda v_1^b - 0.3 = -\frac{3.9\lambda^2 - 5.4\lambda + 2.7}{3(3-\lambda)} < 0$$

for every $\lambda \in (0,1)$, a contradiction. Thus this case cannot arise.

It remains to consider $\lambda v_1^b = 0.3$. Then $v_1^b = 0.3/\lambda$. The fixed-point equations for firms 2 and 3 are

$$v_2^b = \frac{1.2 - \lambda v_3^b}{3} + \frac{2\lambda v_2^b}{3}, \qquad v_3^b = \frac{1.2 - \lambda v_2^b}{3} + \frac{2\lambda v_3^b}{3}.$$

They imply $v_2^b = v_3^b = 1.2/(3-\lambda)$. Firm 1's fixed-point equation then pins down $\pi_2(G) + \pi_3(G)$:

$$\frac{0.3}{\lambda} = \frac{1.5 - 2.4\lambda/(3-\lambda)}{3} + \frac{0.3(\pi_2(G) + \pi_3(G))}{3}. \tag{E.9}$$

At $\lambda = 3/5$, this equation requires $\pi_2(G)+\pi_3(G) = 2$, so $\pi_2(G) = \pi_3(G) = 1$. For $\lambda \in (3/5, 1)$, feasible mixing probabilities exist. For example, the symmetric mixing probabilities satisfy

$$\pi_2(G) = \pi_3(G) = \frac{9 - 18\lambda + 13\lambda^2}{2\lambda(3-\lambda)},$$

which lies in $(0,1)$ for $\lambda \in (3/5, 1)$. Therefore, for $\lambda > 3/5$, the unique v^b is

$$v_1^b(\lambda) = \frac{0.3}{\lambda}, \qquad v_2^b(\lambda) = v_3^b(\lambda) = \frac{1.2}{3-\lambda}.$$

The mixing probabilities need not be unique, but v^b is. Finally, for $\lambda \in [0, 3/5)$, no feasible mixing probabilities satisfy equation (E.9). □

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. (1987): *Perfect Equilibria in Bargaining Models*, London School of Economics, International Centre for Economics and Related ....

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.

Buchanan, J. M. and G. Tullock (1965): *The Calculus of Consent: Logical Foundations of Constitutional Democracy*, vol. 100, University of Michigan press.

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.

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.

Craig, L., L. Harris, and T. Shohfi (2026): "Riskless Principal Trades in Corporate Bond Markets," *The Journal of Fixed Income*, 35, 18–29.

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.

Ederington, L. H., W. Guan, and P. K. Yadav (2015): "Dealer Spreads in the Corporate Bond Market: Agent vs. Market-Making Roles," CFR Working Paper 15-11, Centre for Financial Research.

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.

Goldstein, M. A. and E. S. Hotchkiss (2020): "Providing Liquidity in an Illiquid Market: Dealer Behavior in U.S. Corporate Bonds," *Journal of Financial Economics*, 135, 16–40.

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

Harris, L. (2015): "Transaction Costs, Trade Throughs, and Riskless Principal Trading in Corporate Bond Markets," Working Paper, University of Southern California.

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.

Jacobsen, S. and K. Venkataraman (2025): "Receiving Investors in the Block Market for Corporate Bonds," *Journal of Financial Economics*, 170, 104061.

Kamate, V. and A. Kumar (2024): "Dealer Networks, Client Sophistication and Pricing in OTC Derivatives," *Journal of International Money and Finance*, 140, 102986.

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.

Li, D. and N. Schürhoff (2019): "Dealer Networks," *The Journal of Finance*, 74, 91–144.

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.

O'Hara, M. and X. A. Zhou (2021): "The Electronic Evolution of Corporate Bond Dealers," *Journal of Financial Economics*, 140, 368–390.

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.

Schultz, P. (2017): "Inventory Management by Corporate Bond Dealers," Working Paper, University of Notre Dame.

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.

Sirri, E. R. (2014): "Report on Secondary Market Trading in the Municipal Securities Market," Research Paper, Municipal Securities Rulemaking Board.

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.

Watson, J. (2024): "Contractual Chains," *Econometrica*, 92, 1735–1774.

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

Zitzewitz, E. (2010): "Paired Corporate Bond Trades," Working Paper, Dartmouth College.