Dynamic Bargaining between Hospitals and Insurers

Jacob Dorn*

November 5, 2024

LATEST VERSION HERE.

Abstract

This paper introduces the Nash-in-Kalai model, a new bargaining framework designed to tractably capture forward-looking dynamics in markets with staggered contracting. The Nash-in-Kalai model aligns with the usual Nash-in-Nash model in many leading static cases, but uniquely controls state space growth in dynamic settings. I use novel data on hospital–insurer contracts in West Virginia, which are typically trade secrets. Typical contract durations range from three years to a decade or longer. I use this data to provide the first estimated structural model of vertical market bargaining with staggered contracts. I find that negotiators were substantially forward-looking: I reject the null hypothesis of myopia and estimate an annual discounting rate of $\beta = 0.899$. I quantify a new counterfactual effect of a proposed government-driven reform on private payment dynamics. I find that the reform would increase private spending after nine years by $4.98 billion, while a myopic model lacking forward-looking responses would overestimate the effect by $2.35 billion and miss short-term dynamics, including the possibility of payments decreasing.

1 Introduction

Many vertical markets feature multiyear contracts formed through bilateral bargaining. In industries such as telecommunications, manufacturing, and the $400 billion American

*This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2039656. 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. I am indebted to many conversations with my advisors, Kate Ho, Alessandro Lizzeri, and Ulrich Müller; the generous help of Michael Morris, Stacy Pridemore, Barbara Skeen, and Chessie Short; and helpful conversations with David Ahn, Eduard Boehm, Nick Buchholz, Colleen Carey, Allan Collard-Wexler, Zack Cooper, Janet Currie, Gregory Dobbels, Barton Hamilton, Katherine Hempstead, Heather Howard, Jakub Kastl, Harry Koos, Victoria Larsen, Quan Le, W Bentley MacLeod, Lukas Mann, Tyler Maxey, Mikkel Plagborg-Møller, Eric Qian, Radhika Ramakrishnan, Jackson Reimer, Jesse Silbert, Suren Tavakalov, Yuci Zhou, Esmée Zwiers, and many other generous people whom I omit for brevity. ChatGPT was used in editing. All errors are my own.

hospital–insurer market, contracts often use price indexing schemes to manage payment dynamics over multiple years. Vertical market bargaining models generally assume that contracts are short-lived and formed simultaneously through a Nash equilibrium in Nash bargains. This static framework has allowed for tractable answers to important static questions, but it cannot quantify dynamic effects in which contracts are staggered.

This paper studies vertical market bargaining with the potential for staggered contracting. I use a novel panel dataset on hospital–insurer contracts from West Virginia to demonstrate that staggered contracting is a quantitatively important phenomenon. Typically, vertical market contracts are trade secrets. I show that a typical contract duration could be as short as three years or as long as 15 years. A static model with one-period contracts cannot accurately capture contract dynamics or anticipatory responses to future changes, calling for a new forward-looking model. Multiyear and staggered contracting is likely to be pervasive in many markets (Joskow, 1987, Abbey, 2012, Marcelo, 2021).

I use a simple illustrating example to demonstrate the tractability challenges introduced by staggered contracting. Imagine a monopolist insurer that negotiates unit-care prices with two symmetric hospitals. Consumers become patients who each require one unit of care, and those consumers split equally within the insurer's network of hospitals that have contracted with the insurer. If the insurer disagrees with the first hospital, then all patients will be forced to go to the second hospital. In such a model, an increase in the second hospital's price has a spillover effect of increasing the insurer's cost of providing care to patients if the insurer only contracts with the second hospital. The second hospital's price thus increases the insurer's willingness to pay for the first hospital's agreement.

These bilateral contract spillovers have previously been modeled as static. In a static model, contracts are short-lived and formed at the same time, so that spillovers only occur through anticipated simultaneous agreements. The smaller literature on dynamic Nash-in-Nash still assumes simultaneous period-by-period contracting with no internalized price spillovers. Suppose the illustrating insurer negotiates one-year contracts every six months. When negotiating with the first hospital today, the insurer knows that a higher first hospital agreement price today will lead to more costly disagreement, and a higher negotiated price, when bargaining with the second hospital in six months. If negotiators are forward-looking and care about future profits, staggered contracting price spillovers are internalized and generate new time-dependent bargaining states. I show that even the threat of internalizing spillovers after disagreement can yield a complicated and opaque recursive Nash-in-Nash bargaining solution, often for little empirical gain.

To address the limitations of the existing literature, I propose the Nash-in-Kalai model. The Nash-in-Kalai solution is a Nash equilibrium in Kalai (1977) proportional bargains.

The approach has the exact same predictions as Nash-in-Nash in leading static cases, but is uniquely tractable in dynamic markets. I show that the Nash-in-Kalai model uniquely controls state space growth and yields an intuitive moment on the expected net present value (NPV) payment. The expected NPV payment moment coincides with the Nash-in-Nash moment under myopia, which corresponds to static-type estimation with accurate timing. There is no such moment under forward-looking Nash bargaining.

The tractability advantages of the Nash-in-Kalai solution come from a static utility property called *step-by-step*. The step-by-step property is a way of iteratively constructing the Kalai proportional solution in utility space: first bargaining over a reduced utility frontier, and then using the smaller agreement as the disagreement point in the full bargain. In a dynamic setting with recursively defined disagreement, the step-by-step property can be repeatedly applied to replace the Nash-in-Kalai disagreement point with the value of repeated disagreement. This effective disagreement point corresponds to Binmore et al. (1989)'s impasse point, and involves fewer bargaining states. This makes the Nash-in-Kalai model uniquely suitable for studying settings in which negotiators are forward-looking about long-lived contracts, while retaining the tractability and intuition of the static Nash-in-Nash approach of previous work.

My empirical application leverages the new West Virginia contract panel dataset to provide the first estimated structural model of vertical market bargaining with staggered contracting. My approach is similar to the existing empirical literature, but with two key differences made possible by the panel dataset: I estimate bargaining parameters based on gains from trade available at the moment of bargaining, and I allow negotiators to balance gains from trade across multiple periods. As in the Ho and Lee (2017) model from which I build on, the key inputs to bargaining are hospital and insurer demand. I estimate hospital demand using patients with a complete network, and I estimate insurer demand and bargaining parameters through GMM. I estimate an annual discounting rate of $\beta = 0.899$, strongly rejecting the null hypothesis of myopia. I find that the forward-looking model is more accurate than the myopic model, even for short-term predictions. My estimates suggest that a hypothetical dynamic Nash-in-Nash model would be close to the estimated Nash-in-Kalai model, underscoring the usefulness of Nash-in-Kalai as a tool for tractably capturing key dynamics in dynamic markets whether the Kalai proportional solution is exactly or just approximately true.

I apply my estimated model to quantify the counterfactual impact of proposed increases to Medicare rate inflation on the private insurer market, focusing on Medicare's role as a private payment benchmark and incorporating forward-looking responses. American hospitals have often argued that the payments they receive from Medicare, a large government program,

have failed to keep pace with costs. To simulate a more favorable dynamic, I consider a one percentage point annual increase in Medicare rates, roughly counteracting observed deflation relative to reported costs. The change is implemented as a surprise announcement at the end of 2006 that is fully known in every subsequent period.

I find that persistent Medicare payment increases would have real effects on the payments negotiated by commercial insurers, even after accounting for forward-looking offsets. I estimate that after nine years, this change to Medicare rates would increase spending on behalf of the commercially insurer by 1.319%. This effect is 25.4% of the predicted increase if negotiators did not adjust their contracted multiples. Extrapolated nationally in 2015 and converted to 2019 dollars, the effect corresponds to a $4.98-billion increase in spending. I find that a key mechanism is negotiators' forward-looking offset of anticipated future Medicare-driven price increases by reducing starting prices, to the point that I estimate spending in one year would decrease. Approaches based on the static literature would predict incorrect effects. A myopic model that corresponds to static-type estimation with accurate timing would overestimate the 2015 effect by $2.35 billion and miss short-run dynamics. A static model with single-period contracting would fail to capture any effect at all.

The vertical market bargaining literature has largely focused on static models, with modern empirical work generally using the Nash-in-Nash framework, which corresponds to a Nash equilibrium in Nash bargains. Most Nash-in-Nash models are static and include only a single period (Crawford et al., 2018, Ho and Lee, 2019); a smaller literature considers simultaneous period-by-period contracting with inertia (Lee and Fong, 2013, Liu, 2021, Tiew, 2022, Deng et al., 2022, 2023). This literature cannot accurately capture dynamics of staggered contracting as I show occurs in practice. I propose the Nash-in-Kalai model, a Nash equilibrium in Kalai (1977) proportional bargains. Other bargaining solutions like the Kalai and Smorodinsky (1975) solution lacks the same dynamic advantages. The Nash-in-Kalai model coincides with the Nash-in-Nash model in transferable utility (TU) settings in which negotiations are zero-sum at the margin. TU bargaining includes the microfoundations of period-by-period Nash-in-Nash (Lee and Fong, 2013, Collard-Wexler et al., 2019) as well as common empirical practice (Ho and Lee, 2017, Crawford et al., 2018, Ho and Lee, 2019). Some empirical work has estimated static models with nontransferable utility (NTU) bargaining (Crawford and Yurukoglu, 2012, Grennan, 2013, Gowrisankaran et al., 2015). Even in these NTU models in which Nash-in-Nash and Nash-in-Kalai do not coincide, Nash bargaining is often interpreted as nearly following a Kalai proportional strategy of splitting gains from trade proportionally to bargaining weights (Brügemann et al., 2018, Ghili, 2022). There is empirical work on dynamic bargaining over a single good (Backus et al., 2020, Keniston et al., 2024) and search-on-the-job (Cahuc et al., 2006, Bagger et al., 2014, Bilal et al., 2022, Gottfries, 2022),

typically with TU frameworks. My counterfactual analysis relates to work on price-setting dynamics in the context of market inflation (Taylor, 1980, Calvo, 1983) and pharmaceutical price caps (Abbott, 1995). I build on this literature by offering the first structural estimates of vertical market bargaining with staggered contracts, and quantifying a new dynamic mechanism in an important market. I discuss other related work further in Appendix A.

The remainder of this work is as follows. Section 2 documents that West Virginia contracts were multiyear and formed at different times. Section 3 defines the Nash-in-Nash model, illustrates the state space disadvantage of a dynamic Nash-in-Nash model, and presents the Kalai proportional bargaining solution. Section 4 defines the dynamic Nash-in-Kalai solution and proves key tractability benefits. Section 5 describes the empirical model, data, estimation procedure, parameter estimates, and estimated counterfactual effects. Section 6 concludes.

2 Motivating Descriptives

I begin by motivating the need for a dynamic bargaining model of vertical markets: vertical market contracts are multiyear and formed at different times. I am able to say this definitely using a public record dataset on West Virginia hospital–insurer contracts between 2006 and 2015. Static models are used in part because it is rare to see the data on contract timing (Sorensen, 2003, Reinhardt, 2006, Gudiksen et al., 2019), making the setting an exciting opportunity to understand vertical market dynamics.

I am able to amass a panel dataset on hospital–insurer contracts because West Virginia made the contracts public records. The contract disclosure was a byproduct of a "corridor" rate regulation system. The state mandated a dynamic ceiling on hospital list prices and an average cost floor on private insurer payments. The regulator in charge of certifying the average-cost floor made the regulated contracts public records. The state destroyed the contracts at the end of the corridor system in 2016, but retained scans of contract summary information that I leverage for estimation and counterfactuals.

The contract dataset I use includes Discount Contract Lists (DCLs) and Detailed Contract (DC) forms. The DCLs are a panel dataset of projected discounts relative to list prices by contract for 2006–2015. Appendix Figure 10 gives one example.[1] I use these reports to observe payment rates and infer some negotiation times. The DC reports were supplemental information for contracts that either constituted at least 5% of a hospital's projected utilization or which fell into certain rare exceptions, and were retained beginning midway through

[1]The reports exclude Medicare Advantage contracts. Medicare Advantage is a large and ostensibly commercial insurance product that is funded by Medicare and often included with traditional Medicare (CMS, 2022).

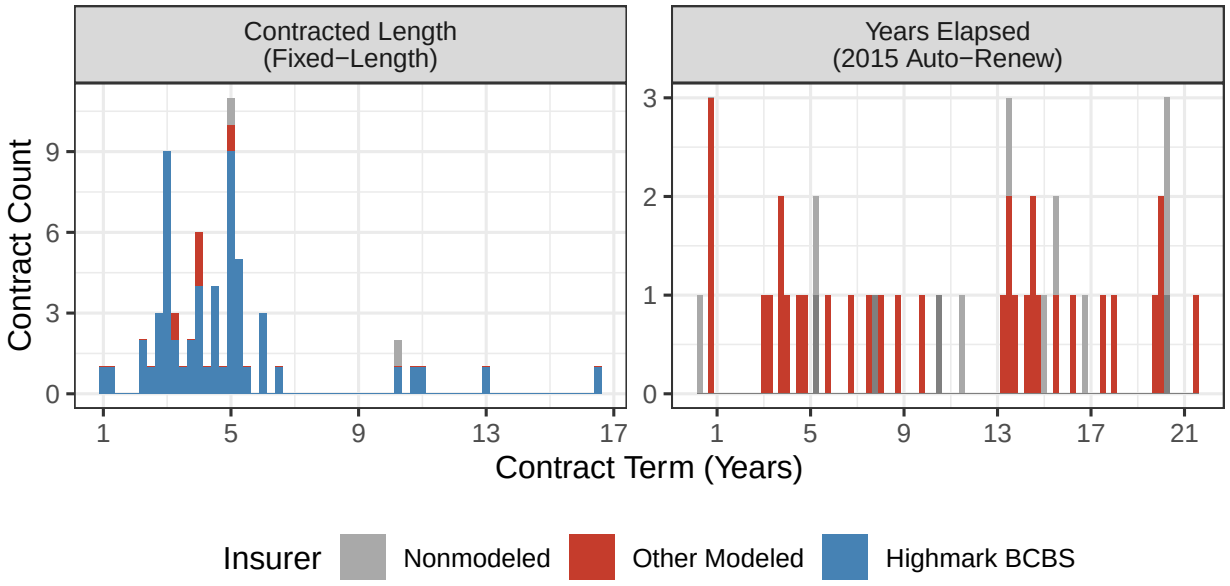


Figure 1: Contracts are multiyear. Distribution of reported contract lengths for fixed-term contracts in general (left panel) and auto-renew contracts with reported formation dates as-of fiscal 2015 (right panel). Colors indicates insurer. The two nonmodeled contracts correspond to Wheeling–Pittsburgh Steel.

2010. Appendix Figure 11 presents one example. The DC data includes detailed information on scale, as well as information on formation and expiration dates (if applicable).[2]

The contract data is public record, so I am able to reference specific firms and provide the data for other researchers to use. I ultimately focus my estimation on six insurers that were more likely to report retrospective formation dates: Highmark BCBS; a regional insurer, HPUOV; and the four largest for-profit firms, Aetna, Carelink, Cigna, and UnitedHealth.[3] I refer to these insurers as "modeled" and group the remaining tail of small insurers into the "nonmodeled" category in this section's plots. The cleaned dataset is available at `https://jacobdorn.info/files/ContractData/ContractData.zip`. For more on the state's rate review system, see Murray and Berenson (2015). For more on the setting and data, see the companion paper Dorn (2024).

Figure 1 demonstrates that hospital–insurer contracts in West Virginia were multiyear. Figure 1(a) includes contracts with reported formal expiration dates. These contracts were associated with Highmark BCBS, and corresponded to standard lengths of three or five years.

[2]There are also reported submission dates. I use contract approval date, rather than submission date, as a measure of contract negotiation and start date because approval was generally quick, while contract submission date is sometimes used to refer to a recent resubmission of an extant contract.

[3]Carelink was a regional subsidiary of Coventry between 1999 and Aetna's acquisition of Coventry at the end of 2014.

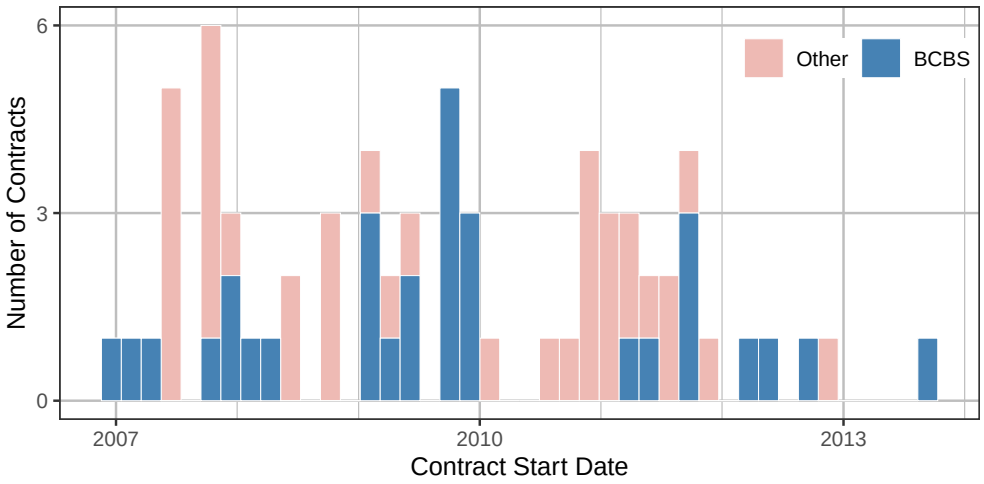


Figure 2: Contracts are formed at different times. Histogram of contract start dates for contracts used in the estimation sample and introduced in 2007–2014 for Highmark BCBS (blue) and other modeled insurers (pink). Vertical lines indicate January 1 of a given year. Contracts were not systematically introduced on the same dates.

The right panel includes "auto-renew" contracts: contracts with an annual commitment that would automatically renew until one side requested termination. Auto-renew contracts were associated with smaller contracts that would rarely be reported in the DC data, and so I can generally only infer a formation dates for larger auto-renew contracts. Figure 1(b) shows that even many of the larger auto-renew contracts with reported expiration dates in 2015 had been formed at least a decade previously. In companion work, Dorn (2024) estimates that the average auto-renew renewal probability in this era was 93.4%.

Figure 2 demonstrates that contracts were formed at different times, even within a given year. The figure plots contract start dates for contracts used in the estimation sample. These are bargains with reliable start and end date information in my sample. Contracts were visually formed in many years and all throughout a given year. This pattern of staggered formation held both within Highmark BCBS and across insurers.

Figures 1 and 2 show that a forward-looking model is needed to accurately predict dynamic effects. Forward-looking negotiators can respond to changes in future conditions when choosing starting terms. The typical static approach to vertical market bargaining assumes all contract structures in place in a market were formed recently and simultaneously. There is no static representation of staggered contracts. Static-type estimation can be applied to staggered contracts through a model that imposes myopia, so that negotiators only think of one period of profits and fail to offset even predictable changes to future conditions.

In Section 5, I use the West Virginia contract data to provide the first structural estimates of a vertical market bargaining model with staggered contracts. The sections leading up

to the empirical estimates develop the necessary forward-looking model, beginning with a static Nash-in-Nash framework for simultaneous contracting and culminating in a dynamic Nash-in-Kalai model with staggered contracting. I will ultimately find that negotiators are forward-looking, and as a result substantially offset my counterfactual change of interest substantially, but incompletely.

3 Static Nash-in-Nash and Nash-in-Kalai

This section describes the existing Nash-in-Nash model and illustrates why I propose a different characterization for dynamic vertical markets. I start with the usual static Nash-in-Nash model, which by construction cannot capture dynamic counterfactuals with staggered contracting. I use a example to illustrate that in the presence of potential staggered contracting, dynamic Nash-in-Nash introduces new tractability challenges, often with limited empirical gain. I then define the Nash-in-Kalai solution in static markets.

3.1 Static Nash-in-Nash

I begin with the static Nash-in-Nash model, which is a tractable solution concept for simultaneous bargaining but turns out to be an inconvenient generalization when contracting can be staggered.

A static Nash-in-Nash model is a Nash equilibrium in bilateral Nash bargains where bargaining outcomes only affect the current period and do not affect any other bargaining outcomes.

I base my notation for a static vertical market on Lee et al. (2021). n players compete in a single period t. Each pair $i < j$ has the simultaneous opportunity to mutually choose a non-null contract $\mathbb{C}_{ijt} \in \mathcal{C}_{ijt}$; if negotiations fail, then ij choose the null contract $\mathbb{C}_{0,ijt}$. A successful agreement comes with a j-to-i price p_{ijt} and requires paying a negotiation cost r_{it} and r_{jt}. (This adds generality, as the literature generally assumes $r_{it} = r_{jt} = 0$.) Once contracts are chosen, I write the set of market contracts as $\mathbb{C}_t = (\mathbb{C}_{ijt})_{i<j}$, the number of new agreements for player i as R_{it}, and the profits of player i as

$$V_{it}(\mathbb{C}_t, R_{it}) = \pi_{it}(\mathbb{C}_t) + T_{it}(\mathbb{C}_t) - r_i R_{it},$$

where the net ijt demand is $D_{ijt}(\mathbb{C}_t)$ and player i's net transfer is $T_{it}(\mathbb{C}_t) = \sum_{j \neq i} p_{ijt} D_{ijt}(\mathbb{C}_t)$. Vertical markets can generally be partitioned into upstream suppliers and downstream retailers, but I write the model more generally.

The Nash equilibrium restriction of Nash-in-Nash is that firms calculate gains from trade taking as given every simultaneous strategy. The restriction on simultaneous strategies corresponds to a 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). I write $(\mathbb{C}_{t,-ij}, \hat{\mathbb{C}}_{ijt})$ for the market contract that replaces the ijt agreement with $\hat{\mathbb{C}}_{ijt}$. If ij reach an agreement, then player i's gains from trade from a contract $\mathbb{C}_{ijt}$ are $V_{it}((\mathbb{C}_{t,-ij}, \mathbb{C}_{ijt}), R_{it}) - V_{it}((\mathbb{C}_{t,-ij}, \mathbb{C}_{0,ijt}), R_{it} - 1)$. It is useful to write the effect of the ij contract on i's non$-j$ transfers and pre-transfer profits as:

$$\begin{aligned}
[\Delta_{ij} T_{it,-j}] &\equiv E\left[\sum_{k \neq i,j} \hat{p}_{ikt}\left(D_{ikt}\left(\hat{\mathbb{C}}_t\right) - D_{ikt}\left((\hat{\mathbb{C}}_{t,-ij}, \mathbb{C}_{0,ijt})\right)\right)\right] \\
[\Delta_{ij} \pi_{it}] &\equiv E\left[\pi_{it}\left(\hat{\mathbb{C}}_t\right) - \pi_{it}\left((\hat{\mathbb{C}}_{t,-ij}, \mathbb{C}_{0,ijt})\right)\right],
\end{aligned} \tag{1}$$

so that the ij gains from trade are $GFT_i\left(\mathbb{C}_{ijt}\right) = [\Delta_{ij}\pi_{it}] + [\Delta_{ij} T_{it,-j}] - r_i + p_{ijt} D_{ijt}(\mathbb{C}_t)$.

I will illustrate the Nash-in-Nash framework with a simple triangular market. A monopolist insurer (player 1) sells insurance at a regulated price of $\phi = \$10,000$. The insurer's sales depends on how many of the two symmetric hospitals (players 2 and 3) agree to be in the insurer's network. A hospital-i agreement is associated with a price per unit of care provided p_{it}, and is formed through Nash-in-Nash bargaining with insurer weight τ. If both hospitals agree, the insurer sells $D^M = 6,000$ units of insurance. If only one hospital agrees, the insurer sells $D^M = 5,000$ units of insurance. Consumers become patients and demand one unit of care, split evenly within the insurer's network and at zero cost to the hospitals. In equilibrium, both hospitals agree to be in the insurer's network. $r_i = 0$ for all players.

The Nash-in-Nash model allows contracts to have spillovers on each other through gains from trade. In the illustrating example, if only hospital i reaches an agreement, then hospital i provides 5,000 units of care. If hospital $5 - i$ also reaches an agreement, then hospital i provides only 3,000 units of care. When the insurer negotiates with hospital $5 - i$, the anticipated price paid to hospital i increases the insurer's gains from a further agreement with hospital $5 - i$ and typically increases the realized price paid to hospital $5 - i$. In a static model, these spillovers in gains from trade occur through anticipated agreements, but the spillovers are not internalized. That lack of internalization is convenient, but inaccurate when forward-looking negotiators choose staggered contracts.

The Nash bargaining component of Nash-in-Nash is that contracts are reached through bilateral Nash bargaining. More generally, a bargaining solution X is a function from a space of feasible agreement values $P \subset \mathbb{R}^2$ and a disagreement value $v^D \in \mathbb{R}^2$ to a point $C \in P \cup v^D$ such that $C_1 \geq v_1^D$ and $C_2 \geq v_2^D$ (Figure 3(a)). For a bargaining solution X, I write that

a set of contracts $\hat{\mathbb{C}}_t = \{\hat{\mathbb{C}}_{ijt}\}$ is a *Nash-in-X* equilibrium if for all ij, the realized value $\left(V_{it}(\hat{\mathbb{C}}_t, \hat{R}_{it}), V_t(\hat{\mathbb{C}}_t, \hat{R}_{jt})\right)$ is equal to a solution X applied to the bargaining problem over feasible gains from trade. If ij reach a non-null agreement, the Nash bargaining solution maximizes an asymmetric product of gains from trade:

$$\hat{\mathbb{C}}_{ijt}^{(Nash)} \in \underset{\mathbb{C}_{ijt} \in \mathcal{C}_{ijt}}{\text{argmax}} \left(GFT_i\left(\mathbb{C}_{ijt}\right)\right)^{1-\tau_{ij}} \left(GFT_j\left(\mathbb{C}_{ijt}\right)\right)^{\tau_{ij}}.$$

This asymmetric product can be viewed as a joint utility function (Nash, 1950, p. 159). If value functions are differentiable in the price p_{ijt} and $\tau_{ij} < 1$, then first-order conditions yield that at distribution of Nash-in-Nash contracts $\mathbb{C}_{ijt}^{(Nash)}$:

$$\frac{E\left[GFT_j\left(\hat{\mathbb{C}}_{ijt}^{(Nash)}\right)\right]}{E\left[GFT_i\left(\hat{\mathbb{C}}_{ijt}^{(Nash)}\right)\right]} = \frac{\tau_{ij}}{1-\tau_{ij}} \times \overbrace{\frac{E\left[D_{ijt} + p_{ijt}\frac{\partial D_{ijt}}{\partial p_{ijt}} - \frac{\partial \pi_{jt}(\hat{\mathbb{C}}_t)}{\partial p_{ijt}} - \frac{\partial T_{jt,-ij}(\hat{\mathbb{C}}_t)}{\partial p_{ijt}}\right]}{E\left[D_{ijt} + p_{ijt}\frac{\partial D_{ijt}}{\partial p_{ijt}} + \frac{\partial \pi_{it}(\hat{\mathbb{C}}_t)}{\partial p_{ijt}} + \frac{\partial T_{it,-ij}(\hat{\mathbb{C}}_t)}{\partial p_{ijt}}\right]}}^{MVR}, \tag{2}$$

where MVR is the negative of the ratio of the marginal profits as a function of p_{ijt}.

I write that a game is *transferable utility* (TU) if the MVR is always equal to one in every subgame.[4] In a TU game, a pair can first decide on their efficient agreement and then choose transfers to split gains from trade as they desire. TU bargaining is an important special case that the microfoundations of the Nash-in-Nash solution. In the TU case, the ratio of gains from trade is equal to the ratio of bargaining weights $\frac{\tau_{ij}}{1-\tau_{ij}}$.

Lemma 1 (Static Nash-in-Nash payments)**.** *Consider a Nash-in-Nash equilibrium of a static TU game in which ij reach an agreement with probability one. Then:*

$$E[p_{ijt}^{(Nash)} D_{ijt}] = -\tau_{ij}\left([\Delta_{ij}\pi_{it}] + [\Delta_{ij}T_{it,-j}]\right) + (1-\tau_{ij})\left([\Delta_{ij}\pi_{jt}] + [\Delta_{ij}T_{jt,-i}]\right) + \tau_{ij}r_i - (1-\tau_{ij})r_j.$$

Proof. Appendix D □

The illustrating example is a TU game. When player 1 negotiates with hospital i at an expected hospital $5-i$ price of $p_{5-i,t}$, the gains from trade correspond to $[\Delta_{i1}\pi_{1t}] = \$10,000(6,000-5,000)$, $[\Delta_{i1}T_{1t,-i}] = -p_{5-i}\left(3,000-5,000\right)$, $[\Delta_{ij}\pi_{it}] = 0$, and $[\Delta_{i1}T_{it,-1}] = 0$. Then the Nash-in-Nash payment is $\frac{D^M}{2}p_{i,t} = (1-\tau)(\phi(D^M - \underline{D}^M) + (\underline{D}^M - \frac{D^M}{2})p_{5-i,t})$, and the equilibrium price is $p_{2,t} = p_{3,t} = \text{“}p^*_{(Static)}\text{”} = \frac{(1-\tau)\phi(D^M-\underline{D}^M)}{\tau D^M/2+(1-\tau)(D^M-\underline{D}^M)} = \$10,000\frac{1-\tau}{1+2\tau}$.

[4]This is a slight abuse of notation: TU is usually used in the context of coalitional games in which utility may be transferred across sets of agents, but I focus my attention on coalitional games in which contracting is only formed bilaterally. I also restrict my attention to games in which utility is transferable through per-unit-demand payments. One could achieve transferable utility through lump-sum agreement payments by writing D_{ijt} to be an indicator of reaching a non-null contract.

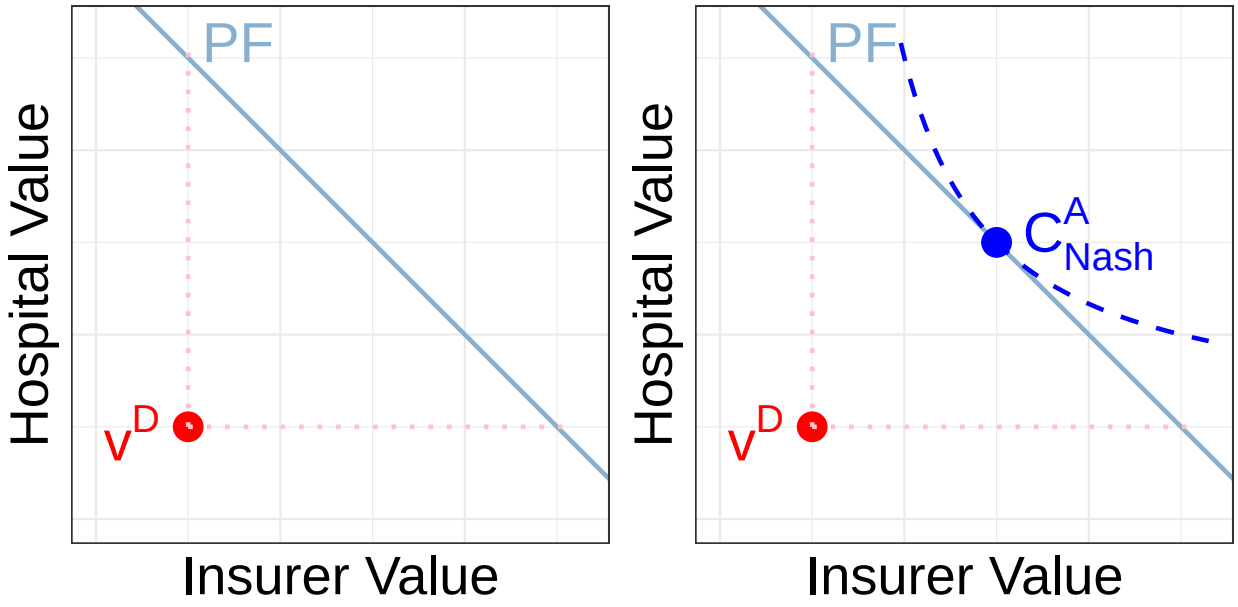


Figure 3: Left: an arbitrary Nash bargaining problem between an insurer (x axis) and hospital (y axis) with Pareto frontier PF in light blue. The disagreement point value is indicated with v^D in red, and implies Pareto improvement constraints (dotted lines). Right: the Nash bargaining solution, indicated by C^A_{Nash} in blue. Dashed blue line corresponds to Nash bargaining indifference curve over joint gains from trade at the Nash bargaining agreement.

The literature has so far been static. In the next subsection, I illustrate the challenges of modeling a dynamic market with Nash-in-Nash.

3.2 Dynamic Nash-in-Nash

The goal of this work is to build a tractable model of forward-looking bargaining with staggered contracts. I now use the running illustration example to demonstrate why a Nash-in-Nash extension to staggered contracting introduces new tractability and interpretability challenges through the possibility of internalized contract spillovers.

There continue to be two symmetric hospitals that can form one-year contracts, but in equilibrium they now reach simultaneous agreements on January 1 of each year with an annual discounting rate of $\beta \in [0, 1)$. A natural starting point is an annual model with one-year disagreement.

However, Section 2 showed that accurate empirical predictions may require acknowledging staggered contracts. In some sense, in the presence of negotiation constraints, simultaneous bargaining is the exception in practice rather than the rule. Whether staggered contracting is an important feature to model depends on the setting and question. When accuracy requires forward-looking negotiators that acknowledge the existence or threat of staggered

contracting, it will turn out that Nash-in-Nash bargaining will introduce new tractability issues.

To consider the simplest case, what would happen if a pair that disagrees on January 1 can limit their losses by returning to the bargaining table on July 1? I use a device I call split-in-half to study whether exclusion lengths matter.

Definition 1 (Split-in-half invariance)**.** *A dynamic Nash-in-X model is* invariant to being split in half *if for every subgame of every Nash-in-X equilibrium of the game, there is a corresponding subgame of a model with subsequent period lengths cut in half (formalized in Appendix Definition 3) and a Nash-in-X equilibrium with the same subgame predictions.*

If some Nash-in-X prediction cannot be sustained after shortening the period length, then I write the model is not invariant to being split in half. The split-in-half process to introduce new equilibria; ruling out new equilibria is more difficult than showing a mapping of existing equilibria to split-in-half equilibria.

Proposition 1 (Running example is not invariant)**.** *The annual illustrating Nash-in-X model with $\tau \in (0,1)$, $\beta > 0$, and $\underline{D}^M \geq D^M/2 > 0$ is not invariant to being split in half among split-in-half subgames with differentiable price responses.*

Proof. Appendix D. □

The logic of Proposition 1 is presented in Figure 4. Construct the value of disagreeing via backwards induction. If the insurer disagrees with hospital i on January 1 in the split-in-half model, they will return to the bargaining table on July 1 and negotiate relative to six more months of exclusion. If the firms negotiate on July 1, they will negotiate with an internalized spillover on the other hospital's next negotiation. This spillover leads Nash bargainers to choose a July 1 disagreement contract is more favorable to the insurer. On January 1, the firms know that disagreement will improve the insurer's leverage, and so the January 1 contract is shifted towards the insurer as well. The full solution involves a system of equations in steady-state payments and spillover rates.

I refer to this phenomenon as *state space growth*. A recursive Nash-in-Nash model will typically define gains from trade relative to future agreements. Those future agreements are recursively defined, and each corresponds to a new bargaining state with different spillovers on future negotiations. An unbounded number of bargaining states may be needed to accurately approximate a bargaining process with short-lived exclusion. Even when sufficient computing resources exist to approximate the dynamic Nash-in-Nash bargaining problem, I show in Appendix B.4 the Nash-in-Nash model does not easily accommodate even ostensibly ignorable uncertainty.

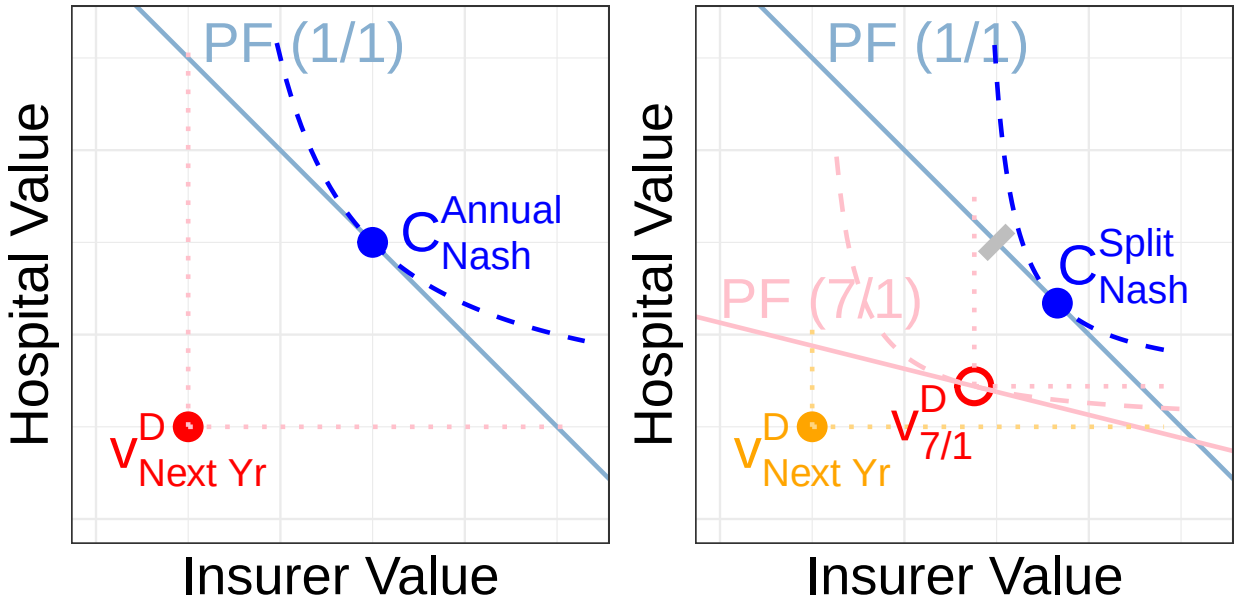


Figure 4: In the illustration example, the predicted values (blue dot) are different in the annual model (left, agreement $C_{\text{Nash}}^{\text{Annual}}$) and split-in-half model (right, agreement $C_{\text{Nash}}^{\text{Split}}$ relative to annual agreement indicated by grey notch). Insurer (x axis) and hospital (y axis) values are presented relative to the value of disagreeing for one year ($v_{\text{Next Yr}^D}$). Full agreement and associated Pareto frontier are indicated in blue, the relevant disagreement point is indicated in red, the split-in-half disagreement point's disagreement point is indicated in orange, dashed lines indicate Nash product indifference curves at the chosen agreement, and dotted lines indicate Pareto improvement constraints.

I derive the split-in-half running example's closed-form solution in Lemma 2 in Appendix B.1, but it is opaque and hard to parse. Yet, for all the theoretical complexity of the recursive Nash-in-Nash process, the empirical effect of incorporating the threat of forward-looking spillovers. This is because while internalized spillovers exist in a dynamic vertical market, they are often small in practice: the hospital contracts only overlap for six months, and disagreement with one hospital only leads one-third of consumers to substitute to the other hospital. The simultaneous price is around 5% smaller than the static price, but is awkward to calculate and opaque. I propose the Nash-in-Kalai approach to arrive at an internally-consistent model of vertical market bargaining that cuts through the complications of state space growth.

3.3 The Kalai Proportional Solution

The Nash-in-Kalai model aligns with Nash-in-Nash in many leading static cases, but will more tractably extend to dynamic settings with staggered contracts. The principal advantage

of this model for dynamic settings is that it will manage the growth of empirically relevant bargaining states, but I begin by introducing the model for static bargaining.

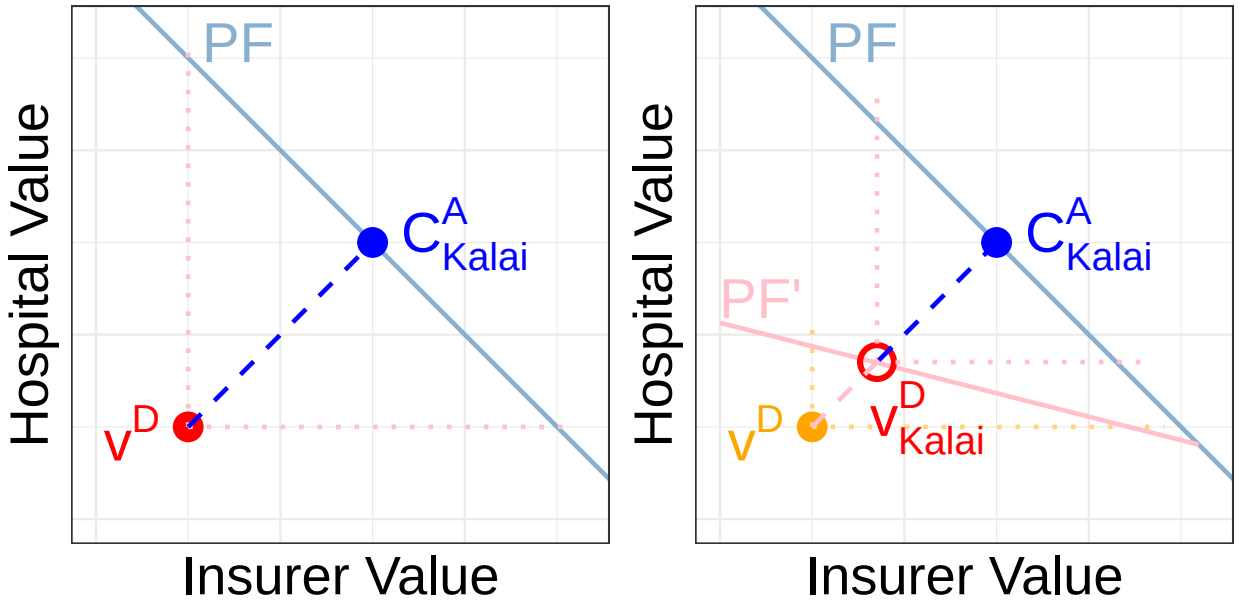


Figure 5: Left: The Kalai proportional solution, as applied to the Nash bargaining problem of Figure 3(a), with proportional split indicated by dashed blue line. Right: the step-by-step property of Kalai proportional bargaining. Kalai proportional bargaining has the same predictions when the outcome of bargaining over the smaller feasible set (PF' with first step outcome v^D_{Kalai} in pink) is used as the disagreement point in bargaining over the full set of feasible values (PF) with agreement C^A_{Kalai} (blue dot).

The Nash-in-Kalai solution is a Nash equilibrium in Kalai proportional bargains. The Kalai proportional bargaining solution with player-j bargaining weight τ_{ij} chooses the contract $\mathbb{C}_{ijt}$ that maximizes gains from trade among contracts that give player j a τ_{ij} share of the joint gains from trade. That is, if Kalai proportional bargaining yields a contract $\mathbb{C}^{(Kalai)}_{ijt}$ and $\tau_{ij} < 1$, then the ratio of gains from trade satisfies:

$$\frac{GFT_j\left(\hat{\mathbb{C}}^{(Kalai)}_{ijt}\right)}{GFT_i\left(\hat{\mathbb{C}}^{(Kalai)}_{ijt}\right)} = \frac{\tau_{ij}}{1-\tau_{ij}}. \tag{3}$$

The solution is presented in Figure 5(a).

The Nash-in-Kalai characterization of payments generalizes the Nash-in-Nash TU payment from Lemma 1 to NTU bargaining.

Corollary 1 (Static Nash-in-Kalai payments)**.** *Let $\mathbb{C}_t$ follow a Nash-in-Kalai equilibrium of*

a static game such that $i < j$ agree almost surely. Then whether or not bargaining is TU,

$$E[p_{ijt}^{(Kalai)} D_{ijt}] = -\tau_{ij}\left([\Delta_{ij}\pi_{it}] + [\Delta_{ij}T_{it,-j}]]\right) + (1-\tau_{ij})\left([\Delta_{ij}\pi_{jt}] + [\Delta_{ij}T_{jt,-i}]]\right) + \tau_{ij}r_i - (1-\tau_{ij})r_j.$$

In the illustrating example, the static equilibrium Nash-in-Kalai price is the Nash-in-Nash $p^*_{(Static)}$. That is because bargaining in that example is TU.

The unique feature of the static Kalai proportional bargaining solution that I exploit is the *step-by-step* (Kalai, 1977) or *path independence* (Roth, 1979) property. I demonstrate this property in Figure 5(b). If f is a bargaining solution over a feasible utility set S with disagreement point v_D, then the step-by-step property is that for any $S' \subseteq S$ that generates a valid bargaining problem, $f(S, f(S', v^D)) = f(S, v^D)$.[5] That is, if one starts with a bargaining problem, reduces the set of feasible values, applies the bargaining solution to the smaller bargaining problem as a first step, and treats the outcome as the disagreement point when negotiating over the full set of feasible values as a second step, then Kalai proportional bargaining will predict the same contract as in the original bargaining problem. Kalai (1977) shows that this feature is unique to Kalai proportional bargaining in games with free disposal, and Roth (1979) shows that this feature is unique to bargaining solutions with "proportional character" in games that may lack free disposal. I discuss some other relative advantages and disadvantages of the Nash and Kalai proportional solutions as bilateral NTU solution concepts in Appendix A.

While the step-by-step property is interesting, Kalai cautions that "one has to be careful in the interpretation." As he accurately notes, from a static perspective, the property "is a condition on the feasible utility levels and not directly on the underlying game."

The key insight of this work is to note that for a dynamic bargaining game, the step-by-step property is an essential tool for tractability. Consider the split-in-half game, but under Nash-in-Kalai bargaining rather than Nash-in-Nash bargaining. Return to the step-by-step Figure 5(b). Relabel the step-by-step graph's ultimate disagreement point v^D as $v^D_{\text{Next Yr}}$, ultimate Pareto frontier PF as $PF(1/1)$, smaller frontier PF' as $PF(7/1)$, smaller frontier agreement v^D_{Kalai} as $v^D_{7/1}$, and ultimate agreement C^A_{Kalai} as $C^{\text{Split}}_{\text{Kalai}}$. The result is an analog of the Nash-in-Nash split-in-half Figure 4(b), but now shows that the split-in-half predictions are the same as in the original game. Any gains from trade available by excluding for six months rather than one year are already split proportionally to bargaining weights, and so can be treated as null. This yields an invariance to splitting the period length in half.

Corollary 2 (Running example, Nash-in-Kalai)**.** *Consider the split-in-half subgame of Lemma 2,*

[5] I modify the notation slightly from Kalai and Roth in order to avoid imposing a disagreement point normalization.

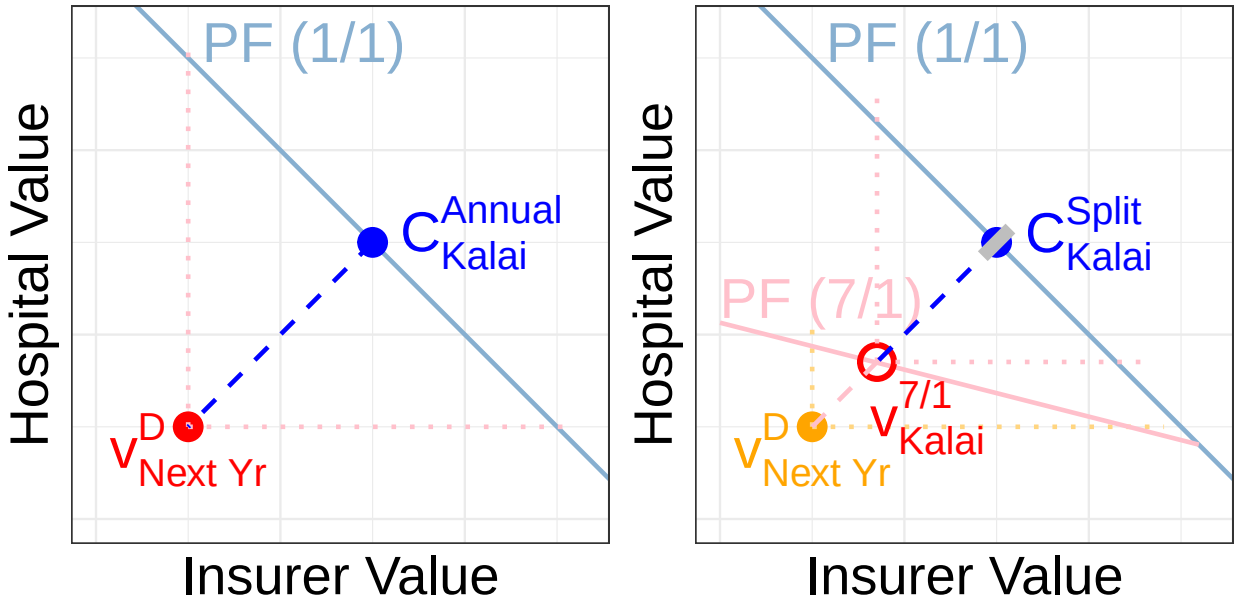


Figure 6: Nash-in-Kalai split-in-half invariance in the illustrating example. Graph is based on Figure 4, but for Kalai proportional bargaining. By the step-by-step property demonstrated in Figure 5(b), any gains from trade available in the off-equilibrium July 1 bargain can be ignored for the purposes of computing the equilibrium January 1 gains from trade.

but under a Nash-in-Kalai model. Then there is an equilibrium where the simultaneous price $p^{(Kalai),*}_{(Simult)}$ *is the static price* $p^{*}_{(Static)} = \frac{(1-\tau)\phi(D^M - \underline{D}^M)}{\tau D^M/2 + (1-\tau)(D^M - \underline{D}^M)} = \$10,000 \frac{1-\tau}{1+2\tau}$.

Proof. Appendix D. □

In a more general forward-looking model with arbitrarily short period lengths, the step-by-step logic of Figure 6 can be used to derive a tractable equivalent disagreement point. The Nash-in-Kalai recursive disagreement point is the value of disagreeing once and then reaching an agreement. By the step-by-step property, the econometrician can replace this disagreement value with the value of disagreeing twice without modifying the predicted contract. By repeatedly applying this argument, the Nash-in-Kalai recursive disagreement point can be replaced with the value of repeated surprise disagreement.

The Nash-in-Kalai equivalent disagreement point constructed via repeated application of the step-by-step property corresponds to Binmore et al. (1989)'s impasse point. In the ij impasse point, i and j continually attempt to reach an agreement, but surprise all actors and continue to disagree without abandoning negotiations. Under Markov strategies, the impasse point is well-defined regardless of how disagreement is reached. The impasse point captures the dynamic intuition of static Nash-in-Nash disagreement: everyone behaves anticipating ij to reach an agreement, but ij always surprise the market and disagree. The impasse

value only depends on ij bargaining states through states at which others respond to an anticipated ij agreement. Most ij bargaining states under impasse have no impact on the value of impasse, and so can be ignored for the purposes of calculating the Nash-in-Kalai bargaining problem.

In the next section, I define a dynamic Nash-in-Kalai model and use this insight to prove that the Nash-in-Kalai has more general tractability advantages.

4 Dynamic Nash-in-Kalai Bargaining

I now propose the Nash-in-Kalai model for dynamic vertical markets. The model has the same predictions as Nash-in-Nash in the important TU class of static models. I show that in a dynamic setting, the Nash-in-Kalai model yields a moment on NPV payments that generalizes the static Nash-in-Nash moment. Further, I prove that the new dynamic term in the Nash-in-Kalai model moment is zero in steady state, and that the model possesses a valuable tractability simplification inside or outside steady state.

4.1 Model Timing

I study dynamic vertical markets of the following form.

Time is indexed by t. One unit of time is divided into $m \geq 1$ periods, with the time of period s written as t_s: $t_1 = 0$, $t_2 = 1/m$, $t_{1+m} = 1$, and so on. The per-period discounting rate is equal to $\beta^{1/m} \in [0, 1)$, with myopia corresponding to $\beta = 0$.

The timing in each period t_s is as follows:

1. *Information is revealed and non-bargaining choices are made.* Information $\mathcal{I}_{t_s}$ is revealed, possibly through multiple sub-periods. The resulting information includes the history of the game, period pre-transfer profit functions $\pi_{t_s i}(\mathbb{C}_{t_s})$, period demand functions $D_{t_s,ij}(\mathbb{C}_{t_s})$, null contracts $\mathbb{C}_{0,ijt_s}$, the set of feasible bilateral contracting states $\mathcal{C}_{ijt_s}$ (either a single renewal contract, a single null contract, or a set of feasible agreements), and any information about future values. Every feasible contract ij includes p_{ijt_s}, the net i-to-j transfer per unit of equilibrium demand $D_{t_s,ij}(\mathbb{C}_{t_s})$ provided in period t_s.

2. *Contracts are bargained.* For $i < j$ pairs with more than one feasible contract, new contracts $\hat{\mathbb{C}}_{ijt_s}$ are chosen through bilateral Kalai proportional bargaining relative to the null contract, with j receiving bargaining weight $\tau_{ij} \in [0, 1]$. I write the number of new contracts formed by firm i as $\hat{R}_{it_s}$.

3. *Flow profits are formed.* Flow profits for agent i at contract state $\hat{\mathbb{C}}_{t_s}$ with associated j-to-i prices $\hat{p}_{ijt_s} = -\hat{p}_{jit_s}$ and network of firms with agreements $\mathcal{G}_{it_s}$ are equal to

$$\pi_{t_s i}(\hat{\mathbb{C}}_{t_s}) + \sum_{j \in \mathcal{G}_{it_s}} \hat{p}_{ijt_s} D_{t_s,ij}(\hat{\mathbb{C}}_{t_s}) - r_i \hat{R}_{it_s},$$

where $r_i \geq 0$ is the cost of validating a new contract.

Definition 2 (Nash-in-Kalai equilibrium)**.** *A* Nash-in-Kalai equilibrium *is a conditional random variable distribution $\mathcal{I}_{t_s} \mid (\mathcal{I}_{t_{s-1}}, \mathbb{C}_{t-1})$, a bilateral contract choice distribution $\hat{\mathbb{C}}_{t_s,ij}(\mathcal{I}_{t_s})$, and recursive value functions (Appendix B.2) such that (i) all strategies are Markov strategies, (ii) any strategies in the formation of $\mathcal{I}_{t_s}$ are in the appropriate equilibrium, and (iii) for every ij and every subgame, the expected NPV profit of the distribution of $\hat{\mathbb{C}}_{t_s,ij}(\mathcal{I}_{t_s}) \mid \mathcal{I}_{t_s}$ when taking other strategies as given solves the Kalai proportional bargaining problem over value functions with bargaining weight $\tau_{ij} \in [0,1]$ on player $j > i$.*

Due to the bilateral nature of contracting, the recursive value functions have only been useful as collections of flow profits.

I discuss each stage in further detail.

4.1.1 Information is Revealed

The information updating stage 1 is written in a generic way in order to accommodate many types of vertical markets.

Stages 1 and 3 can be modified to include separate games. The first stage can include premium-setting, contract renewal, or other strategic phases. The final stage can include demand reallocation in response to upstream prices. I focus on the bargaining stage and abstract from these strategic responses here.

One could just as easily imagine strategic interactions after bargaining rather than before bargaining. Such a change could be accommodated at the cost of additional notation to track the outcome of the end-of-period response. I will eventually imagine a period to be arbitrarily short, so that a single period is almost irrelevant. A more substantive change would be if non-bargaining competition occurred at the same time as bargaining in stage 2. Such simultaneous competition would be insubstantial if it occurred only at specific times, but difficult to model if the strategic response was revisited in every period while a pair remained in impasse.

4.1.2 Contracts are Bargained

The core of my analysis is the bargaining Stage 2.

I assume that a contract must have a price per unit of demand, which corresponds to transfers in the scale of utility. 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; or per-period transfers as a function of each firm's network of realized agreements. However, I will assume that a given period's realized contracts $\mathbb{C}_t$ (and any $\mathcal{I}_t$ information decided earlier) are sufficient statistics for flow profits before transfers and negotiation costs.

I also assume that a bilateral pair chooses a contract taking as given both future strategies and all simultaneous bilateral strategies. 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 in a dynamic market: I will have in mind a model where the length of a period tends to zero, so that the simultaneous negotiations will typically have no importance in the limit.

4.1.3 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 vertical markets.

I implicitly rule out payments under null contracts. One could imagine modifying the game to include payments outside of a contract, for example through out-of-contract purchasing (Prager and Tilipman, 2022). I rule that out for conciseness. One could extend the Nash-in-Kalai model to cover these cases at the cost of yet more notation.

I model a negotiation cost borne after bargaining succeeds. 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.

The flow profit representation is not unique. For example, one can double demand and subtract the new payments from π to yield the same profits. I only assume that a representation of this form exists. The representation here is motivated by static TU games in which π can be written as a function of networks alone, and transfers reflect lump-sum or per-unit payments.

4.2 Key Properties of the Dynamic Nash-in-Kalai Solution

I now show that the dynamic Nash-in-Kalai model yields a moment on NPV payments that generalizes the static Nash-in-Nash moment on current payments. The moment delivers key tractability properties. In particular, shortening the length of disagreement does not change the moment, so that a complex Nash-in-Kalai model defined in continuous time can have the exact same predictions as a particular and more tractabile model defined in discrete time. The existence of a payment moment and a tractable continuous-time foundation for a discrete-time model are key advantages for empirical work.

I maintain some regularity conditions.

Assumption 1. (Regularity conditions)
Players are risk-neutral, share rational expectations, and follow Markov strategies. If ij reach their null contract in t_s and do not negotiate a new contract in t_{s+1} then ij reach their null contract in t_{s+1}. There is a uniform transversality condition: if $\mathcal{F}_{t_r|t_s}(\mathcal{I}_{t_s})$ is the set of feasible information sets in period $t_r \geq t_s$ after information $\mathcal{I}_{t_s}$ is reached in period t_s, then

$$\lim_{h\to\infty} \sup_{\mathcal{I}_{t_s}} \sup_{\mathcal{I}_{t_{s+h}} \in \mathcal{F}_{t_{s+h}|t_s}(\mathcal{I}_{t_s})} \sup_{i} \beta^{h/m} \left| V^{(1)}_{i,t_{s+h}}(\mathcal{I})(\mathcal{I}_{t_{s+h}}) \right| = 0.$$

4.2.1 NPV Payment Moment

The dynamic Nash-in-Kalai bargaining model yields a moment on expected NPV payments, and as a result the econometrician can construct moments on payments for estimation. The moment naturally generalizes the static Nash-in-Nash bargaining moment to incorporate multiple periods of gains from trade.

Theorem 1 (Nash-in-Kalai Moment)**.** *Consider a dynamic Nash-in-Kalai equilibrium that satisfies Assumption 1. Suppose players $i < j$ form a contract in a subgame time t_0 that remains in place through the (potentially random) terminal time t^* with (potentially random) realized prices p^*_{ijt}. Then the expected NPV of realized payments $D_{t_s,ij}p^*_{ijt_s}$ at the moment of contract formation is equal to the sum of the expected NPV of flow period Nash-in-Nash payments, a negotiation cost payment, and an impasse repricing payment term:*

$$\mathbb{E}_{t_0}\left[\sum_{t_0 \leq t_s \leq t^*} \beta^{\frac{t_s - t_0}{m}} D_{t_s,ij} p^*_{ijt_s}\right] = \text{Pay}_{NiN} + \text{Pay}_{NC} + \text{Pay}_{IRT}, \tag{4}$$

where the expected NPV of static Nash-in-Nash payments (Lemma 1) is:

$$\text{Pay}_{NiN} = \mathbb{E}_{t_0}\left[\sum_{t_0 \le t_s \le t^*} \beta^{\frac{t_s - t_0}{m}} \begin{pmatrix} -\tau_{ij} & ([\Delta_{ij}\pi_{it_s}] + [\Delta_{ij}T_{it_s,-j}]) \\ +(1-\tau_{ij}) & ([\Delta_{ij}\pi_{jt_s}] + [\Delta_{ij}T_{jt_s,-i}]) \end{pmatrix}\right], \tag{5}$$

the negotiation cost payment Pay_{NC} *is equal to* $\tau_{ij}r_i - (1-\tau_{ij})r_j$*, and the impasse repricing payment* Pay_{IRT} *is:*

$$\text{Pay}_{IRT} = E_{t_0}\left[\sum_{t_s \ge t_0} \beta^{\frac{t_s - t_0}{m}} \begin{Bmatrix} -\tau_{ij} & \left(\hat{\pi}^A_{it_s} + \hat{T}^A_{it_s} - \hat{\pi}^D_{it_s} - \hat{T}^D_{it_s}\right) \\ +(1-\tau_{ij}) & \left(\hat{\pi}^A_{jt_s} + \hat{T}^A_{jt_s} - \hat{\pi}^D_{jt_s} - \hat{T}^D_{jt_s}\right) \end{Bmatrix}\right], \tag{6}$$

where $\hat{\pi}^A_{t_s,k}$ *and* $\hat{T}^A_{t_s,k}$ *correspond to the (potentially random) path of profits and net transfers if* ij *enter their impasse point beginning in period* t^*+1 *and the* ij *equilibrium contract is replaced with the* ij *null contract in periods* t_0 *through* t^**, and where* $\hat{\pi}^D_{t_s,k}$ *and* $\hat{T}^D_{t_s,k}$ *correspond to those paths if* ij *enter their impasse point beginning in period* t_0*.*

This is a generalization of the static Nash-in-Nash bargaining moment, which corresponds to the myopic case $\beta = 0$. In many empirical applications, the realized payments $D_{t_s,ij}p^*_{ijt_s}$ can be observed, the flow Nash-in-Nash payments in Pay_{NiN} depend on estimable demand functions and a small number of bargaining parameters, and Pay_{NC} depends on only a few parameters. The term Pay_{IRT} accounts for the fact that in future periods, the Nash-in-Nash disagreement point differs from the Nash-in-Kalai impasse point. The static Nash-in-Nash gains used in Pay_{IRT} are calculated relative to disagreement under the non-ij contracts formed in equilibrium, while the Nash-in-Kalai gains are calculated relative to disagreement under the non-ij contracts formed when continually expecting expecting ij to exit impasse.

Proof. Consider a period t_0's stage 2 subgame in which $i < j$ form a new contract with positive probability. For simplicity, I proceed assuming there is a pure strategies equilibrium, and that in every subgame that can be reached after t_0, the expected number of subsequent ij agreements is infinite.[6] Also for simplicity, assume that a realized contract $\mathbb{C}_{t_s,ij}$ will be unchanged until the next period in which ij have the opportunity to negotiate. I write the anticipated non-ij contracts in period t_s as $\mathbb{C}_{t_s,-ij}$.

I consider two off-equilibrium paths: an agreement-followed-by-impasse path (superscripted A), and an immediate-impasse path (superscripted D). I write $t^A_0 = t^D_0 = t_0$, and I write t^A_a and t^D_a for the (potentially random) a^{th} future period in which ij defect from

[6] In estimation, it is infeasible to include an infinite number of bargains without imposing strong assumptions like steady state behavior. I therefore impose a finite horizon model in estimation as an approximation to an infinite-horizon model.

equilibrium and reach the null contract on that path. For conciseness, I leave some scratch work to Appendix D (Page 90).

For $d \geq 0$, define $V^{(2,E)}_{k,(d)^D}$ to be the (potentially random) expected NPV to player k of ij existing impasse in period t^D_d at the other-pair subgame equilibrium contracts $\hat{\mathbb{C}}_{t^D_d,-ij}$. Let the subgame equilibrium have ij choose $\hat{\mathbb{C}}_{t^D_d,ij}$. Also define $V^{(2,I)}_{k,(d)^D}$ to be the expected NPV of remaining in impasse, and define $GFT^{(2,E)}_{(d)^D} \equiv V^{(2,E)}_{i,(d)^D} + V^{(2,E)}_{j,(d)^D} - V^{(2,I)}_{i,(d)^D} - V^{(2,I)}_{j,d}$ to be the joint subgame gains from trade. Because ij follow the Kalai proportional solution, $\hat{C}_{t^D_d,ij}$ is chosen to maximize joint gains from trade subject to:

$$(\tau_{ij})\left(V^{(2,E)}_{i,(d)^D}\left(\left(\hat{\mathbb{C}}_{t^D_d,-ij},\hat{\mathbb{C}}_{t^D_d,ij}\right)\right) - V^{(2,I)}_{i,(d)^D}\right) = (1-\tau_{ij})\left(V^{(2,D)}_{j,d}\left(\left(\hat{\mathbb{C}}_{t^D_d,-ij},\hat{\mathbb{C}}_{t^D_d,ij}\right)\right) - V^{(2,I)}_{j,d}\right).$$

I iteratively expand and obtain (Equation (10)):

$$V^{(2,I)}_{i,(d)^D} = E_{t^D_d}\left[\sum_{t_s \geq t^D_d} \beta^{\frac{t_s - t^D_d}{m}}\left(\hat{\pi}^D_{it_s} + \hat{T}^D_{it_s}\right) + \sum_{h>0}(1-\tau_{ij})\beta^{\frac{t^D_{d+h}-t^D_d}{m}} GFT^{(2,E)}_{(d+h)^D}\right]$$

$$V^{(2,I)}_{j,(d)^D} = E_{t^D_d}\left[\sum_{t_s \geq t^D_d} \beta^{\frac{t_s - t^D_d}{m}}\left(\hat{\pi}^D_{jt_s} + \hat{T}^D_{jt_s}\right) + \sum_{h>0}(\tau_{ij})\beta^{\frac{t^D_{d+h}-t^D_d}{m}} GFT^{(2,E)}_{(d+h)^D}\right].$$

Analogous results hold to characterize $V^{(2,I)}_{i,(d)^A} = V^{(2,E)}_{i,(d)^D}\left(\left(\hat{\mathbb{C}}_{t^D_d,-ij},\hat{\mathbb{C}}_{t^D_d,ij}\right)\right)$ and $V^{(2,I)}_{j,(d)^A} = V^{(2,E)}_{j,(d)^D}\left(\left(\hat{\mathbb{C}}_{t^D_d,-ij},\hat{\mathbb{C}}_{t^D_d,ij}\right)\right)$, except that in periods t_0 through $(t^A_1 - 1)$, ij have a contract on the A path, so that ij profits are $\hat{\pi}^A + \hat{T}^A + [\Delta_{ij}\pi] + [\Delta_{ij}T_{-.}] + pD_{ijt}$ instead of $\hat{\pi}^A + \hat{T}^A$ and ij pay an additional negotiation cost in period t_0.

Substituting these results together, characterizing the surplus split of the $t_0 = t^D_0 = t^A_0$ bargain, and moving around terms, I obtain that for any equilibrium agreement:

$$\begin{aligned}
E_{t_0}\left[\sum_{t_0 \leq t_s < t^A_q} \beta^{\frac{t_s-t_0}{m}} p^*_{ijt_s} D_{ijt_s}\right] &= E_{t_0}\left[\sum_{t_0 \leq t_s < t^A_q} \beta^{\frac{t_s-t_0}{m}} \begin{Bmatrix} -\tau_{ij} & ([\Delta_{ij}\pi_{it_s}] + [\Delta_{ij}T_{it_s,-j}]) \\ +(1-\tau_{ij}) & ([\Delta_{ij}\pi_{jt_s}] + [\Delta_{ij}T_{jt_s,-i}]) \end{Bmatrix}\right] \\
&\quad + (\tau_{ij}r_i - (1-\tau_{ij})r_j) \\
&\quad + E_{t_0}\left[\sum_{t_0 \leq t_s} \beta^{\frac{t_s-t_0}{m}} \begin{Bmatrix} -(\tau_{ij}) & \left(\hat{\pi}^A_{it_s} + \hat{T}^A_{it_s} - \hat{\pi}^D_{it_s} + \hat{T}^D_{it_s}\right) \\ +(1-\tau_{ij}) & \left(\hat{\pi}^A_{jt_s} + \hat{T}^A_{jt_s} - \hat{\pi}^D_{jt_s} + \hat{T}^D_{jt_s}\right) \end{Bmatrix}\right] \\
&= \text{Pay}_{NiN} + \text{Pay}_{NC} + \text{Pay}_{IRT},
\end{aligned}$$

which is the desired equality. □

4.2.2 Other Tractability Advantages

There are a few important corollaries of Theorem 1 for empirical use: the Pay_{IRT} term that reflects off-equilibrium behavior is zero in steady state, and a discrete-time Nash-in-Kalai model can have the exact same predictions as a continuous-time model generated by continually splitting the discrete model's period length in half.

Theorem 1 writes the expected NPV payment at the time of contracting as

$$\mathbb{E}_{t_0}\left[\sum_{t_0 \leq t_s \leq t^*} \beta^{\frac{t_s - t_0}{m}} D_{t_s,ij} p^*_{ijt_s}\right] = \text{Pay}_{NiN} + \text{Pay}_{NC} + \text{Pay}_{IRT}.$$

The term Pay_{IRT} is new to the literature. This new term is zero if negotiators are myopic ($\beta = 0$) or contracts are single-period and have no effect on future outcomes. When negotiators are forward-looking ($\beta > 0$) and balance multiple periods of gains from trade, the term captures the internalized bargaining spillovers. The transfer splits the effect of impasse on disagreement payoffs, relative to the value of disagreement at the equilibrium profit functions and contracts.

The impasse repricing term Pay_{IRT} is zero in also steady state.

Corollary 3 (Steady state Pay_{IRT} is zero)**.** *Let a steady state equilibrium of a Nash-in-X game be a Nash-in-X equilibrium such that the equilibrium profit functions and contracts $\hat{\mathbb{C}}_t$ are constant, and in any subgame with the equilibrium outcome, if a pair ij defects and disagrees for one period, then profit functions and contracts return to the original equilibrium the next period. Then in any steady state equilibrium of a Nash-in-Kalai game,* $\text{Pay}_{IRT} = 0$.

Proof. Appendix D. □

Real-world markets are rarely in steady state. In practice, profit inputs like consumer demand change over the course of an ij contract. As a result, if ij disagree in t_0 and return to the bargaining table at $t_s > t_0$, it is likely that they will agree to a different path of prices than their initial agreement. In turn, non-ij bargainers negotiating at t_s in the ij impasse point negotiate relative to a slightly different anticipated ij contract as would emerge in equilibrium. In empirical practice, Pay_{IRT} will then be non-zero. Where markets are close to steady state, Corollary 3 implies that approximating Pay_{IRT} to zero will produce an approximately valid moment. This is the strategy I follow in my empirical application.

As demonstrated in Figure 6, the Nash-in-Kalai model tractably controls state space growth.

Corollary 4 (State space simplification)**.** *Let $\mathcal{B}$ be the set of Nash-in-Kalai games with full networks in every original and split-in-half equilibrium (formalized in Appendix Assumption 2). Then for every $b \in \mathcal{B}$ and every subgame of b, the Nash-in-Kalai model is invariant to being split in half.*

Proof. Appendix D. □

In my view, Corollary 4 is the principal benefit of the Nash-in-Kalai model for dynamic vertical markets. One can believe in a continuous-time model that is represented by repeatedly splitting a discrete-time model's period length in half. Corollary 4 shows that a discrete-time model can plausibly capture all dynamics in the continuous-time model. The key requirement will be that if ij enter impasse, the discrete-time model will include every moment at which a non-ij bargain is conducted in anticipation of ij exiting impasse. The many bargaining states that may emerge between those moments as ij remain in impasse and return to the bargaining table an hour, minute, or second later have no impact on the Nash-in-Kalai model's predictions. Because the step-by-step property is unique to the Kalai proportional solution, Corollary 4 is unique to the Nash-in-Kalai model.

I have established that a discrete-time Nash-in-Kalai model yields the same predicted payments as a continuous-time Nash-in-Kalai model defined as the limit as period lengths are increasingly short. I also show that in both models, the Nash-in-Kalai model yields an intuitive moment on expected net present value payments that is amenable to empirical work in the presence of staggered contracts, as is common in American healthcare and other vertical markets.

Having established the tractability of the Nash-in-Kalai model in principle, I now proceed to evaluate the model in practice. I require a setting with data on contract timing and gains from trade over multiple years. In the next section, I provide the first estimated structural model of vertical market bargaining with staggered contracts, using the dataset that motivated the forward-looking model in Section 2. In my counterfactual analysis, the forward-looking Nash-in-Kalai model allows negotiators to anticipate and offset future payment increases by reducing starting prices, allowing me to quantify both a significant effect that static models would rule out and key dynamics that a myopic model would overlook.

5 Empirical Application

This section estimates the dynamic Nash-in-Kalai model introduced in Section 4 using the West Virginia hospital–insurer contract data introduced in Section 2. The dataset, made public as a byproduct of a state regulation on hospital list prices, covers ten years of contract

formation, expiration, and payments. The data allows estimation of bargaining parameters using accurate contract dates, while the model allows negotiators to balance gains from trade over the multiple years a contract may remain in place.

Unlike the previous static literature, I estimate bargaining parameters with only new contracts and allow negotiators to balance gains from trade over multiple periods. I find that a forward-looking model is more accurate than the existing static or myopic approaches: I estimate a discounting rate of $\beta = 0.899$ and clearly reject the null hypothesis of myopia that corresponds to the static model's one-period estimation with accurate contract timing. I apply the estimated model to study the counterfactual effects of proposed dynamic government rate reform on private spending benchmarked to government-set rates, where I find both a significant nine-year payment increase that a static model would rule out and a short-run spending reduction that a myopic model would overlook. These are the first structural estimates of a dynamic vertical market with staggered contract formation.

5.1 Empirical Model

My empirical model is motivated by Ho and Lee (2017)'s static Nash-in-Nash model.

Flow profits are adapted from Ho and Lee. Contracts can be formed between hospitals and insurers (also known as managed care organizations, or MCOs). I adapt their notation and refer to hospitals with subscripts i or h and with superscripts H, and refer to insurers with subscripts j or n and with superscripts M. I assume that the period length is one day; by Corollary 4, the daily timing is equivalent to assuming a continuous-time model in which strategic choices are made at most daily and shorter period lengths do not introduce new contracting threats that change equilibrium behavior. I write insurer j's network in period t as $\mathcal{G}^M_{jt} = \{\text{hospitals } i \mid \mathbb{C}_{ijt} \neq \mathbb{C}_{0,ijt}\}$. I write a hospital i's network $\mathcal{G}^H_{it}$ analogously, and write $\mathcal{G}_t = \{ij \mid \mathbb{C}_{ijt} \neq \mathbb{C}_{0,ijt}\}$ for the full network. The insurer charges premiums ϕ_{jt}, which I assume are set simultaneously during stage 1 of every January 1. I write the vector of premiums as ϕ_t. The transfer demand function is hospital inpatient units of care, $D_{ijt}(\mathbb{C}_t) = D^H_{ijt}(\mathcal{G}_t, \phi_t)$, which is zero for ij that reach the null contract.

If i is a hospital and j is an insurer, then the pre-transfer profit functions are:

$$\pi^H_{it}(\mathbb{C}_t) = -\sum_{n \in \mathcal{G}^H_{it}} D^H_{int}(\mathbb{C}_t) c_i$$
$$\pi^M_{jt}(\mathbb{C}_t) = D^M_{jt}(\mathcal{G}_t, \phi_t)(\phi_{jt} - \eta_j),$$

where c_i is the hospital's per-unit cost, D^M_{jt} is the insurer j demand function, η_j is insurer j's per-enrollee noninpatient costs. I describe the demand function specifications in Ap-

pendix C.1. The associated flow pre-transfer gains from trade (Equation (1)) for a hospital i and insurer j at the equilibrium contracts $\hat{\mathbb{C}}_t$, networks $\hat{\mathcal{G}}_t$, and premiums $\hat{\phi}_t$ are:

$$\begin{aligned}
[\Delta_{ij}\pi_{it}^H] &= -c_i D_{ijt}\left(\hat{\mathbb{C}}_t\right) - c_i \sum_{n\in\mathcal{G}_{it}^H}\left(D_{int}\left(\hat{\mathbb{C}}_t\right) - D_{int}\left((\hat{\mathbb{C}}_{t,-ij}, \mathbb{C}_{0,ijt})\right)\right) \\
[\Delta_{ij}\pi_{it}^M] &= \left(D_{jt}^M\left(\hat{\mathcal{G}}_t, \hat{\phi}_t\right) - D_{jt}^M\left(\hat{\mathcal{G}}_t/ij, \hat{\phi}_t\right)\right)\left(\hat{\phi}_{jt} - \eta_j\right),
\end{aligned} \tag{7}$$

where $\hat{\mathcal{G}}_t/ij$ removes ij from the equilibrium networks. I assume that premiums are set annually through Nash-Bertrand competition, not taking into account any effect of premiums on future contract negotiations. The stylized nature of my premium data precludes taking a more precise approach to premium-setting.

I assume that contracts are negotiated as a benchmark (either a Medicare-linked "prospective" price or a list-price-linked "share of charges" price), a fixed benchmark price multiple (or equivalently initial price), and expiration (either a fixed expiration date or an auto-renew contract that can renew on January 1 of each year). I assume that firms that enter impasse continually attempt to reach a contract in good faith to avoid further painful exclusion, formalized as Good-Faith Disagreement in Appendix Assumption 3. I set $r_i = 0$ because it is not clear the parameter is separately-identified from the insurer costs r_j, and I assume the insurer costs r_j are constant for non-Highmark-BCBS insurers to increase statistical power. I approximate Pay_{IRT} to zero based on Corollary 3.

Based on Theorem 1, the initial price $p_{ijt_0}^*$ given expectations over the distribution of benchmark prices $p_{it}^{B_{ijt}}$ and expiration t^* solves:

$$\begin{aligned}
p_{ijt_0}^* \mathbb{E}_{t_0}\left[\sum_{t=t_0}^{t^*}\beta^{t-t_0} D_{ijt}^H(\mathcal{G}_t, \phi_t)\frac{p_{it}^{B_{ijt}}}{p_{it_0}^{B_{ijt}}}\right] &= \mathbb{E}_{t_0}\left[\sum_{t=t_0}^{t^*}\beta^{t-t_0}\left(-\tau_{ij}\left[\Delta_{ij}\pi_{it}^H\right] + (1-\tau_{ij})\left[\Delta_{ij}\pi_{jt}^M\right]\right)\right] \\
&\quad - (1-\tau_{ij})r_j.
\end{aligned}$$

The left-hand side is observed up to the patience parameter β. The right-hand side depends on the insurer's bargaining weight τ_{ij}, negotiation costs r_j, hospital demand D^H, insurer demand D^M, hospital costs c_i, insurer noninpatient costs η_j, and prices per unit of care p.

I parameterize the bargaining weights as an insurer fixed effect with hospital size effects on a logit scale:

$$\log\left(\tau_{ij}/(1-\tau_{ij})\right) = log(\tau_j/(1-\tau_j)) + \tau^{Size} log(HospSize_{i,2006}/MeanHospSize_{2006}), \tag{8}$$

where hospital size is measured as the size of the bargaining system in my first year of 2006 and τ_j is insurer bargaining power measured at the average-sized hospital system. Larger

Table 1: 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

hospitals have more bargaining weight if τ^{Size} is negative. I assume that the insurer bargaining weight parameter τ_j is shared for the for-profit insurers I model (Aetna, Cigna, Carelink, and UnitedHealth). In estimation, I allow τ_j to be below 0 or above 1 as a plausibility test; if that happens, then I take $\tau_{ij} = \tau_j$.

5.2 Data and Estimation

The main estimation dataset is the public record contract data referenced in Section 2. In the interest of brevity, I include only a high-level summary of my estimation procedure here.

I describe the core contract data in Section 2. Table 1 presents count statistics for the full contract data, the contract data involving the six modeled insurers, and the 63 contracts used in bargaining estimation.

I complement the novel contract data with more standard data on hospital and insurer quantities demanded. I group diagnoses into one of six main categories based on Ho (2006)'s International Classification of Diseases (ICD) categories (see Appendix Table 4 for frequencies). I use 2016 uniform billing (UB) data to estimate hospital and insurer quantities demanded for the identified insurers. The UB data is discharge data on every inpatient stay in West Virginia, and identifies primary payor if the primary payor is one of Aetna, Highmark BCBS, or HPUOV. The data is not claims data, because it does not include negotiated or realized payments. I estimate insurance demand for previous years and other modeled insurers using data on annual sales and premiums in the West Virginia fully insured market. The insurer fully insured sales data is digitized from state-level accident and health reports like Offices of the Insurance Commissioner (2016). The reports cover every plan sold in which an insurer is paid a premium to provide comprehensive medical insurance. I also use state financial reports on Medicare rates in supplemental analyses.

My hospital and insurer demand estimation strategy is fairly standard. I use hospital demand as an input to both the insurer demand model and bargaining model, and use insurer demand as an input to the bargaining model. I estimate hospital choice with logistic maximum likelihood as a function of diagnosis category and location. I estimate insurer

demand through the generalized method of moments (GMM). My strategies are similar to existing approaches. The most notable feature is that I leverage Affordable Care Act (ACA) premium restrictions, which prevented insurers from differentiating premiums within geographic rating areas beyond a limited set of homogeneously incorporated factors. I use these rating areas to control for most sources premium variation, and use the correlation of sales with network quality within rating areas to identify the effect of network quality on insurer sales.

I estimate bargaining parameters through GMM on estimated and predicted net present value payments. Unlike existing static approaches, I allow negotiators to balance gains from trade over the multiple years in which their contract remains in place, and I only use new contracts with reliable start and end dates in estimation. I estimate a finite-horizon bargaining model, because West Virginia was nonstationary which precludes estimating a compelling infinite-horizon model. I consider the first $T = 5$ years of gains from trade, which I take as an approximation to an infinite-length model.

My bargaining instruments are insurer and grouped-hospital-size dummy variables. The insurer and hospital size generate variation in contract length and terms, generating variation in how future conditions affect starting prices. Both instruments also identify the bargaining weight heterogeneity. I constrain insurer noninpatient costs η_j using moments on reported medical spending. My main specification use hospital reported costs as a proxy for the opportunity cost of inpatient care, although I consider other cost models as robustness checks. I bootstrap standard errors by resampling inpatient cases and state-level sales.

See Appendix C.1 for more on estimation, Appendix C.2 for sources of identification, Appendix C.6 for key caveats, and Appendix C.7 for a detailed description of the data analysis.

5.3 Parameter Estimates

I estimate an annual patience parameter of $\beta = 0.899$ and overwhelmingly reject the null hypothesis of myopia ($\beta = 0$). I find that most negotiations were close to the TU case in which Nash and Kalai coincide, and I find no evidence that a dynamic Nash-in-Nash model would be more accurate than the model I use.

My estimated hospital and insurer demand systems are generally plausible. The associated parameter estimates are presented in Appendix C.3.

I present bargaining estimates under two models, both of which use the same hospital and insurer demand estimates. The first strategy is a *myopic* approach that estimates bargaining parameters for contracts with confirmed start and end dates, but constrains the discount

rate β to zero to recover a static-type estimation strategy. The second strategy is a *forward-looking* approach described above, which allows the annual discount rate β to take on any value between zero (myopia) and one (discounting after CPI inflation).

	Parameter				
	β	τ_{BCBS}	τ_{HPUOV}	τ_{FP}	$-\tau^{Size}$
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***
($\text{Pay}_{IRT} = 0$)	(0.03)	(0.006)	(0.026)	(0.005)	(0.028)
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 2: Estimated bargaining and patience weights for the myopic (first row) and more general forward-looking (second row) bargaining models. The MCO τ_j bargaining weights are estimated for Highmark BCBS (BCBS), HPUOV, and the modeled for-profit insurers (FP) and represent the insurer's predicted share of gains from trade at the average bargain's hospital bargaining system log 2006 size. I present estimates of noninpatient costs (η) and net negotiation costs (r_j) in Appendix Table 8. Estimates under alternative bargaining models are presented in Appendix Table 9.

The bargaining estimates are presented in Table 2. I find that bargainers are forward-looking. I find excellent model fit (Appendix Figure 12). My estimated model overwhelmingly rejects a null hypothesis of myopia, which corresponds to the case $\beta = 0$. My estimated patience parameter of $\beta = 0.899$ is also below one. These indicate that firms care about future period profits ($\beta > 0$) but value a dollar today less than a dollar tomorrow ($\beta < 1$). The results are qualitatively similar under most other specifications (Appendix Table 9).

The forward-looking model does a better job of predicting initial rates than the myopic model, even though the forward-looking model targets payments in later contract years. For example, the correlation between forward-looking and myopic predicted and real starting share of list prices is 0.521 and 0.452, respectively (Appendix Figure 13). The staggered nature of contracting in Section 2 rejects a truly static model, and the forward-looking estimates reject a myopic model in which negotiators only value one period of gains from trade.

The estimated bargaining weights are generally empirically plausible. I find little heterogeneity in bargaining power across insurers: I estimate that insurers keep 85%–89% of the joint surplus when bargaining with a medium-sized hospital system (Table 2). The estimated weights are somewhat larger than other estimates in the literature based on claims

data (Ho and Lee, 2017, Ghili, 2022). There is no restriction in my estimation procedure that both sides must gain from the negotiated contracts, so it is reassuring that the myopic and forward-looking models estimate bargaining weights τ_{ij} between zero and one.

I find that while bargaining was NTU, most bargains used in estimation were close to the TU case in which Kalai and Nash coincide. I estimate the marginal effect of ij's initial price on NPV ij transfers and the spillover effect on non$-ij$ NPV transfers. In 62 of the 63 cases, both firm's NPV spillovers are at most 10% of the direct effect on their bilateral NPV payment (Appendix Figure 14). In the one exception, a Nash-in-Nash model would predict a higher payment than the Nash-in-Kalai model, but the realized payment was lower than the estimated model predicts (Appendix Figure 15). These results are consistent with the Nash-in-Kalai model providing a tractable approximation in my setting even if negotiators truly followed a more complex Nash-in-Nash model, and I find no evidence that a Nash-in-Nash model would be more accurate than the Nash-in-Kalai model I propose.

5.4 Counterfactuals

I now apply the estimated bargaining model to quantify a novel dynamic mechanism: the effect of government-set Medicare rate dynamics on private insurer payments through Medicare's role as a benchmark. I find that Medicare benchmark dynamics have an important role in private insurer spending, but a myopic model lacking forward-looking responses would both substantially overstate the effects and miss important short-run dynamics.

Medicare, a large American government program, pays hospitals at rates that the American Hospital Association argues have deflated related to costs (AHA, 2022). I ask how a counterfactual change in Medicare rates to roughly track West Virginia hospital reported cost growth would affect spending in the private insurer market.

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

I classify contracts as either likely list-price-linked "share of charges" contracts or likely-Medicare-linked "prospective" contracts. This notation follows Cooper et al. (2019). See

Appendix C.1 for more on the algorithm and Appendix C.5 for more on the logic.[7] I present estimated frequencies Table 3. I classify an estimated 46.7% of inpatient spending as prospective, and find that prospective contracts were especially common for the largest insurer Highmark BCBS, as well as for HPUOV in the regions in which it was a larger actor.

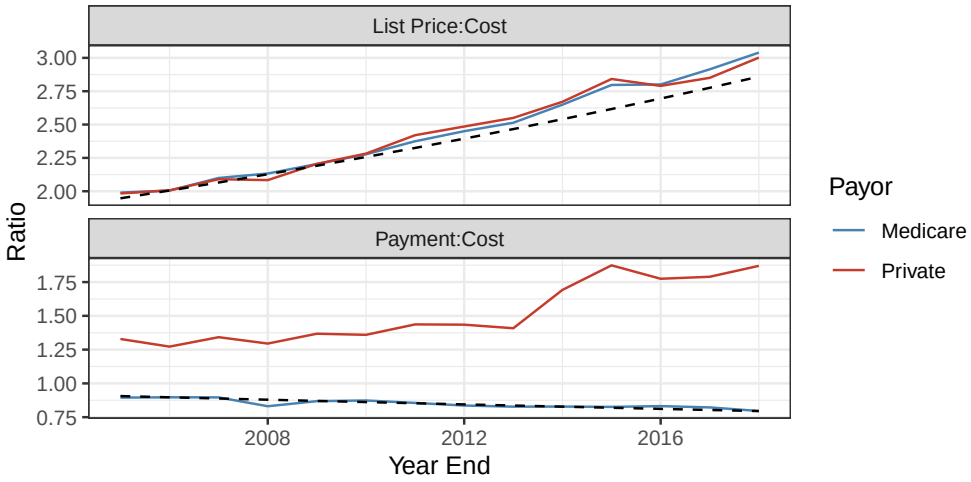


Figure 7: The ratio of list price charges (top) and real payments (bottom) to reported costs by Medicare (red) and private payors (blue) for West Virginia hospitals. Dashed lines represent Medicare 2006 values extrapolated based on three ppt. and negative one ppt. annual changes.

The choice of benchmark had important consequences on how prices evolve after a contract is formed. Figure 7 presents the ratio of charges and payments to hospital reported costs. List price charges increased quickly, roughly three percentage points faster than hospital reported costs. Medicare rates depreciated relative to hospital reported costs by roughly one percentage point annually. As a result, if insurer A formed a contract with a hospital paying a fixed multiple of charges, and insurer B formed a contract at the same time with the same average initial payment calculated as a fixed fraction of Medicare rates, then after five years, insurer A could easily pay 20% more. The associated divergence in payments between negotiations was an important driver of the divergence in payments between Highmark BCBS and other insurers in this era (Dorn, 2024). A counterfactual increase in Medicare benchmark rates over time could have real effects even in the long-run when negotiators discount the future and price-setting is staggered, similar to a (Taylor, 1980) model of staggered price-setting with inflation.

I study the counterfactual effect of an additional one-percentage-point annual increase in hospital prospective prices, announced at the end of 2006 to begin in 2007. The one-percentage-point increase would roughly offset the divergence between West Virginia Medi-

[7]Note that some charges are based on diagnoses, and are therefore also prospective. The key difference is that Medicare-linked prices evolve based on a national formula.

care rates and hospital costs in this era. The change would correspond to a $26.5-billion increase in 2015 Medicare hospital expenditures (CMS, 2022, expressed in 2019 dollars). I do not have access to the DRG weight schemes used in the construction of payments, so I assume all prices imputed as prospective would increase one percentage point faster annually beginning in 2007. My counterfactual can also be interpreted as a world in which these contracts' prices were required to inflate one percentage point faster annually.

Counterfactual payments are calculated as follows. I first calculate counterfactual contract prices holding benchmark choice, contract length, premiums, and NPV residual payment fixed, and using plug-in estimates of counterfactual benchmark price growth. I then estimate back-of-the-envelope downstream premium effects based on a Nash-Bertrand premium competition model assuming Ho (2006)'s estimated own-price elasticity average own-price elasticity of -1.14. I do not estimate a reinforcing effect of higher premiums on prices, which biases my estimates towards zero. I discuss these choices further in Appendix C.1.

I summarize the effects by year under a hypothetical static model, the estimated forward-looking model, and the estimated myopic model in Figure 8. The effects are presented as a percentage of modeled insurer spending, which is roughly three-quarters of all private insurer spending (Dorn, 2024). A static model with one-period contracts (black) would fail to estimate any effect: even if Medicare doubled rates annually, Medicare-benchmarking negotiators could cut their contracted multiples in half and arrive at the same real outcomes. I find that contracts are multiyear, so that the reform will at least have a short-run effect.

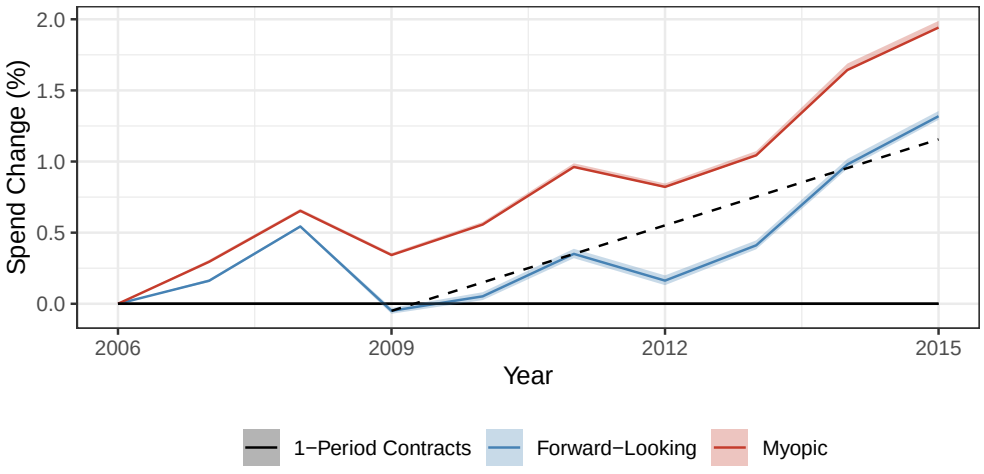


Figure 8: Estimated counterfactual spending effects from a one-percentage-point increase in Medicare rates from a myopic (blue) and dynamic (red) bargaining model. The dashed line indicates 0.20 percentage point additional annual spending increases starting in 2009.

Under the forward-looking model (red), I estimate quantitatively meaningful spending

effects and important short-run dynamics. The estimated effect under the forward-looking model is negative in 2009, when Highmark BCBS and HPUOV negotiated many new prospective contracts and correspondingly would reduce their starting prices relative to the status quo. There is a similar dip in 2012 when Highmark BCBS revised three-year contracts formed in 2009. In later years, contracts remain in place and effects compound. After nine years, the estimated increase in spending is 1.3%. The estimated increase in West Virginia spending in the commercially insured inpatient market would be $7.1 million. The percent change in spending after nine years, if extrapolated to the 2015 national hospital market and inflation-adjusted to 2019 dollars, corresponds to a $4.98-billion effect. I present effects on insurer payments, hospital payments, and premiums in Appendix Figures 16, 17, and 18, respectively.

A myopic model would miss these short-run dynamics and generally overestimate the effects by 45% or more. In the middle of the panel, when the forward-looking model estimates are small or negative due to many negotiations, the myopic model incorrectly predicts a substantial spending increase. The reason the myopic bargaining model overestimates the effects of benchmark price increases is because forward-looking bargainers respond to anticipated future increases, while myopic bargainers only care about accumulated increases so far.

The estimated 2015 spending increase of 1.319% is below the compounded benchmark price increase of 9.37% for three main reasons: many payments were benchmarked to unaffected list prices, the Medicare-benchmarked contracts were renegotiated every few years, and forward-looking bargainers revise starting prices downward based on the anticipated future benchmark price increases. I measure the importance of the first mechanism by considering a same-multiple model, wherein the firms keep their original negotiated benchmark multiples α in place. Appendix Figure 19 presents the different estimated effects. The mechanical model that only accounts for list price contracts would estimate a massively too large nine-year effect that corresponds to nine compounded 0.564 percentage point spending increases. As seen in Figure 8, the myopic model that only includes the first two reasons and equilibrium spillovers, would also substantially overestimate the effects and miss important dynamics like a spending reduction in 2009.

The Nash-in-Kalai model's ability to account for multiyear and staggered contracts allows me to quantify both how the proposed Medicare rate reform would change price dynamics between negotiations, and how forward-looking negotiators would respond to anticipated future Medicare-driven price increases when choosing contract multiples and starting prices. Approaches based on the static literature would have been biased in one of two directs: a static model would be biased towards zero, while the myopic model that corresponds to static-type estimation with accurate timing would be biased away from zero.

6 Discussion

This paper studies dynamic bargaining in vertical markets. I demonstrate a new challenge in capturing dynamic vertical markets: state space growth. I propose the Nash-in-Kalai model, which generalizes the Nash-in-Nash approach in important static settings while uniquely overcoming this new tractability challenge in dynamic markets. I estimate an empirical Nash-in-Kalai bargaining model with data from West Virginia. I find that negotiators are forward-looking and substantially offset proposed dynamic changes to Medicare rates. Nevertheless, I find that a dynamic increase to Medicare rates would still have significant effects on private insurer spending through Medicare's role as a benchmark — a mechanism that existing static models cannot capture.

The empirical analysis here opens up exciting new avenues for future empirical work. In the West Virginia context, my companion paper Dorn (2024) documents that smaller insurers experienced rapid price increases under long-lived contracts. A forward-looking model is needed to quantify the impact of these contracts on competition, spending, and premiums. Further, static approaches to antitrust questions, such as merger effects, typically assume that all contracts in a market were recently formed. In West Virginia, such static models would underestimate smaller insurers' bargaining power by misattributing fast ex post price increases to a lack of ex ante bargaining power. Future work will explore the implications of staggered contracting for antitrust analysis.

The Nash-in-Kalai model is likely to be valuable beyond healthcare. The framework may prove useful for modeling media asset acquisition (Morrissette, 2023), programming disputes with time-bound programming (Hayes, 2023), disagreement with inertia costs (Handel, 2013), staggered team formation in the presence of aggregate wage limits (Mulholland and Jensen, 2019), and supply chain bargaining in the presence of inflation (Davidson, 1988, Baudendistel, 2023). One speculative application could be to extend popular empirical search-on-the-job models (Cahuc et al., 2006, Bagger et al., 2014, Bilal et al., 2022) beyond TU bargaining.

This paper leaves a new methodological question for future work. I show that the Nash-in-Kalai model is uniquely tractable in dynamic markets, but I cannot prove that the Nash-in-Kalai model is always accurate. In the West Virginia application, I find that the Nash-in-Kalai model likely offers similar predictions to a more complex dynamic Nash-in-Nash model. However, not all settings will resemble West Virginia. This raises a critical open question: is there a tractable approximation to dynamic Nash-in-Nash in settings where Nash-in-Kalai and Nash-in-Nash predictions meaningfully differ? Yet even in these challenging new settings, the Nash-in-Kalai model will remain a plausible and useful benchmark and a foundation for others to build on.

References

Abbey, D. C. (2012). *Cost-Based, Charge-Based, and Contractual Payment Systems*. Productivity Press Inc., 1st edition.

Abbott, T. A. (1995). Price regulation in the pharmaceutical industry: Prescription or placebo? *Journal of Health Economics*, 14(5):551–565.

AHA (2022). Fact sheet: Majority of hospital payments dependent on Medicare or Medicaid.

Andreoni, J. and Bernheim, B. D. (2009). Social image and the 50-50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, 77(5):1607–1636.

Atkinson, G. (2009). State hospital rate-setting revisited. *Commonwealth Fund*, 69(1332).

Backus, M., Blake, T., Larsen, B., and Tadelis, S. (2020). Sequential bargaining in the field: Evidence from millions of online bargaining interactions. *The Quarterly Journal of Economics*, 135(3):1319–1361.

Bagger, J., Fontaine, F., Postel-Vinay, F., and Robin, J.-M. (2014). Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics. *American Economic Review*, 104(6):1551–96.

Baudendistel, M. (2023). CPG cost pressure begins to ease. *FresightWaves*.

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

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., Rubinstein, A., and Wolinsky, A. (1986). The Nash bargaining solution in economic modelling. *The RAND Journal of Economics*, 17(2):176–188.

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

Borağan Aruoba, S., Rocheteau, G., and Waller, C. (2007). Bargaining and the value of money. *Journal of Monetary Economics*, 54(8):2636–2655.

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.

Bárcena-Ruiz, J. C. and Casado-Izaga, F. J. (2008). Timing of endogenous bargaining over costs and firms' locations. *Journal of Economics*, 95(2):149–166.

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.

CBO (2020). The federal budget in 2019: An infographic.

Clemens, J. and Gottlieb, J. D. (2017). In the shadow of a giant: Medicare's influence on private physician payments. *Journal of Political Economy*, 125(1):1–39.

Clemens, J., Gottlieb, J. D., and Molnár, T. L. (2017). Do health insurers innovate? Evidence from the anatomy of physician payments. *Journal of Health Economics*, 55:153–167.

Clemens, J. and Ippolito, B. (2019). Uncompensated care and the collapse of hospital payment regulation: An illustration of the Tinbergen Rule. *Public Finance Review*, 47(6):1002–1041.

CMS (2020). Files for FY 2016 final rule and correction notice.

CMS (2022). National health expenditure data: Historical.

CMS (2023). Market rating reforms.

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.

Cooper, Z., Craig, S. V., Gaynor, M., and Reenen, J. V. (2019). The price ain't right? Hospital prices and health spending on the privately insured. *Quarterly Journal of Economics*, 134(1):51–107.

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.

Cromwell, J. (1987). Impact of state hospital rate setting on capital formation. *Health Care Financing Review*, 8(3):69–82.

Davidson, C. (1988). Multiunit bargaining in oligopolistic industries. *Journal of Labor Economics*, 6(3):397–422.

De Fraja, G. (1993). Staggered vs. synchronised wage setting in oligopoly. *European Economic Review*, 37(8):1507–1522.

Deng, S., Hui, S., Jia, D., Leccese, M., and Sweeting, A. (2022). Bargaining, forward-looking buyers and dynamic competition.

Deng, S., Jia, D., Leccese, M., and Sweeting, A. (2023). Bargaining and dynamic competition.

Diamond, P. A. (1982). Wage determination and efficiency in search equilibrium. *The Review of Economic Studies*, 49(2):217–227.

Diamond, P. A. and Maskin, E. (1979). An equilibrium analysis of search and breach of contract, i: Steady states. *The Bell Journal of Economics*, 10(1):282–316.

Diebel, A. S. and Diebel, N. (2017). Welfare effects of using hospital rate setting as an alternative to bargaining.

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

Dorn, J. (2024). Six stylized facts from ten years of vertical market contract data.

Draganska, M., Klapper, D., and Villas-Boas, S. B. (2010). A larger slice or a larger pie? An empirical investigation of bargaining power in the distribution channel. *Marketing Science*, 29(1):57–74.

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, P. K. (2021). Compromise is key in infinitely repeated bargaining with an evergreen clause. *Journal of Economic Theory*, 193:105170.

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

Dutta, R. (2022). Bargaining as a struggle between competing attempts at commitment.

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

Finder, B. (2022). Four reasons Medicare is an inadequate Benchmark for commercial health plans. *American Hospital Association*.

Flamini, F. (2020). Divide and invest: Bargaining in a dynamic framework. *Homo Oeconomicus*, 37(1):121–153.

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

Froeb, L. M., Mares, V., and Tschantz, S. (2021). Modeling bargaining competition: Nash-in-Shapley vs. Nash-in-Nash.

Fuchs, W. and Skrzypacz, A. (2022). Dynamic bargaining with private information. In Karagözoğlu, E. and Hyndman, K. B., editors, *Bargaining: Current Research and Future Directions*, pages 61–86. Springer International Publishing, Cham.

Gertler, M., Sala, L., and Trigari, A. (2008). An estimated monetary dsge model with unemployment and staggered nominal wage bargaining. *Journal of Money, Credit and Banking*, 40(8):1713–1764.

Gertler, M. and Trigari, A. (2009). Unemployment fluctuations with staggered Nash wage bargaining. *Journal of Political Economy*, 117(1):38–86.

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

Goldberg, V. P. (1976). Regulation and administered contracts. *The Bell Journal of Economics*, 7(2):426–448.

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

Gottfries, A. (2022). Bargaining with renegotiation in models with on-the-job search.

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.

Grennan, M. and Swanson, A. (2022). Bargaining in healthcare markets: Applications of Nash-in-Nash and extensions. In Karagözoğlu, E. and Hyndman, K. B., editors, *Bargaining: Current Research and Future Directions*, pages 249–270. Springer International Publishing, Cham.

Gudiksen, K. L., Chang, S. M., and King, J. S. (2019). The secret of health care prices: Why transparency is in the public interest.

Handel, B. R. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, 103(7):2643–82.

Hayes, D. (2023). DirecTV and Nexstar dig in for long-haul carriage dispute during lull in summer programming. *Deadline*.

Hermo, S. (2024). Collective bargaining networks, rent-sharing, and the propagation of shocks.

Highmark West Virginia (2011). Hospital outpatient billing and reimbursement guide: Version 11.01: Provider version.

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.

Hoof, S. (2018). Dynamic voluntary provision of public goods: The recursive Nash bargaining solution. In Petrosyan, L. A., Mazalov, V. V., and Zenkevich, N. A., editors, *Frontiers of Dynamic Games: Game Theory and Management, St. Petersburg, 2017*, pages 13–23. Springer International Publishing, Cham.

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.

Jacquet, L., Lehmann, E., and der Linden, B. V. (2014). Optimal income taxation with Kalai wage bargaining and endogenous participation. *Social Choice and Welfare*, 42(2):381–402.

Jahn, M. (2017). Seven must-dos to ensure charge description master (CDM) success. *Becker's Hospital CFO Report*.

Jarosch, G., Nimczik, J. S., and Sorkin, I. (2019). Granular search, market structure, and wages. Working Paper 26239, National Bureau of Economic Research.

Joskow, P. L. (1987). Contract duration and relationship-specific investments: Empirical evidence from coal markets. *The American Economic Review*, 77(1):168–185.

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

Kalai, E. and Smorodinsky, M. (1975). Other solutions to Nash's bargaining problem. *Econometrica*, 43(3):513–518.

Keniston, D., Larsen, B. J., Li, S., Prescott, J., Silveira, B. S., and Yu, C. (2024). Splitting the difference in incomplete-information bargaining: Theory and widespread evidence from the field. *International Economic Review*, n/a(n/a).

Keniston, D. E. (2011). Bargaining and welfare: A dynamic structural analysis.

Kidder, R. (2013). Best practice chargemaster management.

Kostadinov, R. (2021). Renegotiation of long-term contracts as part of an implicit agreement. *Theoretical Economics*, 16(4):1471–1512.

Lagos, R., Rocheteau, G., and Wright, R. (2017). Liquidity: A new monetarist perspective. *Journal of Economic Literature*, 55(2):371–440.

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.

Levin, J. (2003). Relational incentive contracts. *American Economic Review*, 93(3):835–857.

Liebman, E. (2022). Bargaining in markets with exclusion: An analysis of health insurance networks.

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

MacLeod, W. B. and Malcomson, J. M. (1993). Investments, holdup, and the form of market contracts. *The American Economic Review*, 83(4):811–837.

Malcomson, J. M. and MacLeod, W. B. (1989). Implicit contracts, incentive compatibility, and involuntary unemployment. *Econometrica*, 57(2):447–480.

Marcelo, R. (2021). New Kagan report: Cable TV carriage deals revenue estimates. *S&P Global Market Intelligence*.

Maskin, E., de Clippel, G., and Serrano, R. (2021). Sequential bargaining, coalitions, and externalities.

McDonough, J. E. (1997). Tracking the demise of state hospital rate setting. *Health Affairs*, 16(1):142–149. PMID: 9018951.

Miller, D. A. and Watson, J. (2013). A theory of disagreement in repeated games with bargaining. *Econometrica*, 81(6):2303–2350.

Morrisey, M. A. (2013). *Health Insurance, Second Edition*. Health Administration Press.

Morrissette, C. (2023). Big 12 'totally out maneuvered pac-12' by getting early media rights deal done. *USC Trojans on SI*.

Mulholland, J. and Jensen, S. T. (2019). Optimizing the allocation of funds of an nfl team under the salary cap. *International Journal of Forecasting*, 35(2):767–775.

Murray, R. and Berenson, R. A. (2015). *Hospital Rate Setting Revisited: Dumb Price Fixing or a Smart Solution to Provider Pricing Power and Delivery Reform?* Urban Insitute, Washington, DC.

Myerson, R. B. (1981). Utilitarianism, egalitarianism, and the timing effect in social choice problems. *Econometrica*, 49(4):883–897.

NAIC (2022). U.S. health insurance industry analysis report: 2021 annual results.

Nash, J. F. (1950). The bargaining problem. *Econometrica*, 18(2):155–162.

NBER (2021). ICD-9-CM to and from ICD-10-CM and ICD-10-PCS crosswalk or general equivalence mappings.

New York State Procurement (2014). Agreement for basic medical supplies and equipment (statewide) by and between New York Sstate Office of General Services and Medline Industries. Contract Number PC66415.

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.

Offices of the Insurance Commissioner (2008). *Accident and Health Insurance Market Report for 2008*. West Virginia Offices of the Insurance Commissioner, Charleston, WV.

Offices of the Insurance Commissioner (2016). *Accident and Health Insurance Market Report 2016: Including data from 2010-2015*. West Virginia Offices of the Insurance Commissioner, Charleston, WV.

Pauly, M. and Town, R. (2012). Maryland exceptionalism? All-payers regulation and health care system efficiency. *Journal of Health Politics, Policy and Law*, 37(4):697–707.

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.

Reinhardt, U. E. (2006). The pricing of U.S. hospital services: Chaos behind a veil of secrecy. *Health Affairs*, 25(1):57–69. PMID: 16403745.

Rica, S. D. l. and Espinosa, M. P. (1997). Testing employment determination in unionised economies as a repeated game. *Scottish Journal of Political Economy*, 44(2):134–152.

Ridley, D. B. and Zhang, S. (2017). Regulation of price increases. *International Journal of Industrial Organization*, 50:186–213.

Rivard, R. (2010). Battle looms for hospital rate hike: Clarksburg medical center wants to up prices 21 percent. *Charleston Daily Mail*.

Roberts, E. T., Hatfield, L. A., McWilliams, J. M., Chernew, M. E., Done, N., Gerovich, S., Gilstrap, L., and Mehrotra, A. (2018a). Changes in hospital utilization three years into Maryland's global budget program for rural hospitals. *Health Affairs*, 37(4):644–653.

Roberts, E. T., McWilliams, J. M., Hatfield, L. A., Gerovich, S., Chernew, M. E., Gilstrap, L. G., and Mehrotra, A. (2018b). Changes in health care use associated with the introduction of hospital global budgets in Maryland. *JAMA Internal Medicine*, 178(2):260–268.

Roberts, E. T., Mehrotra, A., and Chernew, M. E. (2018c). Maryland's hospital global budget program. *JAMA*, 320(19):2040–2040.

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

Roth, A. E., Murnighan, J. K., and Schoumaker, F. (1988). The deadline effect in bargaining: Some experimental evidence. *The American Economic Review*, 78(4):806–823.

Serrano, R. (2005). Fifty years of the Nash program: 1953-2003. *Investigaciones Economicas*, 29(2):219–258.

Sestini, R. (1999). Union–firm bargaining as a repeated game and the behaviour of wages over the business cycle. *LABOUR*, 13(4):821–857.

Shapiro, C. (2021). Vertical mergers and input foreclosure lessons from the AT&T/Time Warner case. *Review of Industrial Organization*, 59(2):303–341.

Sharfstein, J. M., Stuart, E. A., and Antos, J. (2018a). Global budgets in Maryland: Assessing results to date. *JAMA*, 319(24):2475–2476.

Sharfstein, J. M., Stuart, E. A., and Antos, J. (2018b). Maryland's all-payer health reform—a promising work in progress. *JAMA Internal Medicine*, 178(2):269–270.

Shimer, R. (2006). On-the-job search and strategic bargaining. *European Economic Review*, 50(4):811–830.

Sims, C. A. (2002). Solving linear rational expectations models. *Computational Economics*, 20(1):1–20.

Sorensen, A. T. (2003). Insurer-hospital bargaining: Negotiated discounts in post-deregulation Connecticut. *The Journal of Industrial Economics*, 51(4):469–490.

Sorger, G. (2006). Recursive Nash bargaining over a productive asset. *Journal of Economic Dynamics and Control*, 30(12):2637–2659.

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.

Thompson, A. (2023). AHA statement on FY 2024 proposed IPPS & LTCH payment rule.

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

Tompkins, C. P., Altman, S. H., and Eilat, E. (2006). The precarious pricing system for hospital services. *Health Affairs*, 25(1):45–56. PMID: 16403744.

Vega, K. B. (2013). Successfully negotiating managedcare contracts. *HFMA*.

Watson, J., Miller, D. A., and Olsen, T. E. (2020). Relational contracting, negotiation, and external enforcement. *American Economic Review*, 110(7):2153–97.

Wayback Machine (2021). The Health Plan: Strategic partners. June 23, 2021.

Weber, E., Floyd, E., Kim, Y., and White, C. (2019). Peering behind the veil: Trends in types of contracts between private health plans and hospitals. *Medical Care Research and Review*.

World Bank (2023). GDP (current US$).

Yu, X. and Waehrer, K. (2019). Recursive Nash-in-Nash bargaining solution.

A Additional Literature

I now discuss some of the many strands of literature that interact with this paper.

There is a substantial literature on vertical market bargaining, most of which is static. There is a theoretical literature in industrial organization on overlapping contracts in triangular markets (De Fraja, 1993, Bárcena-Ruiz and Casado-Izaga, 2008, Do and Miklós-Thal, 2022). Such work is highly stylized and inapplicable for empirical markets like West Virginia with many asymmetric firms. Recent work argues the Nash equilibrium assumption can do a poor job of replicating true responses to disagreement (Ho and Lee, 2019, Yu and Waehrer, 2019, Froeb et al., 2021, Liebman, 2022, Ghili, 2022). An interesting avenue for future work is using a dynamic vertical market model to more realistically capture disagreement. Note also that Lee et al. (2021) refer to the mechanism I call spillovers as contracting externalities. There is also some work on the Nash-in-Nash solution for union negotiations: Davidson (1988) introduced the Nash-in-Nash solution for multiunit bargaining between many employers and many insurers at roughly the same time as Horn and Wolinsky (1988) introduced the solution for vertical markets with bargaining between firms. To my knowledge, the multiunit bargaining literature generally models bargaining as static, although Hermo (2024) is a notable exception.

Nash bargaining has many advantages for empirical use in static settings. Symmetric ($\tau = 1/2$) Nash bargaining can be justified as a bilateral bargaining solution as the unique solution that satisfies Pareto efficiency, independence of irrelevant alternatives, symmetry, and scale invariance (Nash, 1950). It can be microfounded as a bilateral solution through alternating offers (Binmore et al., 1986, Coles and Muthoo, 2003, Collard-Wexler et al., 2019) and as a vertical market solution for TU games (Collard-Wexler et al., 2019).

The Kalai proportional solution also has many advantages for empirical use in static settings. Symmetric ($\tau = 1/2$) Kalai proportional bargaining can be justified as a bilateral bargaining solution as the unique solution that satisfies Pareto efficiency, independence of irrelevant alternatives, symmetry, and monotonicity (Kalai, 1977). The static Nash-in-Kalai solution is microfounded in the TU case in which the Nash-in-Nash model is microfounded (Collard-Wexler et al., 2019). Unlike the Nash bargaining solution, the Kalai proportional solution is scale-varying: doubling player i's utility scale leads to player j having a better outcome. This makes the Nash-in-Kalai solution applicable only to settings in which the scale of utility can be well-defined, and bilateral microfoundations for NTU proportional bargaining require an external factor to fix the relative scale of utility such as arbitration (Bossert, 1994) or revocation costs (Dutta, 2012). Scale-varying solutions can be more accurate in lab settings in which utility scales are well-defined (Nydegger and Owen, 1974, Duffy

et al., 2021). Whether the Nash-in-Nash or Nash-in-Kalai solution is more appropriate for a given setting depends on which model best balances accuracy and tractability.

There is a long literature explicitly using the Kalai bargaining solution outside vertical markets. The Kalai proportional bargaining solution or similar proportional solution concepts have has been applied to study liquidity constraints in monetary economics (Borağan Aruoba et al., 2007, Lagos et al., 2017, Hu and Rocheteau, 2020, Duffy et al., 2021), in social choice theory (Myerson, 1981), and occasionally in labor economics (Sestini, 1999, Jacquet et al., 2014). The Kalai proportional bargaining solution has been critiqued as not being scale invariant, so that it can only be microfounded through von Neumann-Morgenstern utility if bargainers care about outcomes besides the utility being divided (Serrano, 2005). Evidence from the lab suggests the Kalai proportional solution may be more accurate in monetary bargaining when the two disagree and utility can be expressed in dollars (Nydegger and Owen, 1974, Duffy et al., 2021). Andreoni and Bernheim (2009) argue that negotiators often care about being perceived as fair. Keniston et al. (2024) find that split-the-difference bargaining strategies are common, but need not correspond to a proportional solution.

Medicare rates, the topic of my counterfactual analysis, is an active topic of regulatory debate. Medicare inpatient expenditures amounted to $319 billion in 2019 (CMS, 2022), comparable to the gross domestic product of Egypt (World Bank, 2023) and enough to account for 27% of American hospital expenditures, 40% of Medicare expenditures, and 7% of federal government spending (CBO, 2020, CMS, 2022). The American Hospital Association argues that Medicare payment increases have failed to keep up with costs (Thompson, 2023) as policymakers reduce Medicare rates to make up for budgetary shortfalls (Finder, 2022). However, there is limited work on the spillovers of Medicare payment dynamics on private insurer payment dynamics.

I find that forward-looking bargainers reduce starting prices to offset payments in net present value terms. There is a related literature on capping the inflation of pharmaceutical prices with forward-looking responses in price setting (Abbott, 1995, Ridley and Zhang, 2017). However, the dynamics I highlight here are closer to models of price-setting with inflation expectations (Taylor, 1980, Calvo, 1983). My counterfactual model is especially similar to the Taylor model; in both settings, nominal price inflation can have real effects when price setting is staggered and price-setters discount future profits.

This work relates to literature on dynamic bargaining outside vertical markets. Much of this literature is theoretical (Malcomson and MacLeod, 1989, MacLeod and Malcomson, 1993, Levin, 2003, Miller and Watson, 2013, Watson et al., 2020, Dutta, 2021), including dynamic underpinnings of static bargaining solutions (Binmore et al., 1986, Stole and Zwiebel, 1996, Brügemann et al., 2018, Collard-Wexler et al., 2019, Maskin et al., 2021, Dutta, 2022)

and dynamic negotiations over a single product in isolation (Roth et al., 1988, Rica and Espinosa, 1997, Keniston, 2011, Backus et al., 2020).

An important empirical literature on dynamic bargaining involves wage negotiations. In the search on the job literature, contracts formed at the start of employment affect the subsequent path of wages. Early examples include Diamond and Maskin (1979), Diamond (1982). Shimer (2006) pointed out that when workers cannot fully leverage their employer when considering changing firms, then higher wages can prevent a worker from leaving to a less-productive firm, the bargaining utility set can be nonconvex, and the Nash bargaining approach may not imply unique gains from trade absent randomization. As a result, the empirical search-on-the-job literature has used models in which initial wages do not affect the path of employment (Cahuc et al., 2006, Bagger et al., 2014) and has generally assumed transferable utility (Bilal et al., 2022, Gottfries, 2022). In the associated applied models, the Nash and Kalai proportional bargaining solutions agree. Jarosch et al. (2019) independently discovered the implications of path independence of zero-sum bargains in a special case of their model of post-disagreement threats. (Similar results were found for a stationary vertical market by Shapiro (2021).) The TU search-on-the-job approach is inapplicable to vertical markets with asymmetric and time-varying spillovers. Not all search-on-the-job models can be viewed as using Kalai proportional bargaining solution. In more recent work, Gottfries (2022) proposes a search model that nests the TU case, but allows for NTU bargaining. Gertler and Trigari (2009) propose a model of staggered firm Nash bargaining over multi-period wages, in which wages reset (and have no subsequent effect) after the next bargain, workers have quasilinear utility, and the firm has a linear and known externality on wages paid to employees hired before the next bargain. They show that this case yields a tractable solution incorporating the externality and calibrate the model to monthly jobs data. That approach is applied to an empirical setting with overlapping contracts in Gertler et al. (2008). However, in this literature, overlapping contracts interact through aggregate market states rather than based on substitution patterns.

This paper contributes to the empirical literature on vertical market contracting beyond West Virginia hospitals. This project sheds light on the practice of long-term contracting, which has implications for the literature on bargaining in hospital markets — as examined in work like Gowrisankaran et al. (2015), Ho and Lee (2017), Ghili (2022), and Prager and Tilipman (2022) — and may hold in other vertical markets, such as retail supply chains (Draganska et al., 2010, Baudendistel, 2023), medical device supply (Grennan, 2013, New York State Procurement, 2014, Grennan and Swanson, 2022), and content distribution (Crawford et al., 2018, Marcelo, 2021). This work exploits public record contract data, including Sorensen (2003)'s analysis of contracts in Connecticut, and complements Cooper et al. (2019)'s anal-

ysis of five years of claims data from three major national insurers. Their research measures service-level prices and forms a crucial foundation to my work. Weber et al. (2019) complement Cooper et al. by describing prices using claims data that encompasses both large and small insurers in Colorado. My counterfactual mechanism complements Clemens and Gottlieb (2017), who point out that Medicare can also be relevant to the outside option for capacity-constrained physicians. Clemens et al. (2017)'s descriptive analysis of Texas BCBS 2010 claims suggests Medicare may play a similar role as a benchmark for negotiated physician payments.

I study West Virginia's outcomes under a rate review system. Murray and Berenson (2015) offer an excellent description of the West Virginia and Maryland rate review systems from the year before the West Virginia system ended. McDonough (1997) has more on the history of rate review systems generally from after other state systems ended. There is a literature studying rate review systems more generally — for example, Cromwell (1987), Atkinson (2009), Pauly and Town (2012), Diebel and Diebel (2017), Sharfstein et al. (2018b,a), Roberts et al. (2018a,b,c), and Clemens and Ippolito (2019) — largely with a focus around changes to Maryland's rate-setting system in 2014.

The existence of multiyear hospital–insurer contracts begs the question of why the contracts exist in the first place. Early health insurance plans took the form of either indemnity plans without negotiated discounts or hospital-sold prepaid plans that aimed to reduce marginal costs during the Great Depression. In later years, negotiated discounts can be viewed as the insurer offering to steer patients towards a given hospital in exchange for the favorable payment rate,[8] with the insurer using the discounts to sell more plans (Morrisey, 2013). Viewed from this framework, hospital–insurer contracts resolve a potential holdup problem (Goldberg, 1976, MacLeod and Malcomson, 1993). Without a contract, a hospital could drum up business with the promise of discounts, and then renege and raise prices. In this framework, multiyear contracts might exist to mitigate holdup created if in-network hospitals become more valuable to consumers through learning. Such dynamic demand processes or a full analysis of the drivers of contract length are outside the scope of this work.

This work also touches on several strands of theoretical literature that are outside of its scope. There is a small and generally theoretical literature on dynamic recursive bargaining solutions. In that literature, contracts and choices today affect the outcome under the same bargaining solution in the future. My model is a dynamic recursive bargaining solution. There is relevant work under specific bargaining solutions in non-vertical settings (e.g. Sorger, 2006, Hoof, 2018, Flamini, 2020, Dutta, 2021); dynamic foundations of static bargaining solutions

[8] In practice, inpatient steering is typically achieved by treating a hospital as in-network rather than by steering between in-network hospitals.

(e.g. Binmore et al., 1986, Stole and Zwiebel, 1996, Coles and Muthoo, 2003, Brügemann et al., 2018, Collard-Wexler et al., 2019, Backus et al., 2020, Maskin et al., 2021, Dutta, 2022); and relational contracting under potential enforcement and restrictions (Levin, 2003, Miller and Watson, 2013, Watson et al., 2020, Kostadinov, 2021). "Dynamic bargaining" is sometimes used to refer to bargaining in settings that change over time rather than bargaining with contract outcomes that are dynamic (Fuchs and Skrzypacz, 2022), a related concept that to my knowledge has not been applied in vertical market settings except in dynamic bargaining models that lead to simultaneous formation of single-period contracts (Lee and Fong, 2013, Collard-Wexler et al., 2019). It is likely that the price system I describe can be exactly or approximately viewed as a linear rational expectations system. The corresponding literature is outside the scope of this work and generally focuses on different properties. (In the notation of Sims (2002), from my understanding the linear rational expectations literature must solve for Γ_0 and Γ_1, whereas I estimate Γ_0 and Γ_1 from the bargaining model and am interested in counterfactual changes to Γ_0.)

B Additional Theoretical Content

B.1 Nash-in-Nash Split-in-Half

Lemma 2 (Running example, Nash-in-Nash)**.** *Consider the split-in-half subgame of the running example dynamic Nash-in-Nash bargaining model, and suppose $\tau < 1$. Then there is a Markov perfect equilibrium such that the simultaneous negotiation price is:*

$$p_{(Simult)}^{(Nash),*} = (1-\tau)\phi(D^M - \underline{D}^M)\frac{B * 2(1-\beta^{1/2}) + (1+\beta^{1/2})}{A - B * \left(D^M \beta^{1/2}(2 - \tau(1 - \alpha_{(Alt)}^{(Nash)}\beta^{1/2}))\right)},$$

where $\phi = \$10,000$, $D^M = 6,000$, $\underline{D}^M = 5,000$, the marginal effect of a price negotiated on July 1 on the other hospital's price negotiated on the next January 1 is $\alpha_{(Alt)}^{(Nash)} = \frac{(1-\tau)(2\underline{D}^M - D^M)}{(1+\beta^{1/2})D^M + \beta^{1/2}(1-\tau)(2\underline{D}^M - D^M)}$, and the two constants A and B are:

$$A = \tau(1+\beta^{1/2})\frac{D^M}{2} + (1-\tau)(1+\beta^{1/2})(D^M - \underline{D}^M) + \frac{\tau\beta(1+\beta^{1/2})(\alpha_{(Alt)}^{(Nash)})^2 D^M}{2(1-\beta(\alpha_{(Alt)}^{(Nash)})^2} - \frac{\tau\beta^{3/2}\alpha_{(Alt)}^{(Nash)} D^M}{2(1-\beta^{1/2})}$$

$$B = \frac{\frac{\tau\beta D^M}{2}\left(\frac{(1+\beta^{1/2})(\alpha_{(Alt)}^{(Nash)})^2}{1-\beta(\alpha_{(Alt)}^{(Nash)})} - \frac{\alpha_{(Alt)}^{(Nash)}}{1-\beta^{1/2}}\right)}{D^M\left(\tau + \beta^{1/2}(2 - \tau(2 - \alpha_{(Alt)}^{(Nash)}))\right) + 2(1-\beta^{1/2})(D^M - \underline{D}^M)(1-\tau)}.$$

Proof. Appendix D. □

$p^{(Nash),*}_{(Simult)}$ only collapses to $p^{*}_{(Static)}$ when $\beta = 0$, which corresponds to myopic bargaining. Myopic bargaining is the special case in which bargaining over NPV gains from trade corresponds to static-type bargaining over a single period's gains. Myopic models cannot accurately capture forward-looking responses to dynamic changes, so that a forward-looking bargaining may be needed for dynamic settings.

B.2 Dynamic Value Functions

I begin with the recursively-defined expected value function.

Formalizing the expected NPV profit in bargaining requires notation for the future equilibrium. Fix a period t_s in which ij have multiple feasible contracts. By assumption, the null contract $\mathbb{C}_{0,ijt_s}$ is also feasible. I will assume that strategies are Markov, so that future conditions can evolve based on the contracts realized today but not whether a null contract was reached through agreement or disagreement. Define $\pi_{t_r|t_s}(\mathbb{C}_{t_r} \mid \mathbb{C}_{t_s})$ to be the (potentially random) profit function that will occur in period t_r with realized agreements $\mathbb{C}_{t_s}$ this period. Similarly define $D_{t_r|t_s}(\mathbb{C}_{t_r} \mid \mathbb{C}_{t_s})$ for the demand functions. Also define $\mathcal{I}_{t_s}(\mathbb{C}_{t_s})$ to be the realized information, $\mathbb{C}_{0,k\ell t_r|t_s}(\mathbb{C}_{t_s})$ to be the realized null contracts for $k\ell$ in period t_r, $\mathcal{C}_{k\ell t_r|t_s}(\mathbb{C}_{t_s})$ for the set of feasible contracts, and $\hat{\mathbb{C}}_{t_r|t_s}(\mathbb{C}_{t_s})$ for the equilibrium contracts. Define the period t_s ij substitution function as $\left(\hat{\mathbb{C}}_{t_r,-ij|t_s}, \mathbb{C}_{t_s,ij}\right)$, which replaces the ij element of the equilibrium contract $\hat{\mathbb{C}}_{t_s}$ with $\hat{\mathbb{C}}_{t_s,-ij}$. I then extend this contract replacement to periods $t_r > t_s$ by defining $\left(\hat{\mathbb{C}}_{t_r,-ij|t_s}(\mathbb{C}_{t_s,ij}), \mathbb{C}_{t_r,ij}\right)$ to be the contract which replaces the ij element of $\hat{\mathbb{C}}_{t_r|t_s}\left(\hat{\mathbb{C}}_{t_r,-ij|t_s}(\mathbb{C}_{t_s,ij}), \mathbb{C}_{t_r,ij}\right)$ with the arbitrary $\mathbb{C}_{t_r,ij}$.

I will now define the value function of ij agreeing to a contract $\mathbb{C}_{t_s,ij}$ in stage 2 of period t_s. The value function will depend in large part on the one-period-ahead value of beginning with a modified information state. The ex ante expected value to agent i of beginning the period t_s with information $\mathcal{I}_{t_s}$ with (potentially random) realized contracts $\hat{\mathbb{C}}_{t_s}(\mathcal{I}_{t_s})$, number of new contracts $\hat{R}_i(\mathcal{I}_{t_s})$, and otherwise equilibrium response functions of that form is:

$$V^{(1)}_{i,t_s}\left(\mathcal{I}_{t_s}\right) = E_{\mathcal{I}_{t_s}}\left[\pi_{it_s}\left(\hat{\mathbb{C}}_{t_s}(\mathcal{I}_{t_s})\right) + \sum_{j\neq i}\hat{p}_{ijt_s}(\mathcal{I}_{t_s})D_{t_s,ij}\left(\hat{\mathbb{C}}_{t_s}(\mathcal{I}_{t_s})\right) - r_i\hat{R}_i(\mathcal{I}_{t_s}) + \beta^{\frac{1}{m}}V^{(1)}_{i,t_s+1}\left(\mathcal{I}_{t_s+1}\right)\right].$$

Also define $\hat{R}_{ijt_s}$ as an indicator for ij forming a new contract in t_s, and $\hat{B}_{ijt_s} \geq \hat{R}_{ijt_s}$ to be an indicator ij having the opportunity to bargain in a period t_s. Then the ex post expected value of $i < j$ choosing a contract $\mathbb{C}_{t_s,ij} \in \mathcal{C}_{ijt_s}/\mathbb{C}_{0,ijt_s}$ with associated price p_{ijt_s}, including

the negotiation cost, is:

$$V_{i|ijt_s}^{(2,A)}\left(\mathbb{C}_{t_s,ij} \mid \hat{\mathbb{C}}_{t_s,-ij|t_s}\right) = E_{t_s.2}\left[\pi_{it_s}\left(\left(\hat{\mathbb{C}}_{t_s,-ij|t_s}, \mathbb{C}_{t_s,ij}\right)\right) - r_i(\hat{R}_{it_s} - \hat{R}_{ijt_s} - r_i)\right] + E_{t_s.2}\left[\sum_{k\neq i} p_{ikt_s} D_{t_s,ij}\left(\left(\hat{\mathbb{C}}_{t_s,-ij|t_s}, \mathbb{C}_{t_s,ij}\right)\right)\right] + \beta^{\frac{1}{m}} E_{t_s.2}\left[V_{i,t_s+1}^{(1)}\left(\mathcal{I}_{t_s+1}\left(\left(\hat{\mathbb{C}}_{t_s,-ij|t_s}, \mathbb{C}_{t_s,ij}\right)\right)\right)\right].$$

The value to $j > i$ is similar, but with reindexing as needed. The value of ij disagreeing is:

$$V_{i|ijt_s}^{(2,D)}\left(\hat{\mathbb{C}}_{t_s,-ij|t_s}\right) = E_{t_s.2}\left[\pi_{it_s}\left(\left(\hat{\mathbb{C}}_{t_s,-ij|t_s}, \mathbb{C}_{0,ijt_s}\right)\right) - r_i(\hat{R}_{it_s} - \hat{R}_{ijt_s})\right] + E_{t_s.2}\left[\sum_{k\neq i,j} p_{ikt_s} D_{t_s,ij}\left(\left(\hat{\mathbb{C}}_{t_s,-ij|t_s}, \mathbb{C}_{0,ijt_s}\right)\right)\right] + \beta^{\frac{1}{m}} E_{t_s.2}\left[V_{i,t_s+1}^{(1)}\left(\mathcal{I}_{t_s+1}\left(\left(\hat{\mathbb{C}}_{t_s,-ij|t_s}, \mathbb{C}_{0,ijt_s}\right)\right)\right)\right].$$

Given these ex post value functions, I recursively define the ex ante value function for i in terms of the ij contract outcome. The three possibilities are renewing at equilibrium contract $\hat{\mathbb{C}}_{t_s,-ij|t_s}$ ($\hat{B}_{ijt_s} = 0$), successfully negotiating the (potentially random) equilibrium contract $\hat{\mathbb{C}}_{t_s,-ij|t_s}$ ($\hat{B}_{ijt_s} = \hat{R}_{ijt_s} = 1$), and disagreeing to the null contract $\hat{\mathbb{C}}_{0,ijt_s}$ ($\hat{B}_{ijt_s} = 1, \hat{R}_{ijt_s} = 0$). As a result, the ex ante expected NPV profit can be written as:

$$V_{i,t_s}^{(1)}\left(\mathcal{I}_{t_s}\right) = E_{\mathcal{I}_{t_s}}\left[(1 - \hat{B}_{ijt_s}(1 - \hat{R}_{ijt_s}))V_{i|ijt_s}^{(2,A)}\left(\hat{\mathbb{C}}_{t_s,ij} \mid \hat{\mathbb{C}}_{t_s,-ij|t_s}\right) + (1 - \hat{B}_{ijt_s})r_i\right] + E_{\mathcal{I}_{t_s}}\left[\hat{B}_{ijt_s}(1 - \hat{R}_{ijt_s})(1 - \hat{R}_{ijt_s})V_{i|ijt_s}^{(2,D)}\left(\hat{\mathbb{C}}_{t_s,-ij|t_s}\right)\right].$$

B.3 Additional Assumptions

Definition 3 (Formal split-in-half subgame)**.** *Let the original periods be t_{s+k}, with the subgame corresponding to some fixed s. Let the new periods be $\bar{t}_{s+k}$. Then for all $k > 0$, the new times are $\bar{t}_{s+k} = t_s + (t_{s+k} - t_s)/2$. The new per-period discounting rate is $\bar{\beta} = \sqrt{\beta}$. If $k \geq 0$ is even, then flow profit and demand functions are the same in $\bar{t}_{s+k}$ and $\bar{t}_{s+k+1}$. If k is even and the history in stage 1 of $\bar{t}_{s+k}$ is (after appropriate change of notation including changing one-period disagreement to two-period disagreement) in the support of histories in original-game period $t_{s+k/2}$, then (1) the conditional distribution of $\mathcal{I}_t$ is (after appropriate the change of notation) the same in the new $\bar{t}_{s+k}$ and the original $t_{s+k/2}$; (2) the joint conditional-on-$\mathcal{I}_t$ distributions of flow profits and demand functions in $\bar{t}_{s+k+2\ell}$ for all integer $\ell \geq 0$ is (after appropriate change of notation) the conditional distribution of (flow profits) / $(1+\sqrt{\beta})$ and*

(demand) / $(1+\sqrt{\beta})$ in the original $t_{s+k/2}$; and (3) the associated set of feasible contracts in stage 2 for all pairs are unchanged. If k is odd, then $\mathcal{I}_{\bar{t}_{s+k}}$ includes no new information on future profit functions than was available from the end of stage 2 of $\bar{t}_{s+k-1}$.

Assumption 2 (Full networks)**.** Consider a dynamic bargaining game and an associated X bargaining model with a first subgame. I write that the model has full networks if:

1. *Subgame Nash-in-X solutions exist.* For every subgame and any set of restrictions on subsequent strategies, there is a set of strategies such that for every unconstrained response, the response is bilaterally Pareto optimal and every nonconstrained bargain produces the X bargaining solution to the recursive bilateral bargaining problem.

2. *Full network equilibrium.* For every constrained solution of that form and every history that can be reached through equilibrium play or bilateral defections to disagreement, the next subgame equilibrium includes agreement in all unconstrained negotiations.

The first restriction is that the process of splitting the game in half does not make it impossible to form Nash-in-X solutions. The second restriction is that it is impossible to reach a subgame in which the equilibrium prediction is to disagree. This full-networks assumption is a device adapted from Collard-Wexler et al. (2019) to rule out added opportunities for punishment. For example, consider a forward-looking model in which ij prefer to wait a full period before reaching an agreement. It is possible that in a split-in-half model, they will prefer to reach an agreement in the first half-period; other firms, anticipating the faster ij agreement, may change their behavior and yield a split-in-half effect. The full-network assumption may require restrictions on the parameter space. For example, in a subgame in which a retailer agreed to pay a negative price to one manufacturer rather than a usual positive price, the retailer may prefer to not add an additional manufacturer in order to maximize their off-equilibrium payment received.

The good-faith disagreement assumption is as follows.

Assumption 3 (Good-faith disagreement)**.** Let $G^{(3)}_{t:T|t-1}$ be the distribution function of equilibrium networks $\hat{\mathcal{G}}_t, \ldots \hat{\mathcal{G}}_T$ in stage 2 of period t as a function of the history h^{t-1} reached at the end of stage 3 of period $t-1$. Consider some subgame period t_0 in stage 2. Let h^{t_0} be a a draw from the support of the t_0 subgame equilibrium of stage 2. Let $\hat{h}^{t_0}$ be the history formed by replacing the ij contract with the null contract. Then for all finite T, $G^{(3)}_{t:T|t-1}(h^{t_0}) = G^{(3)}_{t:T|t-1}(\hat{h}^{t_0})$.

Assumption 3 essentially says that bargainers attempt to exit impasse as soon as possible: if ij bargaining fails in period t_0, then i returns to j's network as soon as they would

have been in-network in equilibrium. Future networks are assumed to be unaffected by ij disagreement because everyone continues to expect ij to reach a contract.[9] Good-faith disagreement is in line with the empirical rarity of disagreement, which suggests disagreement is painful and consistently avoided. The good-faith assumption is also in line with the Nash-in-Nash bargaining model's Nash equilibrium assumption: under single-period Nash-in-Nash bargaining, contracts are formed assuming other bargains will succeed. I implicitly rule out any effect of current networks on subsequent future demand. If Assumption 3 did not hold and disagreement affected subsequent networks, as in Lee and Fong (2013)'s model, then the form of payments would be modified to incorporate the effect of impasse on subsequent network formation.

B.4 Uncertainty Can Introduce Bias under Nash Bargaining

In this subsection, I use a toy model to show that an empirical researcher must accurately model uncertainty under Nash bargaining even if uncertainty is completely at random.

Imagine a researcher has access to an infinite number of samples from the following bargaining process. A fair coin is flipped. If the coin comes up heads, the insurer's realized gains relative to the status quo of 0 profits are $\widehat{GFT}^M(p) = 30 - p$ and the hospital's realized gains are $\widehat{GFT}^H(p)$. If the coin comes up tails, the insurer's realized gains are $\widehat{GFT}^M(p) = 10 - p$ and the insurer's realized gains are $\widehat{GFT}^H(p) = 2p$. (To emphasize the role of uncertainty, in this section, I write the realized gains from trade as $\widehat{GFT}$ and the gains from trade as $E[\widehat{GFT}]$ rather than GFT.)

After taking first-order conditions, the Nash bargaining solution over price p in terms of gains from trade $E[\widehat{GFT}(p)]$ can be written as

$$0 = (1-\tau)E[\widehat{GFT}^M(p^*)]\frac{\partial E[\widehat{GFT}^H(p)]}{\partial p}\mid_{p=p^*} + \tau E[\widehat{GFT}^H(p^*)]\frac{\partial E[\widehat{GFT}^M(p)]}{\partial p}\mid_{p=p^*}.$$

The game is transferable utility in expectation, so that both the symmetric Nash and Kalai proportional bargaining solutions yield $p^* = 10$.

Note that if price is viewed as a treatment, uncertainty in this process is strongly ignorable. The coin toss outcome is completely independent of the price that is chosen. There are

[9] The assumption only refers to one disagreement, but Assumption 3 rules out continued impasse affecting subsequent networks by inductively applying the following argument to construct $t_0 + 2$ networks under impasse. First draw $t_0 + 1$ networks from $G^{t_0(3),t+1:\infty}(\hat{h}^{t_0(3)})$, which by Assumption 3 is the same as drawing from $G^{t_0(3),t+1:\infty}(h^{t_0(3)})$; then draw $h^{t_0+1(3)}$ from the distribution of t_0+1 histories conditional on $\hat{h}^{t_0(3)}$ and $t_0 + 1$ networks; and then substitute the null contract for $ijt_0 + 1$. This process generates the distribution of $(\mathcal{G}_{t_0+1}, \mathcal{G}_{t_0+2})$ under impasse. By Assumption 3 applied to both draws, this process generates the same distribution as drawing $(\mathcal{G}_{t_0+1}, \mathcal{G}_{t_0+2})$ from $G^{t_0(3),t+1:\infty}(h^{t_0(3)})$.

no unobservables that the negotiators use to choose contracts.

Suppose a naive researcher presumes they can form a moment based on the realized gains. In particular, they exactly observe the realized $\widehat{GFT}$ functions. The naive researcher unfortunately constructs a Nash bargaining moment

$$0 = E\left[Z\left((1-\hat{\tau}^{Nash})\widehat{GFT}^{M}(p)\frac{\partial\widehat{GFT}^{H}(p)}{\partial p}\mid_{p=p^*} + \hat{\tau}^{Nash}\widehat{GFT}^{H}(p)\frac{\partial\widehat{GFT}^{M}(p)}{\partial p}\mid_{p=p^*}\right)\right],$$

where the instrument Z is a dummy. When the coin comes out tails, $\widehat{GFT}^{M}(p^*) = 0$. When the coin comes out heads, $\frac{\partial\widehat{GFT}^{H}(p)}{\partial p}\mid_{p=p^*} = 0$. As a result, the researcher correctly observes the insurer only gains when the hospital is indifferent to payments, and incorrectly concludes that $\hat{\tau}^{Nash} = 0$. Conversely, the Kalai proportional bargaining moment will set

$$0 = E\left[Z\left((1-\hat{\tau}^{Kalai})\widehat{GFT}^{M}(p^*) - \hat{\tau}^{Kalai}\widehat{GFT}^{H}(p^*)\right)\right] = E\left[(1-\hat{\tau}^{Kalai})(20-p^*) - \hat{\tau}^{Kalai}p^*\right],$$

which is satisfied at the correct bargaining weight $\hat{\tau}^{Kalai} = 1/2$.

At a high level, the methodological issue emerges when there can be a correlation between one side's gains and the other side's marginal value based on uncertainty resolved after the negotiation. The Nash bargaining process that just cares about expected values at the moment of negotiation. If the correlation is incurred from uncertainty after negotiation, then a plug-in moment will be biased. If the correlation is observed and incorporated by the negotiators before bargaining, than a plug-in moment will be unbiased.

C Additional Empirical Content

C.1 Estimation Overview, Continued

I infer contract benchmarks based on the contract data. I infer that a payor with the same reported discount of list prices (or a difference of 0.01% after rounding) in consecutive hospital reports was a share of charges (list-price-benchmarked) contract that paid as a fixed discount of list prices. For the first observation of a hospital–payor pair, I infer that a round-number discount followed by a change of payor or a new share of charges contract was the final year of an expiring share of charges contract. I infer all other contracts were prospective and used Medicare as the benchmark.

I estimate a finite horizon model: I consider only the first T years of a contract in calculating gains from trade, where $T = 5$ for my current analysis. I cannot estimate a compelling

infinite horizon model because West Virginia is nonstationary. I take $T = 5$ as an approximation: as the length of available data goes to infinity, T would go to infinity slowly to enable the number of bargains in estimation to go to infinity as well. The finite horizon makes a constant patience parameter β at best an approximation because the fifth year in should principle best approximate later years. The finite horizon model also calls for care in modeling how impasse affects other bargainers near the end of the horizon. I calculate gains from trade through 2016. Contracts in 2016 are calculated by linearly extrapolating list price levels from calendar year 2015 contract reports and extrapolating contract list price shares from 2015.

I estimate hospital demand with maximum likelihood. I adapt the notation of Ho (2006), but similar models are widely used (Capps et al., 2003, Gowrisankaran et al., 2015, Ho and Lee, 2017, Prager and Tilipman, 2022). I assume that to a patient, the utility of a potential hospital is a function of the patient's diagnosis, the hospital's quality, and the patient's location. In particular, I assume the utility of consumer i visiting in-network hospital h with diagnosis ℓ (cancer, cardiac, digestive, labor, neurological, or other) is:

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

where $\delta^H_{h,\ell}$ is a hospital-diagnosis fixed effect, $\nu_{i,h,\ell}$ are patient-hospital characteristics (distance in miles, distance squared, and distance interacted with emergency), and ε is a type 1 extreme value shock. I estimate the model with Blue Cross patients in 2016, as all hospitals are in-network for Blue Cross.

I estimate insurer demand mainly using cross-sectional data from 2016, the year in which I have estimates of local sales. The equation I estimate in 2016 is:

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

where $u_{i,j,c,m}$ is the utility of individual i choosing insurer i in county c within market m, $\tilde{\delta}^M_{j,k,m}$ is an insurer–rating-area fixed effect that includes premium levels, $WTP_{j,k,c}$ (Capps et al., 2003) is the ex ante expected utility of insurer j's network to an individual of age-group k in county c, γ_k are age-group-dependent coefficients on WTP, $\xi_{j,k,c}$ is an age–county unobservable, and ε is a type 1 extreme value shock. Similar models have been used by Ho and Lee (2017) and Ghili (2022). The equation is estimated using the moment $E[WTP_{j,k,c}\xi_{j,k,c}] = 0$, matching observed county-age shares for insurers identified in the inpatient data, and matching state-level sales for all modeled insurers. I assume that the insurer-rating area fixed effects $\tilde{\delta}^M_{j,m}$ are constant across markets for the two insurers that are not identified in the inpatient data, Cigna and UnitedHealth.

The utility equation does not include premiums. Since 2014, the Affordable Care Act (ACA) restricts insurer premium setting substantially (CMS, 2023). Insurers set premiums (outside the large-group market) by geographic rating area defined by the state of West Virginia.[10] Premium variation in 2016 is essentially subsumed into the $\tilde{\delta}^M_{j,m}$ insurer-rating-area fixed effects.

My counterfactual analysis must account for changes in insurer attractiveness and premiums over time. For years before 2016, I include an insurer-time fixed effect $\tilde{\delta}^M_{j,t}$. The fixed effect $\tilde{\delta}^M_{j,t}$ captures systematic changes in insurer value and premiums in previous years. I solve for the values to match state-level sales by year after adjusting for changes in networks and local population. I discuss other estimation details in Appendix C.7.

The bargaining model is estimated based on observed and predicted payments for observed bargains. I use the estimated hospital and insurer demand systems as inputs to gains from trade in bargaining, with other bargaining parameters like the bargaining weights τ_{ij} and discounting rate β estimating through GMM. I define $\omega_{ijt}(\bar{\theta})$ to be the normalized NPV residual payment from ij negotiating in period t at parameters $\bar{\theta}$:

$$\omega^p_{ijt}(\bar{\theta}) = \frac{\sum_{t=t_0}^{t^*} \bar{\beta}^{t-t_0} \left(D^H_{ijt} p^H_{ijt} - \left\{ -\bar{\tau}_{ij}[\Delta_{ij}\bar{\pi}^H_{it}] + (1-\bar{\tau}_{ij})[\Delta_{ij}\bar{\pi}^M_{jt}] \right\} \right) - (1-\bar{\tau}_{ij})\bar{r}^M_j}{\sum_{t=t_0}^{\lfloor mean(t-t_0) \rfloor} \bar{\beta}^{t-t_0} + (mean(t-t_0) - \lfloor mean(t-t_0) \rfloor)\, \bar{\beta}^{\lceil mean(t-t_0) \rceil}}, \tag{9}$$

where a bar denotes a parameter that is estimating in my bargaining model and $mean(t-t_0)$ is the average bargain's number of years elapsed. The denominator is added to express ω_{ijt} in terms of the ij NPV payment and an aggregate normalization to avoid attenuating the estimated patience parameter β. (As seen in Appendix Table 9, I would estimate a similar $\hat{\beta}$ if I instead normalized by the average value of $\sum \beta^t$ across bargains used in model estimation.) My main specification calibrates hospital costs from reported hospital average costs, which should roughly track the outside option of Medicare rates if hospitals are near capacity, and adjust hospital costs in robustness tests. The parameters to estimate are the τ_j insurer bargaining weights, τ^{Size} contribution of size to hospital bargaining weight, β patience parameter, η_j insurer noninpatient costs, and payment-equivalent negotiation costs r_j.

Bargaining moments are constructed as follows. I take the NPV payment residual $\omega^p_{ijt}(\bar{\theta})$ from Equation (9). I define $\omega^M_{jt}(\bar{\theta}) = \sum_{t=2011}^{2016} \frac{\bar{\eta}_j D^M_{jt} + \sum_h D^H_{hjt} p_{hjt}}{\phi_{jt} D^M_{jt}} - MLR_{j,t}$ as the difference between model-implied medical loss ratio and the medical loss ratios $MLR_{j,t}$ reported to CMS for years 2011 and later. My moments are $E[Z^p \omega^p] = 0$ and $E[Z^M \omega^M] = 0$. The

[10] Insurers also have a limited ability to adjust premiums based on tobacco use, family size, and age. In practice, insurers applied at most a small adjustment for tobacco use, a similar adjustment for large families, and the same age multipliers. I hold these multiples fixed in my analysis.

hospital–insurer payment instruments Z^p are insurer dummies and indicators for hospital size in six groups. The insurer medical loss ratio instruments Z^M are insurer dummies.

Similar strategies have been used with various datasets in static models, though there are important differences. Some notable papers with similar identification strategies are Grennan (2013), Gowrisankaran et al. (2015), Ho and Lee (2017, 2019), Ghili (2022), Liebman (2022) and Prager and Tilipman (2022). Gowrisankaran et al. (2015) and Prager and Tilipman (2022) assume insurers maximize a criterion other than profits. Grennan (2013) and Ghili (2022) have a non-zero-sum downstream response to negotiated prices, which could be partially captured in levels by my flexible hospital–insurer bargaining weight specification (Equation (8)). However, such time-varying NTU bargaining cannot fit into my Kalai proportional bargaining model for reasons discussed in Section 3.2. Ho and Lee (2019) and Ghili (2022) consider network formation in response to disagreement, which is at odds with my good-faith disagreement and simultaneous bargaining Assumptions 1 and 3. Many of these works estimate premium responsiveness (Gowrisankaran et al., 2015, Ho and Lee, 2017, 2019, Liebman, 2022, Ghili, 2022). I instead use market premium regulations to estimate demand given observed premiums and focus on price counterfactual effects that can be conservatively bounded without an estimate of premium responsiveness.

I now discuss my construction of counterfactual prices. The choice of benchmark is highly associated with bargaining power and likely to be at most moderately affected by the change in benchmark dynamics. Most fixed-length contracts were formed by Highmark BCBS, so holding lengths constant mostly corresponds to assuming that Highmark BCBS would not change which contract terms were formed on a three-year or five-year basis and that the other insurers would not change their share of charges renewal strategies. Premiums would increase in response to higher bargained prices. I interpret the ω^p NPV payment residuals as counterfactual-invariant unobserved components of gains from trade, and choose starting payments to leave ω^p unchanged. This approach may introduce slight bias if benchmark price uncertainty is first-order relative to the change to benchmark price dynamics. One potential avenue for future work is to apply tools from the time series literature to develop a more precise approach to counterfactuals.

C.2 Identification Intuition

Hospital demand is identified by selection on observables. If consumers are highly likely to choose Charleston Area Medical Center (CAMC) relative to Saint Francis Hospital one mile away, my estimates will infer that CAMC offers more utility to consumers after adjusting for location. The degree to which patients with similar diagnoses choose closer hospitals

identifies the ρ distance coefficients. There are three key assumptions for the hospital demand model to be accurate for bargaining estimation. First, Blue Cross hospital choice should be representative of the generic consumer hospital choice decision conditional on location and diagnosis (i.e., no endogeneity of insurance choice with respect to hospital value). Second, observed choices should identify counterfactual choice probabilities with different hospital choice sets (i.e., unconfoundedness and correct functional form). Third, observed hospital choice should capture the value of hypothetical hospital networks when choosing an insurer.

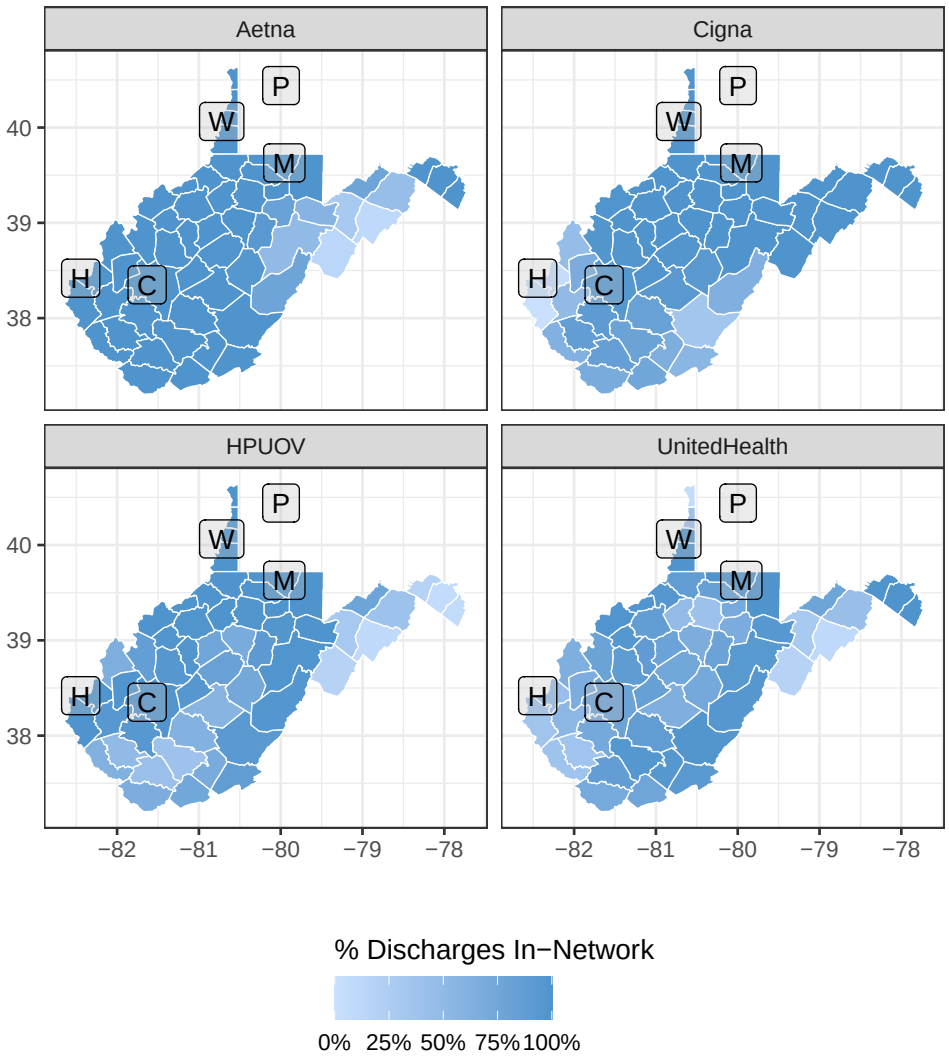


Figure 9: Percent of 2016 inpatient discharges by county of residence that are in the 2015 reported network of (clockwise from top-left) Aetna, Cigna, UnitedHealth, and the Health Plan of the Upper Ohio Valley. Highmark BCBS (omitted) is in-network in all West Virginia hospital reports. The large cities of Charleston, Huntington, Morgantown, Wheeling, and Pittsburgh, PA are indicated by letter labels.

Insurer demand is identified based on variation in network quality conditional on premiums. Insurer regional coverage was heterogeneous within West Virginia's 11 market rating areas (see Appendix Figure 9). The γ_k coefficients are identified by the degree to which consumers are more likely to choose an insurer with better coverage within a rating area that standardizes premiums. The key identification assumption is exogeneity: the market-level unobservables should be uncorrelated with network quality itself. I do not model any potential variation of large-group employer premiums within ACA rating areas. My main analysis holds premiums fixed and calculates back-of-the-envelope downstream effects of counterfactual prices on Nash-Bertrand optimal premiums based matching the average premium elasticity from Ho (2006), which corresponds to a coefficient on premiums of -0.00032. As a result, my analysis does not require estimating premium responsiveness. I discuss these and some other caveats in demand estimation in Appendix C.6.

Identification of bargaining parameters comes from various sources. The η_j insurer noninpatient costs are identified primarily from the CMS medical loss ratio reports but are shifted by the GMM procedure based on observed payments. The flow gains from trade $[\Delta_{ij}\pi]$ are identified from estimated demand, η_j noninpatient costs, and calibrated hospital costs c. The τ_{ij} bargaining weights are identified by the ratio of realized gains from trade at the negotiated price and varying the hospital or insurer. The β bargaining weights are identified by future gains from trade and payments helping explain the realized price conditional on current gains. The r^M negotiation costs are identified from any remaining differences in the levels of payments and the levels of predicted payments. I discuss potential biases in Appendix Appendix C.6.

C.3 Other Empirical Estimates and Results

I present diagnosis category frequencies in Table 4. Labor and cardiac discharges are the largest share of actively defined categories. Medicare patients are more likely to have cardiac discharges and less likely to have labor discharges than the commercially insured sample I use in estimation.

I present estimated hospital demand distance parameters in Table 5. Regardless of diagnosis, consumers prefer closer hospitals and have a diminishing loss from distance. Consumers with neurological conditions are relatively insensitive to distance. Patients travel less far for emergency care outside labor cases. In the 201 hospital-diagnosis fixed effects I estimate but omit for space, consumers place the highest value on Ruby Memorial, the West Virginia University (WVU) Health system's flagship hospital. Hospitals near the state's border like Cabell Huntington and Mon Health Medical Center are generally higher value than their

Table 4: The percent of discharges by diagnosis category for all inpatient discharges (top row), the commercial nonnewborn discharges I use for estimation (middle row), and the Highmark BCBS subset of commercial discharges I use for hospital demand estimation (bottom row).

Discharges	Labor	Cardiac	Digestive	Neurological	Cancer	Other
All	16.86	15.42	8.29	7.47	1.04	50.16
Commercial WV	19.85	12.41	9.97	5.60	1.33	50.11
Highmark BCBS	20.22	12.49	9.84	5.52	1.49	49.72

	Dependent variable:					
	choice					
	Cancer	Cardiac	Digestive	Labor	Neurological	Other
	(1)	(2)	(3)	(4)	(5)	(6)
Distance	-0.115^{***}	-0.113^{***}	-0.117^{***}	-0.121^{***}	-0.077^{***}	-0.108^{***}
	(0.014)	(0.004)	(0.005)	(0.004)	(0.003)	(0.002)
Distance Squared	0.0004^{***}	0.0004^{***}	0.0004^{***}	0.0003^{***}	0.0002^{***}	0.0003^{***}
	(0.0001)	(0.00002)	(0.00002)	(0.0001)	(0.00002)	(0.00001)
Distance x Emergency	-0.010	-0.012^{***}	-0.024^{***}	0.020^{***}	-0.013^{***}	-0.015^{***}
	(0.015)	(0.003)	(0.004)	(0.005)	(0.004)	(0.001)
Observations	284	2,469	2,048	4,143	1,094	10,053
R^2	0.555	0.577	0.615	0.646	0.497	0.555
Log Likelihood	−286.987	−2,722.077	−2,324.572	−3,923.918	−1,297.677	−12,578.030
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 5: Estimated consumer valuation of distance in hospital choice (in utility units) by diagnosis category. Consumers generally are admitted to closer hospitals, have a diminishing loss from travel, and — with the exception of labor cases — are especially unlikely to travel distances for emergency care.

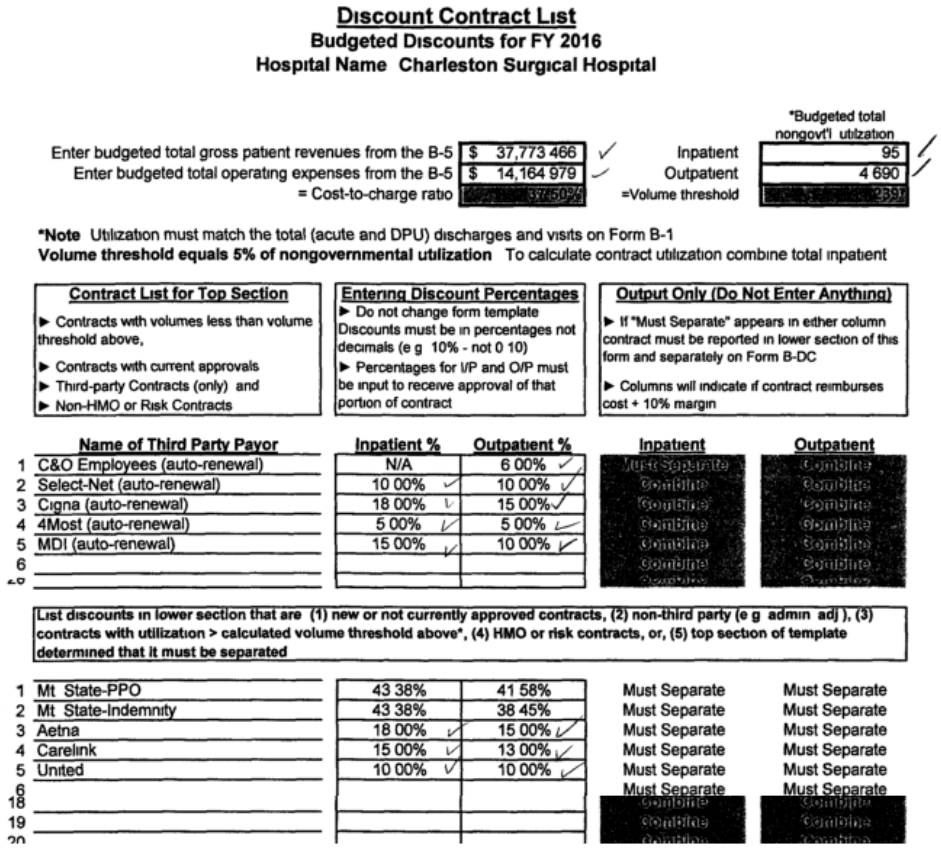
state-level share would suggest, consistent with border hospitals also competing for patients from neighboring states. Hospital fixed effects are comparable for most diagnoses but are smaller for labor discharges.

The insurer demand estimates are presented in Table 6 and Table 7. Table 6 includes insurer mean $\tilde{\delta}^{M}_{j,m}$ fixed effects inclusive of premiums. Consumers are more likely to choose Highmark BCBS than can be explained by the insurer's inpatient network alone, which may reflect a better outpatient network, better perceived quality, inertia, or Highmark BCBS's nonprofit status. The regional HPUOV is also nonprofit and has a larger fixed effect than the three national for-profit insurers, which in part reflects lower premiums. Appendix Table 7 presents the estimated WTP coefficients. Consumers are more likely to purchase insurance with a better network. The network value coefficients differ substantially by age in absolute terms. The scale of network valuation reflects age differences: younger consumers are less likely to get sick and so have a smaller variation in WTP across networks.

Insurer and hospital payment effects are in Figures 16 and 17. In 2009, Highmark BCBS payments would change by -0.42% due to the forward-looking response for Highmark BCBS's many new prospective contracts. Other insurers often see reduced payments in years where they negotiated new prospective contracts with the WVU Health System. Turning to hospital payments, in 2015, most hospitals would have spending increases of 1% to 3%. A fair number of medium-sized hospitals would be almost entirely unaffected.

Figure 18 presents the full path of estimated downstream premium effects. I estimate that in 2015, Highmark BCBS premiums would increase by 0.2%, and HPUOV premiums would increase by 0.51%. The other modeled insurers are less likely to use Medicare-based benchmarks and would see premiums increase by 0.1%. Altogether, premiums would increase by 0.21%, which extrapolates nationally based on inflation-adjusted earned premiums (NAIC, 2022) to $1.33 billion. The extrapolated premium effects are smaller than the extrapolated payment effects because this premium extrapolation holds outpatient payments constant.

Figure 19 presents estimated effects under a mechanical model that accounts for list price contracts not being directly affected, as well as the myopic and forward-looking estimates from the main text. The mechanical model would predict a nine-year effect that corresponds to nine The myopic model, which accounts for the first two responses and equilibrium spillovers, reduces this annualized increase rate by 0.35 percentage point. The forward-looking model, which incorporates all three responses, further reduces the annualized increase by a further 0.068 percentage point and, as seen in Figure 8, yields dynamics that the myopic model misses.



Discount Contract List
Budgeted Discounts for FY 2016
Hospital Name Charleston Surgical Hospital

			*Budgeted total nongovt'l utilization
Enter budgeted total gross patient revenues from the B-5	$ 37,773 466	Inpatient	95
Enter budgeted total operating expenses from the B-5	$ 14,164 979	Outpatient	4 690
= Cost-to-charge ratio	37 50%	=Volume threshold	239

***Note** Utilization must match the total (acute and DPU) discharges and visits on Form B-1
Volume threshold equals 5% of nongovernmental utilization To calculate contract utilization combine total inpatient

Contract List for Top Section
- Contracts with volumes less than volume threshold above,
- Contracts with current approvals
- Third-party Contracts (only) and
- Non-HMO or Risk Contracts

Entering Discount Percentages
- Do not change form template
Discounts must be in percentages not decimals (e g 10% - not 0 10)
- Percentages for I/P and O/P must be input to receive approval of that portion of contract

Output Only (Do Not Enter Anything)
- If "Must Separate" appears in either column contract must be reported in lower section of this form and separately on Form B-DC
- Columns will indicate if contract reimburses cost + 10% margin

	Name of Third Party Payor	Inpatient %	Outpatient %	Inpatient	Outpatient
1	C&O Employees (auto-renewal)	N/A	6 00%	Must Separate	Combine
2	Select-Net (auto-renewal)	10 00%	10 00%	Combine	Combine
3	Cigna (auto-renewal)	18 00%	15 00%	Combine	Combine
4	4Most (auto-renewal)	5 00%	5 00%	Combine	Combine
5	MDI (auto-renewal)	15 00%	10 00%	Combine	Combine
6				Combine	Combine
18				Combine	Combine

List discounts in lower section that are (1) new or not currently approved contracts, (2) non-third party (e g admin adj), (3) contracts with utilization > calculated volume threshold above*, (4) HMO or risk contracts, or, (5) top section of template determined that it must be separated

1	Mt State-PPO	43 38%	41 58%	Must Separate	Must Separate
2	Mt State-Indemnity	43 38%	38 45%	Must Separate	Must Separate
3	Aetna	18 00%	15 00%	Must Separate	Must Separate
4	Carelink	15 00%	13 00%	Must Separate	Must Separate
5	United	10 00%	10 00%	Must Separate	Must Separate
6				Must Separate	Must Separate
18				Combine	Combine
19				Combine	Combine
20				Combine	Combine

Figure 10: Discount Contract List (DCL) scan, with white space and handwritten notes omitted, for Charleston Surgical Hospital in fiscal year 2016 The top panel of contracts lists smaller contracts that are only included in the DCLs, while the bottom panel lists contracts with detailed information reported in the DC scans (Figure 11).

Summary Information of Discount Contracts
Total - FY 2016 Budget
Hospital Name ___ Charleston Surgical Hospital

Fiscal Year Ending 12/31/2016
Submission Date 11/2/2015

1	Name of Purchaser or Third Party Payor	Total	=	Combined Contracts	+	Mt State-PPO	Mt State-Indemnity	Aetna
2	Date of Contract					8/1/2015	8/1/2015	11/1/1994
3	Date Contract Expires					12/31/2018	12/31/2018	Auto Renewal
4	Projected Inpatient Discharges	92		16		50	2	17
5	Projected Gross Inpatient Revenue	2 878 926		417 403		1 627 278	65 091	567 464
6	Inpatient Discount Percent	31 54%		12 00%		43 38%	43 38%	18 00%
7	Projected Amount of Inpatient Discount	908 049		50 088		705 865	28 235	102 144
8	Projected Net Inpatient Revenue	1 970 877		367 315		921 413	36 857	465 320
9	Projected Inpatient Cost	1 079 592		156 525		610 226	24 409	212 798
10	Projected Inpatient Charge per Discharge					32 545 55	32 545 55	33 380 24
11	Projected Inpatient Cost per Discharge					12 204 52	12 204 52	12 517 53
12	**Projected Cost to Charge Ratio**	37 50%				37 50%	37 50%	37 50%
13	Projected Outpatient Visits	3 985		619		2 594	136	388
14	Projected Gross Outpatient Revenue	12 312 629		1 276 012		8 755 454	459 037	1 162 818
15	Outpatient Discount Percent	34 27%		12 00%		41 58%	38 45%	15 00%
16	Projected Amount of Outpatient Discount	4 219 778		153 121		3 640 518	176 500	174 423
17	Projected Net Outpatient Revenue	8 092 852		1 122 890		5 114 936	282 537	988 395
18	Projected Outpatient Cost	4 617 213		478 502		3 283 279	172 138	436 055
19	Projected Outpatient Charge Per Visit					3 375 27	3 375 27	2 996 95
20	Projected Outpatient Cost Per Visit					1 265 72	1 265 72	1 123 85
21	**Projected Cost to Charge Ratio**	37 50%				37 50%	37 50%	37 50%
22	Uncompensated Care Percent of Gross Patient Revenue							
23	Will Contract(s) Provide a Quantifiable Economic Benefit to the Hospital? *Circle*			Yes		Yes	Yes	Yes
24	Is the Discount Amount Below Actual Cost of Service? *Circle*			No		No	No	No
25	Will Cost Be Shifted to Any Other Purchaser of Third Party Payor as a Result of this Contract? *Circle*			No		No	No	No
26	Date contract submitted to HCA					7/8/2015	7/8/2015	10/31/2014
27	the Authority? (If yes please submit revised contracts) *Circle*			No		No	No	No

NOTE This page should include only the total combined and **3 (three)** separate contract columns Use this form in its current version only **Any modifications will be returned**

Figure 11: The first page of detailed contract (DC) data for Charleston Surgical Hospital in fiscal year 2016. The data includes unusual information on contract formation and scale.

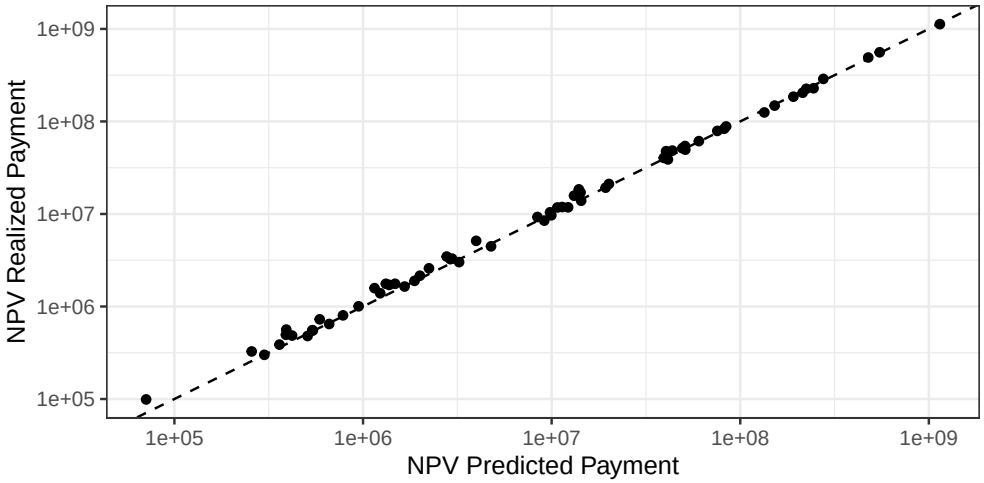


Figure 12: For the 63 estimation bargains, the predicted (x axis) and realized (y axis) NPV payment within the finite horizon used in estimation. Net present values are calculated using the estimated $\beta = 0.899$. Axes are log-scaled for comparability. Perfect prediction fit is indicated with the dashed line. There may be some bias for small contracts (generally non-Highmark-BCBS contracts at small-to-medium hospitals) that get little weight in the estimation procedure, but otherwise the model fit seems to be quite good.

	MCO:				
	Aetna	Highmark BCBS	HPUOV	Cigna	UnitedHealth
	-1.39***	1.33***	-0.8***	-3.54***	-2.43***
	(0.13)	(0.13)	(0.13)	(0.11)	(0.11)
Note:					*p<0.1; **p<0.05; ***p<0.01

Table 6: Estimated average 2016 insurer value including premiums ($\tilde{\delta}^{M}_{j,m}$) after accounting for variation in inpatient network quality.

Table 7: Insurer demand coefficient on network willingness to pay by age group. Consumers are generally more likely to purchase insurance from insurers with better networks. The coefficients are largest for young groups with smaller standard deviations in network quality.

	WTP Coefficient				
	γ_{0-17}	γ_{18-44}	γ_{45-64}	γ_{65-74}	γ_{75+}
	26.6***	4.94***	2.76***	2.79***	2.05***
	(2.65)	(0.67)	(0.33)	(0.27)	(0.15)
Note:					*p<0.1; **p<0.05; ***p<0.01

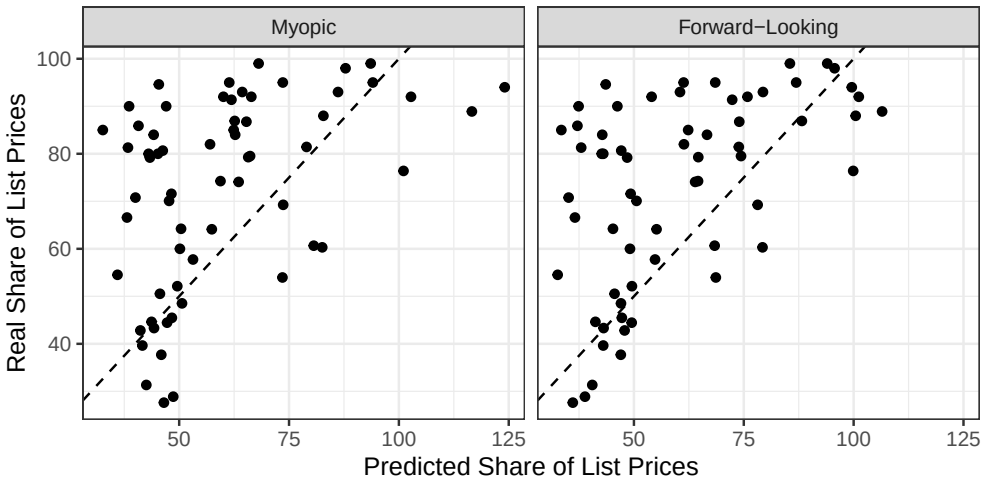


Figure 13: Predicted (x axis) and realized (y axis) bargain starting share of list prices under myopic (left, $R^2 = -0.116$) and forward-looking (right, $R^2 = 0.027$) bargaining models. R^2 can be negative because model predictions are chosen to minimize NPV payment residuals, while this R^2 reflects list price share residuals.

Table 8: Additional estimated bargaining parameters with estimated τ^{Size} hospital size coefficient (top) and with τ^{Size} set to zero (bottom) in calculating hospital–insurer bargaining weights. 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 estimated β would be similar if η were constrained to exactly fit MLR reports (Appendix Table 9). 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}
Only-2015	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
	Parameter ($\tau^{Size} = 0$)							
	η_{BCBS}	η_{HPUOV}	η_{Aetna}	$\eta_{UnitedHealth}$	η_{Cigna}	$\eta_{Carelink}$	r^M_{yBCBS}	r^M_{nBCBS}
Only-2015	3639***	3412***	3660***	2010***	4622***	3139***	10001***	23581***
(Nash/Kalai)	(14)	(28)	(37)	(27)	(26)	(36)	(1143)	(1415)
Myopic	4639***	3412***	3659***	2008***	4624***	6176***	17779***	10000***
(Nash/Kalai)	(14)	(349)	(37)	(28)	(26)	(493)	(904)	(0)
Forward-Looking	4638***	3413***	3659***	2008***	4624***	5972***	10000***	18098***
($\text{Pay}_{IRT} = 0$)	(1546)	(295)	(487)	(447)	(390)	(521)	(761)	(699)
Data	3600	3356	3554	1999	4635	3114		
Note:								*p<0.1; **p<0.05; ***p<0.01

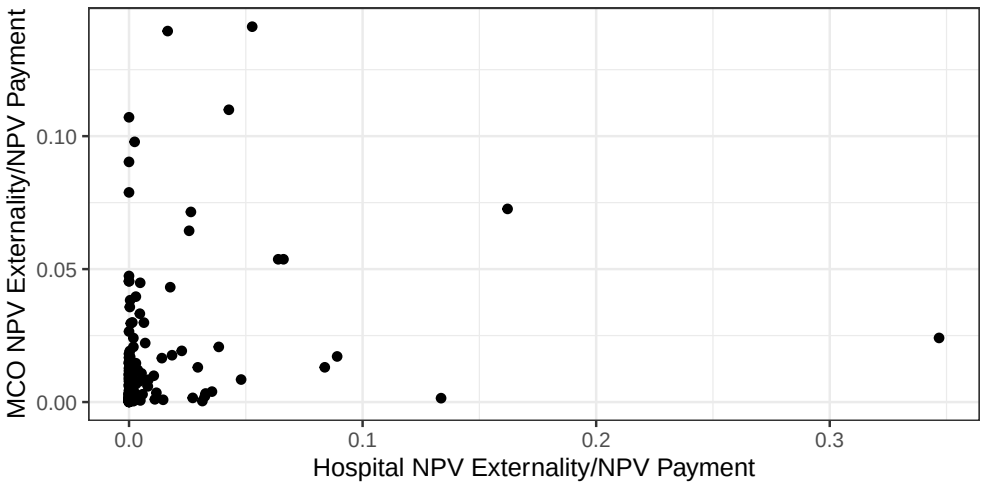


Figure 14: The estimated direct spillovers of each bargained contract on hospital (x axis) and insurer (y axis) payments relative to direct payments under the contract. All estimated insurer spillovers and all but three of the 63 estimated hospital spillovers are less than 10% of payments under the contract, so that the marginal ratio term in Nash bargaining gain from trade split would be close to one.

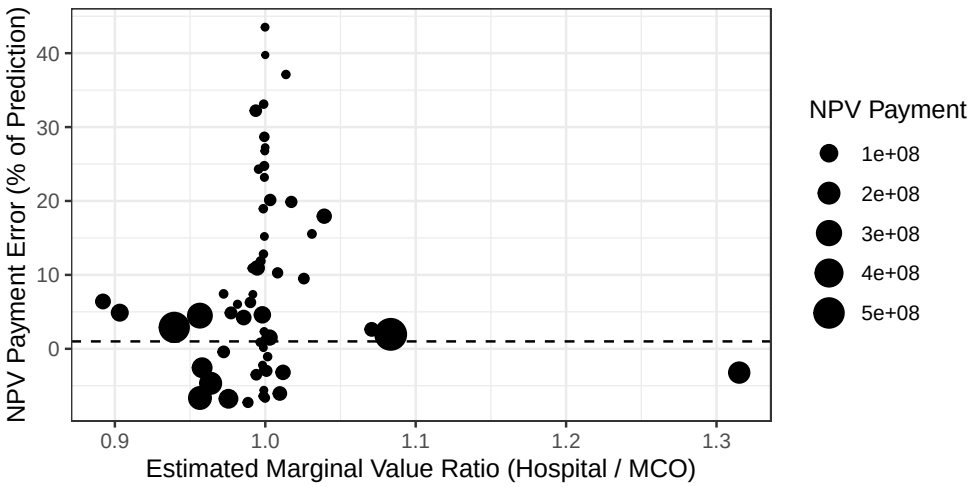


Figure 15: The NPV payment residual as a percentage of predicted payment (y axis) and marginal value ratio (x axis) for estimation bargains. Size corresponds to the NPV payment (weight in bargaining moment). Most marginal value ratios are close to one (x axis near one) corresponding to similar predicted Nash and Kalai proportional splits of gains from trade. A true Nash bargaining model would predict a positive correlation between marginal value ratio (x axis) and realized payments net of Kalai proportional prediction (y axis). The handful of contracts with larger marginal value ratios have residuals close to zero (y values near zero), and if anything exhibits a negative correlation.

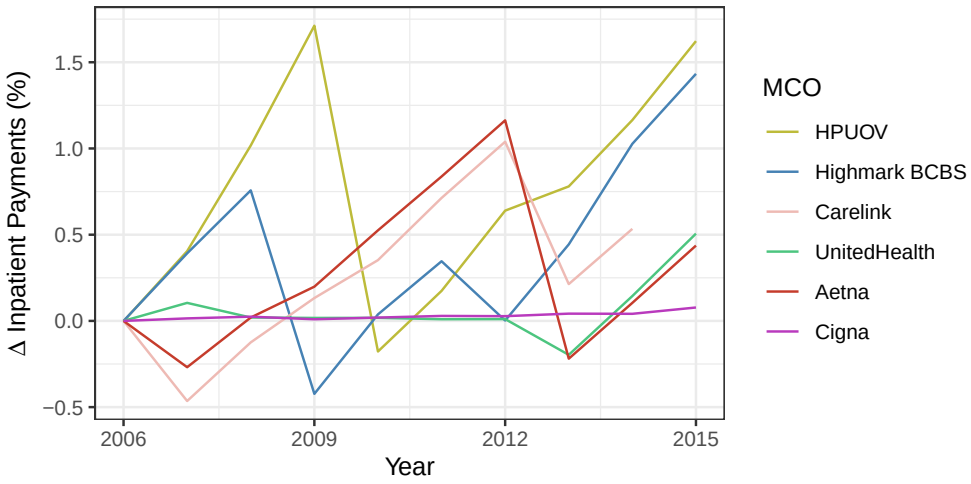


Figure 16: Estimated effects on payments by insurer from a one percentage point annual increase in Medicare rates.

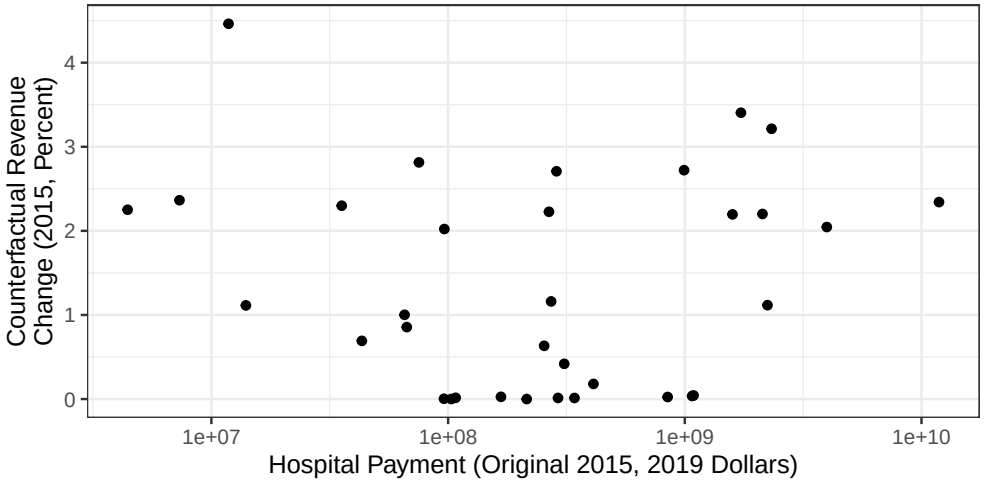


Figure 17: Estimated effects of increased Medicare cost reimbursement on each hospital's received payments in 2015. There is some indication that smaller hospitals would see larger private payment increases.

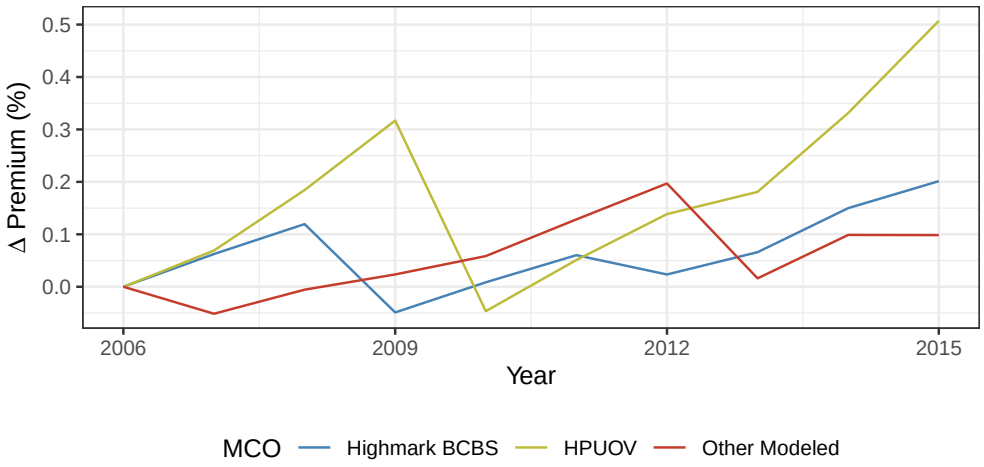


Figure 18: Estimated effects on premiums by insurer for Highmark BCBS (red), HPUOV (green), and the other modeled insurers (blue). Among the other modeled insurers, only Carelink and Aetna have large payment effects, and those effects are reduced after 2012 (see Figure 16).

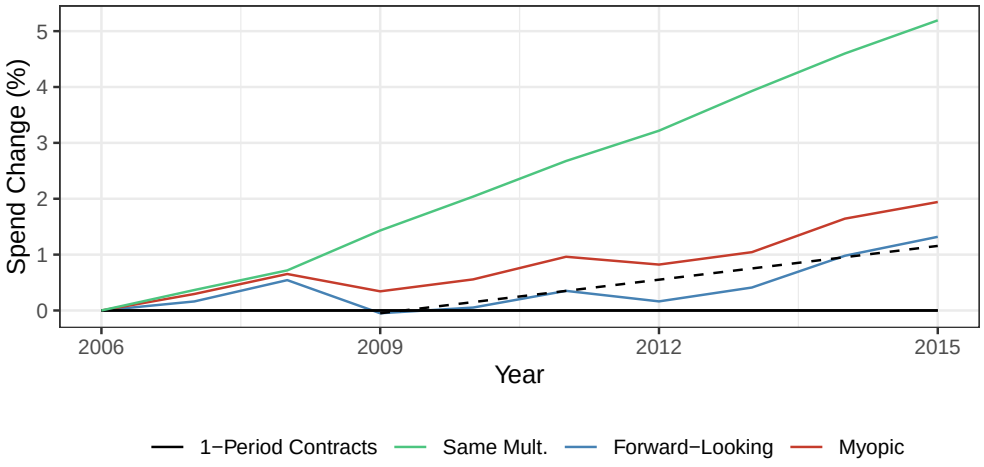


Figure 19: Comparison of estimated effects to a same-multiple model (green) in which bargainers keep their original benchmark multiple in place. The models mainly diverge due to contracts renegotiating a lower initial share of Medicare prices (present myopic model, red) and firms responding to expected future price increases by negotiating lower starting prices (present only in forward looking model, blue). There are also equilibrium spillover effects in both the forward-looking and myopic models, but the magnitudes are smaller than the forward-looking response needed to hold NPV payments constant.

C.4 Other Bargaining Models

The rows in Appendix Table 9 are as follows. The first two rows are the main myopic and forward-looking model estimates, as presented in Table 2. The third row normalizes payments by the average value of $\sum \beta^t$. The fourth row estimated hospital costs as a multiple of list prices rather than calibrating hospital costs. (The estimated multiple is 1.45.) The fifth, sixth, and, seventh rows multiply hospital costs by a fixed scalar. The eighth row multiplies hospital costs by the hospital's reported Medicare payment-to-cost ratio to proxy for the outside option of Medicare patients if hospitals are capacity-constrained. The ninth row takes η values to fit MLR reports rather than calculating them. The tenth row multiplies insurer gains from trade by the hospital's reported share of commercial costs from inpatient care. The eleventh row fixes β. Rows twelve, thirteen, and fourteen fix τ. Rows fifteen and sixteen are versions of the main model estimates, but without τ^{Size} size interactions. Rows seventeen and eighteen are "only-2015" model estimates that use 2015 payments and treat all 2015 contracts as new, with and without hospital size interactions, respectively.

Note also that the estimated discount parameter β is somewhat dependent on the exact split of hospital groups in constructing instruments: bootstrapped β standard errors with reestimated hospital groups increase from 0.03 to 0.109, driven by the 18% of bootstraps in which United Hospital Center's assignment changes.

C.5 Discussion of Cooper et al.'s Work on Prospective Contracts

This work owes a tremendous debt to Cooper et al. (2019). In this section, I discuss how their work on prospective contracts relates to my analysis.

Cooper et al. (2019) estimate 74% of large for-profit insurers' prospective contract cases are paid as a fixed markup over Medicare and find that Medicare benchmarks are associated with larger hospitals. Negotiations of prospective contracts in West Virginia were more likely and more important at larger hospitals, both of which are associated with Medicare benchmark usage in Cooper et al. (2019)'s analysis. Reinhardt (2006) also claims that heterogeneous DRG weights were more typical. I cannot directly compare payment schemes in West Virginia to Medicare rates without access to claims or pricing data by insurer, but Highmark BCBS often used customized DRG weights in inpatient prices disclosed after my dataset ended, but used Medicare rates directly in outpatient calculations at the start of the era I study.

I do not have service-level price disclosures in the era I look at or reliable measures of Highmark BCBS price increases during the post-2021 price disclosure era. I therefore proceed in my main analysis assuming that Highmark BCBS DRG weight increases tracked Medicare

Table 9: Comparison of estimated bargaining parameters with other potential modeling choices. I describe the rows in Appendix C.4.

	Parameter				
	β	τ_{BCBS}	τ_{HPUOV}	τ_{FP}	$-\tau^{Size}$
Myopic	·	0.876***	0.825***	0.861***	1.037***
(Baseline)	(·)	(0.012)	(0.232)	(0.034)	(0.199)
Forward-Looking	0.899***	0.854***	0.877***	0.889***	0.989***
(Baseline)	(0.03)	(0.006)	(0.026)	(0.005)	(0.028)
Forward-Looking	0.925	0.854	0.876	0.89	0.991
(Mean $\sum \beta^t$ normalization)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.497	0.939	0.938	0.942	1.009
(Estimate Hospital Costs)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	1	1	1	1	-0.276
(Hospital Costs * 2)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.931	0.838	0.858	0.875	0.969
(Hospital Costs * 0.9)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	1	0.778	0.781	0.821	0.903
(Hospital Costs * 1/2)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.895	0.834	0.847	0.871	0.913
(Medicare Costs)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.826	0.864	0.874	0.891	0.892
(η from MLR)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.722	0.881	0.905	0.897	0.847
(Inpat. Share GFT Weight)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.99	0.854	0.875	0.881	1
($\beta = 0.99$)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.696	0.001	0.001	0.001	·
(Hospital TIOLI)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.817	0.5	0.5	0.5	·
($\tau = 0.5$)	(·)	(·)	(·)	(·)	(·)
Forward-Looking	0.52	0.999	0.999	0.999	·
(MCO TIOLI)	(·)	(·)	(·)	(·)	(·)
Myopic	·	0.863***	0.845***	0.631***	·
(No Hosp. Size)	(·)	(0.006)	(0.016)	(0.027)	(·)
Forward-Looking	0.714***	0.852***	0.86***	0.685***	·
(No Hosp. Size)	(0.032)	(0.012)	(0.008)	(0.028)	(·)
Only-2015	·	0.487**	-7.54	0.694***	3.354
(Baseline)	(·)	(0.191)	(17.204)	(0.175)	(22.875)
Only-2015	·	0.365***	0.278*	0.16***	·
(No Hosp. Size)	(·)	(0.011)	(0.143)	(0.048)	(·)

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

rates whether Highmark directly used Medicare weights or used customized weights in the era I study. My analysis can also be interpreted as a counterfactual in which Highmark BCBS payment rates were required to increase one percentage point faster annually than in the status quo, regardless of how the payments were calculated. The choice of benchmark only enters the bargaining model through the realized prices. The distinction between Medicare as a benchmark and Medicare-based benchmarks with heterogeneous weights does matter to comparing services (that I aggregate into a generic unit of care) and the interpretation of the counterfactual for policy purposes.

Cooper et al. also argue that Medicare-benchmarked contracts were likely to be boilerplate take it or leave it offers. Cooper et al. do not directly measure boilerplate usage, but large insurers often make such offers to physicians (Abbey, 2012) and Highmark BCBS used shared markups over Medicare for outpatient care at some hospitals (Highmark West Virginia, 2011). That said, stakeholders did not recall boilerplate Highmark BCBS contracts, I have found qualitatively that Highmark BCBS prices disclosed under post-2021 regulations are often calculated as hospital-specific markups over a shared diagnosis-based schedule, and Highmark BCBS contracts were typically implemented at different times (Figure 2). I therefore conclude the use of boilerplate contracts in West Virginia was likely limited.

C.6 Caveats and Limitations

My hospital demand model abstracts from various features to focus on price-setting. Insurers can differ systematically and between plans based on cost-sharing and can put hospitals in separate tiers, but as discussed in Appendix A, the effects of inpatient cost-sharing at common American levels are generally small. I assume 2016 hospital demand was equal to previous hospital demand and as a result do not capture historical hospital investment or changes in patient steering through physician integration. Discussions with stakeholders suggest hospital perceptions were mostly time-invariant. I do not model separate hospital demand by sex, which leads to less precision and could introduce bias by missing premium discrimination before 2014. I do not capture any supply incentives introduced by the choice between prospective Medicare-based payments (which pay based on diagnosis) and list-price-based payments (which pay based on services) or the level of prices, though I hold benchmark choice constant in counterfactuals.

I only model hospital demand by West Virginia residents for West Virginia hospitals. Some degree of state bias at borders is inevitable when data ends at state lines that real humans can cross, and I likely miss some competition in the state's northern panhandle (due to the proximity of Pittsburgh area hospitals), eastern panhandle (due to the proximity

of larger cities in Maryland and Virginia), in Wheeling (due to the proximity of East Ohio Regional Hospital in Ohio), in Huntington (due to the proximity of King's Daughters hospital in Kentucky), and for the Health Plan of the Upper Ohio Valley (which had a comparable line of business in Ohio).

I abstract away from some small potential responses to dropping a contract. I assume there was no out-of-network care by insured patients in West Virginia. Out-of-network care is more common for outpatient care, but can happen for emergency inpatient care and might have become more common if a desirable hospital left an insurer's network. I model consumer substitution to small rural critical access hospitals in response to an insurer dropping a modeled hospital, but treat those payments (which by construction are small) as zero.

The premium data is limited relative to other settings. Large-group premiums were not regulated by the ACA and could reflect different age-based price discrimination or idiosyncratic rating areas than in the ACA-regulated market. It is possible for there to be an unmodeled interaction of age-based premiums with market area that is not captured by market-insurer fixed effects; residual correlation of outpatient and inpatient networks; unmodeled heterogeneity in Cigna, UnitedHealth, and small insurer quality across rating areas; variation in insurer entry in the individual market within rating areas; and variation within rating area in the self-insured market. Aetna premiums in 2015 and 2016 may be mismeasured due to misalignment of premium payments and insurance dates after the insurer acquired Carelink in 2015. I only observe premiums annually. Intrayear premium-setting could be accommodated in the model with appropriate data.

My insurer demand model is highly stylized. I do not have data on choices by family or by employer, so I model the reduced-form selection of insurance by individual based on individual diagnoses probabilities. As discussed in Appendix A, the model of plan choice allows selection only on observables and I treat self-funded and fully insured plans as equally profitable to insurers despite their different business practices. I do not precisely measure outside options across multiple insurers due to lack of identifiers of small insurers. I also do not model the outside option of no insurance or how insurance rates might vary across different areas in West Virginia. I do not model "BlueCard" incentives created by Blue Cross pooling networks: in a border hospital, part of the hospital's value of contracting with Highmark Blue Cross might include the value of additional consumers from CareFirst (Maryland and Virginia) and Anthem (Kentucky and parts of Ohio) Blue Cross in a way that Highmark does not value. There also may be asymmetric value created in other states, but the insurers I model are present on the other sides of West Virginia's more populous Ohio and Pennsylvania borders. That said, the key goal of my model is accurately capturing dynamic bargaining incentives, so the largest concerns reflect any unmodeled changes in

these incentives over time during the era I study.

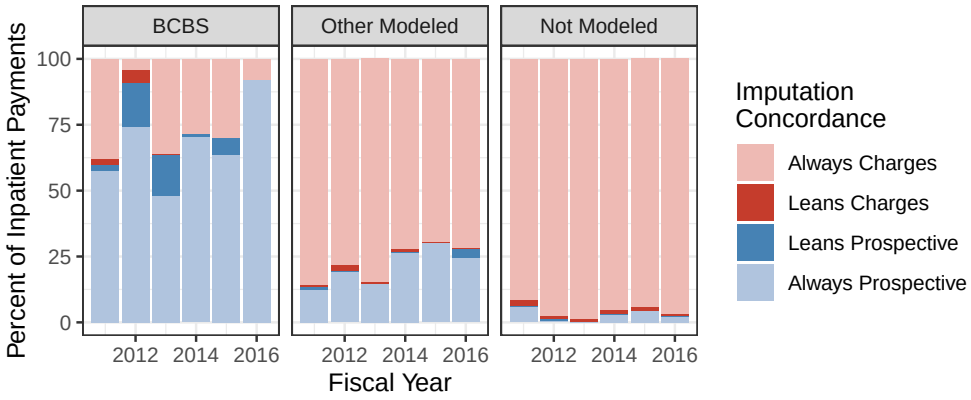


Figure 20: Estimated concordance of benchmark classification by insurer and fiscal year starting with larger contract scale reports in 2011. Blue and red contracts are classified by my algorithm as prospective and share of charges, respectively. Dark colors correspond to round-number discounts reaching the opposite conclusion. An estimated 94.2% of inpatient payments have the same imputation across methods, 1.1% are assigned as share of charges by only repeating non-round-number discounts, and 4.8% are assigned as prospective by only reporting varying round-number discounts (driven by Highmark BCBS-Cabell Huntington in 2011–13).

I capture benchmark usage imperfectly. Both one-off repeated discounts and one-off round-number discounts can reflect coincidence or typographical errors, so each imputation approach has tradeoffs. I generate similar estimated benchmark usage whether I use my main definition or alternatively infer share of charges contracts from round-number discounts. I summarize the concordance across measures in Appendix Figure 20. Both approaches treat as fully prospective any contracts that were benchmarked to list prices with different discounts within inpatient care, which may have included a few Highmark BCBS contracts in early years of my data (Rivard, 2010); any use of per diems for subsets of care like labor cases (Weber et al., 2019); or any use of list prices in outlier payments.

My stylized approach to hospital pricing is standard but abstracts from the relative prices of services. For example, insurers could use list prices for different services to price-discriminate between share of charges payors, though I found no evidence they actually did so. I assume that units of care are proportional to hospital list prices to align with the reported contract data. I inflation-adjust based on CPI which is imperfectly aligned with both hospital care and specific West Virginia conditions. The CPI inflation adjustment may be particularly problematic for noninpatient costs (η) and hospital costs (c), with offsetting effects. It is unclear how these offsetting issues would bias estimates of the patience parameter β, if at all.

I treat benchmark prices as stylized prices per unit of care and treat them as updated annually. It is standard in this literature to aggregate multiple services to a generic unit of care. Firms could apply separate multiples to different care aggregates within a contract, but the average should be a reasonable summary statistic. Firms could commit to time-varying multiples, but to my knowledge they rarely (if ever) do so in practice. List price multiples are theoretically bounded above by one, but generally list prices are intentionally set far above what any contract could reach. I abstract from some other roles of benchmark prices to focus on price negotiation dynamics. I discuss these other mechanisms further in Appendix A. I do not model chargemaster and Medicare timing within a calendar year. List prices governed by the chargemaster could be updated at frequencies other than annually, though such higher frequencies are not standard (Reinhardt, 2006, Tompkins et al., 2006, Jahn, 2017). In principle, chargemasters can be used to price discriminate among insurers that use list prices as a benchmark, but I found no indication that hospitals do so (Reinhardt, 2006, Abbey, 2012, Kidder, 2013).

The model abstracts from many potential aspects of bargaining. I impose a finite horizon model under the view that it is part of an increasingly long-term approximation to a true infinite horizon model. The finite horizon is an approximation — for example, it might be more appropriate to place extra weight on the fifth year as a proxy for subsequent profits. The theoretical arguments deriving estimation moments would not hold if the negotiators had asymmetric patience parameters, risk aversion, or different expectations of the future. Whether these sorts of informational differences can be incorporated in dynamic bargaining is an interesting avenue for future work. It is possible for the bargainers to implicitly or explicitly bargain over nonprice objects like adjudication processes or cooperation, although such objects are generally viewed as secondary (Abbey, 2012, Vega, 2013, Gooch, 2019). I assume that disagreement does not affect subsequent demand functions; disagreement is a dynamic process that affects consumer inertia which I hope to explore in future work.

My estimation procedure implicitly rules out selection of negotiations on unobservable ex ante gains from trade. If, as in the main model estimates, the insurers generally keep most gains from trade, then unobservables leading to higher hospital costs will lead to missing contracts with higher prices, potentially introducing selection bias. That said, networks are substantial in core regions for the insurers I study, so that any bias would likely occur through selection on contract length rather than selection on contract existence. Such hypothetical selection bias is an interesting direction for future work.

The model's timing abstracts from many real-world subtleties around timing in the service of empirical tractability. I model bargaining as succeeding at the start of the day on which it was accepted by the regulator and ending on an unambiguous day of the year. In

practice, contracts are agreed to before they go into effect and occasionally expired contracts would remain in place on a short-term basis while negotiations remained ongoing. Short-term extensions are equivalent to auto-renew contracts in my model, but I will miss extensions that began and ended between contract reports. Impasse is assumed to affect insurer demand in a static process throughout the year; insurance contracts with employers and individuals are reached at staggered times, a process that does not correspond perfectly with this paper's static insurer demand model. I attempt to capture annual patience with an annual patience parameter, using contract dates within a year when estimating the bargaining model but treating profits as equally profitable within a year, which is a convenient but unrealistic simplification. For counterfactuals, I simplify prices and treat bargaining as being conducted at the start of the year to focus on the counterfactual effects at the cost of some precision.

There are also a few other places wherein the bargaining model is currently simplified for convenience. The bargaining model is estimated with a high-dimensional optimization that may not reach a global optimum. The estimation procedure is somewhat affect by initial conditions; most notably, negotiation cost contributions to payments do not always move much from their initial value of $10,000. I use state-level premiums to calculate insurer gains from trade, which may introduce some bias from relative age discrimination between insurers. As discussed, I do not currently estimate an impasse repricing transfer term (with unclear effect), I hold premiums constant in counterfactuals (with conservative effect), and do not have the sufficient variation to statistically test a Nash bargaining model (which likely would produce similar estimates as the model I estimate).

I do not model the network formation process. As argued in the literature referenced in Appendix A, a frictionless Nash-in-Nash bargaining model rarely speaks persuasively to why networks are not complete — for example, an insurer might exclude a hospital to increase their leverage in other negotiations. In my model, networks are theoretically restricted by hospital costs and negotiation costs. I do not view those bargaining frictions as a fully compelling model of network formation. Other disagreement models could be accommodated by adjusting the impasse repricing transfer term to include other disagreement effects. I focus on changes to benchmark price increases that are comparable to normal levels of benchmark price increases. As a result, the counterfactuals would be unlikely to affect network formation substantially.

There are various limitations on counterfactuals. The one percentage point payment increase is based on West Virginia data, but a national equivalent might be closer to 0.7 percentage point annual increase. It is possible that my deterministic stance on uncertainty in counterfactuals introduces bias if the uncertainty over benchmark prices was first-order relative to the changes that would be incurred in the counterfactual I consider. As mentioned

in Appendix A, I do not model any effects on how much care would be reported or how hospital investment might change in counterfactuals due to the limited effects of prices on supply. The set of benchmarks used has changed over time and has shifted towards prospective payments, making out-of-sample extrapolation unclear but potentially meaning national effects of Medicare-based benchmark prices would be larger in years after 2015. West Virginia is a small market, so it is possible that bargaining is less frequent than in large markets that constitute a lot of, but by no means all of, American hospital care. The contract data I use is partial in the earliest and latest years, and as a result I may miss some bargains that were not reported. The downstream estimates of effects on premiums are highly stylized.

C.7 Detailed Data Description

The first step of data processing is cleaning the contract reports. I discuss this cleaning in Dorn (2024). Networks are inferred by calendar year of submission, with missing years inferred from the closest report breaking ties to previous reports. There is a small amount of manual network handling. I drop contracts for nonstandard care like psychiatric care, lab fees, or professional fees. When an insurer reports multiple contracts (for example Highmark BCBS separately reports their indemnity and PPO contracts), I aggregate payments using the closest available reported number of discharges where possible. (I take an unweighted average if I never have estimated number of discharges per contract.) In this paper, I include First Health contracts as HPUOV contracts based on HPUOV's description of First Health as a "strategic partner" (Wayback Machine, 2021).

I focus on the regulated hospitals and treat the remote Critical Access Hospitals (CAHs) that were deregulated after 2000 as negligible. The state also allowed border hospitals to keep their contracts private. I use the fiscal year 2016 report to infer list price payment rates for Weirton Medical Center, and treat the small Williamson Memorial Hospital as equivalent to a CAH.

The West Virginia bargains for estimation were identified manually. The main source was reported contract start and end dates in the panel contract dataset, but regulator contract decisions were also used as a supplement. I mitigate the bias introduced by Aetna's acquisition of Carelink at the start of 2015 by not including any contracts which lasted into 2016 (under the finite horizon) in bargaining estimation. I identify likely bargains for use in counterfactuals but not estimation based on remaining occasions on which either a contract was introduced, a share of charges contract was changed or replaced, a first year after expiration (with a change in discount rates), manual research suggests a change in payments, or

the year after a contract was reported as being expected to expire, so long as the automatic processing identification does not happen in the last period in which I observe the contract. I treating the effective date as January 1 (except for one case in which other data suggests the contract began January 2).

In the inpatient data, I exclude rehab, long-term, and psychiatric hospitals; exclude newborns, residents of other states, and noncommercially insured patients (but including public employees who chose HPUOV to align with the fully insured sales data); take the hospital's main location from Medicare cost reports; and identify systems that reported joint contracts based on personal research. I assign patients locations by county geographic centroid. I estimate probabilities of patients having misclassified insurance status based on reported care frequencies where typos seem likely.

I infer diagnosis categories based on Ho (2006)'s classification of ICD-9 codes. The West Virginia inpatient data lacks ICD-9 codes and only has ICD-10 codes for 59% of discharges, so I convert the data's MS-DRG codes to ICD-10 codes using CMS (2020) and then into ICD-9 codes using NBER (2021). I supplement this conversion with manual research for common DRG codes this method fails to classify. I drop the 2% of cases for which the DRG conversion did not yield an ICD code and I did not reach an active category determination. Where this process maps multiple ICD-9 codes to the same DRG category, I choose the most common ICD-9 code's category.

I calculate Medicare payment-to-cost ratios from state uniform financial reports (UFRs) and linearly interpolate payment-to-cost ratios where missing in the available data.

I convert premiums and other monetary data to 2019 dollars based on consumer price index (CPI) inflation. Federal regulations precluded the state from collecting data on sales in the self-funded market. I infer self-funded sales in 2016 to match estimated combined sales in the UB data and use the fully insured market to infer premiums and insurer values before 2016.

There are some subtleties to my insurer data. Fully insured sales by insurer and self-funded sales estimates come from reports like Offices of the Insurance Commissioner (2008, 2016) for the comprehensive market. Insurer sales are aggregated by group code where possible and outliers are cleaned. In 2008–2009, sales were not reported. As a result, I linearly interpolate the missing lives and inflation-adjusted premiums. I similarly linearly interpolate the sales estimates for the ERISA (self-funded) market for missing years. I calculate MLRs from 2011–2018 CMS reports for West Virginia business in the individual, small-group, and large-group markets. I aggregate MLRs by NAIC company code where available and by name where NAIC codes are not available and take the numerators and denominators from the MLR_NUMERATOR and MLR_DENOMINATOR variables in part 5 (for 2011–13) or

part 3 (for 2014 and later) of the reports. I aggregate medical loss and premium revenue across insurance products by group code. Inflation rates are calculated using World Bank CPI inflation over years relative to 2019 from the `priceR` package: 2017 nominal payments are inflation-adjusted based on the inflation rates reported for 2017 and 2018.

Hospital demand and ex ante WTP are calculated as follows. I calculate the probability of any diagnosis in the inpatient data in 2016 conditional on age, assuming each person has at most one inpatient discharge per year. I then obtain the potential hospitals each Highmark BCBS patient could have visited and run a weighted logit regression of choice on hospital and ν characteristics by diagnosis. The regression is weighted to include probability-of-Blue-Cross-weighted choices at hospitals that misclassified Blue Cross care. I then extrapolate the estimates to calculate the ex post willingness to pay for every conceivable county-age-hospital combination conditional on diagnosis and aggregate the measure into an ex ante WTP measure for every hospital-insurer-age-location-year combination. The WTP measure is calculated as follows:

$$WTP_{j,k,c} = \sum_{\ell} \mathbb{P}(\text{Diagnosis } \ell \mid \text{Age group } k) \log \left(\sum_{h \in \mathcal{G}_j^M} exp(u_{c,h,\ell}^H) \right),$$

where the $u_{c,h,\ell}$ ex ante hospital utility to a consumer in county c with diagnosis ℓ is from the hospital choice model.

Insurer demand estimation is an involved process involving substantial data cleaning for 2016 alone. I first estimate insurer sales based on the fraction of commercially insured inpatient diagnoses from an age group in a county in the inpatient data. The county commercially insured population is taken as the Census intercensal population estimate multiplied by the state estimated fraction of age group with commercial insurance in the inpatient data. I then adjust the inpatient data to ensure every insurer has at least one estimated sale per age-county (taking the needed population from other sales estimates proportionally) and then include non-Highmark (other state) Blue Cross in the outside option. I infer state-level insurer sales in the self-funded market in 2016 for Aetna, Highmark BCBS, and HPUOV based on the difference between state-level sales estimates and state-level fully insured sales. I extrapolated self-funded sales for the two insurers not identified in the inpatient data, Cigna and UnitedHealth, by assuming the sales ratio between the markets is equal to the median estimated ratio. I scale down estimated sales to insure the modeled insurers never exceed 85% of a county-age group's estimated sales individually or exceed 90% in aggregate.

Once sales are estimated, I estimate insurer demand with an outer loop–inner loop algorithm. I An outer loop proposes Cigna and UnitedHealth $\tilde{\delta}^M$ fixed effects (including state-

level premiums) and an inner loop produces county-age-insurer implied values of $\gamma_k WTP_k + \xi$ to fit age-insurer-county ales estimates for the modeled insurers. I then iteratively update the Cigna and UnitedHealth fixed effects based on the current WTP coefficient estimates until conversion. The WTP coefficients are calculated by market-size-weighted regression of $\gamma_k WTP_k + \xi$ (inferred from 2016 sales estimates) on WTP_k (from hospital demand and sickness probabilities). To calculate pre-2016 demand, I calculate pre-2016 WTP in utility by insurer, county, and age. I solve for changes to state-level insurer value (inclusive of state-level premiums) to match state-level sales after adjusting for county-level population changes and ASO market size changes changes. I assume that Carelink's ξ values before its acquisition by Aetna at the end of 2014 were equal to Aetna's values in those same markets.

For bargaining, I estimate the effect of insurer network on hospital and insurer sales as follows. I predict sales under both the observed networks and under counterfactual networks that drop the insurer–system pair at the observed premiums. (Premium changes would be measured in Pay_{IRT}, which estimation currently sets to zero.) I measure the effects for a bargain year as a weighted average of calendar years: if a bargain began 3/4 of the way through 2010, then the first year of gains from trade under the bargain will be a weighted average of 1/4 of the gains from 2010 and 3/4 of the gains from 2011. I calculate the inputs to gains from trade, like the change in hospital costs, for the bargaining estimator. For calculating τ_{ij} hospital heterogeneity, I calculate hospital costs incurred in 2006 as the sum of reported list prices multiplied by the estimated cost-to-charge ratio. I calculate demand estimates from the estimated models to mitigate reporting endogeneity.

The bargaining optimization proceeds as follows. For constructing τ_{ij}, I normalize log hospital system size by the mean log system size in bargaining to report τ_j at the mean. The hospital groups in the price instruments Z^p are chosen to group hospitals by approximate size while ensuring a reasonable number of bargains for each hospital group: the hospitals are first ordered by the quantity of NPV realized payments in estimation bargains if β were equal to 0.8, and then split into six groups based on quantiles of payments taken to the power of 0.3, a quantity which was chosen to balance information with group size. Bargaining parameters are optimized over a simple moment weighting that tries to make the scales roughly comparable across moments: it weights MLR squared moments by 10^5 and normalized payment squared moments by the average NPV payment if β were equal to 0.8 (with the ω-style denominator normalization) by the relevant hospital group or insurer. The optimization attempts to re-optimize 10 times before returning the estimated parameters. Standard errors are calculated by bootstrap by resampling inpatient cases and state-level insurance choices 100 times. The bootstrapped confidence intervals for counterfactuals take the estimated demand functions as fixed and incorporate the uncertainty in bargaining parameters.

The payment residual Appendix Figures 13, 14, and 15 are constructed as follows. The validation in Appendix Figure 13 scales the initial payment rate at the level needed to match the given model's prediction perfectly in NPV terms. This approach intuitively corresponds to holding benchmark price increases constant. In Appendix Figures 14 and 15, I take the marginal effect of a dollar negotiated under a contract on each future contract in the model and then evaluate the change in NPV profits over the finite horizon.

The counterfactual calculation process is as follows. I take the estimated $\hat{\tau}_{ij}$ from the relevant bargaining models, calculate payment multiples to infer the realized counterfactual ratio of starting price to NPV payment, and add 2016 data based on 2015 for computing counterfactual effects on late inferred bargains. There is some further data handling around the Aetna-Carelink acquisition at Davis Medical Center, which had a contract with Carelink but not Aetna before the acquisition. On the rare occasion that a midyear negotiation led to a change of benchmark, I infer the smaller starting price, which increases the forward-looking offsets slightly. I construct realized price transition matrices and matrices of the ex post effect of future prices on bargained prices, calculate the realized residual (including demand-driven components of gains from trade) which is held fixed in counterfactuals, confirm that geometric sum estimates of effects would converge, calculate counterfactuals by matrix inversion, and calculate some summary statistics for later analysis.

The implementation of counterfactuals involves substantial data cleaning. I calculate when contract terms were changed under a modeled or inferred bargain. Counterfactual prices are adjusted at the start the year of negotiation or inferred change. Negotiations inferred from an expected expiration date past the final contract report are implemented at the start of the next calendar year to allow for potential roll-over. There is further data cleaning, for example ensuring that new bargains are not inferred during years that are a part of an estimation sample bargain, ensuring inferred benchmark choice by year is consistent with the inferred negotiation dates around the Aetna/Carelink acquisition, and holding fixed some small hospitals that did not provide inpatient data in 2016 that was used to estimate hospital demand.

I estimate back-of-the-envelope downstream premium effects based on Ho (2006)'s estimated own-price elasticity. I calibrate a coefficient on inflation-adjusted premiums to match the estimated average own-price elasticity of -1.4. I solve for the wedge in insurer-year marginal cost needed to make realized premiums optimal under simultaneous annual Nash-Bertrand premium-setting. I then find the new equilibrium premiums under that change in insurer-year marginal cost and the new predicted negotiated premiums, inclusive of any patient reallocation in response to the new premiums. I do not model how the increased premiums further increase the negotiated prices.

D Proofs

Proof of Lemma 2. Suppose the insurer sells D^M units of insurance with full networks at premiums ϕ and sells $\underline{D}^M$ units with one hospital in their network. I will consider only non-negative anticipated prices, so that joint gains from trade are always positive.

Let the gains from trade to the contracting insurer from agreeing to price p_i with anticipated simultaneous price p_{-i} be $GFT^M_{(Alt)}(p_i \mid p_{-i})$ and $GFT^M_{(Simult)}(p_i \mid p_{-i})$ in alternating and simultaneous bargaining, respectively. Let the simultaneous equilibrium prices be $p^*_{(Alt)}$ and $p^{(Nash),*}_{(Simult)}$, respectively. Let the hospital's gains from trade be GFT^H.

I abuse notation and write the single period discounting rate as $\bar{\beta} = \beta^{1/2}$

I consider equilibria of the form $p^*_i = p^*_{(Alt)} + \alpha^{(Nash)}_{(Alt)}(p_{-i} - p^*_{Alt})$ and $p^*_i = p^{(Nash),*}_{(Simult)} + \alpha^{(Nash)}_{(Simult)}(p_{-i} - p^*_{Simult})$, which I later verify exist.

I first characterize gains from trade with alternating period bargaining. By agreeing at a price p_i, the insurer gains $\phi(D^M - \underline{D}^M)$ in premium revenue this period and reduces its payments to hospital $-i$ from $\underline{D}^M p_{-i}$ to $\frac{D^M}{2}p_{-i}$. It pays a NPV payment to hospital i of $\frac{1}{1-\bar{\beta}}\frac{p^*_{(Alt)}D^M}{2} + \frac{(1+\bar{\beta})(p_i - p^*_{(Alt)})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2}$. Beginning in the next period, it pays a NPV payment to hospital $-i$ of $\frac{\bar{\beta}}{1-\bar{\beta}}\frac{p^*_{(Alt)}D^M}{2} + \frac{\alpha^{(Nash)}_{(Alt)}\bar{\beta}(1+\bar{\beta})(p_i - p^*_{(Alt)})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2}$. By disagreeing, starting in the next period, it pays a NPV payment to hospital i of $\frac{\bar{\beta}}{1-\bar{\beta}}\frac{p^{(Nash),*}_{(Simult)}D^M}{2}$ and pays a NPV payment to hospital $-i$ of $\frac{\bar{\beta}}{1-\bar{\beta}}\frac{p^{(Nash),*}_{(Simult)}D^M}{2}$. Therefore:

$$\begin{aligned}
GFT^M_{(Alt)}(p_i \mid p_{-i}) &= \phi(D^M - \underline{D}^M) + \frac{2\underline{D}^M - D^M}{2}p_{-i} + \frac{\bar{\beta}D^M}{1-\bar{\beta}}(p^{(Nash),*}_{(Simult)} - p^*_{(Alt)}) \\
&\quad - \frac{p^*_{(Alt)}D^M}{2} - \frac{1+\bar{\beta}}{1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}}\frac{D^M}{2}(p_i - p^*_{(Alt)}) \\
&= GFT^M_{(Alt)}(p^*_{(Alt)} \mid p^*_{(Alt)}) + \frac{2\underline{D}^M - D^M}{2}\left(p_{-i} - p^*_{(Alt)}\right) - \frac{1+\bar{\beta}}{1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}}\frac{D^M}{2}(p_i - p^*_{(Alt)}) \\
GFT^H_{(Alt)}(p_i \mid p_{-i}) &= \frac{1}{1-\bar{\beta}}\frac{p^*_{(Alt)}D^M}{2} + \frac{(1+\bar{\beta})(p_i - p^*_{(Alt)})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2} - \frac{\bar{\beta}}{1-\bar{\beta}}\frac{p^{(Nash),*}_{(Simult)}D^M}{2} \\
&= GFT^H_{(Alt)}(p^*_{(Alt)} \mid p^*_{(Alt)}) + \frac{(1+\bar{\beta})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2}(p_i - p^*_{(Alt)}).
\end{aligned}$$

The marginal value ratio between the insurer and hospital is is:

$$\frac{\frac{1+\bar{\beta}}{1-\alpha_{(Alt)}^{(Nash)}\bar{\beta}}\frac{D^M}{2}}{\frac{(1+\bar{\beta})}{1-(\alpha_{(Alt)}^{(Nash)}\bar{\beta})^2}\frac{D^M}{2}} = 1+\alpha_{(Alt)}^{(Nash)}\bar{\beta},$$

so that the equilibrium gain from trade ratio is $\frac{\tau(1+\alpha_{(Alt)}^{(Nash)}\bar{\beta})}{1-\tau}$.

I next characterize the gains from trade from simultaneous bargaining. Generically, I leverage the following useful repeated disagreement fact in the absence of uncertainty. Let $V(d)$ be the value to an agent of disagreeing d times. Then the gains from trade $V(0)-V(1)$ are equal to $V(0)-V(2)-(V(1)-V(2))$, the gains from trade relative to disagreeing twice minus the gains from trade obtained by only disagreeing once. An insurer agreeing to a simultaneous price p_i relative to disagreeing twice, the insurer gains premium revenue $\phi(D^M-\underline{D}^M)(1+\bar{\beta})$ and changes its payments in the next two periods from $(1+\bar{\beta})\underline{D}^M p_{-i}$ to $(1+\bar{\beta})D^M\frac{p_i+p_{(-i)}}{2}$. Similarly, by agreeing relative to disagreeing twice and returning to the simultaneous equilibrium, the hospital increases its profits by $(1+\bar{\beta})\frac{D^M}{2}p_i$. Therefore the gains from trade in simultaneous bargaining can be written as:

$$\begin{aligned}
GFT_{(Simult)}^{M}(p_i \mid p_{-i}) &= \phi(1+\bar{\beta})(D^M-\underline{D}^M)+p_{-i}\underline{D}^M(1+\beta)-(1+\bar{\beta})D^M\frac{p_i+p_{-i}}{2}\\
&\quad-\bar{\beta}GFT_{(Alt)}^{M}(p_{(Alt)}^{*}+\alpha_{(Alt)}^{(Nash)}(p_{-i}-p_{(Alt)}^{*}) \mid p_{-i})\\
&= (\phi-p_{-i})(1+\bar{\beta})(D^M-\underline{D}^M)+(1+\bar{\beta})D^M\frac{p_{-i}-p_i}{2}\\
&\quad-\bar{\beta}GFT_{(Alt)}^{M}(p_{(Alt)}^{*}+\alpha_{(Alt)}^{(Nash)}(p_{-i}-p_{(Alt)}^{*}) \mid p_{-i})\\
&= GFT_{(Simult)}^{M}(p_{(Simult)}^{(Nash),*} \mid p_{(Simult)}^{(Nash),*})\\
&\quad+\bar{\beta}\begin{pmatrix} GFT_{(Alt)}^{M}(p_{(Alt)}^{*}+\alpha_{(Alt)}^{(Nash)}(p_{(Simult)}^{(Nash),*}-p_{(Alt)}^{*}) \mid p_{(Simult)}^{(Nash),*}) \\ - \quad GFT_{(Alt)}^{M}(p_{(Alt)}^{*}+\alpha_{(Alt)}^{(Nash)}(p_{-i}-p_{(Alt)}^{*}) \mid p_{-i}) \end{pmatrix}\\
&\quad+(1+\bar{\beta})\left(\underline{D}^M-\frac{D^M}{2}\right)(p_{-i}-p_{(Simult)}^{(Nash),*})-(1+\bar{\beta})\frac{D^M}{2}(p_i-p_{(Simult)}^{(Nash),*})\\
GFT_{(Simult)}^{H}(p_i \mid p_{-i}) &= (1+\bar{\beta})\frac{D^M}{2}p_i-\bar{\beta}GFT_{(Alt)}^{H}(p_{(Alt)}^{*}+\alpha_{(Alt)}^{(Nash)}(p_{-i}-p_{(Alt)}^{*}) \mid p_{-i})\\
&= GFT_{(Simult)}^{H}(p_{(Simult)}^{(Nash),*} \mid p_{(Simult)}^{(Nash),*})+(1+\bar{\beta})\frac{D^M}{2}(p_i-p_{(Simult)}^{(Nash),*}).
\end{aligned}$$

It is difficult to characterize the best response $\alpha_{(Simult)}^{(Nash)}$, but the GFT function here is linear, so such an $\alpha_{(Simult)}^{(Nash)}$ exists. Note that simultaneous bargaining is TU, so that the equilibrium gain from trade ratio is $\frac{\tau}{1-\tau}$.

An equilibrium is a set of equilibrium prices $p_{(Alt)}^{*}, p_{(Simult)}^{(Nash),*}$ and coefficients $\alpha_{(Alt)}^{(Nash)}$ and

$\alpha^{(Nash)}_{(Simult)}$ such the following system of equations holds:

$$\frac{\tau(1+\alpha^{(Nash)}_{(Alt)}\bar{\beta})}{1-\tau} = \frac{GFT^M_{(Alt)}(p^*_{(Alt)} + \alpha_{(Alt)}{}^{(Nash)}(p_i - p_{-i}) \mid p_{-i})}{GFT^M_{(Alt)}(p^*_{(Alt)} + \alpha_{(Alt)}{}^{(Nash)}(p_i - p_{-i}) \mid p_{-i})}$$

$$\frac{\tau}{1-\tau} = \frac{GFT^M_{(Simult)}(p^{(Nash),*}_{(Simult)} + \alpha^{(Nash)}_{(Simult)}(p_i - p_{-i}) \mid p_{-i})}{GFT^M_{(Simult)}(p^{(Nash),*}_{(Simult)} + \alpha^{(Nash)}_{(Simult)}(p_i - p_{-i}) \mid p_{-i})}$$

and agreement at the equilibrium prices is profitable.

First I characterize $\alpha_{(Alt)}{}^{(Nash)}$ via:

$$\begin{aligned}
\frac{\tau(1+\alpha^{(Nash)}_{(Alt)}\bar{\beta})}{1-\tau} &= \frac{GFT^M_{(Alt)}(p^*_{(Alt)} + \alpha_{(Alt)}{}^{(Nash)}(p_i - p_{-i}) \mid p_{-i})}{GFT^M_{(Alt)}(p^*_{(Alt)} + \alpha_{(Alt)}{}^{(Nash)}(p_i - p_{-i}) \mid p_{-i})} \\
&= \frac{GFT^M_{(Alt)}(p^*_{(Alt)} \mid p^*_{(Alt)}) + \frac{2\underline{D}^M - D^M}{2}\left(p_{-i} - p^*_{(Alt)}\right) - \frac{1+\bar{\beta}}{1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}}\frac{D^M}{2}(p_i - p^*_{(Alt)})}{GFT^H_{(Alt)}(p^*_{(Alt)} \mid p^*_{(Alt)}) + \frac{(1+\bar{\beta})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2}(p_i - p^*_{(Alt)})} \\
&= \frac{GFT^M_{(Alt)}(p^*_{(Alt)} \mid p^*_{(Alt)}) + \frac{2\underline{D}^M - D^M}{2}\left(p_{-i} - p^*_{(Alt)}\right) - \frac{1+\bar{\beta}}{1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}}\frac{D^M}{2}(p_i - p^*_{(Alt)})}{\frac{1-\tau}{\tau(1+\alpha^{(Nash)}_{(Alt)}\bar{\beta})}GFT^M_{(Alt)}(p^*_{(Alt)} \mid p^*_{(Alt)}) + \frac{(1+\bar{\beta})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2}(p_i - p^*_{(Alt)})}.
\end{aligned}$$

Because the equilibrium prices already split the gains from trade at the desired ratio, it must be the case that:

$$\begin{aligned}
\frac{\tau(1+\alpha^{(Nash)}_{(Alt)}\bar{\beta})}{1-\tau} &= \frac{\frac{2\underline{D}^M - D^M}{2}\left(p_{-i} - p^*_{(Alt)}\right) - \frac{1+\bar{\beta}}{1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}}\frac{D^M}{2}(p_i - p^*_{(Alt)})}{\frac{(1+\bar{\beta})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2}(p_i - p^*_{(Alt)})} \\
&= \frac{\frac{2\underline{D}^M - D^M}{2}\left(p_{-i} - p^*_{(Alt)}\right) - \frac{1+\bar{\beta}}{1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}}\frac{D^M}{2}\alpha^{(Nash)}_{(Alt)} * (p_{-i} - p^*_{(Alt)})}{\frac{(1+\bar{\beta})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2}\alpha^{(Nash)}_{(Alt)} * (p_{-i} - p^*_{(Alt)})} \\
&= \frac{\frac{2\underline{D}^M - D^M}{2} - \frac{1+\bar{\beta}}{1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}}\frac{D^M}{2}\alpha^{(Nash)}_{(Alt)}}{\frac{(1+\bar{\beta})}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}\frac{D^M}{2}\alpha^{(Nash)}_{(Alt)}} \\
&= \frac{\frac{2\underline{D}^M - D^M}{D^M} - \frac{1+\bar{\beta}}{1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}}\alpha^{(Nash)}_{(Alt)}}{\frac{(1+\bar{\beta})}{(1-\alpha^{(Nash)}_{(Alt)}\bar{\beta})(1+\alpha^{(Nash)}_{(Alt)}\bar{\beta})}\alpha^{(Nash)}_{(Alt)}}
\end{aligned}$$

$$\begin{aligned}
\frac{\tau}{1-\tau} &= \frac{\frac{2\underline{D}^M - D^M}{D^M} - \frac{(1+\bar{\beta})\alpha_{(Alt)}^{(Nash)}}{1-\alpha_{(Alt)}^{(Nash)}\bar{\beta}}}{\frac{(1+\bar{\beta})\alpha_{(Alt)}^{(Nash)}}{1-\alpha_{(Alt)}^{(Nash)}\bar{\beta}}} \\
&= \frac{\frac{2\underline{D}^M - D^M}{D^M}(1-\alpha_{(Alt)}^{(Nash)}\bar{\beta}) - (1+\bar{\beta})\alpha_{(Alt)}^{(Nash)}}{(1+\bar{\beta})\alpha_{(Alt)}^{(Nash)}} \\
(1+\bar{\beta})\alpha_{(Alt)}^{(Nash)} &= (1-\tau)\frac{2\underline{D}^M - D^M}{D^M}(1-\alpha_{(Alt)}^{(Nash)}\bar{\beta}) \\
\alpha_{(Alt)}^{(Nash)} &= \frac{(1-\tau)\frac{2\underline{D}^M - D^M}{D^M}}{1+\bar{\beta}\frac{D^M+(1-\tau)(2\underline{D}^M - D^M)}{D^M}} \\
&= \frac{(1-\tau)(2\underline{D}^M - D^M)}{(1+\bar{\beta})D^M + \bar{\beta}(1-\tau)(2\underline{D}^M - D^M)},
\end{aligned}$$

which is part of the claim.

Note that $2\underline{D}^M - D^M > 0$ and $2\underline{D}^M - D^M < D^M$, so that $\alpha_{(Alt)}^{(Nash)} \in [0,1)$ and the system is convergent.

Next consider the simultaneous equilibrium price $p_{(Simult)}^{(Nash),*}$. I write $p_{(Next)}^* = p_{(Alt)}^* + \alpha_{(Alt)}^{(Nash)}(p_{(Simult)}^{(Nash),*} - p_{(Alt)}^*)$ for the next-period price with hospital i after disagreement on the simultaneous contract today. Then:

$$\begin{aligned}
\frac{\tau}{1-\tau} &= \frac{(1+\bar{\beta})(\phi - p_{(Simult)}^{(Nash),*})(D^M - \underline{D}^M) - \bar{\beta}GFT_{(Alt)}^M(p_{(Next)}^* \mid p_{(Simult)}^{(Nash),*})}{(1+\bar{\beta})\frac{D^M}{2}p_{(Simult)}^{(Nash),*} - \bar{\beta}GFT_{(Alt)}^H(p_{(Next)}^* \mid p_{(Simult)}^{(Nash),*})} \\
&= \frac{(1+\bar{\beta})(\phi - p_{(Simult)}^{(Nash),*})(D^M - \underline{D}^M)}{(1+\bar{\beta})\frac{D^M}{2}p_{(Simult)}^{(Nash),*} - \bar{\beta}GFT_{(Alt)}^H(p_{(Next)}^* \mid p_{(Simult)}^{(Nash),*})} \\
&\quad - \frac{\left(\frac{\tau}{1-\tau} + \frac{\tau(1+\alpha_{(Alt)}^{(Nash)}\bar{\beta})}{1-\tau} - \frac{\tau}{1-\tau}\right)\bar{\beta}GFT_{(Alt)}^H(p_{(Next)}^* \mid p_{(Simult)}^{(Nash),*})}{(1+\bar{\beta})\frac{D^M}{2}p_{(Simult)}^{(Nash),*} - \bar{\beta}GFT_{(Alt)}^H(p_{(Next)}^* \mid p_{(Simult)}^{(Nash),*})} \\
\frac{\tau}{1-\tau} &= \frac{(1+\bar{\beta})(\phi - p_{(Simult)}^{(Nash),*})(D^M - \underline{D}^M) - \bar{\beta}\frac{\tau\alpha_{(Alt)}^{(Nash)}\bar{\beta}}{1-\tau}GFT_{(Alt)}^H(p_{(Next)}^* \mid p_{(Simult)}^{(Nash),*})}{(1+\bar{\beta})\frac{D^M}{2}p_{(Simult)}^{(Nash),*}} \\
\tau(1+\bar{\beta})\frac{D^M}{2}p_{(Simult)}^{(Nash),*} &= (1-\tau)(1+\bar{\beta})(\phi - p_{(Simult)}^{(Nash),*})(D^M - \underline{D}^M) \\
&\quad - \tau\bar{\beta}^2\alpha_{(Alt)}^{(Nash)}GFT_{(Alt)}^H(p_{(Next)}^* \mid p_{(Simult)}^{(Nash),*}) \\
\tau(1+\bar{\beta})\frac{D^M}{2}p_{(Simult)}^{(Nash),*} &= (1-\tau)(1+\bar{\beta})(\phi - p_{(Simult)}^{(Nash),*})(D^M - \underline{D}^M) \\
&\quad - \tau\bar{\beta}^2\alpha_{(Alt)}^{(Nash)}\left(\frac{D^M(p_{(Alt)}^* - \bar{\beta}p_{(Simult)}^{(Nash),*})}{2(1-\bar{\beta})} + \frac{(1+\bar{\beta})\alpha_{(Alt)}^{(Nash)}(p_{(Simult)}^{(Nash),*} - p_{(Alt)}^*)}{1-(\alpha_{(Alt)}^{(Nash)}\bar{\beta})^2}\frac{D^M}{2}\right)
\end{aligned}$$

$$p_{(Simult)}^{(Nash),*}\left(\begin{array}{ll} & \tau(1+\bar{\beta})\frac{D^M}{2}+(1-\tau)(1+\bar{\beta})(D^M-\underline{D}^M) \\ + & \frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha_{(Alt)}^{(Nash)})^2D^M}{2(1-(\alpha_{(Alt)}^{(Nash)}\bar{\beta})^2)} \\ - & \frac{\tau\bar{\beta}^3\alpha_{(Alt)}^{(Nash)}D^M}{2(1-\bar{\beta})}\end{array}\right)$$

$$=(1-\tau)(1+\bar{\beta})\phi(D^M-\underline{D}^M)+p_{(Alt)}^*\left(\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha_{(Alt)}^{(Nash)})^2D^M}{2(1-(\alpha_{(Alt)}^{(Nash)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^2\alpha_{(Alt)}^{(Nash)}D^M}{2(1-\bar{\beta})}\right).$$

Thus, $p_{(Alt)}^*$ and $p_{(Simult)}^{(Nash),*}$ solve the system of equations:

$$\frac{\tau}{1-\tau}=\frac{2(1-\bar{\beta})(\phi-p_{(Alt)}^*)(D^M-\underline{D}^M)+2\bar{\beta}D^M(p_{(Simult)}^{(Nash),*}-p_{(Alt)}^*)}{(1+\alpha_{(Alt)}^{(Nash)}\bar{\beta})D^M(p_{(Alt)^*}-\bar{\beta}p_{(Simult)}^{(Nash),*})}$$

$$p_{(Simult)}^{(Nash),*}\left(\tau(1+\bar{\beta})\frac{D^M}{2}+(1-\tau)(1+\bar{\beta})(D^M-\underline{D}^M)+\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha_{(Alt)}^{(Nash)})^2D^M}{2(1-(\alpha_{(Alt)}^{(Nash)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^3\alpha_{(Alt)}^{(Nash)}D^M}{2(1-\bar{\beta})}\right)$$

$$=(1-\tau)(1+\bar{\beta})\phi(D^M-\underline{D}^M)+p_{(Alt)}^*\left(\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha_{(Alt)}^{(Nash)})^2D^M}{2(1-(\alpha_{(Alt)}^{(Nash)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^2\alpha_{(Alt)}^{(Nash)}D^M}{2(1-\bar{\beta})}\right).$$

where the $GFT_{(Alt)}$ functions are specified above.

The first equation yields:

$$p_{(Alt)}^*=\frac{p_{(Simult)}^{(Nash),*}D^M\bar{\beta}(2-\tau(1-\alpha_{(Alt)}^{(Nash)}\bar{\beta}))+2\phi(1-\bar{\beta})(D^M-\underline{D}^M)(1-\tau)}{D^M\left(\tau+\bar{\beta}(2-\tau(2-\alpha_{(Alt)}^{(Nash)}))\right)+2(1-\bar{\beta})(D^M-\underline{D}^M)(1-\tau)}.$$

And the second equation yields:

$$p_{(Simult)}^{(Nash),*}=\frac{p_{(Alt)}^*\left(\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha_{(Alt)}^{(Nash)})^2D^M}{2(1-(\alpha_{(Alt)}^{(Nash)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^2\alpha_{(Alt)}^{(Nash)}D^M}{2(1-\bar{\beta})}\right)+(1-\tau)(1+\bar{\beta})\phi(D^M-\underline{D}^M)}{\tau(1+\bar{\beta})\frac{D^M}{2}+(1-\tau)(1+\bar{\beta})(D^M-\underline{D}^M)+\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha_{(Alt)}^{(Nash)})^2D^M}{2(1-(\alpha_{(Alt)}^{(Nash)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^3\alpha_{(Alt)}^{(Nash)}D^M}{2(1-\bar{\beta})}}$$

Note that when $\bar{\beta}=0$, this yields $\frac{(1-\tau)\phi(D^M-\underline{D}^M)}{\tau D^M/2+(1-\tau)(D^M-\underline{D}^M)}$; when D^M, $\underline{D}^M$, and ϕ are taken from the empirical example, this yields the static equilibrium $\$10,000\frac{1-\tau}{1+2\tau}$. When $\tau=0$ so that the hospitals have all bargaining power, this reduces to ϕ, so that the hospitals keep all premium revenue, even if $\bar{\beta}>0$.

Plugging the first result into the second result, I obtain:

$$p^{(Nash),*}_{(Simult)} = \frac{\frac{2\phi(1-\bar{\beta})(D^M-\underline{D}^M)(1-\tau)}{D^M\left(\tau+\bar{\beta}(2-\tau(2-\alpha^{(Nash)}_{(Alt)}))\right)+2(1-\bar{\beta})(D^M-\underline{D}^M)(1-\tau)}\frac{\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2D^M}{2(1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^2\alpha^{(Nash)}_{(Alt)}D^M}{2(1-\bar{\beta})}}{\tau(1+\bar{\beta})\frac{\underline{D}^M}{2}+(1-\tau)(1+\bar{\beta})(D^M-\underline{D}^M)+\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2D^M}{2(1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^3\alpha^{(Nash)}_{(Alt)}D^M}{2(1-\bar{\beta})}}}{1-\frac{D^M\bar{\beta}(2-\tau(1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}))}{D^M\left(\tau+\bar{\beta}(2-\tau(2-\alpha^{(Nash)}_{(Alt)}))\right)+2(1-\bar{\beta})(D^M-\underline{D}^M)(1-\tau)}\frac{\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2D^M}{2(1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^2\alpha^{(Nash)}_{(Alt)}D^M}{2(1-\bar{\beta})}}{\tau(1+\bar{\beta})\frac{\underline{D}^M}{2}+(1-\tau)(1+\bar{\beta})(D^M-\underline{D}^M)+\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2D^M}{2(1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^3\alpha^{(Nash)}_{(Alt)}D^M}{2(1-\bar{\beta})}}}$$
$$+\frac{\frac{(1-\tau)(1+\bar{\beta})\phi(D^M-\underline{D}^M)}{\tau(1+\bar{\beta})\frac{\underline{D}^M}{2}+(1-\tau)(1+\bar{\beta})(D^M-\underline{D}^M)+\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2D^M}{2(1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^3\alpha^{(Nash)}_{(Alt)}D^M}{2(1-\bar{\beta})}}}{1-\frac{D^M\bar{\beta}(2-\tau(1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}))}{D^M\left(\tau+\bar{\beta}(2-\tau(2-\alpha^{(Nash)}_{(Alt)}))\right)+2(1-\bar{\beta})(D^M-\underline{D}^M)(1-\tau)}\frac{\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2D^M}{2(1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^2\alpha^{(Nash)}_{(Alt)}D^M}{2(1-\bar{\beta})}}{\tau(1+\bar{\beta})\frac{\underline{D}^M}{2}+(1-\tau)(1+\bar{\beta})(D^M-\underline{D}^M)+\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2D^M}{2(1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^3\alpha^{(Nash)}_{(Alt)}D^M}{2(1-\bar{\beta})}}}.$$

Define

$$A \equiv \left(\tau(1+\bar{\beta})\frac{D^M}{2}+(1-\tau)(1+\bar{\beta})(D^M-\underline{D}^M)+\frac{\tau\bar{\beta}^2(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2D^M}{2(1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2)}-\frac{\tau\bar{\beta}^3\alpha^{(Nash)}_{(Alt)}D^M}{2(1-\bar{\beta})}\right)$$

$$B \equiv \frac{\frac{\tau\bar{\beta}^2D^M}{2}\left(\frac{(1+\bar{\beta})(\alpha^{(Nash)}_{(Alt)})^2}{1-(\alpha^{(Nash)}_{(Alt)}\bar{\beta})^2}-\frac{\alpha^{(Nash)}_{(Alt)}}{1-\bar{\beta}}\right)}{D^M\left(\tau+\bar{\beta}(2-\tau(2-\alpha^{(Nash)}_{(Alt)}))\right)+2(1-\bar{\beta})(D^M-\underline{D}^M)(1-\tau)}.$$

Then I obtain:

$$p^{(Nash),*}_{(Simult)} = (1-\tau)\phi(D^M-\underline{D}^M)\frac{B*2(1-\bar{\beta})+(1+\bar{\beta})}{A-B*\left(D^M\bar{\beta}(2-\tau(1-\alpha^{(Nash)}_{(Alt)}\bar{\beta}))\right)},$$

which is the desired equality.

□

Proof of Proposition 1. The claim is that for an example of the form claimed in Proposition 1, there is no equilibrium with $p^{(Nash),*}_{(Simult)} = p^*_{(Static)}$.

Proof by contradiction. Suppose not, and $p^{(Nash),*}_{(Simult)} = p^*_{(Static)}$ in some simultaneous bargaining period t_0. To simplify the notation, I assume that the equilibrium is symmetric.

First, I show that it is not an equilibrium response to exclude on July 1. Consider the subgame in which hospital i disagreed on January 1 of a given year. I show that there is a strictly Pareto-improving agreement relative to disagreeing twice. Consider the value of agreeing at the price $p^*_{(Simult)}$. The insurer's expected net present value is at least the value if they disagree with hospital $5-i$ on January 1 of the next year, and then follow the equilibrium (by Markov strategies) response of simultaneous-agreement equilibrium on July

1 of next year. By agreeing at $p^*_{(Simult)}$ on July 1, the insurer gains relative to disagreeing a second time are at least

$$\begin{aligned}
&\phi(D^M - \underline{D}^M) + p^*_{(Simult)}(\underline{D}^M - D^M/2)/(1+\beta^{1/2}) \\
-\ &\beta^{1/2}\left(\phi(D^M - \underline{D}^M) + p^*_{(Simult)}(\underline{D}^M - D^M/2)\right)/(1+\beta^{1/2}) \\
=&(1-\tilde{\beta})\left(\phi(D^M - \underline{D}^M) + p^*_{(Simult)}(\underline{D}^M - D^M/2)\right)/(1+\beta^{1/2}) > 0.
\end{aligned}$$

Similarly, the hospital's value of agreement is at least as high as the value of agreeing on July 1 and then disagreeing next July 1 and returning to equilibrium, in which case they provide $D^M/(2(1+\beta^{1/2}))$ units of care in the first half-period and at least $\min\{D^M/2, \underline{D}^M\}/(1+\beta^{1/2}) = D^M/(2(1+\beta^{1/2}))$ units of care in the second half-period at $p^*_{(Simult)}$; with disagreement, they provide those units of care at $p^*_{(Simult)}$ in the second and third half-periods. Because $\beta \in (0,1)$, $\beta^{1/2} \in (0,1)$ and agreement at $p^*_{(Simult)}$ strictly Pareto-dominates disagreement. By differentiability of prices responses, there is a Nash bargaining solution and it is not an equilibrium response to exclude on July 1.

Next, write $p^*_{(Alt)}$ for an alternating-period equilibrium price, which exists by differentiability of the price response and the weakly increasing nature of gains from trade in previously formed other-firm prices. To save on notation, I abbreviate *Static* and *Simult* to S in the remainder of this proof.

It is useful to write $V^H_{(Alt)}(p_i \mid p_{3-i})$ and $V^H_{(S)}(p_i \mid p_{3-i})$ for the value to the hospital of agreeing to a price p_i, when the other hospital agrees simultaneous or in the last period, respectively, to a price p_{3-i}. Similarly let $V(\cdot)^M(p_i \mid p_{3-i})$ be the value to the insurer. Let $p^*_{(Alt)}(p_{3-i})$ be the alternating period price response function. It is useful to write $p^{(1)}_{(Alt)} = p^*_{(Alt)}(p^*_{(S)})$ and then inductively write $p^{(t+1}_{(Alt)}) = p^*_{(Alt)}(p^{(t)}_{(Alt)})$. Then:

$$\begin{aligned}
V^M_{(Alt)}(p_i \mid p_{3-i}) &= \sum_{u=0}^{\infty} \bar{\beta}^u D^M \phi - V^H_{(Alt)}(p_i \mid p_{3-i}) - p_{3-i} * D^M/2 - \bar{\beta} V^H_{(Alt)}(p^*_{(Alt)}(p_i) \mid p_i) \\
&= \sum_{u=0}^{\infty} \bar{\beta}^u D^M \phi - V^H_{(Alt)}(p_i \mid p^*_{(S)}) - p_{3-i} * D^M/2 - \bar{\beta} V^H_{(Alt)}(p^*_{(Alt)}(p_i) \mid p^*_{(S)}) \\
V^M_{(S)}(p_i \mid p_{3-i}) &= \sum_{u=0}^{\infty} \bar{\beta}^u D^M \phi - (1+\bar{\beta}) D^M \frac{p_i + p_{3-i}}{2} - \bar{\beta}^2 V^M_{(S)}(p^*_{(S)} \mid p^*_{(S)}) \\
&= \sum_{u=0}^{\infty} \bar{\beta}^u D^M - V^H_{(S)}(p_i \mid p_{3-i}) - V^H_{(S)}(p_{3-i} \mid p_i) \\
&= \sum_{u=0}^{\infty} \bar{\beta}^u D^M \phi - V^H_{(S)}(p_i \mid p^*_{(S)}) - V^H_{(S)}(p_{3-i} \mid p^*_{(S)}).
\end{aligned}$$

The value of disagreeing in simultaneous negotiations at anticipated price p_{3-i} is $V^{D,H}_{(S)}(p_{3-i}) \equiv \bar{\beta} V^{H}_{(Alt)}(p^*_{(Alt)}(p_{3-i}) \mid p_{(3-i)})$ to the hospital and is $V^{D,M}_{(S)}(p_{3-i}) \equiv \underline{D}^M(\phi - p_{3-i}) + \bar{\beta} V^{M}_{(Alt)}(p^*_{(Alt)}(p_{3-i}) \mid p_{3-i})$ to the insurer. The value of disagreeing in alternating periods at anticipated p_{3-i} is $V^{D,H}_{(Alt)}(p_{3-i}) \equiv \bar{\beta} V^{H}_{(S)}(p^*_{(S)} \mid p_{(S)^*})$ to to hospital and $V^{D,M}_{(Alt)}(p_{3-i}) \equiv \underline{D}^M(\phi - p_{3-i}) + \bar{\beta} V^{M}_{(S)}(p^*_{(S)} \mid p_{(S)^*})$ to the insurer.

Because simultaneous bargaining is TU, simultaneous Nash-in-Nash bargaining solution at $p^*_{(S)}$ requires:

$$\begin{aligned}
&(1-\tau)\left(V^M_{(S)}(p^*_{(S)} \mid p^*_{(S)}) - V^{D,M}_{(S)}(p^*_{(S)})\right) \\
=&\tau\left(V^H_{(S)}(p^*_{(S)} \mid p^*_{(S)}) - V^{D,H}_{(S)}(p^*_{(S)})\right) \\
\Leftrightarrow\ &(1-\tau)\left(V^M_{(S)}(p^*_{(S)} \mid p^*_{(S)}) - \underline{D}^M(\phi - p^*_{(S)}) - \bar{\beta} V^M_{(Alt)}(p^{(1)}_{(Alt)} \mid p^*_{(S)})\right) \\
=&\tau\left(V^H_{(S)}(p^*_{(S)} \mid p^*_{(S)}) - \bar{\beta} V^H_{(Alt)}(p^{(1)}_{(Alt)} \mid p^*_{(S)})\right).
\end{aligned}$$

Further, by construction of the static Nash-in-Nash solution, the static agreement splits gains from trade relative to two split-in-half disagreements, i.e.:

$$\begin{aligned}
&(1-\tau)\left(V^M_{(S)}(p^*_{(S)} \mid p^*_{(S)}) - (1+\bar{\beta})\underline{D}^M(\phi - p^*_{(S)}) - \bar{\beta}^2 V^M_{(S)}(p^*_{(S)} \mid p^*_{(S)})\right) \\
=&\tau\left(V^H_{(S)}(p^*_{(S)} \mid p^*_{(S)}) - \bar{\beta}^2 V^H_{(S)}(p^*_{(S)} \mid p^*_{(S)})\right) \\
\Leftrightarrow\ &(1-\tau)\left(V^M_{(S)}(p^*_{(S)} \mid p^*_{(S)}) - \underline{D}^M(\phi - p^*_{(S)}) - \bar{\beta} V^{D,M}_{(Alt)}(p^*_{(S)})\right) \\
=&\tau\left(V^H_{(S)}(p^*_{(S)} \mid p^*_{(S)}) - \bar{\beta} V^{D,H}_{(Alt)}(p^*_{(S)})\right).
\end{aligned}$$

Therefore, by subtraction and multiplying by the real number $\bar{\beta}^{-1} \in (1, \infty)$

$$(1-\tau)\left(V^M_{(Alt)}(p^{(1)}_{(Alt)} \mid p^*_{(S)}) - V^{D,M}_{(Alt)}(p^*_{(S)})\right) = \tau\left(V^H_{(Alt)}(p^{(1)}_{(Alt)} \mid p^*_{(S)}) - V^{D,H}_{(Alt)}(p^*_{(S)})\right).$$

But by Nash-in-Nash first-order conditions:

$$\begin{aligned}
\frac{V^M_{(Alt)}(p^{(1)}_{(Alt)} \mid p^*_{(S)}) - V^{D,M}_{(Alt)}(p^*_{(S)})}{V^H_{(Alt)}(p^{(1)}_{(Alt)} \mid p^*_{(S)}) - V^{D,H}_{(Alt)}(p^*_{(S)})} &= \frac{\tau}{1-\tau} \frac{-\frac{\partial V^M_{(Alt)}(p|p^*_{(S)})}{\partial p}\Big|_{p=p^{(1)}_{(Alt)}}}{\frac{\partial V^H_{(Alt)}(p|p^*_{(S)})}{\partial p}\Big|_{p=p^{(1)}_{(Alt)}}} \\
&= \frac{\tau}{1-\tau}\left(1 + \bar{\beta}\frac{\frac{\partial V^H_{(Alt)}(p^*_{(Alt)}(p)|p^*_{(Alt)})}{\partial p}\Big|_{p=p^{(1)}_{(Alt)}}}{\frac{\partial V^H_{(Alt)}(p|p^*_{(S)})}{\partial p}\Big|_{p=p^{(1)}_{(Alt)}}}\right),
\end{aligned}$$

so that

$$\frac{\partial V^H_{(Alt)}(p^*_{(Alt)}(p) \mid p^*_{(Alt)})}{\partial p} \mid_{p=p^{(1)}_{(Alt)}} = 0.$$

But by assumption, $\frac{\partial V^H_{(Alt)}(p|p^*_{(Alt)})}{\partial p} \mid_{p=p^*_{(Alt)}\left(p^{(1)}_{(Alt)}\right)} > 0$, so that it must be the case that $\frac{\partial p^*_{(Alt)}(p)}{\partial p} \mid_{p=p^{(1)}_{(Alt)}} = 0$. As a result of this equality and chain rule:

$$\frac{\partial V^H_{(Alt)}(p^*_{(Alt)}(\bar{p}) \mid p^*_{(S)})}{\partial p} \mid_{\bar{p}=p^{(1)}_{(Alt)}} = \frac{\partial V^H_{(Alt)}(\bar{p} \mid p^*_{(S)})}{\partial p} \mid_{\bar{p}=p^{(2)}_{(Alt)}} * \frac{\partial p^*_{(Alt)}(p)}{\partial p} \mid_{p=p^{(1)}_{(Alt)}} = 0.$$

More generally, the alternating price response $p^*_{(Alt)}(p)$ is assumed to be differentiable around $p^{(t)}_{(Alt)}$ and generally solves:

$$\frac{V^M_{(Alt)}(p^*_{(Alt)}(p) \mid p) - V^{(D,M)}_{(Alt)}(p)}{V^H_{(Alt)}(p^*_{(Alt)}(p) \mid p) - V^{(D,H)}_{(Alt)}(p)} = \frac{\tau}{1-\tau}\left(1 + \bar{\beta}\frac{\frac{\partial V^H_{(Alt)}(p^*_{(Alt)}(\bar{p})|p^*_{(S)})}{\partial p} \mid_{\bar{p}=p^*_{(Alt)}(p)}}{\frac{\partial V^H_{(Alt)}(\bar{p}|p^*_{(S)})}{\partial p} \mid_{\bar{p}=p^*_{(Alt)}(p)}}\right) = \text{“}\rho(p^*_{(Alt)}(p))\text{”}.$$

Substituting in $p^*_{(S)}$ where possible, I obtain:

$$\begin{aligned}
&V^M_{(Alt)}(p^*_{(S)} \mid p^*_{(S)}) - V^{(D,M)}_{(Alt)}(p^*_{(S)}) + \underline{D}^M(p - p^*_{(S)}) \\
&+V^H_{(Alt)}(p^*_{(S)} \mid p^*_{(S)}) - V^H_{(Alt)}(p^*_{(Alt)}(p) \mid p^*_{(S)}) \\
&+\bar{\beta}\left(V^H_{(Alt)}(p^*_{(S)} \mid p^*_{(S)}) - V^H_{(Alt)}(p^*_{(Alt)}(p^*_{(Alt)}(p)) \mid p^*_{(S)})\right) \\
&=\rho(p^*_{(Alt)}(p))\left(V^H_{(Alt)}(p^*_{(Alt)}(p) \mid p_{(S)^*}) - V^{(D,H)}_{(Alt)}(p^*_{(S)})\right).
\end{aligned}$$

I totally differentiate with respect to p at $p = p^{(1)}_{(Alt)}$ (feasible by assumption of differentiable responses on the equilibrium path) and apply $\frac{\partial p^*_{(Alt)}(p)}{\partial p} \mid_{p=p^{(1)}_{(Alt)}} = 0$. The only term is that is non-zero is $\underline{D}^M(p - p^*_{(S)})$. Therefore $\underline{D}^M = 0$. But $\underline{D}^M > 0$. Contradiction. □

Proof of Lemma 1. Because the game is U, the MVR is one. More generally than Equation (2), the gains from trade must satisfy:

$$\begin{aligned}
\tau_{ij} E\left[GFT_i\left(\hat{\mathbb{C}}^{(Nash)}_{ijt}\right)\right] &= (1-\tau_{ij})\tau_{ij} E\left[GFT_j\left(\hat{\mathbb{C}}^{(Nash)}_{ijt}\right)\right] \\
\tau_{ij}\left([\Delta_{it}\pi_{it}] + [\Delta_{ij}T_{it,-j}] + p_{ijt}D_{ijt} - r_i\right) &= (1-\tau_{ij})\left([\Delta_{it}\pi_{jt}] + [\Delta_{ij}T_{jt,-i}] - p_{ijt}D_{ijt} - r_j\right) \\
p_{ijt}D_{ijt} &= -\tau_{ij}\left([\Delta_{it}\pi_{it}] + [\Delta_{ij}T_{it,-j}] + p_{ijt}D_{ijt}\right) \\
&\quad + (1-\tau_{ij})\left([\Delta_{it}\pi_{jt}] + [\Delta_{ij}T_{jt,-i}]\right) \\
&\quad + \tau_{ij}r_i - (1-\tau_{ij})r_j.
\end{aligned}$$

□

Proof of Corollary 1. Since bargaining is TU, the Nash-in-Nash and Nash-in-Kalai solutions coincide and this result follows by Lemma 1. □

Details in proof of Theorem 1. The iterative expansion is:

$$\begin{aligned}
V^{(2,I)}_{i,(d)^D} &= E_{t^D_d}\left[\sum_{t^D_d \le t_s < t^D_{d+1}} \beta^{\frac{t_s - t^D_d}{m}}\left(\hat{\pi}^D_{it_s} + \hat{T}^D_{it_s}\right) + \beta^{\frac{t^D_{d+1} - t^D_d}{m}} V^{(2,E)}_{i,(d+1)^D}\left(\left(\hat{\mathbb{C}}_{t^D_{d+1},-ij}, \hat{\mathbb{C}}_{t^D_{d+1},ij}\right)\right)\right] \\
&= E_{t^D_d}\left[\sum_{t^D_d \le t_s < t^D_{d+1}} \beta^{\frac{t_s - t^D_d}{m}}\left(\hat{\pi}^D_{it_s} + \hat{T}^D_{it_s}\right) + \beta^{\frac{t^D_{d+1} - t^D_d}{m}}\left(V^{(2,I)}_{i,(d+1)^D} + (1-\tau_{ij})GFT^{(2,E)}_{(d+1)^D}\right)\right] \\
&= E_{t^D_d}\left[\sum_{t_s \ge t^D_d} \beta^{\frac{t_s - t^D_d}{m}}\left(\hat{\pi}^D_{it_s} + \hat{T}^D_{it_s}\right) + \sum_{h>0}(1-\tau_{ij})\beta^{\frac{t^D_{d+h} - t^D_d}{m}} GFT^{(2,E)}_{(d+h)^D} + \lim_{h\to\infty}\beta^{\frac{t^D_{d+h} - t^D_d}{m}} V^{(2,I)}_{i,(d+h)^D}\right] \\
&= E_{t^D_d}\left[\sum_{t_s \ge t^D_d} \beta^{\frac{t_s - t^D_d}{m}}\left(\hat{\pi}^D_{it_s} + \hat{T}^D_{it_s}\right) + \sum_{h>0}(1-\tau_{ij})\beta^{\frac{t^D_{d+h} - t^D_d}{m}} GFT^{(2,E)}_{(d+h)^D}\right],
\end{aligned} \tag{10}$$

where the final line uses $\lim_{h\to\infty}\beta^{\frac{t^D_{d+h} - t^D_d}{m}} V^{(2,I)}_{i,(d+h)^D} = 0$ under transversality conditions.

The gain-from-trade expansion is:

$$\begin{aligned}
&(\tau_{ij})E_{t_0}\left[\sum_{t_0 \le t_s \le t^*} \beta^{\frac{t_s - t_0}{m}}\left([\Delta_{ij}\pi_{it_s}] + [\Delta_{ij}T_{it_s,-j}] + p^*_{ijt_s}D_{ijt_s}\right) - r_i\right] \\
&+(\tau_{ij})E_{t_0}\left[\sum_{t_0 \le t_s} \beta^{\frac{t_s - t_0}{m}}\left(\hat{\pi}^A_{it_s} + \hat{T}^A_{it_s} - \hat{\pi}^D_{it_s} - \hat{T}^D_{it_s}\right)\right] \\
&+(\tau_{ij})E_{t_0}\left[\sum_{h>0}(1-\tau_{ij})\beta^{\frac{t^A_{d+h} - t^A_d}{m}} GFT^{(2,E)}_{(d+h)^A} - (1-\tau_{ij})\beta^{\frac{t^D_{d+h} - t^D_d}{m}} GFT^{(2,E)}_{(d+h)^D}\right] \\
&=(1-\tau_{ij})E_{t_0}\left[\sum_{t_0 \le t_s \le t^*} \beta^{\frac{t_s - t_0}{m}}\left([\Delta_{ij}\pi_{jt_s}] + [\Delta_{ij}T_{jt_s,-i}] + p_{jit_s}D_{ijt_s}\right) - r_i\right] \\
&+(1-\tau_{ij})E_{t_0}\left[\sum_{t_0 \le t_s} \beta^{\frac{t_s - t_0}{m}}\left(\hat{\pi}^A_{jt_s} + \hat{T}^A_{jt_s} - \hat{\pi}^D_{jt_s} - \hat{T}^D_{jt_s}\right)\right] \\
&+(1-\tau_{ij})E_{t_0}\left[\sum_{h>0}(\tau_{ij})\beta^{\frac{t^A_{d+h} - t^A_d}{m}} GFT^{(2,E)}_{(d+h)^A} - (\tau_{ij})\beta^{\frac{t^D_{d+h} - t^D_d}{m}} GFT^{(2,E)}_{(d+h)^D}\right].
\end{aligned} \tag{11}$$

□

Proof of Corollary 3. The path of profit functions, contracts, and transfers are the same on the A and D path. Thus, $\text{Pay}_{IRT} = 0$. □

Proof of Corollary 2. Consider an equilibrium bargain at a period t_0. In this example, contracts and payments on the A and D path coincide: in the first year, the paths coincide because the other hospital's contract is fixed; after the first year, the paths correspond to identical subgames by Markov strategies (Assumption 1). In addition, $r_i = r_j = 0$ for all negotiations. Therefore $\text{Pay}_{IRT} = \text{Pay}_{NC} = 0$.

As a result, by Theorem 1, the split-in-half simultaneous equilibrium contract price $p_{(Simult)}^{(Nash),*}$ satisfies:

$$E_{t_0}\left[(1+\sqrt{\beta})\frac{D^M/2 p_{(Simult)}^{(Nash),*}}{1+\sqrt{\beta}}\right] = E_{t_0}\left[(1+\sqrt{\beta})\frac{-\tau_{ij}*0+(1-\tau_{ij})\left(\phi(D^M-\underline{D}^M)+(\underline{D}^M-D^M/2)p_{(Simult)}^{(Nash),*}\right)}{1+\sqrt{\beta}}\right].$$

This is the same equality as the static Nash-in-Nash equilibrium, so that the equilibrium is as described. □

Proof of Corollary 4. I propose strategies in the split-in-half subgame constructed as follows.

The first step is the constraining step. In any future subsubgame that corresponds (after appropriate change of notation) to a subsubgame of the original game, follow the strategies (with appropriate change of notations) in the original game.

The second step fills in the remaining strategies. After reaching the first subsubgame that does correspond to a subsubgame of the original game, follow a set of Nash-in-Kalai strategies of the type assumed to exist in Assumption 2 under the constraints from the first step.

It remains to show that for every choice in the first step is an equilibrium step, and every bargain is a Kalai proportional bargaining. I abstract from stage 1 and focus on the bargaining stage 2. It remains to show that every bargain is Pareto-efficient, splits gains from trade proportionally to ij bargaining weights, and is Pareto-improving relative to the value of disagreement.

Every bargain is Pareto-efficient among the class of feasible contracts. By construction, every ij choice of non-null contract generates the same expected net present value profit as in the original game. Therefore the space of feasible values with non-null agreements is unchanged, and the proposed contracts are Pareto-efficient.

Every bargain splits gains from trade proportionally to bargaining weights. This follows by the Kalai proportional step-by-step property, although it is so far possible for the sides to split negative gains from trade proportionally to bargaining weights.

Finally, I show that the proposed strategies predict weakly positive gains from trade. Consider a pair ij in some subsubgame t_s, and consider the equivalent strategies when not restricting the pair and following the previous strategies in every subsequent subsubsubgame $t_r > t_s$. By the full network assumption, the new Nash-in-Kalai equilibrium produces some distribution over non-null ijt_s contracts that produce weakly positive gains from trade, are Pareto-efficient, and split gains from trade proportionally to bargaining weights. But the originally proposed strategies for ijt_s produce weakly positive gains from trade, are Pareto-efficient, and split gains from trade proportionally. By Pareto efficiency and the fact that the value functions for a given ijt_s contract are the same due to the restriction on t_r strategies, the value of the two strategies to the pair must be the same, and the proposed ijt_s contract produces weakly positive gains from trade. □