Tractable Staggered Bargaining via the Step-by-Step Property

Jacob Dorn*

March 23, 2026

Abstract

Many supply markets feature staggered bargaining, yet prevailing models assume simultaneous contract formation. Assuming simultaneity risks misattributing price dynamics to bargaining weights, but explicitly modeling staggering introduces explosive state space growth. I show that if bilateral shares are fixed, which corresponds to a property called step-by-step, then one can eliminate future states via a finite dependence-type cancellation. I propose a model of Nash equilibrium in Kalai proportional bargains, the only bargaining family satisfying the step-by-step property in general. I establish that the resulting model, which I call Nash-in-Kalai, yields a tractable GMM estimation framework that nests transferable-utility Nash-in-Nash models. I test the importance of timing using results from a companion paper studying hospital–insurer contracting in West Virginia. I find that contracts are pervasively staggered; negotiators are substantially forward-looking; and a static model produces biased and sometimes incoherent bargaining weights, whereas models with accurate timing produce coherent bargaining weights.

1 Introduction

Staggered contracting is prevalent in firm-to-firm bargaining. In cable carriage, YouTube TV dropped Disney programming for weeks due to constraints from Disney's preexisting

*This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2039656 and by grant T32 HS026116 from the Agency for Healthcare Research and Quality. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation or the Agency for Healthcare Research and Quality. The results here represent a preliminary analysis, and structural estimates of counterfactual behavior should not be interpreted as a direct measure of firm conduct or intent. I am indebted to many conversations with my committee members Alessandro Lizzeri and Ulrich Müller; and helpful conversations with David Ahn, Eduard Boehm, Nick Buchholz, Colleen Carey, Allan Collard-Wexler, Adam Dearing, Gregory Dobbels, Gautam Gowrisankaran, Philip Haile, Barton Hamilton, Adam Harris, Catherine Ishitani, Jakub Kastl, Quan Le, W Bentley MacLeod, Tyler Maxey, Mikkel Plagborg-Møller, Alexander Olssen, Eric Qian, Jackson Reimer, Jesse Silbert, Tong Wang, Frank Yang, and Yihao Yuan; seminar participants at the International Industrial Organization Conference and the University of Pennsylvania; and many other generous people whom I omit for brevity. Artificial intelligence was used in the editing process. This work was previously circulated as "The Nash-in-Kalai Model for Estimation with Dynamic Bargaining and Nontransferable Utility" and "A Microfoundation for the Nash-in-Kalai Model." This work is profoundly indebted to my advisor, Kate Ho; I hope she would have been proud of its final form. All errors are my own.

agreements (Lerner, 2025). In property management, shopping malls leverage multidecade anchor tenant agreements to attract smaller stores (Bell, 2024). In healthcare, a single hospital–insurer negotiation can govern subsequent prices for more than a decade (Dorn, 2025a).

Prevailing models of firm-to-firm bargaining assume that all contracts are formed at the same time. When contracts are truly staggered, a period-by-period model must misspecify gains from trade: a late contract price is discounted at the time of bargaining, but a period-by-period model assigns that price fully to bargaining weights. Simple corrections fail. For example, aggregating multiple years within a period implicitly rules out time-discounting, and adjustments that reduce bias for a specific contract can increase the bias for other agreements. At a fundamental level, no fixed time aggregation can correctly capture gains for all staggered contracts.

The natural response is to explicitly model staggering, but doing so introduces explosive state space growth. Firm-to-firm contracting features contracting externalities: a given pair's price affects their leverage when bargaining with other players (Lee et al., 2021). When contracts overlap, externalities are internalized: if i and j form a three-year contract the year before h and j form a three-year contract, then j must forecast how their current price with i will affect their future price with h, how that future price with h will affect the next agreement with i, and so on. Each pair's remaining contract length then becomes a relevant bargaining state, so that a market with N bilateral agreements lasting T periods generates T^N market configurations. For a monthly model with nine pairs forming five-year contracts, this process yields $60^9 \approx 10$ quadrillion configurations.

This paper shows that a single property of a bargaining solution, called step-by-step, controls the state space growth introduced by staggering. Loosely speaking, under step-by-step, a given pair always splits surplus in fixed proportions.

Under step-by-step, it is as though a two-period agreement is chosen relative to disagreeing for both periods. To see why, consider a two-period, two-player example. Suppose

players 1 and 2 choose a two-period agreement at price p^A, generating period-1 value $V_1^i(p^A)$ for player i. The agreement in period 1 is chosen relative to disagreeing in period 1 and then bargaining over a single-period price p^D with period-2 value $V_2^i(p^D)$. If each agreement gives player 1 a share τ of the joint gains, then the agreement price p^A satisfies

$$\tau\left(V_1^2(p^A) - \beta V_2^2(p^D)\right) = (1-\tau)\left(V_1^1(p^A) - \beta V_2^1(p^D)\right), \tag{1}$$

and the off-equilibrium disagreement price p^D satisfies

$$\tau V_2^2(p^D) = (1-\tau)V_2^1(p^D). \tag{2}$$

Multiplying Equation (2) by β and adding the result to Equation (1) yields that p^A solves

$$\tau V_1^2(p^A) = (1-\tau)V_1^1(p^A). \tag{3}$$

That is, the off-equilibrium bargain (2) already splits gains proportionally to τ, and so it can be replaced by a second disagreement without changing the equilibrium prediction. This is the cancellation that will make estimation tractable.

The two-period cancellation applies inductively for the full contract duration, generating a finite dependence-type cancellation of future states. States between the current agreement and the next agreement are canceled by step-by-step, and states after the next agreement enter the simplified model in the same way. Applying this property reduces the state space bound from T^N under generic staggered bargaining to at most $\min\{T, N\}^N$ under step-by-step, and produces a much tighter bound when equilibrium contracting dates are sparse.

Unfortunately, Nash bargaining loses the step-by-step property in precisely the staggered environments in which the property is needed. Step-by-step requires that a pair's surplus shares are constant in every subgame. Nash bargaining satisfies step-by-step in the transferable-utility (TU) models that cover most dynamic empirical Nash-in-Nash models

(Lee and Fong, 2013; Liu, 2021; Tiew, 2022) and much of modern static practice (Ho and Lee, 2017; Crawford et al., 2018; Collard-Wexler et al., 2019). However, staggering introduces time-varying contracting externalities that tilt the Pareto frontier as future agreements come closer in time and then are passed. Nash bargaining responds to the resulting change of marginal values at the Pareto frontier, breaking the fixed-proportion requirement for step-by-step. As a result, state space control under staggered contracting requires a bargaining solution other than Nash.

This paper instead studies the Kalai (1977) proportional solution. The Kalai proportional solution chooses the best agreement along a fixed-surplus-share ray. Unlike Nash bargaining, the surplus-split ray continues to apply even if internalized contracting externalities change marginal values on the Pareto frontier. Kalai proportional bargaining is the only solution concept to preserve the step-by-step property in general (Kalai, 1977; Roth, 1979).

I formalize a model of Nash equilibrium in Kalai proportional bargains, which I call Nash-in-Kalai. I extend Dutta (2012)'s static bilateral model to microfound multilateral Nash-in-Kalai in terms of delegate revocation costs. The leading microfoundation for Nash-in-Nash (Collard-Wexler et al., 2019) only covers the simultaneous-TU case in which Nash and Kalai coincide. The resulting multilateral solution concept is amenable to empirical practice.

The Nash-in-Kalai model yields a closed-form moment for GMM estimation that nests existing TU estimation strategies. The expected net present value (NPV) payment under Nash-in-Kalai equals the expected NPV of the period-by-period flow payments that would emerge under TU Nash-in-Nash, plus a dynamic adjustment that captures the effect of repeated bilateral disagreement on other pairs' future prices. In steady state, this adjustment term is zero, so that existing TU estimation strategies carry over without modification. Outside steady state, one can iteratively capture the dynamic adjustment through increasingly higher-order bounded rationality approximations. A level-1 approximation appears adequate in my empirical application.

I apply the Nash-in-Kalai framework to data on hospital–insurer contracting from West Virginia, adapting the empirical results of Dorn (2025a) to illustrate staggered timing's empirical implications. West Virginia made data on hospital–insurer contracts public records. Contracts were pervasively staggered across and within years, with lengths running from three years to more than a decade.

I compare three specifications: a static model with misspecified timing, a myopic model with accurate timing but no weight on future years, and the preferred forward-looking Kalai model. The static model produces biased and sometimes incoherent bargaining weights, including an estimated weight of -7.54 for one insurer. The myopic and forward-looking models produce similar and plausible bargaining weights, suggesting that the static model's bias primarily arises from misspecified timing rather than misspecified discounting. However, the forward-looking model's estimated discount factor of $\hat{\beta} = 0.899$ overwhelmingly rejects myopia, so that only the forward-looking model can accurately predict responses to dynamic counterfactuals.

This paper presents the first empirical model of multilateral bargaining that accommodates staggered bargaining. Existing models either assume simultaneous contract formation in a single period (Horn and Wolinsky, 1988; Crawford and Yurukoglu, 2012; Grennan, 2013; Gowrisankaran et al., 2015; Ho and Lee, 2017, 2019; Crawford et al., 2018; Ghili, 2022) or allow multiple periods under a maintained assumption that contracts are reformed in each period (Lee and Fong, 2013; Liu, 2021; Tiew, 2022; Deng et al., 2024; Barwick et al., 2025). Staggered bargaining has been studied in the context of firm-to-firm bargaining (Do and Miklós-Thal, 2022; Zhou and Sibley, 2026), workers searching across employers (Cahuc et al., 2006; Bilal et al., 2022; Jarosch et al., 2024), and employers contracting with multiple workers (Brügemann et al., 2018; Stole and Zwiebel, 1996), but this work has, to my knowledge, remained either theoretical or calibrated to aggregate outcomes. This paper bridges these literatures, proposing a model of staggered bargaining that can be taken to data on individual contracts for the first time.

The rest of this paper is as follows. Section 2 presents the challenges introduced by staggering: period-by-period models introduce bias, but staggered models introduce explosive state space growth. Section 3 illustrates the cancellation from step-by-step, explains why Nash loses step-by-step under staggering, and discusses the Kalai proportional solution's advantages and limitations. Section 4 formalizes the Nash-in-Kalai model of Nash equilibrium in Kalai proportional bargains and presents the revocation costs microfoundation. Section 5 provides the practical upshot of the model for empirical practice. Section 6 outlines the empirical application and key results. Section 7 concludes.

2 Bias and State Space Growth from Staggering

When contracts are meaningfully staggered, a period-by-period model misspecifies gains from trade, risking bias in estimates of bargaining weights. However, explicitly modeling staggered contracts requires tracking an internalized price externality that depends on all contract lengths in a market, introducing explosive state space growth.

2.1 Bias From Misspecified Timing

A simple model shows that when dynamics are meaningful and contracts are staggered, a period-by-period model risks biasing some, if not all, bargaining weights. The problem is that any way of fitting staggered contracts into a series of shared periods will distort the time-weighted gains from trade for every agreement that is not actively targeted.

A single monopolist insurer forms two-period agreements with two hospitals. For each period with an agreement with a given hospital, the insurer gains G and pays a flow payment p, while the hospital receives p and pays a marginal cost $c = 0.9G$. A two-period agreement specifies a first-period price p_1, with the second-period price $p_2 = (1 + \pi)p_1$ following from some known exogenous inflation rate π. Both sides value future periods at a rate $\beta > 0$, so that NPV gains from an agreement are $(1+\beta)G-p_1-\beta p_2$ for the insurer and $p_1+\beta p_2-(1+\beta)c$

for the hospital. The bargaining solution allocates a share $\tau = 0.5$ of joint NPV gains to each side, yielding a starting price of

$$p_1^* = (1+\beta)\left(c + (1-\tau)(G-c)\right)/(1+\beta(1+\pi)).$$

Suppose a researcher wishes to build a structural model of this market. They have done months of research, and know all structural parameters of the model, including the true insurer bargaining weight $\tau = 0.5$, from interviews. However, they wish to validate their understanding by recovering τ from the observed data. If the researcher could study each agreement in isolation, they could recover τ from the split of NPV gains from trade.

Unfortunately, the researcher wishes to follow the prevailing practice of modeling period-by-period bargaining, and they can see that contracts are staggered. In odd-numbered periods $t = 1, 3, \ldots$, hospital O negotiates its starting price p_1, and in even-numbered periods $t = 2, 4, \ldots$, hospital E negotiates its starting price p_1. In multilateral markets, it will be important to model substitution across suppliers under disagreement, so the researcher wishes to fit both hospitals into the same period-by-period model.

The researcher might apply a period-by-period model by estimating bargaining weights separately in each period. In period $t = 1$, they observe $p^O = p_1^*$ and $p^E = p_2^* = (1+\pi)p_1^*$. They plug in these prices and infer bargaining weights

$$\begin{aligned}\hat{\tau}^O &= \frac{G - p_1^*}{G - c} = \tau + \frac{9.5\beta\pi}{1+\beta(1+\pi)} \\ \hat{\tau}^E &= \frac{G - p_2^*}{G - c} = \tau - \frac{9.5\pi}{1+\beta(1+\pi)}.\end{aligned}$$

This estimation approach is essentially static. Unless $\pi = 0$, so that the underlying bargaining problem is essentially static, the period-by-period approach will introduce bias. If I calibrate these numbers based on the list price-based contracts in Section 6.1 below, I might obtain a ten-year-plus-lived contract with an annual discount factor of 0.9 and an annual

inflation rate of 6%, yielding a five-year approximate rates of $\beta = 0.5$ and $\pi = 1/3$.[1] In this case, the econometrician estimates $\hat{\tau}^O = 1.45$ and $\hat{\tau}^E = -1.4$. This is impossible — everyone knows bargaining weights should be between zero and one. The problem is that the period-by-period approach incorrectly treats new contract periods (hospital O in period 1) and old contract periods (hospital E in period 1) as representative of all periods, generating bias in opposite directions.

Stricken with bargaining weights below 0 and above 1, the researcher guesses that the problem is that they have lost the multiperiod structure of contracting. The researcher could aggregate gains across contracts within a period by averaging across the two prices in place in the market, the too-large p_1^* and the too-small p_2^*. Alternatively, the researcher could create synthetic periods that match the true two-period contract length, averaging between p_1^* and p_2^* across the two underlying periods of each contract. Both approaches incorrectly impose complete patience ($\beta = 1$) when aggregating, and both approaches yield the same estimate of $\hat{\tau}^{Shared} = 0.025$, exactly halfway between the two separate estimates.

The researcher has a breakthrough moment: they need to get discounting right. When the firms agree to p_1^*, they discount p_2^* by a factor of β. None of these approaches have adjusted for discounting, so of course they risk introducing bias. The researcher looks back thankfully on their months of interviews that allowed them to know with certainty that the discount factor is $\beta = 0.5$. They take their synthetic periods and downweight the second half: $t = 2$ gets half as much weight as $t = 1$, $t = 4$ gets half as much weight as $t = 3$, and so on. They go to estimate the bargaining weights, and they obtain:

$$\hat{\tau}^O = \frac{(1+\beta)G - p_1^* - \beta p_2^*}{(1+\beta)(G-c)} = \tau.$$

Finally, the researcher has gotten the correct bargaining weight of 1/2! After a dutiful round of champagne, the researcher returns to their computer to confirm the other bargaining

[1] The low-margin calibration $c = 0.9G$ is chosen above to dramatize the outcome.

weight is as it should. They find:

$$\hat{\tau}^E = \frac{(1+\beta)G - p_2^* - \beta p_1^*}{(1+\beta)(G-c)} = \tau - \frac{9.5\pi(1-\beta)}{1+\beta(1+\pi)}.$$

The odd-period hospital's bargaining weight is now correctly specified at $\hat{\tau}^O = 1/2$, but the even-period hospital's bargaining weight is now misspecified even worse as $\hat{\tau}^E = -0.45$. The problem is that while applying an intra-$\tilde{t}$ discount factor correctly specifies the gains from trade for a contract whose start and end date is captured by the model, the intra-$\tilde{t}$ discount factor actually flips the discount factors applied at bargaining for a contract formed in the middle.

The fundamental problem is that no single time aggregation can accurately capture gains from trade for all contracts in a market with staggering. Discounting from hospital O's contract start date gets hospital O's gains right but applies the wrong time discounting to hospital E. An annual model starting at the same time as hospital O's start date does even worse, getting hospital O's time discounting wrong and implying a model in which hospital E only values its last contract period. Any period-by-period approach must choose one window, and staggering ensures that at least one contract will be misaligned with any window.

Eventually, the researcher decides that their only option is to explicitly write down a model that allows contracts to be staggered. Unfortunately, they find that when that model includes both staggering and strategic interactions, they lose much of the tractability of a period-by-period model.

2.2 State Space Growth Under Staggered Contracting

The previous subsection showed that when contracts are staggered, a period-by-period model risks introducing bias. A natural response is to write down a model that explicitly incorporates staggering. I now use a simple example to show why adding staggering to

multilateral bargaining models will introduce explosive state space growth. The core reason is that the contracting externalities that emerge in period-by-period models are internalized across overlapping contracts when bargaining is staggered.

Some time later, the researcher sets off to model a market with important strategic interactions. The core of the market remains one insurer and two hospitals, with hospitals still named O and E. In each day, the insurer receives $4G$ in revenue if both hospitals have an agreement in place, receives $3G$ in revenue if only one hospital has an agreement, and receives no revenue if their network of agreements is empty. However, payments and costs now depend on the network. Revenue corresponds to units of care demanded: patients split equally within the insurer's network, so that each hospital provides 2 units of care if both hospitals have an agreement in place, an agreeing hospital provides 3 units of care if it is the only hospital with an agreement, and both hospitals provide zero care in the absence of an agreement. Inpatient care costs $c = 0.9G$ per unit, so that if price p_h is in place with hospital h, then the hospital gains $2(p_h - c)$ per day if both hospitals have an agreement, and gains $3(p_h - c)$ if it is the only hospital in the insurer's network; the insurer's profit is $4G - 2p_O - 2p_E$ if both hospitals have an agreement, is $3G - 3p_h$ if only hospital h has an agreement, and is zero with an empty network.

This is a setting with contracting externalities (Lee et al., 2021): the insurer's gains from trade with hospital O depend on its negotiated price with hospital E. Consider a day in which the insurer knows it will pay p_E per unit of care to hospital E. The insurer's profit is $4G-2p_O-2p_E$ with an agreement at price p_O, and its profit is $3G-3p_h$ if it only has a single agreement. Thus, the hospital-O agreement generates gains to the insurer of $G + p_E - 2p_O$: by agreeing with hospital O, the insurer increases its revenue by G and reduces the amount of care it sends to hospital E by one unit. The more the insurer agrees to pay hospital E, the more the insurer is willing to pay hospital O, and vice versa.

Remembering the lessons from Section 2.1, the researcher sets off to model this market while allowing staggering. The researcher obtains data on contract timing and finds the

patterns that I see in West Virginia in Section 6.1 in other markets: some contracts are three-year contracts, some contracts are multidecade contracts, and contracts are staggered even within a year. The researcher chooses an arbitrary annual period length; with an eye towards the two-year contracts this illustration focuses on, the researcher splits each year into $T/2$ time periods.

In equilibrium, hospital O reaches agreements on January 1 of odd-numbered years, and hospital E reaches agreements on January 1 of even-numbered years. Mercifully, the contracts have no inflation: an agreement in period t_0 corresponds to a starting price p_{h,t_0} that will remain in place for T periods, with the next contract formed in period $t_0 + T$. In equilibrium, hospital O forms agreements in periods $1, T+1, 2T+1, \ldots$, and hospital E forms agreements in periods $T/2+1, 3T/2+1, 5T/2+1, \ldots$. For illustration purposes, there is no uncertainty.

In equilibrium, both hospitals reach an agreement. The researcher conjectures that the hospital-h price negotiations yield a price-response function $p_h(p_{-h}, \ell)$, where $\ell = T/2$ is the remaining length of the hospital-h agreement. (In fact, the agreement will need to represent a Pareto improvement over disagreement in order to be accepted, but it will be simplest to ignore disagreement constraints in this illustration.) The researcher also conjectures that because the hospitals are symmetric, their price response functions will be symmetric. Then, whatever the price-response function is, the value functions for the marginal agreement for the insurer (M for Managed Care Organization) and Hospital (H for Hospital) take the form of

$$\begin{aligned} V^M(p_h, T, p_{-h}, \ell) &= \sum_{t=0}^{\ell} \beta^t \left(4G - 2p_h - 2p_{-h}\right) + \beta^\ell V^M(p_h(p_h, T-\ell), T, p_h, T-\ell) \\ V^H(p_h, T, p_{-h}, \ell) &= \sum_{t=0}^{T-1} \beta^t \left(2(p_h - c)\right) + \beta^T V^H(p_h(p_h(p_h, T-\ell), \ell), T, p_h(p_h, T-\ell), \ell), \end{aligned} \tag{4}$$

where $p_h(p_h, T-\ell)$ is the price that hospital $-h$ will reach in the next agreement and $p_h(p_h(p_h, T-\ell), \ell)$ is the price that hospital h will reach in the next agreement in response.

Incorporating staggering has already introduced a new type of bargaining state because contract externalities are now internalized. In the hospital-O negotiation, the value functions include an infinite chain of future price responses. The 2025 p_O will affect the 2026 p_E negotiation; that next p_E negotiation will affect the price that hospital O negotiates in 2027, and so on to infinity. These internalized contracting externalities are new: in period-by-period models, contracting externalities occur through anticipated simultaneous agreements, and are never internalized. But in practice, these on-equilibrium externalities are not so bad, and for example only require one bargaining state per year for these two hospitals.

So far, the researcher has only had to account for one path: the equilibrium path. The troublesome state space growth will come from capturing the value of disagreement. Bargaining solutions model an agreement as being chosen based on the value of agreeing (described in Equation (4) above) relative to disagreement. In principal, disagreement could involve a simple response like permanent exclusion, or a carefully tailored response chosen to control state space growth. In practice, disagreement is often painful, with disagreeing firms both losing out on direct profits and losing leverage with competitors. As a result, in many settings, the most natural model of disagreement is disagreeing for one period followed by a new negotiation.

Consider what happens if hospital O and the insurer fail to agree and there are six-month periods, so that $T = 4$. In equilibrium, hospitals O and E continue to negotiate in alternating years. But if hospital O disagrees, it will bargain again on July 1, with the hospital-E negotiation closer in time. That is, the equilibrium price response $p_h(p_{-h}, \ell = T/2)$ is recursively defined in terms of the next-period price response $p_h(p_{-h}, \ell - 1)$. In turn, the July-1 price response $p_h(p_{-h}, \ell - 1)$ is recursively defined in terms of waiting another period and negotiating in yet another new bargaining state. In the end, this process continues T times, with each step introducing a new bargaining state with a new internalized contracting externality.

If the researcher is to consider more players, then there will be increasingly more price-

response functions. With three hospitals, the first hospital's price has an internalized externality on future prices at both other hospitals. This generates two separate effects, each with T potential price-response functions. The other two contracts will also interact in future periods, so that in combination, there are T^2 market configurations for the other pairs. More broadly, with N symmetric agreements in steady state, the full recursive bargaining model involves T^{N-1} price-response functions for each pair. Aggregating over the full market, there are NT^{N-1} price-response functions across all pairs, and as many as T^N market configurations for when each pair negotiates. If $T = 5$ captures an annual model with $N = 9$ for three retailers negotiating with three suppliers, then there are 2 million configurations in the recursively defined bargaining problem. If the researcher tries a more accurate monthly model, there would be some 10 quadrillion market configurations. The researcher can control some of this growth by noting that multiple market configurations correspond to the same relative timing, but even after canceling market configurations with the same underlying timing, there will still be $NT^{N-1} = 1.6$ quadrillion price-response functions. The step-by-step property will control this growth.

3 The Step-by-Step Property and State Space Tractability

The step-by-step property from static axiomatic game theory turns out to be precisely the condition to avoid explosive state space growth from staggered contracting. Under step-by-step, the value of disagreeing once can be replaced with the value of disagreeing for the duration of the equilibrium agreement, generating an equivalent path with finite dependence-type cancellation without modifying the equilibrium agreement. This section defines the property, demonstrates the cancellation, and shows that Nash bargaining loses the property in precisely the staggered settings in which the step-by-step property is most needed. Finally, the section describes the Kalai proportional solution, the only family that

preserves the step-by-step property in general.

3.1 The Step-by-Step Property in Static Games

A bargaining solution satisfies step-by-step if splitting a negotiation into stages — first bargaining over a subset of feasible surplus, and then bargaining over any remaining feasible surplus — yields the same outcome as bargaining over everything at once. This property is roughly equivalent to splitting gains from trade in fixed proportions.

To translate the step-by-step intuition into mathematical terms, view a bargaining solution as a mapping f from feasible outcomes S and a disagreement payoff v^D to a selected outcome $s \in S \cup \{v^D\}$ with $f(S, v^D) \geq v^D$ pointwise. I allow a bargaining solution to satisfy step-by-step for a set of games, where a set of games G is a set of pairs (S, v^D) of feasible outcomes and disagreement values.

Definition 1 (Step-by-step)**.** *A bargaining solution f satisfies step-by-step for a set of games G if for all $(S, v^D) \in G$ and all T for which $f(T, v^D) \in S$, $f(S, f(T, v^D)) = f(S, v^D)$.*

The step-by-step property has previously been understood in static games, where the property is difficult to interpret in practical terms.[2] Kalai (1977) motivates the property through Kissinger's "step-by-step" approach to diplomacy, but cautions that the property is a condition on "feasible utility levels and not directly on the underlying game." After all, why would negotiators commit to bargaining over a restricted set of outcomes before negotiating over the full Pareto frontier?

In dynamic games, step-by-step will have a key practical interpretation. Suppose a pair bargains relative to disagreeing for one period. The value of disagreeing once can be viewed as the outcome of a first-step bargain over the smaller feasible set that commits to disagreeing in the first period.[3] The equilibrium agreement can then be viewed as a second-step bargain

[2] Kalai (1977) and Roth (1979) require S to contain T; most of the results below would follow after this substitution. However, I only require S to contain the first-stage outcome $f(T, v^D)$ in order to relax the conditions when constructing first-stage bargains in the proof of Lemma 1 in Appendix B.1 below.

[3] This logic relies on some implicit assumptions to ensure that a current bargain over the next-period

over any remaining feasible surplus relative to that first-step bargain. By step-by-step, the first-step value of disagreeing once can be replaced with the value of disagreeing twice without changing the equilibrium agreement. That will yield a cancellation of states in Section 3.2.

The step-by-step property will be equivalent to being proportional. To state this equivalence precisely, I first define what it means for a bargaining solution to be proportional. Only Kalai proportional bargaining will be proportional in all games, but a bargaining solution can be proportional for a specific set of games.

Definition 2 (Proportional bargaining solution)**.** *A bargaining solution f is* proportional *for a set of games G if there is a $\tau \in [0,1]$ such that for all $(S, v^D) \in G$ with $f_1(S, v^D) + f_2(S, v^D) > v_1^D + v_2^D$, $\tau = (f_1(S, v^D) - v_1^D)/(f_1(S, v^D) - v_1^D + f_2(S, v^D) - v_2^D)$.*

In bargaining games with $f(S, v^D) = v^D$, there are no gains from trade. An equivalent version of the claim that is less useful constructively, but avoids a reference to positive gains, is to require that in all bargaining games, $\tau(f_2(S, v^D) - v_2^D) = (1 - \tau)(f_1(S, v^D) - v_1^D)$.

I show in Appendix B.1 that under relatively lax conditions, satisfying step-by-step for all games in G is equivalent to being proportional for all games in G. If a bargaining solution is proportional, then one can verify step-by-step by algebra. If a bargaining solution is not proportional, then it is possible to construct a first-step agreement that splits gains in different proportions to the full agreement. Scaling up that first-step agreement eventually produces a second-step problem whose original solution would require one side to accept a worse outcome than the adjusted disagreement point. The formal argument requires the bargaining solution to be homogeneous and translation-invariant, and the set of games to be sufficiently rich to construct such a scaled-up first step agreement. The result relates to results from Kalai (1977) and Roth (1979), but my argument only requires comprehensiveness in one game.

The proportional representation of the step-by-step property helps clarify why step-by-

agreement will produce the same outcome as bargaining over the same agreement next period; I discuss weaker sufficient conditions to apply step-by-step in Section 5.

step will control the state space growth introduced by staggered bargaining.

3.2 State Space Cancellation From Step-by-Step

By applying step-by-step to a recursive bargaining problem, one can replace the value of disagreeing once with the value of disagreeing twice without changing the equilibrium agreement. Inductively applying this property for the duration of an agreement yields a finite dependence-type cancellation of future states.

The step-by-step cancellation requires notation to track the path of outcomes, and the resulting value functions, after the insurer and hospital O disagree once, twice, or d times before reaching the next agreement. Recall from Section 2.2 that in equilibrium, the insurer and hospital O form a $T = 4$-period agreement in period t_0. Write $V_t^{Dd,j}$ for the net present value that player j will receive starting in period t if the insurer and hospital O disagree d times and then reach the resulting subgame's equilibrium agreement. Also write S^{Dd} for the set of feasible period-$t_0 + d$ value functions available after disagreeing d times. It is convenient to write $V_t^{A,j} = V_t^{D0,j}$ for the agreement path; the equilibrium agreement yields value functions $V_{t_0}^{A,M}$ and $V_{t_0}^{A,H}$. In the d-disagreement subgame, the agreement $V_{t_0+d}^{Dd}$ is recursively chosen through:

$$V_{t_0+d}^{Dd} = f\left(S^{Dd}, (3G - 3p_{E,t_0+d}, 0) + \beta V_{t_0+d+1}^{Dd+1}\right),$$

where f is the bargaining solution that the insurer and hospital O follow, and p_{E,t_0+d} is the price that the sides anticipate hospital E will reach in the subgame that follows d disagreements. The disagreement point has two components: a flow payoff in the initial period in which the insurer operates with only hospital E in its network, where the insurer receives $3G - 3p_{E,t_0+d}$ and hospital O earns nothing, plus the discounted continuation value $\beta V_{t_0+d+1}^{Dd+1}$ from bargaining again next period.

If the bargaining solution satisfies step-by-step, then replacing the value of disagreeing

once with the value of disagreeing twice does not change the equilibrium outcome. Recall that by Lemma 1, if f satisfies step-by-step, then there is a weight $\tau \in [0,1]$ such that at every opportunity to bargain,

$$\tau\left(V_{t_0+d}^{Dd,H} - \beta V_{t_0+d+1}^{Dd+1,H}\right) = (1-\tau)\left(V_{t_0+d}^{Dd,M} - (3G - 3p_{E,t_0+d}) - \beta V_{t_0+d+1}^{Dd+1,M}\right).$$

Applying this equality for $d = 0$ and $d = 1$, I obtain by addition that for the equilibrium agreement,

$$\tau\left(V_{t_0}^{A,H} - \beta V_{t_0+1}^{D1,H}\right) = (1-\tau)\left(V_{t_0}^{A,M} - (3G - 3p_{E,t_0}) - \beta V_{t_0+1}^{D1,M}\right) \tag{5}$$

$$\tau\left(V_{t_0+1}^{D1,H} - \beta V_{t_0+2}^{D2,H}\right) = (1-\tau)\left(V_{t_0+1}^{D1,M} - (3G - 3p_{E,t_0+1}) - \beta V_{t_0+2}^{D2,M}\right). \tag{6}$$

Scaling Equation (6) by β and adding the result to Equation (5) yields:

$$\tau\left(V_{t_0}^{A,H} - \beta^2 V_{t_0+2}^{D2,H}\right) = (1-\tau)\left(V_{t_0}^{A,M} - (3G - 3p_{E,t_0}) - \beta(3G - 3p_{E,t_0+1}) - \beta^2 V_{t_0+2}^{D2,M}\right). \tag{7}$$

One state has been removed. In order to link this result to step-by-step, interpret the value of disagreeing for one period as the value of a first-stage bargain next period relative to disagreeing for a second period; by step-by-step (in combination with shared time preferences and beliefs), replacing the value of disagreeing once with the value of disagreeing twice does not change the equilibrium agreement.

The step-by-step property is most powerful when applied inductively for the duration of an agreement. The hospital-O agreement remains in place for T periods, so apply step-by-step T times. The result is

$$\tau\left(V_{t_0}^{A,H} - \beta^T V_{t_0+T}^{DT,H}\right) = (1-\tau)\left(V_{t_0}^{A,M} - \sum_{d=0}^{T-1} \beta^d (3G - 3p_{E,t_0+d}) - \beta^T V_{t_0+T}^{DT,M}\right).$$

That is, the recursive disagreement payoff has been replaced by the expected flow of payoffs over T periods of disagreement, plus a continuation value from returning to the bargaining table in period $t_0 + T$.

This repeated step-by-step application yields a finite dependence-type cancellation. The $T - 1$ states that emerge to capture the value of returning to the bargaining table at an off-equilibrium time are canceled via step-by-step. The continuation value at $t_0 + T$ only involves the next on-equilibrium bargaining state. The researcher already must track the continuation-value state on the equilibrium path, so no new states are introduced. Further, this cancellation applies under uncertainty, so long as negotiators share expectations and a share time preferences.

While the step-by-step property substantially simplifies staggered bargaining, the bilateral intuition can mask the remaining complexity introduced by the equilibrium order of multilateral contracting. The gains from trade in 2025 are evaluated based on the series of p_{E,t_0+d} values. If the insurer and hospital O actually disagree for T periods, then it is possible that the hospital-E negotiation in 2026 will lead to a different price based on the anticipated off-equilibrium outcome for hospital O. As a result, the researcher still needs to model at least one off-equilibrium bargaining state. Put differently, step-by-step will control the state space growth introduced by staggering, but not the state space growth introduced by the ordering of equilibrium agreements. I return to this point later in Section 4.2 below. Still, step-by-step will represent a crucial advantage in modeling staggered bargaining, but the prevailing Nash bargaining solution will lose this property precisely when contracts are staggered.

3.3 Nash Bargaining Loses Step-by-Step Under Staggering

Nash bargaining satisfies step-by-step only under homotheticity. The internalized contracting externalities introduced by staggering are time-varying, so that homotheticity fails and Nash bargaining loses the step-by-step property in precisely the cases in which the

property is most needed.

The Nash bargaining solution is illustrated in Figure 1(a). As a reminder, the Nash bargaining solution with player-1 weight $\tau \in [0,1]$ chooses an outcome f to solve $f(S, v^D) = \text{argmax}_{s \geq v^D \in S}(s_1 - v_1^D)^\tau (s_2 - v_2^D)^{1-\tau}$. When the Pareto frontier of $s \in S$ can be parameterized as a differentiable function $s(p)$ for contract characteristics p, then first-order conditions on p yield $f(S, v^D) = s(p^*)$, where the Nash-bargained characteristics p^* satisfy $\tau s_1'(p^*)(s_2(p^*) - v_2^D) = -(1-\tau)s_2'(p)(s_1(p^*) - v_1^D)$. If $f(S, v^D) \neq v^D$ and $\tau s_1'(p^*) \neq (1-\tau)s_2'(p^*)$,[4] then the solution concept predicts

$$(s_1(p^*) - v_1^D)/(S_1(p^*) - v_1^D + S_2(p^*) - v_2^D) = \tau s_1'(p^*)/(\tau s_1'(p^*) - (1-\tau)s_2'(p^*)). \quad (8)$$

Nash bargaining will satisfy step-by-step when Pareto frontiers are homothetic. By the proportionality argument in Section 3.1 and inspection of Equation (8), Nash bargaining will satisfy step-by-step if for all achievable (S, v^D) with $f(S, v^D) \neq v^D$, $(s_1(p^*) - v_1^D)/(s_1(p^*) - v_1^D + s_2(p^*) - s_2^D)$ is constant. This requirement generally corresponds to homothetic Pareto frontiers (Coles and Muthoo, 2003) in which every pair of games $(S, v^D), (T, u^D) \in G$ that a model generates satisfy $S - v^D = \gamma(T - u^D)$ for some $\gamma \geq 0$. Under homotheticity, Nash bargaining will predict the same surplus share in every game. If a recursive bargaining model only generates homothetic Pareto frontiers, then Nash bargaining will satisfy step-by-step, and the resulting bargaining problem can be simplified accordingly. A leading case is transferable utility (TU) games (Lee and Fong, 2013; Collard-Wexler et al., 2019; Ho and Lee, 2017; Crawford et al., 2018), where the negotiated price has a zero-sum effect on profits for the players involved. The TU case is also a leading case in the dynamic search-on-the-job bargaining literature, where models often assume that workers can fully leverage their employers' surplus when leaving, so that the initial wage does not have a nontransferable

[4]The case $\tau s_1'(p^*) = (1-\tau)s_2'(p^*)$ is only achievable when both are zero, and corresponds to take-it-or-leave-it offers where the no-bargaining-power player i having $s_i'(p^*)$ but with positive gains $s_i(p^*) - s_i^D > 0$. In those cases, step-by-step must be verified from the aggregate gains and not from the Nash first-order condition.

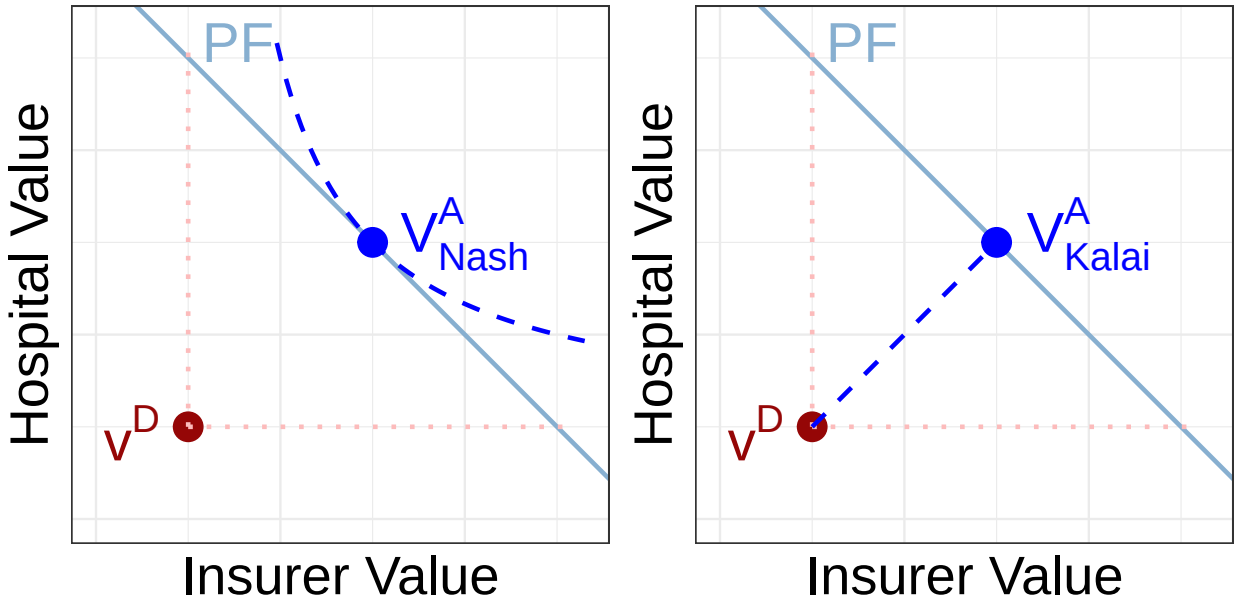


Figure 1: Left: the Nash bargaining solution to the same game. The dashed line indicates the Nash bargaining indifference curve, and the resulting agreement is indicated by v^A_{Nash}. Right: the Kalai proportional bargaining solution for bargaining between an insurer (x axis) and hospital (y axis) with Pareto frontier PF in light blue and disagreement value v^D in red, and Pareto constrained indicated by dotted lines. The desired-surplus-split ray $v = v^D + x(1-\tau)/\tau$ is indicated by the dashed blue ray, and the resulting agreement is indicated by v^A_{Kalai}.

utility (NTU) effect on post-departure wages (Cahuc et al., 2006; Bilal et al., 2022; Jarosch et al., 2024).

Unfortunately, Nash bargaining will fail step-by-step in precisely the staggered contracting environments that motivated step-by-step. To see why, recall the odd-and-even two-year-contract example of Section 2.2. Suppose that in the equilibrium alternating-year bargaining state, the equilibrium price response has a slope of $\lambda_1 > 0$. When bargaining in alternating years, an increase to the initial hospital-h payment $p_{t,h}$ will lead the insurer to pay the other hospital an additional λ_1 in period $t+1$, tilting the Pareto frontier away from TU.[5] No

[5] After working through the full process, the marginal value ratio turns out to be $-1/(1+\beta\lambda_1)$. To verify this, note that a marginal increase to the $t+1$ price formed with hospital h increases the price formed with hospital $-h$ in period $t+1$ by λ_1, increases the $(h, t+2)$ price by λ_1^2, increases the $(-h, t+3)$ price by λ_1^3, and so on. By collecting terms, a marginal current price increase with hospital h increases the insurer's spending proportionally to $\sum_{t=0}^{\infty}(\lambda_1\beta)^t$, but only increases the hospital's payments received proportionally to $\sum_{t=0}^{\infty}(\lambda_1\beta)^{2t}$.

such internalized price externality exists in the off-equilibrium simultaneous bargaining state. As a result, Nash fails step-by-step. More broadly, under staggered agreements, the slope of the Pareto frontier depends on internalized contracting externalities that are inherently time-dependent as future bargains come closer in time and then are passed.

Nash's step-by-step property can still be recovered if contracts are simultaneously bargained in every state, as is the usual practice that Section 2.1 argues can introduce bias, or if negotiators are myopic and do not value future leverage effects, ruling out forward-looking responses from the start. Section 6 finds that both options would be wrong in hospital–insurer contracting: contracts are pervasively staggered and I estimate that firms are substantially forward-looking.

One might hope to find some other bargaining solution that preserves step-by-step in a specific staggered model. But staggering introduces a time-varying nonhomotheticity that makes case-by-case verification circular: the shape of gains from trade depends on the bargaining solution itself, so that verifying step-by-step in a given model requires first solving the recursive problem that step-by-step is needed to simplify. A more viable path is to follow a bargaining solution that satisfies step-by-step in every bargaining model, and not just the model at hand. There is only one such family of bargaining solutions.

3.4 Kalai Proportional Bargaining Preserves Step-by-Step in General

Kalai proportional bargaining selects the best feasible agreement along the ray that splits surplus in proportion to a set of fixed bargaining weights. The step-by-step property's first-step agreement corresponds to moving the value of disagreement upwards along the same ray, and therefore does not change the equilibrium agreement.

The Kalai (1977) proportional solution f_τ^{Kalai} with player-1 weight $\tau > 0$ and for S is

defined as:

$$f_{\tau,1}^{(Kalai)}(S, v^D) = \sup s_1 \quad : \quad s \in S \cup \{v^D\},\ s \geq v^D,\ \tau(s_2 - v_2^D) = (1-\tau)(s_1 - v_1^D)$$
$$f_{\tau,2}^{(Kalai)}(S, v^D) = (1-\tau)/\tau f_{\tau,1}^{(Kalai)}(S, v^D).$$

The bargaining solution with player-1 weight $\tau = 0$ is defined as $f_{\tau,1}^{(Kalai)}(S, v^D) = 0$, $f_{\tau,2}^{(Kalai)}(S, v^D) = \sup s_2 : s \in S \cup \{v^D\}, s_1 = v_1^D, s_2 \geq v_2^D$.

The Kalai proportional solution for bargaining between a hospital (player 2) and insurer (player 1) is illustrated in Figure 1(b). Draw a dashed ray to indicate the values $v \in V$ that yield positive gains from trade and the desired surplus split $\tau(v_2 - v_2^D) = (1-\tau)(v_1 - v_1^D)$. The solution V_{Kalai}^A is the best agreement along this desired-surplus-split ray.

Intuitively, Kalai proportional bargaining differs from Nash bargaining in how the solution concepts reflect marginal and aggregate gains from trade. Both Nash and Kalai proportional bargaining attempt to maximize gains, and both include a bargaining weight that captures relative surplus-splitting. To see the remaining difference, imagine an initial game in which the two solution concepts agree. Now imagine tilting the Pareto frontier more vertically, so that at the old agreement, marginal gains for player 2 come at a smaller cost to player 1. Nash bargaining will predict to choose an agreement that is more favorable to player 2 and less favorable to player 1, while Kalai proportional bargaining will predict the old agreement.

Which stance on marginal values is more accurate is context-dependent, but the Kalai stance is the right stance for recursive bargaining tractability. Under Nash bargaining, disagreement changes the internalized contracting externalities faced by a pair as leverage effects become more and then less important. As a result, repeated disagreement under staggering changes the shape of the Pareto frontier, changing the ratio of marginal values and therefore the surplus-split proportions predicted by Nash bargaining. Kalai proportional bargaining avoids this recursive dependence by selecting gains based on a fixed-proportion ray applied to aggregate gains rather than responding to marginal values at the frontier.

As a result, tilting the Pareto frontier does not change the Kalai proportional bargaining prediction, and the recursive complexity disappears.

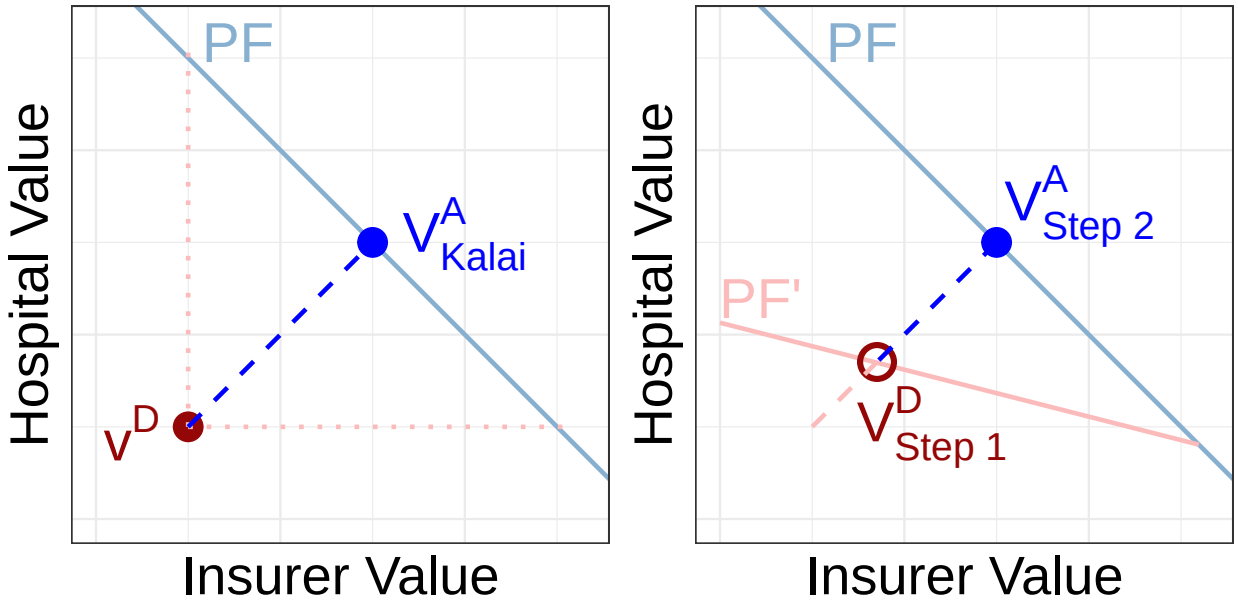


Figure 2: Left: The Kalai proportional bargaining solution for choosing an agreement V^A_{Kalai} (blue) on a Pareto frontier PF relative to the value of disagreement v^D (red), with proportional split illustrated by dashed blue line. Right: illustrating the step-by-step property for a final agreement $V^A_{\text{Step 2}}$ (blue) recursively defined through bargaining relative to a first-step agreement $V^D_{\text{Step 1}}$ (red), chosen through applying the bargaining solution to negotiations over some smaller Pareto frontier PF' (pink, with the dashed pink line indicating the proportional split) relative to the full disagreement point v^D (omitted).

Figure 2 illustrates why Kalai proportional bargaining satisfies the step-by-step property. Recall the intuition of the Kalai proportional solution as choosing the best agreement along a surplus-splitting ray. In Figure 2(b), instead of negotiating over the full agreement, the sides first negotiate over a smaller Pareto frontier (pink) to reach a first-step agreement. The first-step agreement then becomes the disagreement point in the second-step bargain over the full surplus, leading to a final agreement. This recursive process generates the same outcome as one-shot bargaining over the full surplus Figure 2(a), because the first-step agreement corresponds to moving the disagreement value upwards along the same surplus-splitting ray.

In math, if $f(T, v^D) \leq f(S, v^D)$, then the set of feasible S in the $(S, f(T, v^D))$ bargaining game is the set of $s \in S \cup \{f(T, v^D)\}$ satisfying $s \geq v^D$ and $\tau(s_2 - f_2(T, v^D)) = (1-\tau)(s_1 -$

$f_1(T, v^D)$). But by construction, $\tau(f_2(T, v^D) - v_2^D) = (1-\tau)(f_1(T, v^D) - v_1^D)$. Therefore the remaining elements of s are a subset of the feasible elements of the (S, v^D) game. $f(S, v^D)$ is feasible, so by inspection, $f(S, f(T, v^D)) = f(S, v^D)$.

The Kalai proportional solution can be justified from axiomatic, empirical, and microfoundational perspectives. Among bargaining solutions that are homogeneous (so that negotiating in dollars and cents produce the same underlying outcome) and weakly Pareto-optimal in comprehensive games (so that the solution concept does not leave gains from trade on the table in standard games), only Kalai proportional bargaining satisfies step-by-step in general. Further, the solution concept is the unique family of bargaining solutions satisfying weak Pareto-optimality, independence of irrelevant alternatives, strong individual rationality, and a monotonicity axiom that Pareto frontier expansions should be weakly Pareto-improving (Kalai, 1977; Roth, 1979). Empirically, Kalai proportional bargaining fits the outcomes of laboratory games more closely than Nash (Nydegger and Owen, 1974; Duffy et al., 2021). Structurally, the solution concept can be microfounded as the outcome of multi-stage arbitration (Bossert, 1994), a revocations-cost game (Dutta, 2012, 2024), or small-bundle bargaining under liquidity constraints (Hu and Rocheteau, 2020). In Section 4.3, I offer a microfoundation for multilateral Nash-in-Kalai bargaining in terms of delegate revocation costs, with bargaining weights as the ratio of first revocation costs.

Before applying the Kalai proportional solution to generic multilateral games, it is important to note the solution concept's leading limitations. Kalai proportional is scale-varying, so that changing player 1's scale of utility changes player 2's outcome. This is empirically surmountable when payoffs are defined in monetary terms; changing player 1's scale from dollars to cents can be undone by changing the bargaining weight. However, scale-variance means that Kalai proportional bargaining is inappropriate in settings that lack some scale in which to compare relative utilities. Further, scale-variance means that Kalai proportional bargaining cannot be microfounded by von Neumann-Morganstern utility alone, since utility levels cannot be meaningful; the microfoundation in Section 4.3 uses delegate revocation

costs to fix relative utility scales. Further, in games in which one player's most favored offer gives the other player positive gains, then the favored offer corresponds to interior proportions. When an important part of upstream price bargaining is the downstream quantity consequence, then the Pareto frontier may fail comprehensiveness and turns in on itself, and Kalai proportional bargaining may no longer be able to capture supplier take-it-or-leave-it offers with a bargaining weight of one.

This paper does not, and should not, argue that Kalai proportional bargaining is a superior solution concept in the abstract; when Nash and Kalai diverge, the choice should be driven by context and data. This paper establishes that Kalai is the unique solution concept that controls state space growth introduced by staggering, and that the associated solution concept can be justified from revocation costs. When bargaining remains close to TU, the step-by-step property can be applied as either a maintained assumption or an approximation, as Nash and Kalai will nearly coincide. When bargaining is substantially NTU, the two models diverge due to the dependence of Nash bargaining shares on marginal values. Appendix B.3 develops intuitive conditions under which Nash and Kalai will deliver similar predictions either immediately or after appropriate reindexing of bargaining weights (when bargaining is NTU but close to homothetic), and clarifies the degree to which dynamic Nash-in-Nash becomes intractable when the approximation argument does not apply.

4 A Tractable Model of Staggered Bargaining

This section formalizes the arguments of Section 3 in a general framework of forward-looking multilateral bargaining. I first extend previous Nash-in-Nash models to allow for staggering and arbitrary bargaining solutions, a family that I call Nash-in-f. I characterize state space growth in the general framework and show that when the bargaining solution aligns with Kalai proportional bargaining, a case that I call Nash-in-Kalai, the step-by-step property yields a substantially tighter bound on state space growth. I close by offering a

microfoundation for Nash-in-Kalai in terms of delegate revocation costs. Researchers mainly interested in estimation may wish to skip to Section 5, which develops the Nash-in-Kalai moment using the intuition from Sections 2 and 3.

4.1 Potential Recursive Complexity

In a multilateral Nash-in-Nash model with staggered contracts, the internalized price externalities from Section 2.2 mean that every remaining contract duration enters the bargaining state. In a market with N agreements that remain in place for T periods, this results in up to T^N relevant market configurations.

I will offer a framework for repeated bargaining with a generic bargaining solution. The solution concept will be a Nash equilibrium in some forward-looking bargaining solution. I call this generic framework Nash-in-f, for a Nash equilibrium in the bargaining solution f; Nash-in-Nash and Nash-in-Kalai will be specific cases.

I study dynamic markets of the following form.

There are n players, indexed $i = 1, \ldots, n$. The world is infinitely-lived with periods indexed by t. The per-period discount factor is $\beta \in [0, 1)$. Myopia corresponding to $\beta = 0$.

The timing in each period t is as follows:

1. *Potential pre-bargaining actions are taken.* Given the start-of-period state s_t, each player i chooses actions $a_{it}^{(1)}$ in some action space $A_{it}^{(1)}(s_t)$. I write the resulting vector of actions as $a_t^{(1)}$.

2. *Bargaining and potential simultaneous actions are taken.* The actions from stage 1 are announced to all players. Then, at the same time:

   (a) For $i < j$ pairs with more than one feasible contract, new contracts $\mathbb{C}_{ijt}$ in the feasible new-contract spaces $\mathcal{C}_{ijt}\left(a_t^{(1)}\right)$ are chosen through bargaining relative to the null contract $\mathbb{C}_{0,ijt}$. I write the resulting contract state (including which

agreements are newly-formed) as $\mathbb{C}_t$, and the number of new contracts formed by firm i as R_{it}.

(b) Each player i chooses actions $a_{it}^{(2b)}$ in some action space $A_{it}^{(2b)}(s_t)$. I write the resulting vector of actions as $a_t = \left((a_{1t}^{(1)}, a_{1t}^{(2b)}), \ldots\right)$.

3. *Flow profits are formed.* Player i earns flow profits of

$$u_i\left(\mathbb{C}_t, a_t, x_t, \varepsilon_t\right) + \sum_{j \neq i} \text{Pay}_{ji}\left(\mathbb{C}_t, a_t, x_t, \varepsilon_t\right).$$

4. *Information is revealed.* The new state s_{t+1} is revealed to all players. This state includes the chosen action and contract states a_t and $\mathbb{C}_t$, which are observed by the econometrician; s_t; the next period's non-bargaining action spaces A_{it+1}, null contracts $\mathbb{C}_{0,ijt+1}$, contract space functions $\mathcal{C}_{ijt+1}(a_t^{(1)})$ (either a single renewal contract or a set including the null contract), and other observable states in x_{t+1}, which is observed by the econometrician; and the unobservable states $\varepsilon_{t+1}, i_{t+1}$.

Within such a game, strategies are a set of functions $\sigma = \left\{\sigma_i^{(1)}, \sigma_{ij}^{(2a)}, \sigma_i^{(2b)}\right\}$ such that $\sigma_i^{(1)}\left(s_t\right) \in A_{it}^{(1)}(s_t)$, $\sigma_{ij}^{(2a)}(a_t^{(1)}, s_t) \in \mathcal{C}_{ijt}\left(a_t^{(1)}\right)$, and $\sigma_i^{(2b)}\left(a_t^{(1)}, s_t\right) \in A_{it}^{(2b)}(a_t^{(1)}, s_t)$.

Definition 3 (Nash-in-f equilibrium)**.** *A* Nash-in-f equilibrium *is a game and strategies σ of the form of this section such that (i) value functions are calculated under these strategies (see Appendix C.1 for formal definition), (ii) the strategies $\sigma_i^{(1)}$ and $\sigma_i^{(2b)}$ satisfy the appropriate notion of equilibrium induced by this game and the value functions generated by these strategies, and (iii) the players follow the class of bargaining solutions f in the sense that there are a set of bargaining solutions f_{ij} in the class f such that for every feasible pair of $a^{(1)}, s$, the strategy $\sigma_{ij}^{(2a)}(a^{(1)}, s)$ taking other strategies as given and the time-of-bargaining values functions $V_i(\mathbb{C}_{ij}, a^{(1)}, s)$ generate feasible outcomes and disagreement values $S_{ij}\left(a^{(1)}, s\right) = \left\{\left(V_i(\mathbb{C}, a_{-i}, a^{(1)}, s), V_j(\mathbb{C}, a^{(1)}, s)\right)\right\}_{\mathbb{C} \in \mathcal{C}_{ijt}\left(a^{(1)}\right)}$ and disagreement values*

$v_{ij}^D\big(a^{(1)}, s\big) = \big(V_i(\mathbb{C}_{0,ijt}, a^{(1)}, s), V_j(\mathbb{C}_{0,ijt}, a^{(1)}, s)\big)$ *that satisfy*

$$\sigma_{ij}^{(2a)}(a^{(1)}, s) = f_{ij}\left(S_{ij}\big(a^{(1)}, s\big), v_{ij}^D\big(a^{(1)}, s\big)\right).$$

For generality, I allow the σ_i strategies to follow notions of equilibrium besides Markov-perfect. As a result, the notion of equilibrium is broad enough to include punishment for disagreement. In my empirical application, I assume that disagreement is short-lived.

I discuss each stage in further detail in Appendix C.2.

The full state space can reflect non-Markov strategies and non-bargaining actions, so to focus on the state space growth from staggering, consider a natural generalization of the one-insurer, two-hospital model of Section 2.2. There are now R retailers and S suppliers, generating $N = SR$ bilateral agreements; in Section 2.2, $N = 2$. I make several assumptions to control state space growth: there are no underlying market dynamics for players to track; each agreement remains in place for T periods, so there is no state space growth from different contracting frequencies; and price responses are linear, so that the econometrician only has to model length-dependent slopes.[6] I make one assumption in the interest of realism: if a pair disagrees, then in the resulting subgame, they reach a new agreement in the next period. Multiperiod exclusion can control state space growth, but it requires a pair to commit to forgoing the joint profits and leverage that motivated their agreement, and in many settings will not be a credible threat.

In the N-pair, T-period version of the illustrating example from Section 2.2, the number of relevant market configurations grows as T^N and the number of price-responses grows as NT^{N-1}. The former follows by the T^N combinations of pairwise disagreement that can emerge. The price-response bound follows by fixing one pair ij. The argument of Section 2.2

[6]To be more formal, I assume that the subgame price-response strategies for a pair ij in response to the anticipated other-pair contract outcome $\mathbb{C}_{-ij}$ follow the linear Markov form $p_{ij}(\mathbb{C}_{-ij}) = p_{ij}^*(\{\ell\}) + \sum_{hk \neq ij} \lambda_{ij|hk}(\{\ell\})(p_{hk} - p_{hk}^{(}\{\ell\}))$, where $\{\ell\}$ is the anticipated remaining periods (including this period) for each hk pair and $p_{hk}^*(\{\ell\})$ is a long-run outcome that the hk pair will reach in the long-run if ij bargain at this set of contract lengths remaining, and where the long-run outcome can be expressed as a function of the known underlying gains from trade and the unknown length-dependent slopes.

shows that the way ij bargain depends on when the pair expects every other pair to next negotiate: the unique combination of remaining lengths dictates the effect of the currently negotiated price on every subsequent non-ij price, as well as how subsequent non-ij price change will then affect the next ij price negotiation. This yields T^{N-1} states in the ij bargain: every state that could be reached by a combination of agreements and disagreements. There are N pairs, yielding NT^{N-1} relevant price-response functions.

To put the growth in concrete terms, imagine that all S suppliers negotiate bilaterally with all R retailers over two-year contracts. The illustrating example controlled state space growth by restricting to two hospitals ($S = 2$) and one insurer ($R = 1$) with two-year agreements in an annual model, so that there are at most 4 price-response functions. But now suppose there are three hospitals and three insurers, so that $N = 9$, and suppose the players form five-year agreements. As outlined in Section 2 that there would be 3.5 million functions in an annual model, 900 million functions in an every-six-month model, and 1.6×10^{15} functions in a monthly model. The econometrician can only hope to defer explosive state space growth by assuming symmetry: in the extreme case with $N = 9$ symmetric agreements, there are 495 functions in the annual model, $24,310$ functions in the every-six-month model, and 6.5 billion functions in the monthly model.[7]

The step-by-step property will substantially control this state space growth. As noted in Section 3.4, Kalai proportional bargaining is the only solution that preserves the step-by-step property in general.

4.2 State Space Control Under Step-by-Step

When every pair follows a fixed Kalai proportional solution in every bargaining subgame, then the step-by-step property from Section 3 applies to every bargaining, allowing the econometrician to bypass the explosive state space growth introduced by staggering. I call

[7] To derive these figures, the number of weakly-ordered assignments of $N-1$ period lengths from 1 through T is $\binom{N+T-2}{N-1}$ by a stars-and-bars argument.

such a model a Nash-in-Kalai model for a Nash equilibrium in Kalai proportional bargains. In a market with N agreements that remain in place for T periods, a Nash-in-Kalai model results in up to N^N market configurations under fully staggered timing, and reduces to as few as N states if all contracts form simultaneously in equilibrium.

A Nash-in-Kalai model is a Nash-in-f model in which all pairs follow a fixed Kalai proportional solution.

Definition 4 (Nash-in-Kalai equilibrium)**.** *A* Nash-in-Kalai *equilibrium is a Nash-in-f equilibrium in which for every ij, there is a $\tau_{ij} \in [0,1]$ such that for every feasible pair of $a^{(1)}, s$, the strategy $\sigma_{ij}^{(2a)}(a^{(1)}, s)$ satisfy*

$$\sigma_{ij}^{(2a)}(a^{(1)}, s) = f_{\tau_{ij}}^{(Kalai)}\left(S_{ij}\left(a^{(1)}, s\right), v_{ij}^{D}\left(a^{(1)}, s\right)\right).$$

In this definition, I sidestep any issues that could be introduced by noncomprehensiveness. If a model is noncomprehensive, then it could be that there is a subgame and a τ such that every feasible contract that splits gains proportionally to τ is Pareto-inefficient. I implicitly assume that the Nash-in-Kalai weights τ_{ij} and game ensure that no agreement would be Pareto-dominated.

I characterize state space growth in a restricted context to focus on the role of staggered timing.

Proposition 1 (Step-by-step state space control)**.** *Consider a Nash-in-Kalai equilibrium in which N bilateral pairs form T-period contracts, and regularity conditions that prevent state space growth from non-contracting elements hold (formalized in Assumption 1 in Appendix B.2). There is an $z \leq \min\{T, N\}$ and a Nash-in-Kalai equilibrium (other than possible non-equilibrium behavior at periods with no contracting in equilibrium) with at most z^N market configurations, at most $N \min\{z^N, T^{N-1}\}$ price-response functions, and the same equilibrium price-response functions as the original model.*

Importantly, the state space control from step-by-step does not require changes to re-

searchers' estimating equations. In previous TU bargaining models, Nash-in-Kalai immediately reduces to Nash-in-Nash, so that a researcher who has specified a TU Nash-in-Nash model can apply the step-by-step cancellation without modifying the bilateral bargaining solution. More generally, I establish in Corollary 1 below that the Nash-in-Kalai NPV payment reduces to the expected NPV of period-by-period TU Nash-in-Nash flow payments, with a dynamic adjustment that is zero in steady state and otherwise intuitively small. I follow this approach in my empirical application, where the Nash-in-Kalai moment requires only slight modification from the previous static estimation strategy.

The number z is the number of discrete periods at which contracts are formed within a given T-period interval: at most the number of players if $N \leq T$, at most the number of periods if $T \leq N$, and strictly less than T and N if there are periods with no contracting in equilibrium. By step-by-step, the Nash-in-Kalai model can remove every off-equilibrium state at which no pair bargains in equilibrium, leaving z potentially relevant periods within each T-period interval. Each remaining pair may then negotiate at any of the z potentially relevant periods, leaving z^N market configurations. The non-equilibrium caveat is that the equivalent model will require players to disagree in periods in which no players negotiate in equilibrium; this may violate rationality requirements of the Nash-in-Kalai model. The leading simplifying restriction of Assumption 1 is that there are no non-bargaining actions or payoff-relevant states; actions like price-setting can be accommodated by increasing the number of discrete periods to include those actions, and payoff-relevant states directly expand the relevant state space. The key regularity condition on behavior in Proposition 1 is that disagreement is short-lived. Short-lived disagreement ensures that no states are needed to capture the effect of a pair's post-disagreement exclusion on subsequent bargains.

The gains from step-by-step are largest when equilibrium contract timing is sparse. Recall from Section 4.1 that a Nash-in-f model will have NT^{N-1} price-response functions, so that there are 3.5 million states in an annual model, 900 million in an every-six-month model, and 1.6 quadrillion in a monthly model. If the N firms each negotiate in the same period, then

Proposition 1 shows that the econometrician can capture even very short off-equilibrium time periods, or even infinitesimally short exclusion, with only $N = 9$ price response functions. If each pair negotiates in a different period, then Proposition 1 still shows that the necessary state space can never require more than $9^{10} = 3.5$ billion price-response functions: large, but immune to explosive growth in T. Further, Section 5.3 below argues that even when the equilibrium contract timing is not sparse so that a full Nash-in-Kalai model will have a large number of price responses, bounded rationality offers a feasible and principled way to approximate the full process with many fewer states.

4.3 Microfoundation

The previous subsection shows that the Kalai proportional solution controls state space growth. The remaining theoretical question is what noncooperative behavior would give rise to Nash-in-Kalai outcomes. I offer a demands game adapted from Dutta (2012) that microfounds Nash-in-Kalai in staggered NTU settings, which compares favorably to the leading Nash-in-Nash microfoundation (Collard-Wexler et al., 2019). Researchers mainly interested in estimation may wish to skip to Section 5.

The Nash-in-Kalai microfoundation uses revocation costs to fix relative scales. Recall from Section 3.4 that Kalai proportional bargaining is scale-varying, so it cannot be microfounded based on von Neumann-Morganstern utility alone. I follow Dutta (2012) and microfound the solution based on revocation costs, where firms negotiate via delegates; the delegates negotiate after making public contract price commitments, and then face concession costs for prices lost relative to the public commitment. This structure of concession costs fixes relative firm value functions in terms of relative per-dollar-price revocation costs, enabling a microfoundation for the scale-varying Kalai proportional solution concept. The same idea would hold if revocation costs were expressed per payment lost; the key is that concession costs must be expressed on the same units, so that comparing revocation costs corresponds to comparing firm value functions. I make a stronger simplifying assumption

that revocation costs should have an infinite first-derivative, which ensures that revocation costs dominate value function costs for small price deviations.[8]

In the microfoundation, as the period length tends to zero, the revocations-cost equilibrium tends to a discretely-timed Nash-in-Kalai solution. In the equilibrium, every pair of negotiators demands the same agreement: demanding a better agreement would introduce a revocation cost that dominates any profit gained, and demanding a worse agreement would introduce a revocation cost while losing profit for the employer. If the first revocation costs tend to a fixed bilateral ratio $t_{ij} = w_i/w_j$, then the bilateral solution tends to a Kalai proportional solution with ij bargaining weight solving $\tau_{ij} = w_i/(w_i + w_j)$. I then apply the step-by-step property to show that the limit of this continuous-time revocations cost game is a discrete-time Nash-in-Kalai outcome. I formalize this microfoundation in Appendix C.3.

The leading Nash-in-Nash microfoundation (Collard-Wexler et al., 2019) produces Nash outcomes as the limit of frequent alternating offers. However, the Collard-Wexler et al. result only covers bargaining over lump-sum transfers with simultaneous equilibrium contract formation, so that Pareto frontiers are TU. As a result, the Collard-Wexler et al. (2019) model can also be described, with tongue somewhat in cheek, as a microfoundation of Nash-in-Kalai. The Nash-in-Kalai microfoundation I offer also covers the case that motivated this paper, staggered bargaining with NTU Pareto frontiers, so that the Nash-in-Kalai model can also be described, with tongue firmly in cheek, as strictly more microfounded than Nash-in-Nash.

The Nash-in-Kalai microfoundation requires firms to negotiate via delegates who face concession costs, whereas the Collard-Wexler et al. microfoundation is based on alternating offers from internally aligned firms. A natural question is whether a Rubinstein (1982)-type alternating offers model could also microfound Nash-in-Nash in a staggered NTU case. Coles and Muthoo (2003) show that an alternating offers model only microfounds Nash under homotheticity; Section 3.3 argues that staggered bargaining inherently introduces nonhomo-

[8] Dutta (2012) makes the slightly weaker assumption that the sequence of right derivatives tends to infinity, rather than being infinite from the jump. A discussion of alternative revocation cost structures that could also yield a Nash-in-Kalai outcome is outside the scope of this work.

theticity, so that alternating offers tends to a differential equation at the margin rather than an aggregate surplus-splitting rule. As a result, extending the Collard-Wexler et al. (2019) alternating-offers game to a staggered NTU setting would reintroduce the recursive complexity that motivated this paper.

5 Taking the Model to Data

The Nash-in-Kalai model will yield a tractable moment for GMM estimation. The expected NPV payment can be expressed in closed-form in terms of the expected NPV of period-by-period TU Nash-in-Nash flow payments plus an adjustment for forward-looking dynamics, so that previous TU estimation strategies carry over with only slight modification. The adjustment is zero in steady state, and can be captured through a bounded rationality approximation that iteratively builds in higher-order effects. I close the section with a discussion of important considerations before taking the model to data.

I express the Nash-in-Kalai moment in terms of the flow payments that would emerge under period-by-period TU Nash-in-Nash bargaining. Recall that each player i's flow profits can be decomposed into a pre-transfer component u_{it} (capturing production costs, downstream revenue, and so on) and a net transfer $\text{Pay}_{it} = \sum_{j \neq i} \text{Pay}_{jit}$. (I formalize these and other moment-relevant concepts in Appendix C.4.) Under period-by-period TU Nash bargaining, the j-to-i payment needed to split joint gains proportionally to τ_{ij} bargaining weights is:

$$\text{Pay}_{NiN,ijt} = -\tau_{ij}\left([\Delta_{ij}u_{it}] + [\Delta_{ij}\text{Pay}_{it,-j}]\right) + (1-\tau_{ij})\left([\Delta_{ij}u_{jt}] + [\Delta_{ij}\text{Pay}_{jt,-i}]\right), \quad (9)$$

where $[\Delta_{ij}Z_t]$ denotes the effect of the ij agreement on the variable Z_t under equilibrium actions and non-ij contracts, and where $\text{Pay}_{it,-j} = \text{Pay}_{it} - \text{Pay}_{jit}$ collects non-ij payments.[9]

[9]This notation is adapted from Ho and Lee (2017): $[\Delta_{ij}Z_t]$ denotes the effect of the ij agreement on the variable Z under the maintained assumption that non-$\mathbb{C}_{ijt}$ actions and states are unchanged. In the appendix, I also add notation to accommodate cross-period dynamics, where $[\Delta_{ijt_0}Z_t^D]$ denotes how the variable Z in period t is affected if non-ij actions and states follow the path in which ij reach an agreement and then enter impasse in t_0 at the next equilibrium bargain, relative to ij entering immediate impasse, and

5.1 Closed-Form Moment Under Staggered Contracting

The Nash-in-Kalai moment adds a dynamic adjustment term to correct for the forward-looking effects of disagreement.

Theorem 1 (Nash-in-Kalai Moment)**.** *Consider a dynamic Nash-in-Kalai equilibrium that satisfies Appendix Assumption 6 in Appendix C.4. Suppose players $i < j$ form a contract in a subgame time t_0 that remains in place through the (potentially random) terminal time t^*. Let Pay_{ijt} be the payment from j to i in period t. Then the expected NPV Nash-in-Kalai payment is the expected NPV Nash-in-Nash flow payment plus a dynamic adjustment term $\text{Pay}_{DA,ijt}$ that captures the change in ijt disagreement payoffs between entering impasse after agreeing at t_0 and entering impasse after disagreeing in t_0 (formalized in Appendix C.4):*

$$E_{t_0}\left[\sum_{t=t_0}^{t^*}\beta^{t-t_0}\text{Pay}^*_{ijt}\right] = E_{t_0}\left[\sum_{t=t_0}^{t^*}\beta^{t-t_0}\text{Pay}_{NiN,ijt}\right] + E_{t_0}\left[\sum_{t=t_0+1}^{\infty}\beta^{t-t_0}\text{Pay}_{DA,ijt}\right]. \tag{10}$$

The dynamic adjustment term emerges because the Nash-in-Nash disagreement value does not capture the right disagreement value in a forward-looking model. Recall from Section 3.2 that under step-by-step, Nash-in-Kalai gains from trade can be computed relative to the value of sustained disagreement, which is formally called impasse. If ij entering impasse would cause non-ij pairs to depart from their equilibrium prices, then gains from trade calculated at equilibrium prices will misstate the value of disagreement. The dynamic adjustment term corrects for this difference. The dynamic adjustment term operates on a different horizon than the Nash-in-Nash flow payment: impasse only affects actions starting in $t_0 + 1$, but compounding effects of non-ij responses can lead impasse to have an effect beyond the current contract's terminal date t^*.

The dynamic adjustment in Equation (10) reflects that the static Nash-in-Nash disagreement point is almost, but not quite, the right disagreement point for a forward-looking model. In a static Nash-in-Nash model's disagreement point, ij surprise all players by failing to agree

where $\mathbb{C}_{ijt}$ is enforced to be the null contract $\mathbb{C}_{0,ijt}$.

for the game's single period, and no other pair updates its behavior in response. However, the static disagreement point is the wrong disagreement point for a forward-looking problem. Suppose that ij bargain a two-period agreement in t_0 that affects retailer j's leverage with supplier h in t_0+1. Suppose further that i offers high gains to j in t_0 that will not be available in future periods. Then ij know that if they disagree in t_0, then hj will anticipate a lower off-equilibrium ij price in t_0+1; the Nash-in-Nash disagreement point at equilibrium prices will underpredict j's future leverage with h, and therefore overpredict the off-equilibrium p_{hj,t_0+1}.

The dynamic analog of Nash-in-Nash disagreement extends the logic of surprise across periods, which corresponds to Binmore et al. (1989)'s *impasse point*. This is the hypothetical path that emerges if ij would like to reach an agreement, and every player expects them to, but ij continually and surprisingly fail to agree without abandoning negotiations. In any period t, no one updates their behavior based on a belief that i and j will fail to agree going forward, just as in the static case, but now all pairs do incorporate the way in which impasse has led to different outcomes in prior periods. The specific form of the dynamic adjustment term in Equation (10) compares the value of entering impasse starting in t_0 to the actions and prices that emerge when entering impasse starting in t^*+1.[10]

5.2 Nesting Previous Models

In steady state, the Nash-in-Kalai model reduces to the expected NPV of period-by-period Nash-in-Nash flow payments, so that prevailing TU estimation strategies are nested as a special case.

Corollary 1 (Steady state Pay_{DA} is zero)**.** *Suppose that in every subgame in period t in which ij form an agreement, if ij instead disagree in t, then the joint distribution of subsequent*

[10]Section 3.2 develops intuition for the Kalai cancellation by comparing the equilibrium outcome to the value of disagreeing for t^*+1-t_0 periods, which has a more direct analog to finite dependence. One could derive a dynamic adjustment based on this model that would also be valid, but the resulting moment would be explicitly expressed in terms of many ij off-equilibrium bargains, and the result would fail if ij punish disagreement with sustained exclusion.

actions $a_{t'}$ *and non-*ij *contracts* $\mathbb{C}_{t',-ij}$ *is unchanged. Then in all* $t \geq t_0$, $\text{Pay}_{DA,ijt} = 0$.

The intuition is straightforward: if ij disagreement has no predictable payoff-relevant effect on non-ij pairs, then the Nash-in-Nash disagreement value will be equal on average to the Nash-in-Kalai impasse value. This condition holds whenever a market is in a stationary equilibrium, so that the same non-ij contracts emerge when ij deviates. This result also trivially holds in many bilateral-monopoly settings, which then have no non-ij pairs to respond to ij impasse.

As a corollary, the Nash-in-Kalai model immediately nests static and myopic TU Nash-in-Nash models. In static models, one can assume myopia ($\beta = 0$) without changing predictions. In myopic models, $\beta^{t-t_0} = 0$ for all periods $t > t_0$, so that the dynamic adjustment term is zero. By applying that argument to the Nash-in-Kalai payment and Nash-in-Nash flow payment terms, one obtains that $\text{Pay}^*_{ijt_0} = \text{Pay}_{NiN,ijt_0}$ in static or myopic TU models.

The Nash-in-Kalai model also nests dynamic TU Nash-in-Nash models. A leading example is Lee and Fong (2013), in which contracting occurs per-period over TU lump-sum transfers, but networks have externalities across periods through contract formation. By applying Theorem 1 to their setting, one obtains a formal characterization of the transfers that they leave implicit: transfers should be the sum of a Nash-in-Nash term to split in-period games and a dynamic adjustment term to compensate for the effect of immediate impasse on future disagreement values.

More broadly, the Nash-in-Kalai model only requires outcomes that are consistent with Kalai proportional bargaining, not that outcomes arise from applying the Kalai proportional solution. In TU models, Nash-in-Nash outcomes already satisfy proportionality for a given pair, so that the Nash-in-Kalai framework nests previous results without changing their predictions. Staggered bargaining will typically prevent Nash-in-Kalai from perfectly aligning with Nash-in-Nash, but in many settings, the difference will be small, especially after appropriately reindexing bargaining weights. I discuss this point in further detail in Appendix B.3.

5.3 Bounded Rationality State Space Approximation

Outside steady state, the dynamic adjustment can be nonzero when ij disagreement can predictably shift non-ij prices. When the resulting dynamic adjustment may be important, a bounded-rationality approximation provides a tractable way forward.

For a concrete example, suppose supplier i and retailer j form a T-period contract in period t_0, and j will negotiate in period $t_0 + 1$ under the threat of being forced to divert more supply to i. If one period of ij disagreement would predictably lead to a lower ij price in $t_0 + 1$, then the ij disagreement in t_0 would lead hj to bargain in $t_0 + 1$ with greater anticipated retailer leverage, leading hj to lower their price on the impasse path. As a result, the Nash-in-Nash disagreement point in $t_0 + 1$ would reflect the equilibrium p_{hjt_0+1} that is too pessimistic for the retailer, and the Nash-in-Kalai payment must be adjusted downward through the dynamic adjustment to compensate.

A level-1 or level-2 bounded rationality approach is likely to capture gains from trade well. A *level*-1 model assumes that the agree-then-impasse and disagree-then-impasse paths generate the same subsequent actions and non-ij contracts, so that $\text{Pay}_{DA,ijt} = 0$. This is the strategy I adopt, and later test, in my empirical application. A *level*-2 model calculates the dynamic adjustment assuming that future contract responses follow a level-1 model. In principle, one could test whether level-k models beyond $k = 2$ improve GMM model fit, but I expect corrections beyond level 2 to be negligible in practice: the level-k adjustment corresponds to ij incorporating how their impasse affects how hk impasse affects $g\ell$'s dynamic adjustment term. Further, most cross-pair contracting externalities can be approximated to zero, reducing the state space.

Bounded rationality also offers a way to nest dynamic extensions of TU Nash-in-Nash models. Recall from Section 3.3 that simultaneous-pricing models like Ho and Lee (2017) and Crawford et al. (2018) become NTU when contracts are staggered, because of time-varying internalized contracting externalities. In this case, a level-1 staggered Nash-in-Nash model

accounts for the direct effect of an ij price on ij payments, but not the indirect effect on future actions or non-ij prices via leverage. As a result, a level-1 staggered Nash-in-Kalai model would have the same predictions as a level-1 staggered Nash-in-Nash model.

5.4 Considerations for Empirical Practice

Theorem 1 offers a moment for GMM estimation. In practice, implementation requires data on contract timing, identification of the discount factor, an approach to handle the sample's end date, and assumptions on expectations. I discuss each in turn.

The key additional data requirement to estimate a Nash-in-Kalai model is data on contract start and end dates. This is required to correctly specify NPV gains from trade with multiperiod contracts. In my empirical application, I can only establish reliable start and end dates for a subset of contracts; I maintain the assumption that any selection of these contracts on data quality is exogenous to the instruments I apply. Otherwise, estimation generally follows the standard strategy of estimating downstream demand as a function of upstream network quality, and then recovering bargaining parameters from the model-derived moment on expected NPV payments, aggregated over observable instruments.

Relative to a static model, the forward-looking model introduces the discount factor β as a new parameter. The econometrician can calibrate β from domain knowledge, but in many settings it is possible to do better. The discount factor can be identified from predictable variation in the distribution of gains from trade over time. In my empirical application, I exploit the steeper price growth of list price-based contracts, which are disproportionately adopted by larger hospitals and smaller insurers, to identify the discount factor.

A practical challenge for estimating discount factors is that a given panel will end, so that it will be impossible to estimate an infinite-horizon model without extrapolating under strong stationarity assumptions. In practice, a researcher must choose a finite horizon that balances two competing concerns: a shorter horizon risks excluding contracts whose duration exceeds the window, while a longer horizon requires extrapolating gains from trade beyond

the panel. Any choice will attenuate NPV gains for long-lived contracts, with the degree of attenuation depending on the discount factor, panel length, and distribution of contract lengths. A more principled approach to estimation with finite panels is an important avenue for future work.

It is also important to note that the Nash-in-Kalai moment assumes that negotiators share rational expectations and a constant per-period discount factor. Shared expectations are necessary to negotiate over an expected transfer, rational expectations are necessary to construct a valid moment based on the realized transfer, and a constant per-period discount factor is necessary for the players to bargain in terms of consistently defined net present values. Empirical models with asymmetric information or irrational expectations are frontier topics (Larsen and Zhang, 2021; Chen et al., 2024) that are outside the scope of this work, and alternative shared time preferences could be accommodated with additional notation.

6 Empirical Application

My empirical application uses a panel dataset from West Virginia to study hospital–insurer contracting. I document that staggering is pervasive, so that a forward-looking Nash-in-Nash model would face explosive state space growth. The estimated discount factor of $\hat{\beta} = 0.899$ corresponds to strong forward-looking behavior, and overwhelmingly rejects controlling state space control through myopia. Finally, I find that the timing misspecification introduced by the usual static approach would substantially bias bargaining weights: a flexible specification produces an estimated bargaining weight of -7.54, and a more restricted specification would bias weights in the direction that Section 2.1 predicts. Note that the results are adapted from the companion paper (Dorn, 2025a), which studies proposed changes to Medicare reimbursement dynamics; I present some important empirical details in Appendix D.2, and leave further implementation details to that companion paper.

6.1 Setting and Data

Hospital–insurer contracting is a classic setting for studying contracting externalities (Gowrisankaran et al., 2015; Ho and Lee, 2017). West Virginia followed the unusual practice of making the underlying contracts public records, allowing me to document real-world timing.

Hospitals and insurers contract over the payments that the hospitals will receive for providing care to consumers who purchase a plan from the insurer. The hospital offers the insurer a discount relative to list prices; in exchange, the insurer steers its consumers to the hospital. This generates contracting externalities through substitution: if an insurer fails to reach an agreement with hospital B, then some of its patients will be steered to hospital A, so that an increase in p_A reduces the insurer's leverage with B, and increases the negotiated p_A. If contracts are also staggered and negotiators are forward-looking, then the state space growth arguments of this paper will apply.

I leverage a panel dataset on hospital–insurer contracts between 2006 and 2015 that allows me to speak to contract timing. The data was produced as part of a system of regulating hospital list price increases. The data includes reported payment levels, which are typically available, as well reports on contract start and end dates. The data allows me to document staggering and contract duration, and confirm the natural intuition that contracts are formed at different times. I focus on six modeled insurers: Highmark BCBS, the largest insurer; HPUOV, a regional non-profit insurer; and four for-profit insurers. After data cleaning, I have 63 reliable bargains that I use in estimation that involve the six insurers and roughly 30 hospital systems (Table 3).

Figure 3 shows that contracts were multiyear. Figure 3(a) presents reported contract lengths for contracts with a reported end date. Highmark BCBS typically formed three- or five-year agreements. Other insurers mostly formed auto-renew agreements. Figure 3(b) presents retrospective contract length for auto-renew agreements in 2015, the last year of my

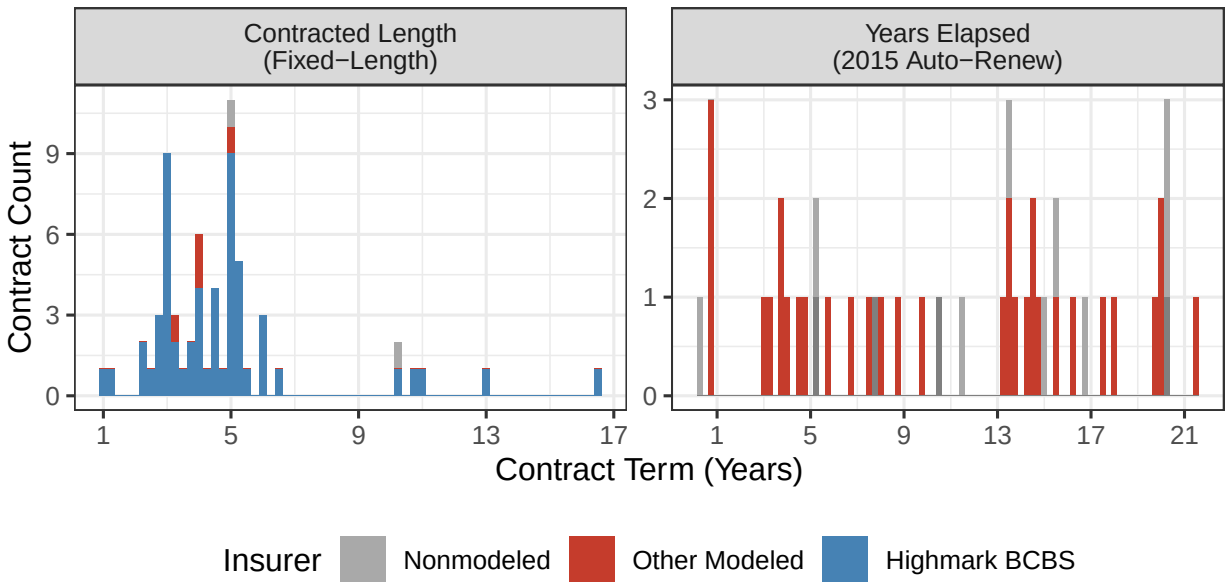


Figure 3: Contracts are multiyear. Distribution of reported contract lengths (including agreements that are not used in estimation) for fixed-term contracts in general (left panel) and auto-renew contracts with reported formation dates as-of fiscal year 2015 (right panel). Colors indicates insurer. The two nonmodeled fixed-length contracts correspond to Wheeling–Pittsburgh Steel.

sample. Most of these auto-renew contracts, in place in 2016 and large enough to report start dates, had been formed in 2002 or earlier. I estimate that the broader average auto-renew renewal probability was 93.4% in each year (Figure 7). The variation in contract length implies at least staggering across years, and suggests a high risk of bias from incorrectly applying a model of annual contracting.

Contracts were meaningfully staggered across and within years. Figure 4 presents the distribution of contract start dates for agreements used in the estimation sample. Even within a given year, Highmark BCBS would form agreements all throughout the year. The prevalence of staggering leaves a capacity for internalized contracting externalities. A Nash-in-Nash model is at the risk of explosive state space growth: a monthly model aimed at capturing intra-year contracting externalities would have $T = 120$ time periods and an average of $N = 148$ ties per year, yielding roughly $NT^{N-1} = 6.5 * 10^{307}$ price-response functions.

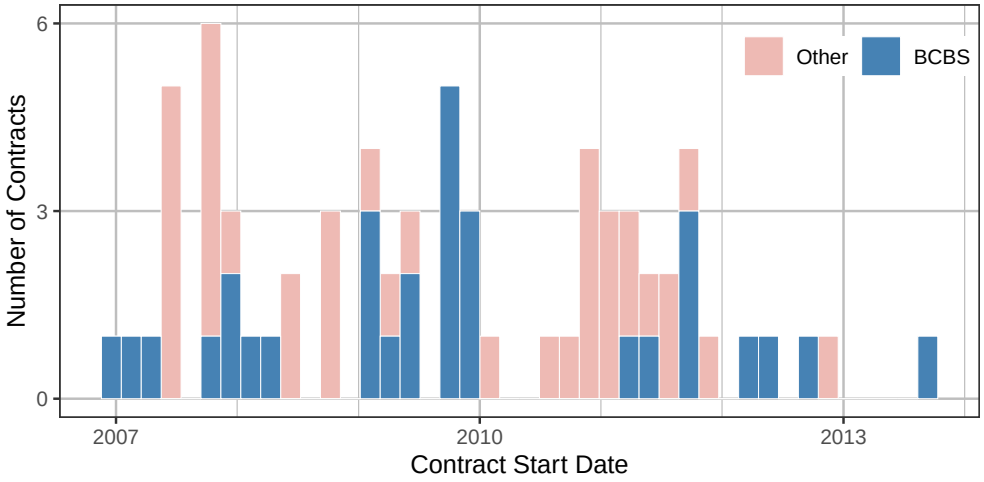


Figure 4: Contracts are formed at different times. Histogram of contract start dates for contracts used in the estimation sample and introduced in 2007–2014 for Highmark BCBS (blue) and other modeled insurers (pink). Vertical lines indicate January 1 of a given year. Contracts were not systematically introduced on the same dates.

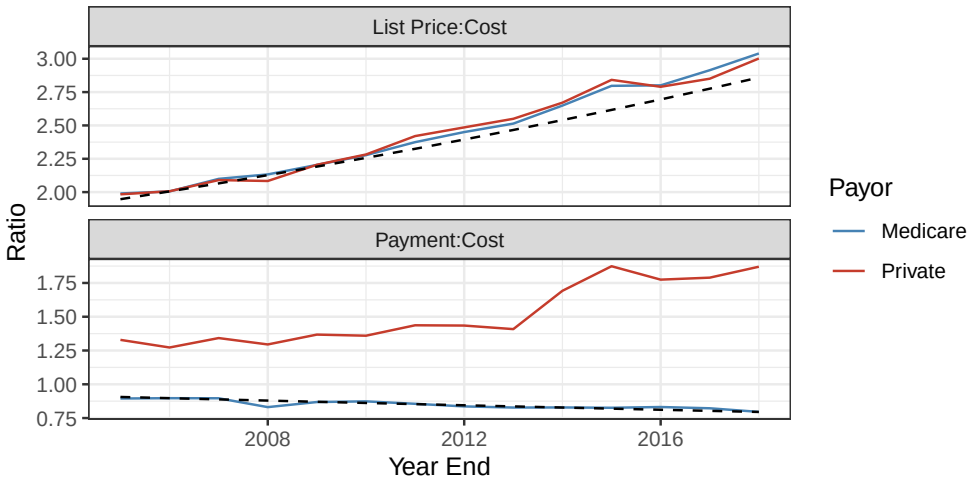


Figure 5: The ratio of list price charges (top) and real payments (bottom) to reported costs by Medicare (red) and private payors (blue) for West Virginia hospitals. Dashed lines represent Medicare 2006 values extrapolated based on three ppt. and negative one ppt. annual changes.

Finally, different firms reached contracts with predictable different price dynamics. I classify agreements into list-price-based contracts whose payments moved proportionally to list prices, or other contracts that were likely negotiated relative to Medicare-based rates. Figure 5 establishes that list prices grew roughly four percentage points faster than Medicare

payments annually. In Table 4, I estimate that roughly 53% of payments were negotiated as a fixed fraction of list prices, and that these list-price-based contracts were contracted among small- and medium-sized insurers. As a result, the logic of Section 2.1 suggests that a period-by-period model will overestimate Highmark BCBS's relative bargaining power by misattributing smaller insurers' high prices in late contract years to bargaining power, even though those high prices were discounted at the time of contract formation.

6.2 Estimation Strategy

I estimate an annual model that is a dynamic extension of the Ho and Lee (2017) model. The key departure is to allow multiperiod contracts with accurate timing under the Nash-in-Kalai moment from Theorem 1.

I follow the standard two-stage approach of estimating hospital demand and insurer demand. I assume that losing a bilateral hospital–insurer agreement leads to substitution through two channels: when choosing insurance plans, consumers derive more value from an insurer with a broader network; after becoming sick, patients who stay with the insurer are forced to reallocate to a different hospital within the insurer's network. I follow the standard practice of estimating the model in reverse order of the consumer's decision, so that previous decisions can be taken as fixed. I estimate hospital demand using a logit model with diagnosis-specific willingness to travel and hospital fixed effects, estimated with Highmark BCBS inpatient discharges in 2016. I then estimate insurer demand, where consumers trade off network quality against premiums, using 2016 data in which Affordable Care Act rating area fixed effects absorb premium variation. I calibrate premium sensitivity from Ho (2006)'s estimated own-price elasticity of -1.4, and back-cast insurer demand in previous years based on changes in state-level fully-insured sales. I present more details on the demand models, identification, and estimation in Appendix D.2.

The demand estimates enter the bargaining model through flow gains from trade. When hospital i and insurer j have an agreement in place, the insurer gains premium revenue

(ϕ_{jt} per life insured net of η_j non-inpatient costs) from additional sales through its broader network and reduces spending by diverting patients from other hospitals, but the insurer and its patients pay the hospital $D^H_{ijt}p_{ijt}$ for the D^H_{ijt} units of inpatient care provided in period t. I follow Ho and Lee (2017) and write $[\Delta_{ij}\pi^M_{jt}]$ as the effect on jt profits from the ij agreement before accounting for the $D^H_{ijt}p_{ijt}$ bilateral transfer. I similarly write $[\Delta_{ij}\pi^H_{it}]$ for the effect on hospital profits before receiving the bilateral transfer: operational costs c_{it} per additional unit of care provided, and lost revenue from patients who switch to insurer j and then still visit hospital i. I also allow insurers to incur a per-contract negotiation cost r^M_j, which partially captures the reduced-form effect of hospital negotiation costs. With this notation, $Pay_{NiN,t} = (1-\tau_{ij})[\Delta_{ij}\pi^M_{it}] - \tau_{ij}[\Delta_{ij}\pi^H_{it}] + (1-\tau_{ij})r^M_j 1\{t = t_0\}$ is the latent Nash-in-Nash flow payment in period t, and Theorem 1 implies a parametric model of expected NPV payments over the life of a contract.

All bargaining models that I estimate imply a moment of the same form. If hospital i and insurer j form a contract starting in period t_0 and ending in the (potentially random) period t^*, then the realized payments $D^H_{ijt}p_{ijt}$ satisfy

$$E_{t_0}\left[\sum_{t=t_0}^{\min\{t^*,t_0+4\}} \beta^{t-t_0}\left(D^H_{ijt}p_{ijt} - (1-\tau_{ij})[\Delta_{ij}\pi^M_{jt}] + \tau_{ij}[\Delta_{ij}\pi^H_{it}] - (1-\tau_{ij})r^M_j 1\{t=t_0\}\right)\right] = 0. \tag{11}$$

The latent Nash-in-Nash flow payment $Pay_{NiN,t}$ is the latent Nash-in-Nash flow payment depends on substitution patterns that are estimated from hospital and insurer demand, hospital marginal costs c_{it} that are calibrated, and insurer non-inpatient costs η_j and negotiation costs r^M_j that are estimated through GMM.

I allow the bargaining weights τ_{ij} to differ by hospital and insurer. I estimate bargaining weights τ_{ij} on a logit scale:

$$\log\left(\tau_{ij}/(1-\tau_{ij})\right) = log(\tau_j/(1-\tau_j)) + \tau^{Size} log(HospSize_{i,2006}/MeanHospSize_{2006}), \tag{12}$$

where hospital size is measured as the size of the bargaining system in my first year of 2006, τ_j is insurer bargaining power measured at the average-sized hospital system, and more negative τ^{Size} corresponds to larger hospitals having more bargaining weight. I allow the τ_j fixed effects to correspond to weights outside the interval $(0, 1)$, in which case a logit specification becomes inappropriate and I impose $\tau_{ij} = \tau_j$.

I estimate three bargaining models. A `static` model captures the usual static approach of collecting data over a short period and misspecifying timing by treating all agreements as recently-formed. This model estimates bargaining parameters (other than β) based on 2015 data, treating each ij pair with an agreement as a new negotiation. A `myopic` model incorporates accurate timing, but controls Nash-in-Nash state space growth by assuming $\beta = 0$. The `forward-looking` model is the Nash-in-Kalai model proposed by this paper, which has accurate timing and allows negotiators to balance gains over multiple periods using a discount factor. The `forward-looking` model uses a five-year finite horizon and level-1 bounded rationality approximations for simplicity. The level-1 approximation sets the dynamic adjustment term of Equation (10) to zero, so that the `myopic` model is a special case of the `forward-looking` model that imposes $\beta = 0$, and the `static` is a variant of the `myopic` model that also misspecifies timing. This yields a simple test of potential bias introduced by the `static` model's misspecified timing and `myopic` model's assumed discounting, but does not attempt to quantify the precise level or effect of internalized contracting externalities.

The key new parameter for the `forward-looking` model to estimate is the discount factor β. Identification comes from predictable variation in the time path of payments across firms. Smaller insurers and larger hospitals were more likely to reach agreements with list-price-based contracts, which Figure 5 establishes leads to predictably faster price growth after a contract is formed. As a result, increases to β increase NPV payments and gains from trade for all firms, but increase NPV payments disproportionately for smaller insurers and larger hospitals. I therefore use insurer dummies and binned hospital-size dummies as instruments

to identify β.

It is important to note that identification of β via firm and insurer instruments relies on restricting heterogeneity in the bilateral τ_{ij} bargaining weights. If τ_{ij} were restricted to be constant for all pairs, then one could identify β by comparing moments across pairs with different payment paths but the same shared bargaining weight. However, I find important heterogeneity in bargaining weight by hospital, precluding this approach. If I allowed τ_{ij} to be fully flexible across pairs, then I could match expected NPV payments for all ij pairs, so that there would be no identifying variation for β by marginalizing by insurer or hospital. Instead, I rely on an assumption that a lower-dimensional representation of τ_{ij} suffices to capture heterogeneity in surplus-splitting and the estimates of hospital and insurer demand suffice to capture heterogeneity in flow gains from trade. Then, the remaining association of starting prices with future price inflation comes from the predictable choice of benchmark. Further, identification of β would still fail if the dynamic path of payments was absorbed by the estimated gradient of τ_{ij}. I mitigate this risk by grouping hospital systems into bins that allow for different forms of association of hospital size with bargaining weights and payment inflation.

6.3 Results

The fully flexible `static` model predicts a bargaining weight of -7.54, and after shutting down heterogeneity, predicts biased bargaining weights that line up with the predictions of Section 2.1. The `myopic` model produces fairly accurate bargaining weights, but incorrectly rules out forward-looking responses to dynamic counterfactuals.

Table 1 presents the estimated bargaining parameters across the three models. I first focus on the `forward-looking` model, the preferred specification.

The `forward-looking` model produces plausible estimates of surplus shares and time discounting. The estimated discount factor of $\hat{\beta} = 0.899$ implies that negotiators value an inflation-adjusted dollar next year equivalently to roughly 90 cents today, and the standard

	Parameter				
	β	τ_{BCBS}	τ_{HPUOV}	τ_{FP}	$-\tau^{Size}$
Static	·	0.487**	-7.54	0.694***	3.354
(Nash/Kalai)	(·)	(0.191)	(17.204)	(0.175)	(22.875)
Myopic	·	0.876***	0.825***	0.861***	1.037***
(Nash/Kalai)	(·)	(0.012)	(0.232)	(0.034)	(0.199)
Forward-Looking	0.899***	0.854***	0.877***	0.889***	0.989***
(level-1 rationality)	(0.03)	(0.006)	(0.026)	(0.005)	(0.028)
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 1: Estimated bargaining and patience weights for the static (first row), myopic (second row), and preferred forward-looking (third row) bargaining models. The MCO τ_j bargaining weights represent the estimated bilateral share τ_{ij} evaluated for a hospital i with size equal to the average bargain value, and include estimated heterogeneity between Highmark BCBS (BCBS), the regional insurer HPUOV, and the modeled for-profit insurers (FP). The estimated forward-looking model overwhelmingly rejects myopia, and the static model produces implausible estimated bargaining weights. For additional parameter estimates, see Appendix Table 5.

error (0.03) overwhelmingly rejects myopia. The insurer bargaining weights, represented as the insurer's share of NPV expected surplus when bargaining with a medium-sized system, is strikingly similar across insurers: Highmark BCBS, the largest insurer, obtains a share of $\hat{\tau}_j = 0.854$; the for-profit insurers, which are smaller in West Virginia, obtain a share of $\hat{\tau}_j = 0.889$; and HPUOV, a regional insurer, is in the middle at $\hat{\tau}_j = 0.877$. Instead, the main source of heterogeneity in shares comes from hospital size: the estimated coefficient $\tau^{Size} = -0.989$ indicates that larger hospital systems negotiate substantially higher-paying agreements. Figure 6 presents the estimated shares across the distribution of hospital size, illustrating that small systems receive near-take-it-or-leave-it offers, while large systems receive 50%–65% of the shared surplus.

The data shows that a `static` model with misspecified timing would infer the wrong bargaining weights. Recall from Section 2.1 that a researcher who interprets multiyear contracts through a static model risks getting incoherent bargaining weights, and in particular under-

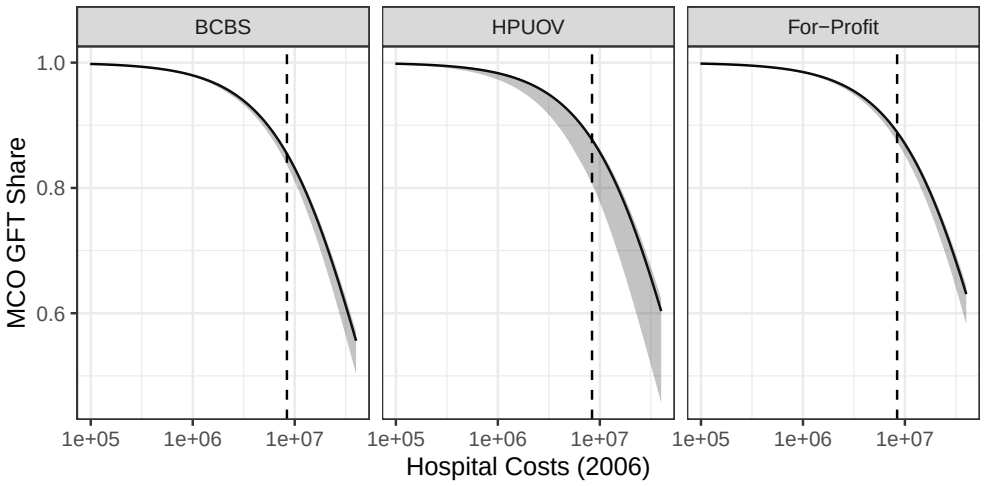


Figure 6: Estimated insurer shares of expected NPV gains from trade under the `forward-looking` model. The model estimates substantial heterogeneity by hospital system size and limited heterogeneity by insurer.

estimating bargaining weights with fast price growth leading to high payments in discounted late contract years. The misspecified `static` model would estimate a bargaining weight of -7.54 for HPUOV, indicating that HPUOV loses \$7.54 for each \$1 of surplus an agreement generates. Some of the extreme parameter estimate comes from an association of hospital size with contract structure coupled with a large estimate of $\tau^{Size} = -3.354$. In an alternative specification in which I fix $\tau^{Size} = 0$, the estimated bargaining weights for Highmark BCBS, HPUOV, and the for-profit insurers would change to the interior weights of 0.365, 0.278, and 0.16, respectively; the `static` model underestimates all bargaining weights, but in particular underestimates bargaining weights more for insurers with more list-price-based contracts with faster price growth.

The `myopic` model suggests that the `static` model's bias come from misspecified timing rather than a failure to allow time discounting. Recall that this model captures accurate contract start and end dates but imposes $\beta = 0$, so that negotiators only value the first year of an agreement. Despite this restriction, the estimated bargaining weights are quite close to the forward-looking model, as seen when comparing Figure 8 to Figure 6. Quantitatively, the `myopic` model estimates that the modeled insurers retain 85%–80% of surplus when negotiat-

ing with an average-sized system, and the τ^{Size} coefficient falls from -0.989 to -1.037. These results have an important practical upshot: for fundamentally static counterfactuals like mergers or price passthrough in which the key structural parameters are bargaining weights, the key empirical requirement is knowing which prices reflect recently-formed agreements and which prices reflect later dynamics; a myopic model applied to accurately timed data may be enough to get bargaining weights right. On the other hand, the `forward-looking` model's strong rejection of myopia implies that the `myopic` model would mischaracterize forward-looking responses to anticipated changes in future market conditions, such as the proposed dynamic Medicare reform studied by Dorn (2025a). For dynamic counterfactuals like the target of Dorn (2025a) — the equilibrium effect of predictable benchmark price inflation after accounting for forward-looking responses — a forward-looking model will be necessary.

The `forward-looking` model allows more flexible gains from trade than previous static or period-by-period models, but the level-1 bounded rationality approximation and five-year finite-horizon window are important restrictions. The five-year finite horizon is chosen to include Highmark BCBS's five-year contracts while avoiding extrapolation beyond the panel's 2015 end date, following the general tradeoff discussed in Section 5.4. At the estimated discount factor of $\hat{\beta} = 0.899$, the five-year window captures the full NPV of fixed-term contracts but only 63% of NPV gains for a contract that remains in place for a decade; the resulting attenuation bias is likely to be larger for smaller insurers with more long-lived auto-renew agreements, and so represents a potential source of residual bias from remaining misspecified timing.

I find in Appendix D.3 that the level-1 approximation is likely to be reasonable in this setting. Recall that the approximation corresponds to firms considering the direct effect of prices on bilateral payments, and not the indirect effect on other agreements through leverage; equivalently, it approximates the full forward-looking problem as TU. For 97% of the estimation sample, the ratio of marginal values is within 10% of the TU benchmark. A

more demanding test would need to confirm that a level-2 correction would not substantially shift the estimates; this test requires estimating how the slope would change across off-equilibrium bargains, which I leave to future work.

7 Conclusion

This paper shows that one can model the prevalent real-world feature of staggered bargaining without introducing bias from misspecifying timing or introducing uncontrolled explosive state space growth from internalized contracting externalities. The step-by-step property yields a much simpler equivalent representation of the bargaining solution, implying a tractable moment for GMM estimation that generalizes previous TU Nash-in-Nash frameworks. I apply the natural multilateral model of generic step-by-step, the Nash-in-Kalai model, to data from West Virginia, where I find that a static model introduces substantial bias to bargaining weights.

The framework enables dynamic counterfactuals that were previously infeasible. The companion paper Dorn (2025a) uses Nash-in-Kalai to study how proposed changes to Medicare reimbursement would propagate across staggered hospital–insurer renegotiations. The framework provides a natural tool for studying how classic macroeconomic dynamics like inflation propagate through staggered contracting, connecting the empirical literature to the macroeconomic literature on staggered price-setting (Taylor, 1980; Calvo, 1983). The framework is broadly applicable when contracts interact and timing matters; for example,the framework could be used to study how hospital list price-setting affects endogenous renegotiation timing for associated auto-renew contracts (Dorn, 2025b). Several extensions remain for implementation: ameliorating boundary bias from finite panel lengths, formally testing whether higher-order bounded rationality approximations improve model fit, and relaxing the maintained assumptions of shared rational expectations and constant discounting.

References

Barwick, P. J., Kwon, H.-S., Li, S., and Zahur, N. B. (2025). Drive down the cost: Learning by doing and government policies in the global ev battery industry. Working Paper 33378, National Bureau of Economic Research.

Beier, E. (2020). 5 steps to better payer contract negotiation. *MediGain*. Year unknown.

Bell, B. (2024). What is an anchor tenant in a shopping mall?

Bilal, A., Engbom, N., Mongey, S., and Violante, G. L. (2022). Firm and worker dynamics in a frictional labor market. *Econometrica*, 90(4):1425–1462.

Binmore, K., Shaked, A., and Sutton, J. (1989). An outside option experiment. *The Quarterly Journal of Economics*, 104(4):753–770.

Bossert, W. (1994). Disagreement point monotonicity, transfer responsiveness, and the egalitarian bargaining solution. *Social Choice and Welfare*, 11(4):381–392.

Brügemann, B., Gautier, P., and Menzio, G. (2018). Intra firm bargaining and Shapley values. *The Review of Economic Studies*, 86(2):564–592.

Cahuc, P., Postel-Vinay, F., and Robin, J.-M. (2006). Wage bargaining with on-the-job search: Theory and evidence. *Econometrica*, 74(2):323–364.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398.

Capps, C., Dranove, D., and Satterthwaite, M. (2003). Competition and market power in option demand markets. *RAND Journal of Economics*, 34(4):737–763.

Chen, X., Hansen, L. P., and Hansen, P. G. (2024). Robust inference for moment condition models without rational expectations. *Journal of Econometrics*, 243(1):105653.

Coles, M. G. and Muthoo, A. (2003). Bargaining in a non-stationary environment. *Journal of Economic Theory*, 109(1):70–89.

Collard-Wexler, A., Gowrisankaran, G., and Lee, R. S. (2019). "Nash-in-Nash" bargaining: A microfoundation for applied work. *Journal of Political Economy*, 127(1):163–195.

Crawford, G. S., Lee, R. S., Whinston, M. D., and Yurukoglu, A. (2018). The welfare effects of vertical integration in multichannel television markets. *Econometrica*, 86(3):891–954.

Crawford, G. S. and Yurukoglu, A. (2012). The welfare effects of bundling in multichannel television markets. *The American Economic Review*, 102(2):643–685.

Deng, S., Jia, D., Leccese, M., and Sweeting, A. (2024). Bargaining and dynamic competition. Working Paper 32360, National Bureau of Economic Research.

Do, J. and Miklós-Thal, J. (2022). Opportunism in vertical contracting: A dynamic perspective. Technical Report DP16951, CEPR Press Discussion Paper.

Dorn, J. (2025a). Dynamic bargaining between hospitals and insurers.

Dorn, J. (2025b). Six stylized facts from ten years of contract data.

Duffy, J., Lebeau, L., and Puzzello, D. (2021). Bargaining under liquidity constraints: Nash vs. Kalai in the laboratory. Technical Report 2113, Federal Reserve Bank of Dallas.

Dutta, R. (2012). Bargaining with revoking costs. *Games and Economic Behavior*, 74(1):144–153.

Dutta, R. (2024). Bargaining as a struggle between competing attempts at commitment. *Review of Economic Studies*, 91(5):2771–2805.

ECG (2020). A plan and guideline for managed care contracting negotiations. Technical report, ECG Management Consultants.

Fletcher, T. A. (2020). Payer contract negotiations: How to improve your negotiating skills. *ICD10 monitor*.

Ghili, S. (2022). Network formation and bargaining in vertical markets: The case of narrow networks in health insurance. *Marketing Science*.

Gooch, K. (2019). Tips for getting the most out of payer negotiations. *Becker's Hospital Review*.

Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 105(1):172–203.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review*, 103(1):145–77.

Harsanyi, J. C. (1956). Approaches to the bargaining problem before and after the theory of games: A critical discussion of zeuthen's, hicks', and nash's theories. *Econometrica*, 24(2):144–157.

Ho, K. (2006). The welfare effects of restricted hospital choice in the US medical care market. *Journal of Applied Econometrics*, 21(7):1039–1079.

Ho, K. and Lee, R. S. (2017). Insurer competition in health care markets. *Econometrica*, 85(2):379–417.

Ho, K. and Lee, R. S. (2019). Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review*, 109(2):473–522.

Ho, K. and Lee, R. S. (2024). Contracting over rebates: Formulary design and pharmaceutical spending. *Econometrica*.

Horn, H. and Wolinsky, A. (1988). Bilateral monopolies and incentives for merger. *The RAND Journal of Economics*, 19(3):408–419.

Hu, T.-W. and Rocheteau, G. (2020). Bargaining under liquidity constraints: Unified strategic foundations of the Nash and Kalai solutions. *Journal of Economic Theory*, 189:105098.

Jarosch, G., Nimczik, J. S., and Sorkin, I. (2024). Granular search, market structure, and wages. *The Review of Economic Studies*, 91(6):3569–3607.

Kalai, E. (1977). Proportional solutions to bargaining situations: Interpersonal utility comparisons. *Econometrica*, 45(7):1623–1630.

Larsen, B. and Zhang, A. L. (2021). Quantifying bargaining power under incomplete information: A supply-side analysis of the used-car industry. Working Paper 21-058, SIEPR.

Lee, R. S. and Fong, K. (2013). Markov-perfect network formation: An applied framework for bilateral oligopoly and bargaining in buyer-seller networks.

Lee, R. S., Whinston, M. D., and Yurukoglu, A. (2021). Chapter 9 - structural empirical analysis of contracting in vertical markets. In Ho, K., Hortaçsu, A., and Lizzeri, A., editors, *Handbook of Industrial Organization, Volume 4*, Handbook of Industrial Organization, pages 673–742. Elsevier.

Lerner, D. (2025). YouTube TV says Disney dispute over ABC, not ESPN. *Awful Announcing*.

Liu, T. (2021). Bargaining with private equity: Implications for hospital prices and patient welfare. SSRN 3896410.

Nydegger, R. V. and Owen, G. (1974). Two-person bargaining: An experimental test of the Nash axioms. *International Journal of Game Theory*, 3(4):239–249.

Olssen, A. L. and Demirer, M. (2022). Drug rebates and formulary design: Evidence from statins in medicare part d.

O'Neill, B., Samet, D., Wiener, Z., and Winter, E. (2004). Bargaining with an agenda. *Games and Economic Behavior*, 48(1):139–153.

PMMC (2019). Best practices for negotiating payer contracts as cost shifts from the payer to the patient. *Xtelligent Healthcare Media*.

Prager, E. and Tilipman, N. (2022). Regulating out-of-network hospital payments: Disagreement payoffs, negotiated prices, and access.

Roth, A. E. (1979). Proportional solutions to the bargaining problem. *Econometrica*, 47(3):775–778.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–110.

STD TAC and Moss, R. (2014). Ten steps to contracting with insurance companies. *STD TAC*.

Stole, L. A. and Zwiebel, J. (1996). Intra-firm bargaining under non-binding contracts. *The Review of Economic Studies*, 63(3):375–410.

Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *Journal of Political Economy*, 88(1):1–23.

Tiew, A. (2022). Flailing firms and joint operating agreements: An application to U.S. local daily print newspapers from 1932 to 1992.

Zhou, Y. and Sibley, D. (2026). Interdependent bargains. *Journal of Mathematical Economics*, 122:103213.

A Additional Content

B Additional Bilateral Content

B.1 Step-by-Step Implies Proportional

I begin with some additional requirements for the bargaining solution.

Definition 5 (Bargaining solution properties)**.** *A bargaining solution is* homogeneous *if for all $t > 0$ and (S, v^D), $f(tS, tv^D) = tf(S, v^D)$. A bargaining solution is* translation-invariant *if for all (S, v^D) and all $u \in R^2$, $f(S + u, v^D + u) = f(S, v^D) + u$. A bargaining solution is* weakly Pareto-optimal *for a set G of bargaining games if for all $(S, v^D) \in G$, there is no $s \in S$ with $s > f(S, v^D)$.*

If f is homogeneous, then bargaining in dollars and bargaining in cents produces the same underlying outcome. If f is translation-invariant, then taking account of fixed outcomes outside the bargaining game does not change the outcome of bargaining. The translation-invariance ensures that I can map a second-stage agreement whose disagreement point is $f(S, v^D)$ back to a first-stage agreement whose disagreement point is v^D. The weak Pareto optimality condition rules out a bargaining solution satisfying step-by-step by always predicting v^D.

I also introduce a few properties of bargaining games for reducing step-by-step to being proportional. I require some generic properties for the class of bargaining games.

Definition 6 (Set of bargaining games properties)**.** *A set of bargaining games G is* invariant to linear transformations *if $(S, v^D) \in G$, $t > 0$, and $u \in R^2$ implies $(tS + u, tv^D + u) \in G$.*

This condition is fairly standard.

I also add some conditions on a specific bargaining game in G, which rules out certain pathological cases. It is useful to write $P(S, v^D)$ as the set of points $s \in S$ that are weakly Pareto-optimal within the set of $s \geq v^D$ in S.

Definition 7 (Single bargaining game properties)**.** *A single game* (S, v^D) *is* non-trivial, comprehensive, and convex *if (i) there is an* $s^A, s^B \in P(S, v^D)$ *such that* $s_1^A > s_1^B$ *and* $s_2^A < s_2^B$*; (ii) if* $s \in S$*,* $s' \geq v^D$*, and* $s' \leq s$*, then* $s' \in S$*; and (iii) if* $s, s' \in S$ *and* $\alpha \in [0, 1]$*, then* $\alpha s + (1 - \alpha)s' \in S$*.*

These conditions are unnecessarily strong, but simplify the proof that step-by-step implies proportional. Comprehensiveness is the strongest assumption. Without comprehensiveness, step-by-step could hold vacuously, when the feasible set remains too thin, and the set of games remains too sparse, to yield a first-step agreement that generates a contradiction. Comprehensiveness is a simple way to ensure that sufficiently scaled-up first-step agreements are always feasible, allowing the proof to proceed.

If these properties hold, then every bargaining solution satisfying step-by-step is proportional.

Lemma 1 (Step-by-step implies proportional)**.** *Suppose* G *is invariant to linear transformations and contains at least one bargaining game that is sufficiently non-trivial, comprehensive, and convex. Suppose further that a bargaining solution* f *is homogeneous and translation-invariant in general and is weakly Pareto-optimal for that one bargaining game. Then if* f *satisfies step-by-step for games in* G*, then* f *is proportional for* G*.*

Note that Kalai proportional bargaining can fail Pareto optimality when a bargaining game is not comprehensive. However, weak Pareto optimality is only required for the bargaining game that is sufficiently non-trivial, comprehensive, and convex, so that in practice the constraint is extremely lax.

Proof. Appendix E. □

This is an if-and-only if: under much weaker conditions, any bargaining solution that is proportional also satisfies step-by-step.

Lemma 2 (Proportional implies step-by-step)**.** *Take a set of bargaining games G and a bargaining solution f. For any S, T, v^D, define $G(S, T, v^D) = \{(S, v^D), (T, v^D), (S, f(T, v^D))\}$. Suppose that for every $(S, v^D) \in G$, if $f(T, v^D) \in S$, then f is proportional and weakly Pareto-optimal on $G(S, T, v^D)$. Then f satisfies step-by-step on G.*

Proof. Let $(S, v^D) \in G$, and let $f(T, v^D) \in S$. Then there is a $\tau \in [0, 1]$ such that

$$\begin{aligned}
\tau(f_2(S, v^D) - v_2^D) &= (1 - \tau)(f_1(S, v^D) - v_1^D) \\
\tau(f_2(T, v^D) - v_2^D) &= (1 - \tau)(f_1(T, v^D) - v_1^D) \\
\tau(f_2(S, f(T, v^D)) - f_2(T, v^D)) &= (1 - \tau)(f_1(S, f(T, v^D)) - f_1(T, v^D)) \\
\Rightarrow \tau(f_2(S, f(T, v^D)) - v_2^D) &= (1 - \tau)(f_1(S, f(T, v^D)) - v_1^D).
\end{aligned}$$

Further, bother $f(S, v^D)$ and $f(S, f(T, v^D))$ are weakly Pareto-optimal within S. Therefore, $f(S, v^D) = f(S, f(T, v^D))$, so that f satisfies step-by-step on G. □

B.2 State Space Control Regularity Conditions

Assumption 1 (Regularity conditions for Proposition 1)**.** There is no uncertainty in equilibrium: x_t is constant and $\varepsilon_t = 0$. An ij contract takes the form of a T-period constant price p_{ij}. The players follow Markov strategies that imply agreement in every subgame reached through a series of equilibrium agreements and off-equilibrium disagreements. The only payoff-relevant part of s is the current contract state. The pre-bargaining action space $A_{it}^{(1)}(s_t)$ and the same-time-as-bargaining action space $A_{it}^{(2b)}(s_t)$ is empty. Let $\mathcal{R} \subseteq \{1, \ldots, T\}$ be a set such that in equilibrium, agreements are only reached in periods $t = r + kT$ for some $r \in \mathcal{R}$. The only relevant notion of equilibrium other than proportionality restrictions is weak Pareto efficiency: the players cannot do jointly better by bilaterally changing agreements, and the players cannot do separately better by deviating to disagreement. The cardinality of $\mathcal{R}$ is z.

One can allow for multiperiod exclusion after disagreement, but additional price-response

functions will be needed if disagreement leads to exclusion beyond the next equilibrium agreement. Uncertainty over market states would also introduce state space growth through the market states. If there is a relevant market state x, then the number of relevant states must also be scaled by the cardinality of x.

B.3 When Kalai Will Be The Right Tool, and When Kalai Approximates Nash

Kalai proportional bargaining produces the same predictions as Nash when bargaining is TU, and produces the same space of predictions after reindexing of bargaining weights when Pareto frontiers are homothetic. When Pareto frontiers are meaningfully nonhomothetic, recursive Nash bargaining departs from the prior empirical literature and quickly becomes intractable, while Kalai proportional bargaining retains the tractability outlined in the main text.

One may believe in the Kalai model as an accurate model from first principles. Kalai proportional bargaining is likely to emerge in settings driven by negotiator incentives (Dutta, 2012), fairness considerations (Nydegger and Owen, 1974), or dynamic incentive considerations (Kalai, 1977), while Nash bargaining is likely to emerge in settings driven by negotiator patience (Rubinstein, 1982; Collard-Wexler et al., 2019), the possibility of concession (Harsanyi, 1956), or in which there is no natural scale of relative utilities.

If a researcher truly believes in Nash, then as Section 3.3 discusses, Nash bargaining will satisfy step-by-step under homotheticity. If every pair of games $(S, v^D), (T, u^D) \in G$ that a model generates satisfy $S - v^D = \gamma(T - u^D)$ for some $\gamma \geq 0$, then every homogenous bargaining solution like Nash bargaining will predict the same surplus share in every game. A leading case is TU games in which upstream input price-setting occurs at the same time as downstream retail price-setting. Homotheticity naturally arises when a retailer reevaluates downstream prices in each period based on the same demand function, so that disagreement in one period yields a smaller but identically shaped Pareto frontier in the next period.

When bargaining is NTU, but Pareto frontiers remain approximately homothetic, then Kalai proportional bargaining can be used to derive a tractable representation of Nash after reindexing bargaining weights. Under homotheticity, Nash bargaining with weight τ yields the same outcome as Kalai proportional bargaining with player-1 share $\tilde{\tau} = \tau s_1'(p^*)/(\tau s_1'(p^*) - (1-\tau)s_2'(p^*))$. As a result, even if the researcher believes in an NTU Nash model, one can apply the math of this paper to construct an equivalent Kalai model with favorable tractability properties.

It is important to note that if one believes the Nash model is the true model, then the Kalai effective weight $\tilde{\tau}$ will be an equivalent representation rather than a structural parameter. If a counterfactual changes the shape of the Pareto frontier while leaving the true Nash weight τ unchanged, then the Kalai effective weight $\tilde{\tau}$ will shift. In addition, if the NTU model fails comprehensiveness — for example, if a supplier's favored input price gives the retailer positive gains from inducing sales — then take-it-or-leave-it offers that correspond to exterior Nash weights τ may no longer correspond to exterior effective $\tilde{\tau}$ bargaining weights. This concern does not apply in my empirical application, as I extend the TU Ho and Lee (2017) model which is comprehensive on a period-by-period level.[11]

As a practical heuristic, the degree of approximation depends on how much of a negotiated price's effect lays outside of the resulting bilateral payment. If the main role of a price is a direct payment, then Nash and Kalai proportional bargaining will produce similar predictions and the choice is one of convenience. If a negotiated price meaningfully affects bilateral quantities through demand responses or cost pass-through, but the generated curve is consistent across periods and the leverage effect is small, then Kalai proportional bargaining will approximate Nash well after an appropriate change of bargaining weights. If an important part of multiperiod contracting is a forward-looking leverage effect, or if the single-period bilateral problem inherently involves time-varying Pareto frontiers, then the resulting nonhomotheticity breaks the approximation.

[11] Technically, forward-looking premium effects could introduce noncomprehensiveness in higher-order bounded rationality approximations that are not estimated here.

When the approximation does break down, the Nash bargaining solution will face serious difficulties. For example, consider a researcher studying how a retailer leverages two suppliers against each other. The first supplier to agree has a leverage effect on the retailer's negotiation with the second supplier. Under Nash, the earlier firm will agree to lower prices in response to the leverage incentive, but the change will be constrained by the threat of returning to the bargaining table with the leverage incentive removed. (Kalai sidesteps the resulting complexity by imposing constant surplus shares regardless of timing.) The resulting nonhomothetic Nash problem is new to the empirical literature, and would reject all previous static NTU and dynamic TU Nash models. Yet even here, this paper offers a path forward for building a future nonhomothetic Nash model. For practical purposes, the empirical question for capturing forward-looking Nash will be whether the change in frontier curvature between adjacent opportunities to bargain is small relative to the modeled gains from trade. Once that point is reached, the remaining differential Nash bargaining equation can be approximated by some Kalai proportional solution, and the remaining dynamics can be simplified by the resulting approximate step-by-step property.

C Additional Multilateral Content

C.1 Nash-in-f Bellman Functions

For conciseness, I write $\sigma_{-i}(2)(a^{(1)}, s)$ for the set of $\sigma_{jk}^{(2a)}(a^{(1)}, s)$ strategies for $j, k \neq i$ and the set of $\sigma_j^{(2b)}$ strategies for $j \neq i$ and abuse notation by writing $a_{-i}^{(2)} = \left(\{\mathbb{C}_{jk}\}_{j,k\neq i}, a_{-i}^{(2b)}\right)$ for the resulting outcome, then value functions satisfy the Bellman equations

$$
\begin{aligned}
V_i^{(1)}\left(a_i^{(1)}, a_{-i}^{(1)}, s\right) &= E\left[V_i^{(2)}\left(\left\{\sigma_{jk}^{(2a)}(a^{(1)}, s)\right\}_{j,k}, \sigma_i^{(2b)}(a^{(1)}, s), \sigma_{-i}^{(2)}(a^{(1)}, s), a^{(1)}, s\right)\right] \\
V_i^{(2)}\left(\{\mathbb{C}_{ij}\}, a_i^{(2b)}, a_{-i}^{(2)}, a^{(1)}, s\right) &= E\left[\begin{array}{ll} & u_i\left(\mathbb{C}, a, x_t, \varepsilon_t\right) + \sum_{j\neq i} \mathrm{Pay}_{ji}\left(\mathbb{C}, a, s\right) \\ + & V_i^{(4)}\left(\left(\mathbb{C}, a, s, x', \varepsilon', i'\right)\right)\end{array}\right] \qquad (13)
\end{aligned}
$$

$$V_i^{(4)}(s') = \beta E\left[V_i^{(1)}\left(\sigma_i^{(1)}(s'), \sigma_{-i}^{(1)}(s'), s'\right)\right].$$

For conciseness in the main text, I also define the time-of-bargaining value functions

$$V_i(\mathbb{C}_{ij}, a^{(1)}, s) = V_i^{(2)}\left(\mathbb{C}_{ij}, \sigma_i^{(2b)}(\ldots), \sigma_{-ij}^{(2)}(\ldots), a^{(1)}, s\right), \tag{14}$$

where strategies are evaluated at their equilibrium values given $a^{(1)}$ and s.

C.2 Further Discussion of Nash-in-f Timing

C.2.1 Potential Pre-Bargaining Actions

I allow actions to be taken before bargaining in Stage 1. This may include decisions like whether or not to allow a contract to renew. Whether actions are allowed before or after bargaining is typically not substantively important: the framework could allow both at the cost of further notation.

C.2.2 Contracts are Bargained

The core of the analysis is the bargaining Stage 2a.

I assume that contracts play a role in payments $\text{Pay}_{ijt} = -\text{Pay}_{jit}$ from player j to player i. A contract can have other characteristics: a fixed end date or an auto-renew clause; a benchmark rule for updating prices while the contract remains in place; restrictions on future actions; or per-period transfers as a function of realized agreements (Olssen and Demirer, 2022; Ho and Lee, 2024).

I also assume that a pair chooses a contract taking as given both future strategies and all strategies for other simultaneous actions. The assumption on future strategies is to my knowledge innocuous. The restriction on simultaneous strategies is the standard passive beliefs assumption that if i or j defects and changes their behavior, neither party adjusts their behavior in, or expectations about, other strategic decisions taken at the same time (Lee

et al., 2021). This decision is relatively innocuous for dynamic models in which short-lived disagreement effectively constrains the feasible value of contract deviations.

C.2.3 Potential Simultaneous Actions are Taken

In Stage 2b, I allow strategic actions to be made at the same time as bargaining. This allows the Nash-in-f model to nest Nash-in-Nash models in which Nash bargaining happens at the same time as other strategic decisions like price-setting (Ho and Lee, 2017; Crawford et al., 2018). It is often most convenient to include occasional actions like premium-setting in this stage, so that repeated disagreement does not lead to new internalized effects of future actions.

C.2.4 Flow Profits are Formed

The representation of profits in Stage 3 as the sum of a pre-transfer flow profit function and negotiated transfers generalizes many, but not all, conceivable dynamic markets. For example, latent utility shocks ε can be captured in the unobserved information i_{t-1}. In my empirical application, I assume that payments under the null contract are zero. I write the model in this section more generally to include disagreement transfers, for example through out-of-contract purchasing (Prager and Tilipman, 2022).

I allow for a negotiation cost to be borne after bargaining succeeds, which I leave implicit in the $u_i(\cdot)$ pre-transfer utility functions. Real negotiation costs are borne both at the stage of preparing for negotiations (Gooch, 2019; ECG, 2020; Fletcher, 2020; Beier, 2020) and at the stage of carefully checking the terms of a potential agreement (STD TAC and Moss, 2014; PMMC, 2019; Fletcher, 2020). I model the bargaining friction as only the ex post cost to validate a potential agreement. Some work includes a sunk negotiation cost (Prager and Tilipman, 2022). Sunk costs can prevent firms from forming Pareto-efficient contracts. In a static model, sunk costs have the advantage of not entering into payments. In a forward-looking model, future sunk costs enter into current payments in a challenging way.

C.2.5 Information is Revealed

The information updating Stage 4 is written in a generic way in order to accommodate many types of markets. Note that I include the $\mathbb{C}_{ijt}$ contract outcome in s_{t+1} but not how off-equilibrium contracts are reached. As a result, I implicitly rule out asymmetric punishments for contract deviations based on the source of the deviation.

C.3 Formal Nash-in-Kalai Microfoundation

C.3.1 The Demands Game

The sequence of discretely timed games is indexed by m. I avoid notation for period lengths by instead showing that as gains from trade tend to zero in m, the implied solution approaches Nash-in-Kalai uniformly. Taking period lengths to zero generally implies increasingly small games from trade, and therefore implies a Nash-in-Kalai limit.

In game m, time is discrete and runs from $t = 0$ to infinity. In period t, a contract structure $\mathbb{C}_t$ emerges. In period t, the contract structure which emerges is a set of lengths hospital i-insurer j lengths $\ell_{ijt,(m)}$ and a set of $i - j$ prices $p_{ijt,(m)} \subseteq \mathcal{P}$, where $\mathcal{P}$ is a closed, convex subset of $\mathbb{R}$.[12] If i and j do not reach a contract, then $\ell_{ijt,(m)} = p_{ijt,(m)} = 0$. I write the set of contracts that emerge as $\mathbb{C}_t$. I assume that at every stage of negotiations, upstream hospitals prefer higher prices while downstream insurers prefer lower prices. A contract corresponds to a constant per-unit rate in place for the full length of the contract: for example, lump-sum payments would fit in the model as a price per period amortized over the contract.

There is no uncertainty. The set of insurer and hospital indices remains the same in every game. I assume that if i and j contract in period t, the agreement will be in place for the known, exogenous value $\ell^*_{ijt,(m)}$.

Timing in period t is as follows:

[12] The game immediately generalizes to other vertical markets by treating hospitals as an upstream market, insurers as a downstream market, and prices as a real-valued numeraire the sides bargain over.

1. The board of directors of every hospital i and insurer j meet with their delegates, who will simultaneously bargain with every potential partner with whom they do not have an agreement.

    - Hospital delegates and insurer delegates choose a price demand to state publicly. The demand is chosen to maximize a weighted average of their employer's net present value profits and a personal concession cost they will face if they agree to a contract that is worse than their demand. The hospital delegates demand a minimum price $\underline{p}^H_{Demand}$ and the insurer delegates demand a maximum price $\bar{p}^M_{Demand}$.

2. The corresponding delegates for each ij pair without a contract in place from the previous period simultaneously meet with their authorized demands. If an ijt pair has jointly feasible demands $\underline{p}^H_{ijt,Demand} \leq \bar{p}^M_{ijt,Demand}$, the delegates reach a jointly feasible contract by Nash bargaining over firm profits, treating their demands as disagreement points and taking equilibrium strategies of other pairs as given. (If this is not well-defined, the pair choose the average demand.)

3. Each delegate of an ij pair without a contract has the opportunity to concede to the other side's demands. Conceding means adopting the other delegate's demand. Without loss of generality, I write concession costs in units of employer net present value profits. Conceding has an associated cost of $c^H_{ijt,(m)}(\underline{p}^H_{Demand} - \bar{p}^M_{Demand})$ and $c^H_{ijt,(m)}(\underline{p}^H_{Demand} - \bar{p}^M_{Demand})$. The concession cost functions are continuous, strictly increasing for positive concessions, equal to zero and have an infinite right-differentiable at zero, and are uniformly lower-bounded by functions with these properties.

4. Each ij pair without a contract meets simultaneously.

    - If the new demands are jointly feasible, they reenter the same joint bargaining process as in Stage 2.

- If the new demands are jointly infeasible, ij do not form a contract in period t.

5. The hospitals and insurers obtain flow profits: $v_{it}^H(\mathbb{C}_t, R_t) = \pi_{it,(m)}^H(\mathbb{C}_t) + r_{i,(m)}^H R_{ijt}$ for hospital i where r_i^H is any new contract negotiation cost, and analogously $v_{jt}^M(\mathbb{C}_t, R_t) = \pi_{jt,(m)}^M(\mathbb{C}_t) + r_{j,(m)}^M R_{ijt}$ for insurer j.

The outcome of each stage is always immediately announced to all parties.

When the delegates arrive with jointly-unachievable demands, they play a one-shot game with a simultaneous payoff matrix adapted from Dutta and depicted in Appendix Table 2. In Appendix Table 2, I write the value functions with agreement (at the anticipated concession decisions) as V^H and V^M and the value with disagreement as V_D^H and V_D^M. The agreement value functions include any effect of the negotiated contract on any other agreements reached through concession in the period. I later show that there is no concession in equilibrium, so that the relevant value functions are also the value functions at the equilibrium simultaneous agreements. If concession costs are too high, neither delegate will be willing to concede to a contract that only improves their employer's profits slightly. The infinite right-derivative at zero ensures that if the demands are close enough, then both delegates will prefer to concede despite the incurred concession cost.

Table 2: After incompatible demands in Stage 2 ($\underline{p}^H > \bar{p}^M$), payoffs in Stage 3 depending on whether hospital delegate (rows) and insurer delegate (columns) concedes or sticks to their initial demands. Table is adapted from Dutta (2024). I omit the ijt subscripts for brevity. p^* is the hypothetical Nash bargained price if both delegates concede and the demands are reversed.

	Concede (C)	Stick (S)
C	$(V^H(p^*) - c^H(\underline{p}^H - \bar{p}^M), V^M(p^*) - c^M(\underline{p}^H - \bar{p}^M))$	$(V^H(\underline{p}^H) - c^H(\underline{p}^H - \bar{p}^M), V^M(\underline{p}^H))$
S	$(V^H(\bar{p}^M), V^M(\bar{p}^M) - c^M(\underline{p}^H - \bar{p}^M))$	(V_D^H, V_D^M)

The concession costs constrain the contracts that can emerge in equilibrium. Suppose jointly compatible demands will lead to the hospital getting the most favorable deal: a take it or leave it outcome that gives the hospital all surplus and leaves the insurer with their disagreement value. Consider a subgame in which the insurer demands a slightly

better deal. The insurer's delegate will not pay a concession cost to concede and obtain the firm's disagreement value. On the other hand, if the new demand is close enough to the disagreement value, the hospital's delegate will prefer to concede (an arbitrarily small cost) to avoid disagreement (a fixed cost to the hospital given m). The same logic can be applied to lopsided deals: the concession costs constrain how much surplus each side can obtain. The higher one side's concession costs, the better of a deal they guarantee themselves in equilibrium. In the limit as the game becomes instantaneous and gains from trade tend to zero, the constraint becomes driven by the derivative of the cost functions at zero.

The model could be generalized in a few directions at the cost of additional notation. I use Nash bargaining in Stage 2 to be tongue-in-cheek. Any other bargaining solution would work. I make contract lengths exogenous to ensure that the space of feasible contract values is convex. The model could likely be extended to enable endogenous contract lengths. Under this game and Dutta (2024)'s game, a delegate pays the same concession cost whether or not the other side concedes; in Dutta (2012)'s original game, the concession cost is paid based on the difference between demanded and realized price, with the same result.

I now write out value functions for the realized contract state. I will assume players follow Markov strategies, so that value functions only depend on realized contracts and negotiation costs (i.e. realized contracts and the previous period's realized contracts). Suppose the players follow strategies $\hat{\sigma}$ which generate period $t+1$ contracts $\hat{\sigma}_{t+1}(\mathbb{C}_t)$. I define the corresponding value functions as:

$$V^H_{it,(m)}(\mathbb{C}_t \mid \mathbb{C}_{t-1}) = \frac{\pi^H_{it,(m)}(\mathbb{C}_t) - \sum_j r^H_{i,(m)} R_{ijt} + \beta_{(m)} V^H_{it+1,(m)}(\hat{\sigma}_{t+1}(\mathbb{C}_t) \mid \mathbb{C}_t)}{1-\beta}$$
$$V^M_{it,(m)}(\mathbb{C}_t \mid \mathbb{C}_{t-1}) = \frac{\pi^M_{jt,(m)}(\mathbb{C}_t) - \sum_i r^M_{j,(m)} R_{ijt} + \beta_{(m)} V^M_{jt+1,(m)}(\hat{\sigma}_{t+1}(\mathbb{C}_t) \mid \mathbb{C}_t)}{1-\beta}.$$

In that equation, R_{ijt} is an indicator for ij forming a new contract in period t.

C.3.2 Continuous-Time Results

I will assume some structure on the value functions which I expect to hold in many vertical market bargaining models. I will make use of the value of ij deviating to a new contract p while holding fixed the outcome of other bargains. I write these unilateral deviation value functions as $V^H_{ijt}(p_{ijt} \mid \hat{\sigma}, \mathbb{C}_{t-1})$ and $V^M_{ijt}(p \mid \hat{\sigma}, \mathbb{C}_{t-1})$. It is not obvious at this stage that unilateral-deviation value functions are the right deviation value in this bargaining game. I show there is no concession in equilibrium, so that these are the relevant value functions.

Assumption 2 (Monotone and differentiable value function)**.** All strategies are Markov. When bargaining in Stage 2 or choosing whether or not to concede in Item 3, hospitals strictly prefer higher prices and insurers strictly prefer lower prices inclusive of any response through subsequent concession decisions and negotiations in period t. The expected value functions of a bargained initial price p at the expected other equilibrium contracts in the same period is written as $V(p \mid \hat{\sigma}, \mathbb{C}_{t-1})$ and is differentiable with bounded derivatives as follows:

$$0 < \varepsilon B \leq \frac{-\partial V^H_{ijt}(p_{ijt} \mid \hat{\sigma}, \mathbb{C}_{t-1})}{\partial p_{ijt}}, \frac{\partial V^H_{ijt}(p_{ijt} \mid \hat{\sigma}, \mathbb{C}_{t-1})}{\partial p_{ijt}} \leq B$$

for uniformly bounding constants $\varepsilon, B > 0$.

Since one side strictly prefers higher prices and the other side strictly prefers lower prices, Assumption 2 allows me to write the value concession game as a value-based price concession game. Assumption 2 could be relaxed to a Lipschitz continuity-type assumption.

The substantive idea of this assumption is that prices have monotonic effects. The first half that includes downstream effects rules out the delegates choosing to form a contract through jointly feasible demands in order to sustain an equilibrium in other simultaneous demands. Under Assumption 2, if a hospital and insurer arrive with jointly feasible demands, they could do better by deviating to the other's demand despite any effects on the downstream contracts.

The value function derivative component of Assumption 2 ensures the hold-fixed contract deviation value functions are differentiable. As a result, the value functions are invertible in bargained prices. For example, if the price domain $\mathcal{P}$ includes only weakly positive prices, then equilibrium hospitals generally prefer strictly higher prices and insurers prefer strictly lower prices. The property is likely to hold in other settings if which higher prices have positive spillovers on other prices for appropriately defined price domains $\mathcal{P}$.

Lemma 3. *Under Assumption 2, for every game m, subgame $\mathbb{C}_{t-1}$, triplet ijt without a contract in place under that subgame, and associated equilibrium strategies $\hat{\sigma}_{(m)}$, there are prices $p^H_{ijt,(m),D}$ and $p^M_{ijt,(m),D}$ that make the hospital and insurer, respectively, indifferent between agreement at that price and disagreement under the expected contracts formed by other pairs in equilibrium.*

I add an assumption to rule out certain nuisance behavior.

Assumption 3. If hjt do not reach a contract in a period t subgame and ijt do reach a contract through negotiation after initial jointly-feasible demands, then hnt continue to not reach a contract if either i or j strengthens their demand.

Assumption 3 rules out a certain edge case in which pairs reach a contract through Nash bargaining between jointly feasible contracts, but neither side can make a stronger demand because it would lead to an anticipated contract that changes other pairs' concession decision. Without Assumption 3, there is no concession in equilibrium, but the contract outcome may be driven by the effect on others' concession decisions in the same period. The content is minimal if, as in my setting, equilibrium networks are fairly complete.

The following lemma shows that there is no concession in equilibrium. As a result, in any equilibrium the firms must negotiate over an individual contract in a way that is optimal taking the outcome of other bargains as fixed. The best deviation over all demands is at least as good as the best deviation over a single demand.

Lemma 4. *Under Assumptions 2 and 3, for every game m with a pure strategy Markov perfect equilibrium $\hat{\sigma}_{(m)}$, every subgame contract is formed through equal demands.*

I obtain Kalai bargaining strategies as the game tends to instantaneous offers and the ratio of first marginal costs tends to some proportion.

Assumption 4 (Sequence of PSMPE tending to instantaneous)**.** As the game index m tends to infinity, bargaining becomes instantaneous in the sense that $\beta(m) \to 1$ and the effect of disagreement becomes negligible: $\max\{V^H(p_D^M) - V^H(p_D^H), 0\}, \max\{V^M(p_D^H) - V^M(p_D^M), 0\} = o_m(1)$ uniformly in subgames and potential bargainers.

This is plausible in many games: the difference between agreeing to a contract now and waiting a second and agreeing to essentially the same contract should be essentially nil.

Assumption 5 (First marginal costs tend to proportional)**.** As the game index m tends to infinity, the ratio of first marginal costs tend to a fixed proportion in the sense that there are finite $w_i^H, w_j^M > 0$ and a sequence of $\delta_n \to 0$ such that $\max_{ij} \sup_{x \in (0, \max\{p_{ijt,D}^M - p_{ijt,D}^H, \delta_n\}]} \frac{c_{ijt,(m)^M}(x)}{c_{ijt,(m)^H}(x)} - \frac{w_i^H}{w_j^M} = o_m(1)$.

The cost proportionality around zero is important to extending the Kalai proportional results from instantaneous to discrete time. It will be important for the microfoundation that as the contract shifts under impasse, the ratio of first marginal costs between the hospital and insurer retain the same proportions. It is not important that the costs be proportional on a price scale in particular, so long as the costs are in the same fixed units — costs could be formed under a payment scale and I would obtain the same result.

The following result follows from an adapted version of Dutta (2012)'s argument.

Proposition 2. *Suppose Assumptions 2, 3, 4, and 5 hold. Then the bargains tend to an instantaneous Kalai proportional solution:*

$$\sup_{\mathbb{C}_{t-1}, R_{ijt,(m)}=1, p_D^M > p_D^H} \left| \frac{V_{ijt}^M\left(\hat{p}_{ijt}(\mathbb{C}_{t-1}) \mid \hat{\mathbb{C}}_{t-ij,(m)}\right) - V_{ijt}^M\left(p_{ijt,D}^M \mid \hat{\mathbb{C}}_{t-ij,(m)}\right)}{V_{ijt}^H\left(\hat{p}_{ijt}(\mathbb{C}_{t-1}) \mid \hat{\mathbb{C}}_{t-ij,(m)}\right) - V_{ijt}^H\left(p_{ijt,D}^H \mid \hat{\mathbb{C}}_{t-ij,(m)}\right)} - \frac{w_i^H}{w_j^M} \right| \to^m 0.$$

I offer the following intuition. The concession costs ensure that in every pure strategy Nash equilibrium, there is no concession and both sides get sufficiently more than their disagreement value that they cannot guarantee a better outcome by demanding more. As the game tends to instantaneous, the gains from trade relative to waiting a period become small and the infinite first marginal costs become binding. The particular form of the constraint is that the ratio of gains from trade tend to the ratio of first marginal costs. This is a scale varying solution concept because the concession costs are made based on prices rather than profits; fixing the ratio of concession costs fixes the relative value of profits.

C.3.3 Discrete-Time Results

The result in Proposition 2 gives a result about disagreeing over an ignorable period of time. Kalai proportional bargaining has a special path independence property that makes this instantaneous-bargaining limit extend to discrete time.

Proposition 3. *Suppose Assumptions 2, 3, 4, and 5 hold. Let* $\tilde{V}^H_{ijt,(m)}(0 \mid \mathbb{C}_{t-1})$ *and* $\tilde{V}^M_{ijt,(m)}(0 \mid \mathbb{C}_{t-1})$ *be the expected values if* ij *disagree in period* t *and remain in impasse until the next period where another pair forms a contract. Suppose the value of agreement relative to impasse is bounded. Then the bargains tend to an* discrete-*time Kalai proportional solution:*

$$\sup_{\mathbb{C}_{t-1}, R_{ijt,(m)}=1, V^M(p)>V^M(0)} \left| \frac{V^H_{ijt,(m)}(\hat{\sigma}(\mathbb{C}_{t-1})) - \tilde{V}^H_{ijt,(m)}(0 \mid \mathbb{C}_{t-1})}{V^M_{ijt,(m)}(\hat{\sigma}(\mathbb{C}_{t-1})) - \tilde{V}^M_{ijt,(m)}(0 \mid \mathbb{C}_{t-1})} - \tau_{ij} \right| \to^m 0.$$

Proposition 3 justifies using a discrete-timed dynamic Kalai proportional bargaining model even when real bargaining is conducted in continuous time: both the value of agreement and of impasse are defined in discrete time. This justifies theoretical and empirical analysis in discrete time.

Only Kalai proportional bargaining justifies estimating a discrete-time bargaining model with a continuous timed underlying microfoundation in general nonstationary games. A

bargaining solution that generally returns the same contract after adding an additional post-disagreement chance to bargain must have proportional character (Roth, 1979). As a result, a comparable result for Nash-in-Nash bargaining would generally yield a differential equation at the margin (Coles and Muthoo, 2003; O'Neill et al., 2004). Nash-in-Nash bargaining in nonstationary environments with access to lump-sum transfers might be able to be microfounded, but only because Nash bargaining with access to lump-sum transfers is transferable utility, and therefore has the same predictions as Kalai proportional bargaining.

C.4 Nash-in-Kalai Moment Details

In the characterization of the Nash-in-Kalai moment, I assume that player i's flow value in period t, v_{it}, is the sum of a pre-payment utility $u_i(\mathbb{C}_t, a_t, x_t, \varepsilon_t)$ and net payments received $\text{Pay}_i(\mathbb{C}_t, a_t, x_t, \varepsilon_t) = \sum_{k \neq i} \text{Pay}_{ki}(\mathbb{C}_t, a_t, x_t, \varepsilon_t)$, where $\mathbb{C}_t$ is the period contract outcome, a_t are period actions (including potential negotiation costs from new contracts), x_t are econometrically observable characteristics that follow an exogenous law of motion, and ε are econometrically unobserved shocks. $[\Delta_{ij} Z_t]$ is the effect of the ij agreement on Z_t at the equilibrium $\mathbb{C}_t$, a_t, and x_t and where the non-ij payment $\text{Pay}_{it,-j}] = \text{Pay}_i(\cdot) - \text{Pay}_{ji}(\cdot)$.

I maintain some regularity conditions.

Assumption 6. (Regularity conditions)
For any state h, the distribution of $(x', i', \varepsilon') \mid h, a_t, \mathbb{C}_t$ follows the law of motion of $(x', i', \varepsilon') \mid (x, i, \varepsilon)$. There is a uniform transversality condition: if $\mathcal{F}_r(h_{t+1})$ is the set of feasible states in period $t + r \geq t$ after state h_{t+1} is reached at the end of period t, then

$$\lim_{r \to \infty} \sup_{h_{t+1}, h_{t+r} \in \mathcal{F}_r(h_{t+1})} \sup_i \beta^r \left| V_i^{(4)}(h_{t+r}) \right| = 0.$$

To define $\text{Pay}_{AD,ijt}$, I use AI to represent the agree-then-impasse path that will replace agreement value functions, and I use DI to represent the disagree-then-impasse path that will replace the disagreement value. I use $[\Delta_{ijt_0} u_{it}^D]$ and $[\Delta_{ijt_0} Pay_{it}^D]$ to describe the difference

Table 3: Count statistics for all hospital–insurer years (All Contracts), hospital–insurer years with modeled insurers (Modeled Contracts), and hospital–insurer years used in bargaining estimation (Estimation Bargains).

Data	Hospitals	Hosp. Systems	MCOs	System-MCO Pairs	System-MCO Years	Bargain Count
All Contracts	38	33	168	613	5108	
Modeled Contracts	35	30	6	159	1482	
Estimation Bargains	32	27	6	53	289	63

between player i's period-t disagreement value if actions follow the AI path (with $\mathbb{C}_{ijt}$ set to the disagreement contract) relative to actions along the DI path, and define the operators applied to u_{jt}^D and Pay_{jt}^D similarly.

Then the dynamic adjustment term is defined as

$$\text{Pay}_{DA,ijt} = E_{t_0}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(\left[\Delta_{ijt_0}u_{it}^D\right]+\left[\Delta_{ijt_0}Pay_{it}^D\right]\right) \\ +(1-\tau_{ij})\left(\left[\Delta_{ijt_0}u_{jt}^D\right]+\left[\Delta_{ijt_0}Pay_{jt}^D\right]\right)\end{array}\right\}\right]. \tag{15}$$

D Additional Empirical Content

D.1 Additional Tables and Figures

MCO	Prospective	Share of Charges
All	46.74	53.26
Modeled MCOs	60.20	39.80
Highmark BCBS	72.27	27.73
HPUOV	56.24	43.76
Other Modeled	13.14	86.86
Nonmodeled	3.03	96.97

Table 4: The estimated share of inpatient payments by benchmark type for fiscal years 2011–16. Prospective contracts were common, especially for Highmark BCBS.

D.2 Key Details for Empirical Analysis

The dataset includes annual panel data reported by hospitals to West Virginia between 2006 and 2015. The state required hospital–insurer agreements to be profitable as part of

Table 5: Additional estimated bargaining parameters. BCBS parameters correspond to Highmark BCBS. "Data" corresponds to average difference between MLR-implied costs per life and estimated average inpatient payments per life insured, and would exactly set the MLR moment to zero for the myopic and forward-looking models. The r^M net negotiation costs are close to their starting point of $10,000 and may weakly identified or unidentified.

	Parameter (τ^{Size} Estimated)							
	η_{BCBS}	η_{HPUOV}	η_{Aetna}	$\eta_{UnitedHealth}$	η_{Cigna}	$\eta_{Carelink}$	r^M_{yBCBS}	r^M_{nBCBS}
Static	3657***	3404***	3658***	2008***	4627***	3139***	10000***	9999***
(Nash/Kalai)	(45)	(85)	(116)	(29)	(32)	(39)	(2614)	(1441)
Myopic	4640***	4036***	3659***	3197***	4624***	3139***	10000***	10000***
(Nash/Kalai)	(14)	(650)	(37)	(374)	(26)	(463)	(1444)	(1)
Forward-Looking	4638***	3631***	3660***	3284***	4626***	3140***	9999***	9999***
($\text{Pay}_{IRT} = 0$)	(130)	(302)	(37)	(69)	(30)	(45)	(29)	(65)
Data	3600	3356	3554	1999	4635	3114		

Note: *p<0.1; **p<0.05; ***p<0.01

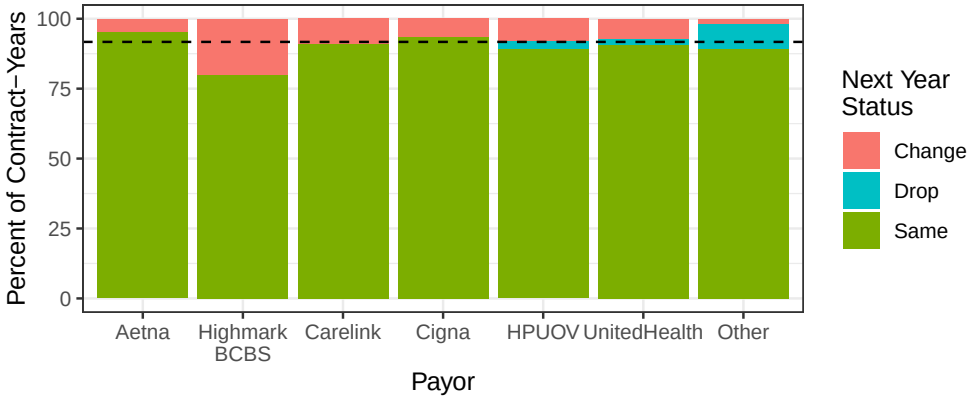


Figure 7: Percentage of auto-renew share of charges contract-years that remained in place (green), were renegotiated (red), or were dropped (blue) for each insurer. The insurers all generally allowed auto-renew contracts to renew, with the average renewal probability of 93.4% indicated by dashed line.

a list price regulation system; the state then required hospitals to report their agreements, so that the state's regulator could sign off on the average-cost floor. The main data I have obtained includes two reports: I have the 2006–15 Discount Contract Lists that report each contract's forecast payment rate as a fixed fraction of list prices for inpatient and outpatient care, as well as the end date for most smaller agreements (auto-renew or fixed end dates). Starting in 2011, the state retained Detailed Contract forms that report start dates, end

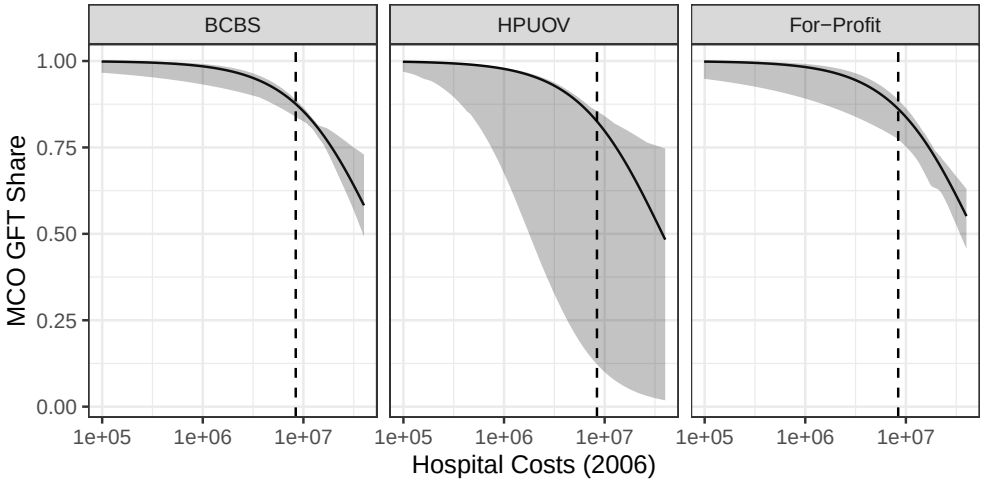


Figure 8: Estimated insurer shares of expected NPV gains from trade under the `myopic` model. The model estimates substantial heterogeneity by hospital system size and limited heterogeneity by insurer, and the estimates are qualitatively similar to the `forward-looking` estimates in Figure 8.

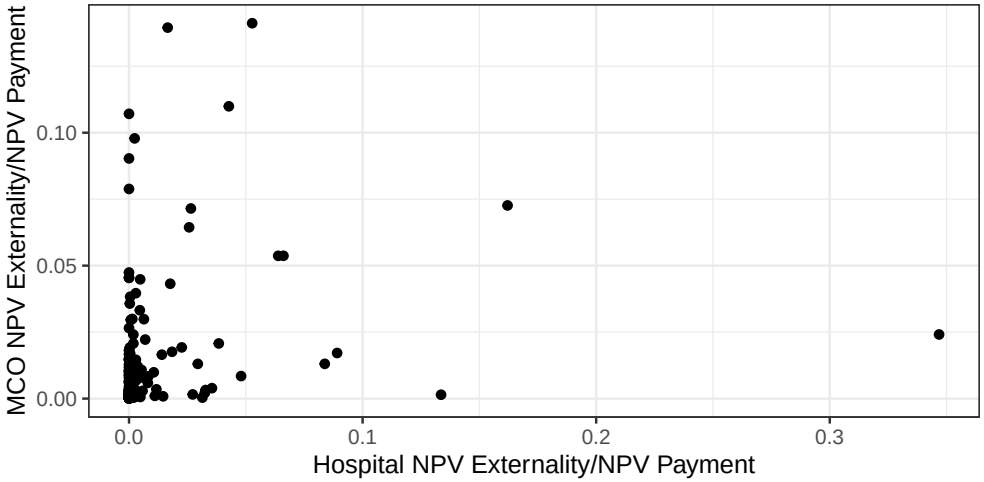


Figure 9: The estimated direct spillovers of each bargained contract on hospital (x axis) and insurer (y axis) payments relative to direct payments under the contract. All estimated insurer spillovers and all but three estimated hospital spillovers are less than 10% of payments under the contract, so that the marginal ratio term in Nash bargaining's gain from trade split would be close to one.

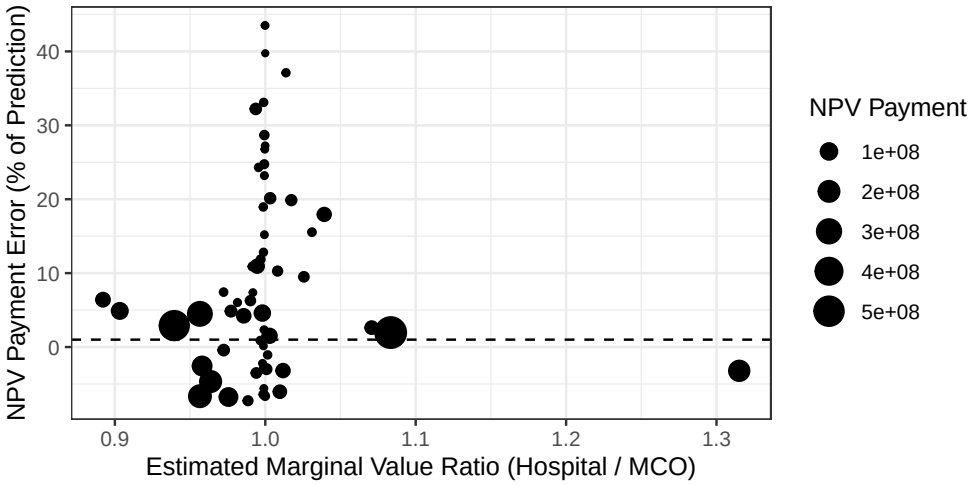


Figure 10: The ratio of realized and model predicted net present value payments (y axis) and marginal value ratio (x axis) for modeled bargains. Size corresponds to the net present value payment, the relevant weight in calculating moments. Most marginal value ratios are close to one (x axis near one) corresponding to similar Nash and Kalai proportional splits of gains from trade. The handful of contracts with larger marginal value ratios have residuals close to zero (y values near horizontal line), so the different equilibrium split models do not seem to be driving the model residuals.

dates, and forecasted quantities for larger agreements. The contract data is public record, and a cleaned version is available at https://jacobdorn.info/files/ContractData.zip. I supplement this data with inpatient discharge data for 2016, as well as various demographic datasets that are described in the companion papers Dorn (2025a) and Dorn (2025b).

I focus my analysis on six insurers that I call modeled: Highmark BCBS, the largest insurer; HPUOV, a nonprofit insurer with a substantial market share in specific regions of the state; and Aetna, Carelink, Cigna, and United, for-profit insurers that were medium-sized players in the state. The hospital competition included a mix of large urban players like Charleston Area Medical Center and WVU Health's flagship Ruby Memorial hospital, mid-sized regional hospitals, and small rural facilities. (Small rural facilities that were declared Critical Access Hospitals by Medicare were exempted from the system I study.)

After cleaning, I have 63 reliable bargains that I use in estimation. Auto-renew contracts often lack reported start dates, but the contracts consistently are negotiated as a fixed fraction of list prices, allowing me to detect contract changes from breaks in fixed long

term discount rates. Importantly, I miss three-year agreements that start and end before 2011 and multiyear agreements that end after 2016. After dropping these and agreements with inconsistent reported start and end dates and focusing on six large- and medium-sized insurers, I obtain a reliable dataset of agreements. I present summary statistics on the hospital–insurer network data and estimation sample in Table 3.

I assume that premiums and demand are negotiated on an annual basis, and each contract is negotiated based on synthetic gains from trade that calculate gains over a weighted average of calendar years from the time of contract formation, taking other prices from the hospital annual reports. This strategy captures the leading bias introduced by a misspecified static model, and leverages Theorem 1 to justify firms negotiating over expected NPV gains from trade. However, this strategy abstracts from intra-year contracting externalities as part of a level-1 bounded rationality approximation in structural estimation. (In Appendix D.3, I find suggestive evidence that this approximation is small in practice.) A finer-grained model that precisely captures intra-year contracting externalities would require the full force of Proposition 1.

The key departure from Ho and Lee (2017) is to allow contracts to be multiyear and staggered. In each pre-bargaining stage 1, insurers and hospitals decide whether to allow their auto-renew contracts to renew. In stage 2a, pairs ij without a contract in place have the opportunity to negotiate a new contract: a commitment length ℓ_{ij}, a starting price per unit of care p_{ij}, and a benchmark b_{ij} that governs price evolution while the contract is in place. At the same time, in stage 2b, insurers set premiums from the year, abstracting from any externality on future price negotiations in the same year. As in a standard static Nash-in-Nash model, the players take as given simultaneous actions; unlike a static model, players take as given future strategies rather than actions.

I follow standard models of hospital and insurer demand. Consumers choose insurance plans based on an insurer's local network quality and premiums. Conditional on receiving a diagnosis ℓ, consumer i chooses a hospital h within the insurer's network of agreements to

maximize

$$u^H_{i,h,\ell} = \delta^H_{h,\ell} + \nu_{i,h,\ell}\rho + \varepsilon_{i,h,\ell},$$

where $\delta^H_{h,\ell}$ is a hospital-diagnosis fixed effect, $\nu_{i,h,\ell}$ capture patient-hospital characteristics (distance, distance squared, and distance interacted with emergency), and ε is a type 1 extreme value shock. I estimate hospital demand using Highbmark BCBS patients in 2016, as all hospitals are in-network for Highmark BCBS; the implicit assumption is that conditional on diagnosis, insurer choice is independent of hospital utility. In the insurer demand stage, the utility of insurer j to consumer i in county c within ACA rating area m in year t is

$$u^M_{i,j,c,m,t} = \gamma_k WTP_{j,k,c,t} + \tilde{\delta}^M_{j,m,t} - \alpha\phi_{j,t} + \xi_{j,k,c,t} + \varepsilon_{i,j,c,m,t},$$

where $WTP_{j,k,c}$ is the ex ante expected utility of the insurer's network in ex post terms (Capps et al., 2003), γ_k are age-dependent coefficients that translate expected ex-post hospital utility into ex ante insurer plan utility, $\tilde{\delta}^M_{j,m,t}$ are insurer–rating-area–year fixed effects that absorb premiums in 2016, and $\xi_{j,k,c,t}$ are age–county unobservables. Insurer demand is identified from variation in network quality within rating areas, and I calibrate a premium sensitivity from Ho (2006)'s estimated own-price elasticity of -1.4.

I allow for insurer negotiation costs r^M_j to be different for BCBS and other insurers. If I allowed i and j to have separate negotiating costs r^H_i and r^M_j, then the net NPV payment contribution would be $(1-\tau_{ij})r^M_j - \tau_{ij}r^H_i$. It is not clear that r^H_i and r^M_j are separately identified, so I assume $r^H_i = 0$ for the purposes of estimating other bargaining parameters. I follow Ho and Lee (2017) by also using moments on reported medical loss ratios to better identify η.

I will highlight some other practical choices. Insurer sales data is limited, so I extrapolate insurer sales from the discharge data's reported residence county and insurer choice. I only use 2016 discharge data, so I extrapolate pre-2016 insurer demand from state-level sales. I

make a correction for attenuation bias from β: if one rephrases Equation (11) as $\sum \beta^t (D_t p_t - \widehat{Dp}_t)$ for predicted payments $\widehat{Dp}$, then larger values of β mechanically have larger absolute errors; I therefore downweight errors by $E[\sum_{t=0}^{mean(t)} \beta^t]$, and find that downweighting by $Mean(\sum_{t=0}^{t^*} \beta^t)$ would lead to similar estimates. I do not have a formal proof that the insurer negotiation costs r_j^M are separately identified from the non-inpatient costs η_j; the estimated values are close to starting points, suggesting that identification of these parameters is a real worry, but at least suggests that any overfitting worries from the negotiation costs are not so bad. The finite horizon assumption allows me to consider the compact interval $\beta \in [0, 1]$, with $\beta = 1$ a limit that is only proper due to the finite-horizon assumption. The back-casting of insurer demand is a limitation of the data that means I can only identify changes in gains from trade driven by prices, premiums, and aggregate demand, and as a result limits the precision of year-to-year variation across pairs. I bootstrap standard errors by resampling inpatient cases and state-level sales, holding constant the set of agreements.

D.3 Is The Level-1 Approximation Strong?

The level-1 bounded rationality approximation sets Pay_{DA,ijt_0} to zero, treating the bargaining problem as approximately TU. This appendix tests this approximation by measuring how far the estimated bargaining problem departs from TU, finding that 97% of bargains are within 10% of TU. Given that benchmark, I ask whether the departures from TU are consistent with Nash bargaining offering better predictions than Kalai proportional bargaining; if anything, the rare departures from TU go in the opposite directions from hypothetical Nash bargaining predictions.

I summarize the degree of NTU through the marginal value ratio for an ij negotiation at time t_0:

$$MV_{ijt_0} = \left(\frac{\partial V_{it_0}^H(p_{ijt_0}^*)}{\partial p_{ijt_0}} \right) \Bigg/ \left(-\frac{\partial V_{jt_0}^M(p_{ijt_0}^*)}{\partial p_{ijt_0}} \right),$$

where $V^H_{it_0}$ and $V^M_{jt_0}$ are the value functions for hospital i and insurer j when negotiating their starting price p_{ijt_0}. Under TU, a dollar increase in p_{ijt_0} increases the hospital's value function by exactly the amount it lowers the insurer's value function, so that $MV_{ijt_0} = 1$. Departures from $MV = 1$ reflect internalized contracting externalities: if the ij price increases subsequent ik or hj prices formed with other parties, then the dollar increase in p_{ijt_0} has a larger gain for the hospital and loss for the insurer, respectively.

I estimate MV_{ijt_0} based on direct price effects. For each non-ij agreement formed by an hk pair during periods t over the five-year finite horizon from t_0, I calculate the current-price response $\frac{\partial p_{hkt}}{\partial p_{ijt}}$ based on the estimated bargaining solution and the direct price effect $\frac{\partial p_{ijt}}{\partial p_{ijt_0}}$ based on observed benchmark-driven price growth. I then calculate how p_{ijt_0} directly affects $V^H_{it_0}$ and $V^M_{jt_0}$, yielding a direct estimate of marginal values. This estimate abstracts from second-order interactions for consistency with the level-1 approximation being tested.

I decompose the marginal value ratio into a direct payment and an externality component. For the hospital, the externality share is

$$ExternShare^H_{it_0} = \frac{\frac{\partial V^H_{it_0}}{\partial p_{ijt_0}} - \sum_{t=t_0}^{\min\{t^*, t_0+4\}} \beta^{t-t_0} D^H_{ijt}}{\sum_{t=t_0}^{\min\{t^*, t_0+4\}} \beta^{t-t_0} D^H_{ijt}}.$$

The formula is similar for the insurer. The numerator removes the direct payment effect from the total marginal value, isolating the indirect effect through internalized contracting externalities relative to the direct effect. A value of zero indicates that the price has no direct effect on an agent's other negotiated transfers.

Figure 9 presents a scatterplot of externality shares. In 89% of estimation bargains, both players have externality effects that are at most 10% as large as the direct payment effect. Further, the contracting externalities push the overall marginal value ratios towards unity. In 97% of the estimation bargains, the marginal value ratio remains within 10% of unity. That is, the bargaining problem is close to TU for nearly all agreements in the estimation sample, which offers reassuring evidence that a level-1 approximation will be adequate for

leading-order questions like the bias in bargaining weights introduced by misspecification.

Given these results, Figure 10 tests whether a Nash bargaining model is likely to do a better job of predicting payments for agreements that depart from TU. As discussed in Appendix B.3, when prices are more valuable to the hospital ($MV_{ijt_0} > 1$), Nash bargaining predicts a higher equilibrium price than Kalai, because Nash responds to the tilted Pareto frontier by shifting the agreement towards the player with the higher marginal value. As a result, if Nash were the true model and Kalai were the estimated model, one would expect to see a positive correlation between MV_{ijt_0} and the NPV payment residual

$$NPVPayErr_{ijt_0} = \sum_{t=t_0}^{\min\{t^*,t_0+4\}} \beta^{t-t_0} D_{ijt}^H (p_{ijt} - \hat{p}_{ijt}),$$

where $\hat{p}_{ijt}$ is the Kalai-predicted price. The statistical power of this test is limited because so many agreements are close to TU, but if anything, the correlation in Figure 10 is the opposite of what Nash would predict. While this test is only heuristic and I have not estimated a forward-looking Nash model, I at least do not find evidence that a more complicated Nash model would offer better fit.

E Lemmas and Proofs

E.1 Lemmas

Lemma 5 (Dutta (2012), Proposition 2)**.** *Suppose $\hat{\sigma}$ is a pure strategy Markov perfect equilibrium of the game I have described in Proposition 2 under Assumption 2 but not necessarily 4 and 5. Suppose i and j could form a strictly Pareto-improving contract in period t. For brevity, I omit the $ijt, (m)$ subscripts. Then there is a unique $y_1, y_2 \in (0,1)$ that satisfies $y_1 + y_2 \geq 1$ and the following property about gains relative to disagreement*

$$V^H\left(y_2 p_D^H + (1-y_2)p_D^M\right) - V^H(p_D^H) = c^H\left((p_D^M - p_D^H)(y_1 + y_2 - 1)\right)$$

$$V^M\left(y_1 p_D^M + (1-y_1)p_D^H\right) - V^M(p_D^M) = c^M\left((p_D^M - p_D^H)(y_1+y_2-1)\right),$$

then the pair (y_1, y_2) *is unique.*

Intuitively, y_1 and y_2 both decrease the left-hand sides to zero but increase the right-hand sides, so there should be a fixed point. In addition, only one of y_1 and y_2 appear on the left-hand side of any given equation. Consider the function $\hat{y}_1(y_2)$ that chooses a $\hat{y}_1$ to hold the first equation with equality at any given y_2. As y_2 increases, the left-hand side of the first equation decreases so $\hat{y}_1 + y_2$ must decrease. Therefore $\hat{y}_1(y_2)$ must decrease faster than y_2. Applying a similar argument to the other equation ensures any fixed point is unique. The next proposition shows the fixed point constrains the equilibrium bargain.

Lemma 6 (Dutta (2012), Proposition 3)**.** *Suppose* $\hat{\sigma}$ *is a pure strategy Markov perfect equilibrium of the game I have described in Proposition 2 under Assumption 2 but not necessarily 4 and 5. Suppose* i *and* j *could form a Pareto-improving contract in period* t*. Then their equilibrium demands are equal and are bounded above and below by* $y_1 p_D^M + (1-y_1)p_D^H$ *and* $y_2 p_D^H + (1-y_2)p_D^M$*, respectively, where* y_1 *and* y_2 *come from Lemma 5.*

Lemma 7. *For a given* ij *pair, let* $y_{1,(m)}$*,* $y_{2,(m)}$ *be the* y_1, y_2 *corresponding to Lemma 5 in game* m *for a given* ij *pair. (If* ij *does not have a strictly Pareto-improving pair, choose some* $y_1, y_2 \in (0,1)$ *satisfying* $y_1 + y_2 = 1$*.) Under the conditions of Proposition 2,* $y_{1,(m)} + y_{2,(m)} \to 1$ *with a convergence rate that is uniform in* (i,j)*.*

E.2 Proofs

Proof of Lemma 1. Proof by contradiction. Suppose f satisfies step-by-step, but is not proportional.

Let $(S, v^D) \in G$ be non-trivial, comprehensive, and convex. By Definition 7, $P(S, v^D)$ contains points s^A and s^B with $s_1^A > s_1^B \geq v_1^D$ and $s_2^B > s_2^A \geq v_2^D$. Write $s^* = f(S, v^D)$. By weak Pareto optimality, $s^* \in P(S, v^D)$. Note that every point $s \in P(S, v^D)$ must have

$s_1 + s_2 > v_1^D + v_2^D$, or else the constraint $s \geq v^D$ from the definition of $P(S, v^D)$ would imply $s = v^D$, but then $(s^A + s^B)/2$ would be an element of S (by convexity) that strictly dominates s (by construction of s^A and s^B), which would imply $s \notin P(S, v^D)$. As a result, the following are uniquely defined:

$$\tau^* = (s_1^* - v_1^D)/(s_1^* + s_2^* - v_1^D - v_2^D)$$
$$\tau^A = (s_1^A - v_1^D)/(s_1^A + s_2^A - v_1^D - v_2^D)$$
$$\tau^B = (s_1^B - v_1^D)/(s_1^B + s_2^B - v_1^D - v_2^D).$$

Because $\tau^A > \tau^B$ by construction, either $\tau^* > \tau^B$ or $\tau^* < \tau^A$. Without loss of generality, suppose $\tau^* > \tau^B$. Note that $s_1^* > s_1^B$: if not, then $\tau^B < \tau^*$ (which was assumed in the construction) would imply $s_2^* < s_2^B$, so that a convex combination of s^A and s^B would dominate s^* and contradict weak Pareto optimality.

Because f is not proportional, there is a game $(\tilde{S}, \tilde{v}^D)$ such that $\tau^*(f_2(\tilde{S}, \tilde{v}^D) - \tilde{v}_2^D) \neq (1 - \tau^*)(f_1(\tilde{S}, \tilde{v}^D) - \tilde{v}_1^D)$. Let $\tilde{\tau}$ solve $\tilde{\tau}(f_2(\tilde{S}, \tilde{v}^D) - \tilde{v}_2^D) = (1 - \tilde{\tau})(f_1(\tilde{S}, \tilde{v}^D) - \tilde{v}_1^D)$. Because $\tau^*(f_2(\tilde{S}, \tilde{v}^D) - \tilde{v}_2^D) \neq (1 - \tau^*)(f_1(\tilde{S}, \tilde{v}^D) - \tilde{v}_1^D)$, at least one side must receive positive gains, so that $\tilde{\tau}$ is uniquely defined.

I now construct a T and a $\tau' < \tau^*$ satisfying $\tau'/(1-\tau') = f_1(T, 0)/f_2(T, 0)$ for $f_2(T, 0) > 0$ based on two cases on $\tilde{\tau}$, which recall by construction is not equal to τ^*.

If $\tilde{\tau} < \tau^*$, take $T = \tilde{S} - \tilde{v}^D$ and $\tau' = \tilde{\tau}$. The desired conditions hold because player 2 must receive positive gains.

If $\tilde{\tau} > \tau^*$, the construction is more involved. Write $\tilde{g} = f(\tilde{S}, \tilde{v}^D) - \tilde{v}^D$. There exists a $t > 0$ small enough such that $t\tilde{g} \leq (s_1^B - v_1^D, s_2^* - v_2^D)$. Take $T = S - v^D - t\tilde{g}$. By convexity and comprehensiveness, $v^D + t\tilde{g} \in S$. By invariance to linear transformations, the game $(v^D + t(\tilde{S} - \tilde{v}^D), v^D)$ is in G. Therefore, by translation invariance and step-by-step applied

to the point $f(v^D + t(\tilde{S} - \tilde{v}^D), v^D) = v^D + t\tilde{g} \in S$:

$$f(T,0) = f(S, v^D + t\tilde{g}) - t\tilde{g} - v^D = f(S, f(v^D + t(\tilde{S} - \tilde{v}^D), v^D)) - t\tilde{g} - v^D = f(S, v^D) - t\tilde{g} - v^D.$$

Then, by construction of t,

$$f_1(T,0) = f_1(S, v^D) - t\tilde{g}_1 - v_1^D \geq s_1^* - s_1^B + v_1^D - v_1^D = s_1^* - s_1^B > 0.$$

Therefore (and because $f_2(T,0) \geq 0$), τ' can be uniquely defined as $f_1(T,0)/(f_1(T,0) + f_2(T,0))$.

I claim that in this case, $\tau' < \tau^*$. Because $f_1(T,0) > 0$, $\tau' > 0$ so that $f_2(T,0) = f_1(T,0)(1-\tau')/\tau'$. Note that $f_1(S, v^D) > v_1^D$, so that $\tau^* > 0$ as well. Note also that since $\tilde{\tau} > \tau^*$ by construction of this case, $\tilde{\tau} > 0$ as well. Then:

$$\begin{aligned}
\frac{1-\tau'}{\tau'} f_1(T,0) = f_2(T,0) = s_2^* - t\tilde{g}_2 - v_2^D &= \frac{1-\tau^*}{\tau^*}(s_1^* - v_1^D) - t\frac{1-\tilde{\tau}}{\tilde{\tau}}\tilde{g}_1 \\
&= \frac{1-\tau^*}{\tau^*}(s_1^* - v_1^D - t\tilde{g}_1) + t\tilde{g}_1\left(\frac{\tilde{\tau} - \tau^*}{\tau^*\tilde{\tau}}\right) \\
&> \frac{1-\tau^*}{\tau^*} f_1(T,0),
\end{aligned}$$

with the final line following because by construction of $\tilde{g}$, in this case $\tilde{g}_1 > 0$ and $\tilde{\tau} > \tau^*$. Therefore $\frac{1-\tau'}{\tau'} > \frac{1-\tau^*}{\tau^*}$, so that $\tau' < \tau^*$.

Thus, I have shown that if f satisfies step-by-step but is not proportional, then there is an (S, v^D) and a $(T, 0)$ such that under without loss of generality conditions, $f(S, v^D) = s^*$, $\tau^*(s_2^* - v_2^D) = (1-\tau^*)(s_1^* - v_1^D)$, $\tau' f_2(T,0) = (1-\tau')f_1(T,0)$, and $\tau' < \tau^*$.

Now, I create the contradiction. Let $\mathcal{C}$ be the convex combination of the points s^*, s^B, and (v_1^D, s_2^B). By comprehensiveness applied to s^B, $(v_1^D, s_2^B) \in S$. By convexity of S, $\mathcal{C} \subseteq S$. By construction, every $s \in \mathcal{C}$ has $s_2 \geq s_2^*$. Also note that every point $s \in \mathcal{C}$ has s_2 equal to a convex combination of s_2^* and s_2^B, so that $s \in \mathcal{C}$ with $s_2 = s_2^*$ implies $s = s^*$.

Now write $s(t) = v^D + tf(T,0)$ and consider the ray for $t \geq 0$. The angle of this ray is

$f_2(T,0)/f_1(T,0) = (1-\tau')/\tau'$, and $\tau' > 0$ so that this angle is finite. The angle of the ray from v^D through s^* is $(1-\tau^*)/\tau^* < (1-\tau')/\tau'$, and the angle of the ray from v^D through (v_1^D, s_2^B) is infinite, so that there is some point $\hat{s}$ on the line segment between s^* and (v_1^D, s_2^B) such that the angle from v^D to $\hat{s}$ is $(1-\tau')/\tau'$. Therefore there is some $\hat{t} > 0$ such that $s(\hat{t}) = \hat{s} \in \mathcal{C}$. And $\hat{s} \neq s^*$, so that by construction of $\mathcal{C}$, $s_2(\hat{t}) = \hat{s}_2 > s_2^*$.

Now I obtain the contradiction. Note that $f(v^D + \hat{t}T, v^D) = v^D + \hat{t}f(T,0) = s(\hat{t}) = \hat{s} \in \mathcal{C} \subseteq S$ by construction of $\mathcal{C}$. By the invariance to linear transformations of G, $(v^D + \hat{t}T, v^D) \in G$. Then, by the maintained assumptions of the step-by-step property on G, translation invariance, and homogeneity,

$$s_2^* = f_2(S, v^D) = f_2(S, f(v^D + \hat{t}T, v^D)) = f_2(S, v^D + \hat{t}f(T,0)) = f_2(S, \hat{s}).$$

But by the constraint $f(S, \hat{s}) \geq \hat{s}$, $f_2(S, s(\hat{t})) \geq \hat{s}_2 > s_2^*$. Contradiction. □

Proof of Lemma 3. Intermediate value theorem applied to the continuous GFT functions. □

Proof of Lemma 4. First, I show there is no concession in equilibrium. Suppose ij concede in equilibrium. If both sides concede, then one delegate could do better by improving their demand and this is not an equilibrium. Suppose one side concedes. Consider that side instead deviating at the demands stage to demand the contract reached. This does not change the demands by any other delegate. This deviation also does not change the expected profit for any other concession decision. Therefore since the strategies are Markov, all other concession decisions are unaffected. Therefore resulting firm profits are unaffected, but the delegate avoids the concession cost and the demand is strictly dominated. Therefore there is no concession in equilibrium.

Now I show by contradiction that there is never a subgame agreement reached by a hospital demanding a strictly lower price than the insurer they negotiate with. By Assumptions 2 and 3, both parties could strictly improve their profits by demanding the contract the other

side demands. Contradiction. Therefore demands are equal in equilibrium of any pair that successfully reaches a contract.

□

Proof of Lemma 5. I am proceeding assuming there is at least one y_1 and y_2 such that $y_1 + y_2 > 1$ and:

$$V^H\left(y_2 p_D^H + (1-y_2)p_D^M\right) - V^H(p_D^H) = c^H\left((p_D^M - p_D^H)(y_1+y_2-1)\right)$$
$$V^M\left(y_1 p_D^M + (1-y_1)p_D^H\right) - V^M(p_D^M) = c^M\left((p_D^M - p_D^H)(y_1+y_2-1)\right).$$

Since $y_1 + y_2 > 1$ and c is increasing for values above 0, the right-hand side of both equations is positive. Therefore the left-hand side is positive, i.e. $y_1, y_2 < 1$.

Now consider more generally the function $\hat{y}_1(y_2) : [0,1] \to [0,1]$ to solve $V^H\left(y_2 p_D^H + (1-y_2)p_D^M\right) - V^H(p_D^H) = c^H\left((p_D^M - p_D^H)(\hat{y}_1(y_2)+y_2-1)\right)$, i.e:

$$\hat{y}_1(y_2) = \frac{(c^H)^{-1}\left(V^H\left(y_2 p_D^H + (1-y_2)p_D^M\right) - V^H(p_D^H)\right)}{p_D^M - p_D^H} + 1 - y_2$$

As pointed out by Dutta in the differentiable case, $\hat{y}_1$ is a continuous function, $\hat{y}_1(1) = 0$, $\hat{y}_1(0) > 1$, and $\hat{y}_1(y_2)$ decreases strictly faster than y_2 since an increase in y_2 by a unit and a decrease in y_1 by one unit leaves $c^H((p_D^M - p_D^H)(y_1+y_2-1))$ unchanged but reduces $V^H(y_2 p_D^H + (1-y_2)p_D^M)$. The function $\hat{y}_2(y_1)$ has the same properties.

By the intermediate value theorem, there there is a fixed point to the function $\hat{y}_1(\hat{y}_2(y_1))$. Since increasing y_1 by ε increases $\hat{y}_1(\hat{y}_2(y_1))$ by strictly more than ε, that fixed point is unique. Since the fixed point (y_1^*, y_2^*) is in $(0,1)$, it must generate positive left-hand sides so that $y_1^* + y_2^* > 1$. □

Proof of Lemma 6. This claim is almost exactly Dutta (2012)'s Proposition 3. Suppose the hospital delegate demands price at least $z_1 p_D^M + (1-z_1)p_D^H$ and the insurer delegate demands price at most $z_2 p_D^H + (1-z_2)p_D^M$. Since there is no concession in equilibrium (Lemma 4), it

must be $z_1 = (1 - z_2)$, so that $z_1 + z_2 = 1$. $z_1 = 1$ and $z_2 = 0$ corresponds to the hospital getting all of the surplus, whereas $z_2 = 1$ and $z_1 = 0$ corresponds to the insurer getting all of the surplus.

After appropriate notation changes, the claim is almost in the setup of Dutta (2012). There is a change to concession costs if both concede, but since that requires bilateral deviation it is irrelevant to the equilibrium and the same result holds. □

Proof of Lemma 7. I proceed for some ij and a sequence of games m satisfying $p_D^M > p_D^H$ and $y_1 + y_2 > 1$; the claim is immediate for the other m.

Recall that $V^H(p_D^M) - V^H(p_D^H), V^M(p_D^H) - V^M(p_D^M) \to_n 0$ by Assumption 4 and $V' \geq \varepsilon B > 0$ by Assumption 2, so that $p_D^M - p_D^H \to_n 0$.

Recall that $y_{1,(m)}$, $y_{2,(m)}$ are defined as the solution to:

$$V^H\left(y_{2,(m)}p_D^H + (1 - y_{2,(m)})p_D^M\right) - V^H(p_D^H) = c^H\left((p_D^M - p_D^H)(y_{1,(m)} + y_{2,(m)} - 1)\right)$$
$$V^M\left(y_{1,(m)}p_D^M + (1 - y_{1,(m)})p_D^H\right) - V^M(p_D^M) = c^M\left((p_D^M - p_D^H)(y_{1,(m)} + y_{2,(m)} - 1)\right).$$

Costs go to zero quickly enough that the infinite right-derivative at zero dominates. The cost functions must tend to zero because $p_D^M - p_D^H \to_n 0$ and $y_{1,(m)} + y_{2,(m)} - 1$ is bounded. The cost functions are also lower-bounded by a function with an infinite right-derivative (Assumption 4). As a result, there is a sequence of $\epsilon_n \to_n 0$ such that $c\left((p_D^M - p_D^H)(y_{1,(m)} + y_{2,(m)} - 1)\right) > B(p_D^M - p_D^H)(y_{1,(m)} + y_{2,(m)} - 1)/(2\epsilon_n)$ for all m large enough and all i, j with $y_{1,(m)} + y_{2,(m)} > 1$.

Note that by construction, $V^H\left(y_{2,(m)}p_D^H + (1 - y_{2,(m)})p_D^M\right) - V^H(p_D^H) \leq B(1 - y_{2,(m)})(p_D^M - p_D^H)$ and $V^M\left(y_{2,(m)}p_D^H + (1 - y_{2,(m)})p_D^M\right) - V^H(p_D^H) \leq B(1 - y_{1,(m)})(p_D^M - p_D^H)$. As a result:

$$B(p_D^M - p_D^H)(y_{1,(m)} + y_{2,(m)} - 1)/\epsilon_n < B(2 - y_{1,(m)} - y_{2,(m)})(p_D^M - p_D^H)$$
$$y_{1,(m)} + y_{2,(m)} - 1 < 2\epsilon_n \to_n 0.$$

Since $1 \leq y_{1,(m)} + y_{2,(m)} \leq 1 + o_m(\epsilon_n)$ for ϵ_n independent of i, j, this completes the proof.

□

Proof of Theorem 1. I define the AI and DI paths and $[\Delta_{ijt_0} u_{it}^D]$ and $[\Delta_{ijt_0} Pay_{it}^D]$ as in Appendix C.4.

Consider a period t_0 Stage 2a subgame in which $i < j$ form a new contract with positive probability. Condition on the event in which they form a contract. Let t^* be the (potentially random) final period in which ij remain in their current contract. Note that $\mathbb{C}_{-ijt}^{AI} = \mathbb{C}_{-ijt}$ and $a_t^{AI} = a_t$ for periods $t = t_0, \ldots, t^*$. It is convenient to write u_{it}^A for the path of u_{it} in equilibrium, u_{it}^{AI} for the path if ij agree for periods t_0 through t^* and then enter impasse, and so on.

It is useful to rewrite the claim in a more compact, but less interpretable, way. The claim can be written as:

$$0 = E_{t_0}\left[\sum_{t=t_0}^{t^*} \beta^{t-t_0} \left\{ \begin{array}{r} -\tau_{ij}\left(u_{it}^A + Pay_{it}^A - u_{it}^{DI} - Pay_{it}^{DI}\right) \\ +(1-\tau_{ij})\left(u_{jt}^A + Pay_{jt}^A - u_{jt}^{DI} - Pay_{jt}^{DI}\right) \end{array} \right\}\right]$$
$$+ E_{t_0}\left[\sum_{t=t^*+1}^{\infty} \beta^{t-t_0} \left\{ \begin{array}{r} -\tau_{ij}\left(u_{it}^{AI} + Pay_{it}^{AI} - u_{it}^{DI} - Pay_{it}^{DI}\right) \\ +(1-\tau_{ij})\left(u_{jt}^{AI} + Pay_{jt}^{AI} - u_{jt}^{DI} - Pay_{jt}^{DI}\right) \end{array} \right\}\right].$$

Recall that by definition of the AI path, the A and AI paths are the same in periods t_0 through t^*. Therefore the claim reduces to showing

$$0 = E_{t_0}\left[\sum_{t=t_0}^{\infty} \beta^{t-t_0} \left\{ \begin{array}{r} -\tau_{ij}\left(u_{it}^{(A,0)} + Pay_{it}^{(A,0)} - u_{it}^{(D,\infty)} - Pay_{it}^{(D,\infty)}\right) \\ +(1-\tau_{ij})\left(u_{jt}^{(A,0)} + Pay_{jt}^{(A,0)} - u_{jt}^{(D,\infty)} - Pay_{jt}^{(D,\infty)}\right) \end{array} \right\}\right].$$

I now add some notation for this proof alone to keep track of paths. Let $t^{(A,d)}$ be the (potentially random) period in which ij reach their $d+1^{\text{st}}$ agreement in equilibrium. Let the superscript (A, d) refer to the path in which all players remain in equilibrium until Stage 2 of period $t^{(A,d)}$, then ij defect to disagreement, and then all players follow equilibrium play.

Note that the subscripts A and (A,∞) refer to the same path, and the subscripts AI and $(A,0)$ refer to the same path. Similarly, let $t^{(D,d)}$ be the period in which ij reach their $d+1^{\text{st}}$ agreement in the subsubgame in which they disagree at the first d opportunities to negotiate, and let (D,d) refer to the path in which ij disagree at every opportunity to negotiate until period $t^{(D,d)}$. Note that the subscripts DI and (D,∞) are the same.

The Kalai proportional bargaining solution requires that:

$$0 = E_{t_0}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(u_{it}^{(A,\infty)}+Pay_{it}^{(A,\infty)}-u_{it}^{(D,0)}-Pay_{it}^{(D,0)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(A,\infty)}+Pay_{jt}^{(A,\infty)}-u_{jt}^{(D,0)}-Pay_{jt}^{(D,0)}\right)\end{array}\right\}\right]$$

and for any finite $t^{(A,d)}$ and $t^{(D,d)}$,

$$\begin{aligned}0 &= E_{t^{A,d}}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(u_{it}^{(A,d+1)}+Pay_{it}^{(A,d+1)}-u_{it}^{(A,d)}-Pay_{it}^{(A,d)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(A,d+1)}+Pay_{jt}^{(A,d+1)}-u_{jt}^{(A,d)}-Pay_{jt}^{(A,d)}\right)\end{array}\right\}\right]\\ &= E_{t^{D,d+1}}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(u_{it}^{(D,d)}+Pay_{it}^{(D,d)}-u_{it}^{(D,d+1)}-Pay_{it}^{(D,d+1)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(D,d)}+Pay_{jt}^{(D,d)}-u_{jt}^{(D,d+1)}-Pay_{jt}^{(D,d+1)}\right)\end{array}\right\}\right].\end{aligned}$$

By the Transverality condition,

$$\begin{aligned}0 &= E_{t_0}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(u_{it}^{(A,0)}+Pay_{it}^{(A,0)}-u_{it}^{(D,\infty)}-Pay_{it}^{(D,\infty)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(A,0)}+Pay_{jt}^{(A,0)}-u_{jt}^{(D,\infty)}-Pay_{jt}^{(D,\infty)}\right)\end{array}\right\}\right]\\ &+\sum_{d=0}^{\infty}E_{t_0}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(u_{it}^{(A,d+1)}+Pay_{it}^{(A,d+1)}-u_{it}^{(A,d)}-Pay_{it}^{(A,d)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(A,d+1)}+Pay_{jt}^{(A,d+1)}-u_{jt}^{(A,d)}-Pay_{jt}^{(A,d)}\right)\end{array}\right\}\right]\\ &+\sum_{d=0}^{\infty}E_{t_0}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(u_{it}^{(D,d+1)}+Pay_{it}^{(D,d+1)}-u_{it}^{(D,d)}-Pay_{it}^{(D,d)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(D,d+1)}+Pay_{jt}^{(D,d+1)}-u_{jt}^{(D,d)}-Pay_{jt}^{(D,d)}\right)\end{array}\right\}\right]\\ &= E_{t_0}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(u_{it}^{(A,0)}+Pay_{it}^{(A,0)}-u_{it}^{(D,\infty)}-Pay_{it}^{(D,\infty)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(A,0)}+Pay_{jt}^{(A,0)}-u_{jt}^{(D,\infty)}-Pay_{jt}^{(D,\infty)}\right)\end{array}\right\}\right]\\ &+\sum_{d=0}^{\infty}E_{t_0}\left[E_{t^{A,d}}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r} -\tau_{ij}\left(u_{it}^{(A,d+1)}+Pay_{it}^{(A,d+1)}-u_{it}^{(A,d)}-Pay_{it}^{(A,d)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(A,d+1)}+Pay_{jt}^{(A,d+1)}-u_{jt}^{(A,d)}-Pay_{jt}^{(A,d)}\right)\end{array}\right\}\right]\right]\end{aligned}$$

$$
\begin{aligned}
&+\sum_{d=0}^{\infty} E_{t_0}\left[E_{t^{D,d+1}}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r}-\tau_{ij}\left(u_{it}^{(D,d+1)}+Pay_{it}^{(D,d+1)}-u_{it}^{(D,d)}-Pay_{it}^{(D,d)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(D,d+1)}+Pay_{jt}^{(D,d+1)}-u_{jt}^{(D,d)}-Pay_{jt}^{(D,d)}\right)\end{array}\right\}\right]\right]\\
&=E_{t_0}\left[\sum_{t=t_0}^{\infty}\beta^{t-t_0}\left\{\begin{array}{r}-\tau_{ij}\left(u_{it}^{(A,0)}+Pay_{it}^{(A,0)}-u_{it}^{(D,\infty)}-Pay_{it}^{(D,\infty)}\right)\\ +(1-\tau_{ij})\left(u_{jt}^{(A,0)}+Pay_{jt}^{(A,0)}-u_{jt}^{(D,\infty)}-Pay_{jt}^{(D,\infty)}\right)\end{array}\right\}\right],
\end{aligned}
$$

which is the desired claim.

□

Proof of Proposition 1. I apply Theorem 1, which has already been proved in this document.

Let each of the N pairs ij negotiate in periods t_{ij}, $t_{ij}+T$, ... Recall that by definition, there are z distinct values of t_{ij}. Call the set of z ordered contracting periods $\boldsymbol{P}$. By construction, $z \leq \min\{N, T\}$.

Let $t_{ij}^{(d)}$ be the periods in which ij bargain after disagreeing in d periods that have a bargain in equilibrium. More formally, if t_{ij} is the h_{ij}-th element of $\boldsymbol{P}$, let $t_{ij}^{(d)}$ be the $(h_{ij}+d)$-th element of $\boldsymbol{P}$, indexed cyclically. (That is, if $h_{ij}=7$, $z=9$, and $d=5$, then $t_{ij}^{(5)}$ will advance 5 steps from 7 to $h+d=12$, but then wrap around to the $h+d-z=3$-rd element of $\boldsymbol{P}$.)

With some abuse of notation, let $s(\boldsymbol{d})$ denote the bargaining state in which each pair ij will bargain at periods $t_{ij}^{(d_{ij})}+kT$. The equilibrium path corresponds to $s(0,\ldots,0)$. If a pair ij disagrees at $s(\boldsymbol{d})$, then the resulting subgame is $s(\boldsymbol{d}+e_{ij})$, where e_{ij} is the unit vector for the pair ij. By step-by-step applied inductively, I will show that any period t that is not in $\boldsymbol{P}$ (indexed cyclically) can be assumed to result in disagreement for any pair that bargains in that period in equilibrium.

The subgame $s(\boldsymbol{d})$ determines a set of price-response functions $p_{ij}^{*}(\{p_{hk}\}, \boldsymbol{d})$ that map anticipated hk prices that emerge in periods $t_0 = t_{ij}^{(\boldsymbol{d}_{ij})}+kT$ into price-responses p_{ijt_0}. By construction, p_{ij}^{*} is constant in p_{ij}; the pair's own anticipated outcome does not directly affect their chosen outcome.

Let the equivalent Nash-in-Kalai model be a model in which each pair ij commits to dis-

agree for every period $t \notin \boldsymbol{P}$, indexed cyclically, but otherwise follow the subgame strategies of the original Nash-in-Kalai model. Then it remains to show that:

(a) This equivalent model has at most z^N market configurations.

(b) This equivalent model has at most $N \min\{z^N, T^{N-1}\}$ price-response functions.

(c) This equivalent model generates the same equilibrium price-response functions as the original model.

(d) This equivalent model is a Nash-in-Kalai model (other than the off-equilibrium Pareto-inefficiency caveat).

The first three items are straightforward. Item (a) follows by counting: each ij pair can disagree up to z times, yielding z^N feasible values of $\boldsymbol{d}$. Item (b) then follows by counting: each market configuration generates one price-response function per pair, yielding Nz^N price-response functions; the relative-positioning argument of the generic Nash-in-f model yields a bound of at most NT^{N-1} price-responses as well. Item (c) follows by construction: this model is constructed to have players follow the original subgame strategies, so the players follow the equilibrium price-response functions in the equilibrium subgame.

It remains to verify Item (d). It is clear that the model remains in the Nash-in-f structure. It remains to show that the solution concept satisfies the constraints of a Nash-in-Kalai model:

(i) There are value functions calculated under these strategies.

(ii) If $t \in \boldsymbol{P}$ (indexed cyclically), the strategies satisfy the appropriate notion of equilibrium generated by these value functions.

(iii) For each ij, there is a $\tau_{ij} \in [0, 1]$ such that each ij pair's outcome is the outcome of $f_{\tau_{ij}}^{(Kalai)}$ applied to the subgame problem.

The statement of Proposition 1 specifically excludes checking equilibrium behavior for $t \notin \boldsymbol{P}$.

Item (i) follows by constructing value functions as follows. If $t \in \boldsymbol{P}$ or $t \notin \boldsymbol{P}$ but no players bargain in this subgame in the original model, then take the value functions from the original model. If $t \notin \boldsymbol{P}$ and at least one pair bargains in period t in this subgame in the original model, then take the value functions from assuming that each pair disagrees until the next $t \in \boldsymbol{P}$ and then consulting the original model's subsequent value functions. By Item (i) applied to the original Nash-in-Kalai model and inspection of behavior applied to other states, the resulting value functions are the value functions calculated under these strategies.

To verify Item (ii), fix some subgame $t \in \boldsymbol{P}$ in which at least one pair negotiates, and fix one such pair ij. Let t' be the next $t \in \boldsymbol{P}$; if $t = \boldsymbol{P}_z$, let $t' = \boldsymbol{P}_1 + T$.

I verify weak Pareto optimality: ij could not do jointly better by deviating to an off-equilibrium agreement, and i and j could not do unilaterally better by deviating to disagreement. Deviating to an off-equilibrium ij agreement generates the same value of that agreement as in the original model, because the subsequent price-response functions coincide with the original model by Item (c). Since the original model satisfies Item (ii) in all subgames, the equivalent model satisfies Item (ii) in this subgame. For the disagreement value, deviating to disagreement generates the value of ij disagreeing in t'. If $t' = t + 1$, the constraint is automatic by Item (ii) applied to the original model, so suppose $t' > t + 1$. In the original model, deviating to disagreement generates the value of ij disagreeing in t and agreeing again in $t + 1$. By weak Pareto efficiency of Kalai proportional bargaining, the $t + 1$ outcome is weakly better than disagreeing until $t + 2$. Applying weak Pareto efficiency of Kalai proportional bargaining to the original model in the intermediate periods $t + 1, \ldots, t' - 1$, I obtain that the original model's value of bargaining at $t + 1$ is weakly better for both players than the equivalent model's value of disagreeing until t'. Weak Pareto efficiency then follows by Item (iii) for the original model.

Item (iii) follows by step-by-step. Let $\{\tau_{ij}\}$ be as in the original model. Fix some subgame

$t \in \boldsymbol{P}$ such that at least one pair negotiates, fix one such pair ij, and recall the definition of t'. I must show that in the equivalent model, the ij agreement at t splits gains from trade in proportion τ_{ij} and both sides do weakly better than disagreeing. For proportionality, the original model contains a series of off-equilibrium bargains between t and t', each of which splits gains proportionally to τ_{ij}. By applying step-by-step inductively along these bargains, one obtains that the agreement in t splits gains (defined relative to the equivalent model's disagreement until t') proportionally to τ_{ij}. Weak Pareto efficiency follows by the inductive argument of Item (ii). Finally, by Item (iii) applied to the original model (and recalling that value functions are unchanged by Item (i)), there is no better agreement that splits further gains proportionally to the τ_{ij}. Therefore the outcome is weakly Pareto efficient and splits gains proportionally to the τ_{ij} ray, so that the outcome of the ij bargain in t is the outcome of $f^{(Kalai)}_{\tau_{ij}}$ applied to the subgame problem. □

Proof of Corollary 1. Under the assumptions of Corollary 1, $\text{Pay}_{DA,ijt} = 0$ in all periods by inspection. □