

HEALTH INSURANCE PRE-EXCHANGES
A MARKET DESIGN APPROACH TO
INSURER-PROVIDER NETWORK FORMATION*

PART I: THEORY[†]

KANE SWEENEY [‡]

Abstract

I propose a centralized marketplace, a *pre-exchange*, for health insurer-provider network formation and examine its welfare impact. Current network formation is highly decentralized and contracting is bilateral within each insurer-provider pair. The key innovation of the pre-exchange is the introduction of explicitly multilateral contracting. In particular, providers set prices that depend both on the identity of the insurer as well as the identities of the other hospitals in the insurer's network. I refer to this as *network-specific pricing*. I argue that the introduction of network-specific pricing enhances efficiency and dramatically lowers equilibrium health care prices. I first introduce a theoretical model that illustrates the general tension that results from the key features of the current market structure: network access, bilateral contracting and per-patient pricing. I then demonstrate how the pre-exchange results in lower equilibrium provider prices. The main theoretical insight is that bilateral pricing produces a strategic externality that providers "price out" with network-specific prices. In a companion paper, I evaluate the impact of introducing a pre-exchange using recent data from Massachusetts.

*December 1, 2012

[†]The companion paper, PART II: EMPIRICS, can be found at my website listed below.

[‡]k-sweeney@kellogg.northwestern.edu, <http://www.kellogg.northwestern.edu/faculty/sweeney/>, Northwestern University, Kellogg School of Management. I would like to thank Rakesh Vohra, David Dranove, Michael Whinston, Mallesh Pai, Chris Ody, Umberto Garfagnini, Aviv Nevo, Jeff Ely, Alex Frankel, Pablo Montagnes.

Contents

Introduction	1
1 A Simple Example	3
2 Related Literature	12
3 Brief Overview of the Health Care Market	16
Theory	18
4 Model: Current Market Structure (Bilateral Bidding)	18
5 Market Design: Pre-Exchanges (Network-Specific Bidding)	30
Discussion	36
6 Conclusion	37
Appendix – Theory	41

Introduction

Christie Romer in the *New York Times*¹:

Here's a frightening thought: Despite the recent Supreme Court decision upholding the Affordable Care Act, serious work on more health care legislation is still needed. [...] Even with the law, health care spending is still projected to rise rapidly over coming decades, so more steps to contain costs will have to be taken.

OECD 2011 Report²:

Overall, the evidence suggests that prices for health services and goods are substantially higher in the United States than elsewhere. This is an important cause of higher health spending in the United States.

Health care spending in the United States is higher than any other nation, and the rate of spending increases has been widely acknowledged as unsustainable. As such, “cost containment” has received considerable attention by policy makers, and recent legislation calls for sweeping reform at both the state and national levels. The pro-competitive aspects of these reforms typically focus on the incentives of health *insurers* through the introduction of health insurance exchanges.³

In this paper, I propose a novel mechanism, a health insurance *pre-exchange*, that focuses on the incentives of health care *providers*. I first construct a model to explain why the current market structure encourages excessively high hospital prices (and profits). I then illustrate how a pre-exchange addresses the concerns by fostering increased hospital price competition without harming efficiency.

The key innovation of the pre-exchange is the introduction of explicitly multilateral contracting between the hospitals and insurers. Bilateral contracts are standard in current network formation, as each insurer-hospital pair meets independently to discuss terms and finalize agreements. One

¹Christie Romer’s *New York Times* article “Only the First Step in Containing Health Care Costs”
<http://www.nytimes.com/2012/07/22/business/health-care-law-and-cost-containment-economic-view.html>

²OECD’s Health at a Glance 2011: “Why is Health Spending in the United States So High?”
<http://www.oecd.org/unitedstates/49084355.pdf>

³The private health insurance market is often cited as a classic example of a market failure due to adverse selection, but the failure is effectively separate from the concerns regarding cost containment. Provisions in the recent US reform (the Patient Protection and Affordable Care Act) address these selection issues by simultaneously mandating insurance purchases and outlawing various types of price discrimination and screening. The list of reforms is, however, much broader than simply insurance exchanges, and some of the other issues are mentioned below.

message of this paper is that these bilateral contracts are not rich enough for hospitals to fully “price out” the relevant network externalities. More precisely: the contracts *are* rich enough to price out the network externalities, but *only* in the absence of strategic behavior.

The central message of this paper, however, is stronger than that: after accounting for strategic behavior, the structure of bilateral contracting produces a strong strategic complementarity among the hospital prices. The best response behavior leads to equilibrium hospital profits well in excess of their contribution to society, and the overall “cost” of health care is excessive.

Pre-exchanges address precisely this issue. In a pre-exchange, hospitals submit different prices for the various networks that include them. I refer to this as *network-specific pricing*. In the simplest terms, hospitals are bidding to exclude other hospitals from an insurer’s network. Equivalently, hospitals are offering the insurer an explicit per-patient discount for dropping a rival or substitute hospital. While this may initially seem counterproductive, permitting the hospitals to target their bids to specific networks entirely decouples the problematic strategic complementarity. This feedback of best responses never starts, and equilibrium prices remain at the desired levels.

Before presenting the formal model, the next section provides some intuition for my results in the context of a simple example. I then review the literature, and describe some institutional details that motivate the the structure and assumptions of the formal model. The formal model allows me to conclude that hospital prices (and profits) are objectively “too high.” With this result in mind, I then outline pre-exchanges and demonstrate how network-specific pricing contains costs without reducing welfare.

In a companion paper (available online), I introduce an empirical framework. The goal is to take these qualitative theoretical insights and quantify the impact in an economically meaningful setting. Specifically, I use recent health care data from Massachusetts to estimate patient demand for insurance in a discrete choice framework. I then simulate equilibrium hospital prices in two settings: (1) a baseline industry model meant to describe the current market structure and (2) the baseline model augmented with a pre-exchange and network-specific pricing.

Before proceeding to the example, I want to emphasize that the current paper addresses pricing incentive issues at the network formation stage, an area untreated by other reform efforts. In fact, it is not clear a priori that the network formation process is malfunctioning in any way or contributing to the bloated cost of care. I view this as one central contribution of my work: by utilizing tools from market design and industrial organization, I am able to highlight concerns that may otherwise go unnoticed.

For instance, in the following motivating example, there will be no “obvious” problems in the baseline scenario. It is only in the light of a better solution that the problem becomes evident. In that sense, the problems associated with network formation are particularly inconspicuous, and

one goal of the current work is to demonstrate the presence and magnitude of the problem.

I defer a proper discussion of other reforms until the literature review, after the example clarifies my point. The present work should be viewed a complement to other reform efforts, and not a substitute for them.

1 A Simple Example

Consider a simple environment with a single insurer and two hospitals h_1 and h_2 . The insurer wants to secure access to a “network” of hospitals, and then offer access to the network in the downstream insurance market. A network is simply a subset of hospitals, so there are four possible networks: $\{h_1, h_2\}, \{h_1\}, \{h_2\}, \{\}$. The insurer has the following valuation v over networks:

$$v(\{h_1, h_2\}) = 10 \quad v(\{h_1\}) = v(\{h_2\}) = 9 \quad v(\{\}) = 0.$$

To motivate the values of v listed above, consider a downstream insurance market with two consumers, p_1 and p_2 . The “ p ” suggests that these consumers will be patients that need hospital treatment. The patients have the following individual valuations over hospitals:

v_{ph}	h_1	h_2
p_1	5	4
p_2	4	5

If the insurer selects the full network $\{h_1, h_2\}$, then patient p_1 visits hospital h_1 and patient p_2 visits hospital h_2 . If only $\{h_1\}$ is selected for the network, then both patients visit h_1 . Similarly, if only $\{h_2\}$ is selected, both patients visit h_2 . If the insurer’s network has no hospitals, then the patients go elsewhere. The insurer’s value for a network can then be interpreted as the revenue it can extract in this downstream insurance market; i.e. the sum of the patient values.

Each hospital has zero marginal cost and enough capacity to treat both patients. Hospitals compete to be included in the insurer’s network by setting prices, and – given the prices – the insurer selects a network. For this example, consider two possible types of price competition: (i) linear, per-patient prices and (ii) lump-sum (total revenue) prices. The former is a familiar piece-rate schedule: the hospital is paid the submitted price for each patient it treats. The latter is a two-part tariff: the hospital gets an upfront fee and then treats patients at cost, which is zero. Two-part tariffs are useful for understanding the problem at hand, and I follow the example with a discussion of the relationship between two-part tariffs and network-specific pricing.

I now describe the insurer’s network choice problem under each type of competition. In general, the insurer will select the network that yields him the highest surplus: his value v less any payments made to the hospitals.

Linear per-patient prices If the hospitals set per-patient prices (bids) of b_1 and b_2 , the insurer selects from the following network options with the corresponding surpluses:

- $\{h_1, h_2\}$ If the insurer includes both hospitals in the network, then p_1 visits h_1 and p_2 visits h_2 , and h_1 is paid b_1 while h_2 is paid b_2 . This yields a surplus of $(5 + 5) - (b_1 + b_2)$.
- $\{h_1\}$ Now both patients visit h_1 , and h_1 is paid $2b_1$ for a surplus of $(5 + 4) - 2b_1$.
- $\{h_2\}$ Now both patients visit h_2 for a surplus of $(4 + 5) - 2b_2$.

Lump-sum (total revenue) prices If the hospitals set lump-sum prices (bids) of B_1 and B_2 , the insurer selects from the following options and surpluses:

- $\{h_1, h_2\}$ Similarly p_1 visits h_1 and p_2 visits h_2 . Hospital h_1 is paid its lump sum B_1 and h_2 is paid B_2 for a surplus of $(5 + 5) - (B_1 + B_2)$.
- $\{h_1\}$ Both patients visit h_1 , but h_1 is paid its lump-sum B_1 for a surplus of $(5 + 4) - B_1$.
- $\{h_2\}$ Both patients visit h_2 , h_2 is paid B_2 for a surplus of $(4 + 5) - B_2$.

The goal of this lengthy example is to emphasize two points: (i) the difference in the insurer’s network selection across the competition formats and (ii) the effect on hospital pricing incentives. Before analyzing the insurer’s network choice decision, consider a simpler problem: when does the insurer prefer the full network $\{h_1, h_2\}$ to the h_2 -only network $\{h_2\}$? That is, when does the insurer want to add h_1 to the network?

With lump-sums, the insurer’s cost-benefit analysis is straightforward: the incremental benefit of adding h_1 is 1 ($= 10 - 9$), and the incremental cost is simply B_1 . Therefore the insurer (weakly) prefers $\{h_1, h_2\}$ to $\{h_2\}$ exactly when $B_1 \leq 1$.

Per-patient pricing is more subtle: the incremental benefit remains $1 = (10 - 9)$, but the incremental cost of adding h_1 is now $b_1 - b_2$. This reflects the fact that one patient will switch from h_2 to h_1 , and therefore the insurer will trade one payment of b_2 for one payment of b_1 ; i.e. $+b_1 - b_2$. The insurer then (weakly) prefers $\{h_1, h_2\}$ to $\{h_2\}$ exactly when the benefit outweighs the cost, $b_1 - b_2 \leq 1$ or $b_1 \leq 1 + b_2$. Since the direct benefit of adding h_1 is constant ($=1$), the threshold is increasing in b_2 because the insurer’s cost-savings are increasing in b_2 . An important corollary is that the insurer is more willing to accept higher prices from h_1 when h_2 ’s prices are higher.

Figure 1 limits attention to this simplified choice between $\{h_1, h_2\}$ and $\{h_2\}$ to illustrate precisely that point: the insurer is increasingly willing to add h_1 for higher prices by h_2 . I find that examining this simplified tradeoff is helpful for gaining intuition, but the insurer’s actual choice problem is richer than this. In particular, the insurer also has the option to select h_1 -only network, $\{h_1\}$, and


Figure 1: Insurer’s Preference between $\{h_1, h_2\}$ and $\{h_2\}$ 

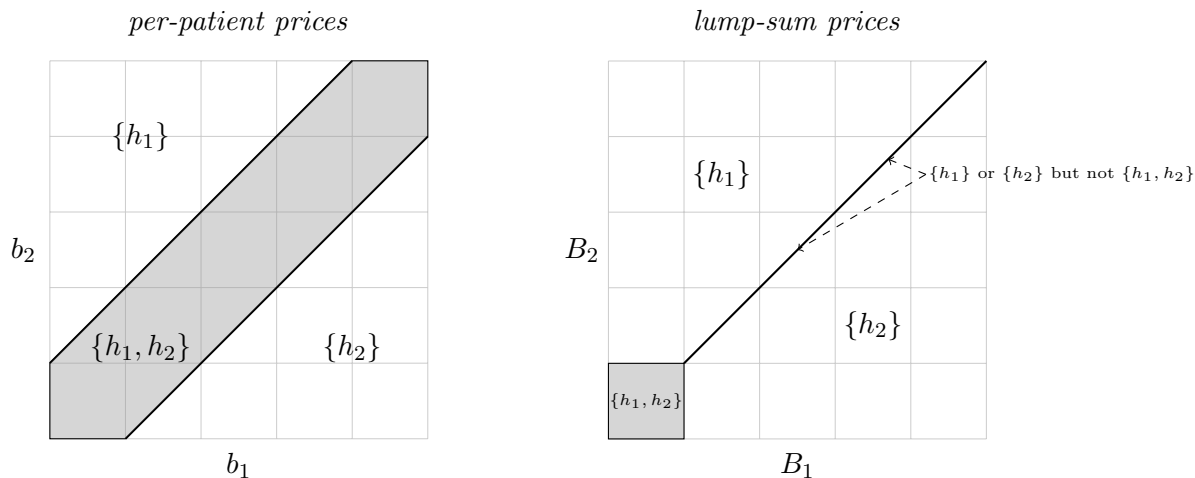
Figure 2 below is a depiction of the mapping from prices to fully optimal network choices.

The left panel of Figure 2 is the insurer’s optimal network choice when hospitals set per-patient prices. If one hospital is sufficiently cheaper than the other, only the cheap alternative is selected. If, however, the prices are close enough, $|b_1 - b_2| \leq 1$, both are selected. This $\{h_1, h_2\}$ tranche along the diagonal reflects the cost-savings argument above: if both hospitals set the same price, for instance, then adding a hospital to the network increases the insurer’s value (+1) without any associated cost savings ($b_1 - b_2 = 0$). Thus, for roughly similar prices, both are included in the network. This simple point will be key to the analysis below.

The right panel of Figure 2 depicts network choices given lump-sum prices. It has a similar general structure, with a notable difference along the diagonal. In particular, if $B_1 = B_2 > 1$, then the insurer strictly prefers either smaller network, $\{h_1\}$ or $\{h_2\}$, to the full network, $\{h_1, h_2\}$. This follows from the same logic as above: if the insurer includes, say, hospital h_2 , then it is not willing to pay $B_1 > 1$ to add h_1 . Importantly, the insurer includes both hospitals if and only if they each ask for less than their incremental value ($= 1$).

By design, per-patient prices and lump-sum prices are not perfectly comparable, but – for illustration purposes – consider a (b_1, b_2) profile above $(1, 1)$ with $b_1 \approx b_2$ and a (B_1, B_2) profile such that $(B_1, B_2) = (b_1, b_2)$. If the insurer selects the full network in both cases, the outcomes are fully equivalent: patients go to the same hospitals and everyone gets the same payoffs. But, with lump-sum prices, the insurer opts for a smaller network. In that sense, the insurer is “more selective” with lump-sum prices, i.e. if h is selected in the left panel, then h is selected in the right panel.

Figure 2: Insurer’s Network Choice



Since the full network is efficient, lump-sum pricing competition is pointwise less efficient. Yet, after accounting for strategic behavior, this “less efficient” market mechanism tends to be – in equilibrium – at least as efficient and, importantly, generates dramatically lower prices.

Given the structure of the insurer’s network choice, now consider the hospital pricing incentives. Figure 3 depicts hospital h_1 ’s best response in each environment.

With per-patient prices, the best response is discontinuous: when h_2 bids low, h_1 bids as high as possible without getting dropped, $b_1 = b_2 + 1$. If, however, b_2 is high enough, h_1 prefers to undercut to $b_1 = b_2 - 1 - \epsilon$.⁴ In this case, the insurer drops h_2 and the substantially cheaper h_1 treats both patients. Note that, for moderate prices, h_1 is extracting both his increase (+1) in value as well the cost-savings he provides (+ b_2). Importantly, the best response is “upward sloping” since h_1 ’s cost-savings is increasing in h_2 ’s price.⁵

Contrast this with lump-sum pricing. If h_2 bids low, then h_1 bids as high as possible without getting dropped, but this is now simply $B_1 = 1$ for all $B_2 \leq 1$. If $B_2 > 1$, then h_1 prefers to undercut slightly to $B_1 = B_2 - \epsilon$. For moderate bid levels, the complementarity is entirely shut down. If the symmetric best response h_2 is overlaid, the equilibrium lump-sum prices are (1, 1) and each hospital is extracting exactly its incremental value to the insurer (= 10 – 9).

There are no pure strategy Nash equilibria in the per-patient pricing environment, but the minimum *rationalizable* bid profile is (2, 2). Although this toy example has discontinuities that I would expect

⁴The dashed line indicates a “slight undercut.”

⁵Note that the discontinuity is not coming into play until hospital h_2 is bidding *three times* its marginal value to the insurer.

Figure 3: Best Responses for h_1





Figure 4: Nash Equilibrium



to be “smoothed out” in practical applications, the message is clear: prices are systematically higher with per-patient pricing than with lump-sum pricing. This is precisely because a hospital wants to extract both his increase in value as well as his cost-savings. With lump-sum pricing, adding a hospital yields no cost savings, and so the hospital simply wants to extract its added value. With per-patient pricing, adding a hospital adds value as well as saves costs, and the cost-savings are increasing in the price of the other hospital. This produces the strong strategic complementarity and the resulting systematically higher prices.

From Two-Part Tariffs to Network-Specific Pricing

Among the simplifying assumptions made in the previous example, I abstracted away from the details of the downstream competition between insurers. The focus of the example, and of the paper more generally, is capturing the strategic pricing incentives of the hospitals, and including only a single insurer isolates that effect.


The presence of downstream competition, however, adds an important, unmodeled constraint which renders two-part tariffs ill-suited to this environment. Namely, with two-part tariffs, price competition in the downstream insurance market becomes too intense. Once the insurers pay the hospitals the lump-sum fees, the *marginal* cost to an insurer of an additional enrollee is substantially below the *average* cost of an additional enrollee. As insurers compete à la Bertrand, rational insurance pricing leads to negative profits. Literature, including Milliou, Petrakis, and Vettas (2004), refer to this as a “prisoner’s dilemma,” and this is the argument in favor of per-patient pricing: it insulates the supply side from this type of unsustainable downstream competition.⁶ See Crawford and Yurukoglu (2012) for further discussion.

The desire to introduce richer contracting while maintaining per-patient pricing leads to the innovation of this paper: *network-specific pricing*. The impact of network-specific pricing can be thought of in two related ways: (1) synthesizing the desirable effects of two-part tariffs (while still satisfying the unmodeled constraint presented by downstream competition: per-patient pricing) or (2) incentivizing more aggressive pricing by making “undercutting” more attractive; i.e. lower risk, higher reward. For emphasis, I sketch the idea of each below.

To illustrate the synthesis, reconsider the example above with network-specific prices: each hospital posts two prices, one for each network that includes it. Specifically, hospital h_1 now sets $b_1(\{h_1, h_2\})$ and $b_1(\{h_1\})$. The interpretation is that $b_1(\{h_1, h_2\})$ is the per-patient price the insurer pays h_1 if the full network is chosen and $b_1(\{h_1\})$ is the per-patient price if only $\{h_1\}$ is chosen. Informally,

⁶There are other institutional reasons why contracts are linear: insurance contracts typically require the patient to pay a percentage of the (per-patient) hospital fee. Another important concern regarding two-part tariffs – that is outside the model – is moral hazard by the hospitals: once the insurers have paid them the fixed fee, hospitals have very low incentives for choosing high effort.

Figure 5: Insurer's Network Choice



think of the hospital as typically offering a per-patient discount if the insurer excludes h_2 so that $b_1(\{h_1\}) \leq b_1(\{h_1, h_2\})$. Similarly, hospital h_2 sets $b_2(\{h_1, h_2\})$ and $b_2(\{h_2\})$.

Network-specific prices If the hospitals set network-specific prices (bids) of $b_1(\{h_1\})$, $b_1(\{h_1, h_2\})$, and $b_2(\{h_2\})$, $b_2(\{h_1, h_2\})$, then the insurer selects from the following options and surpluses:

- $\{h_1, h_2\}$ Patient p_1 visits h_1 and p_2 visits h_2 . Hospital h_1 is paid $b_1(\{h_1, h_2\})$ and h_2 is paid $b_2(\{h_1, h_2\})$ for a surplus of $(5 + 5) - (b_1(\{h_1, h_2\}) + b_2(\{h_1, h_2\}))$.
- $\{h_1\}$ Both patients visit h_1 , and h_1 is paid $2b_1(\{h_1\})$ for a surplus of $(5 + 4) - 2b_1(\{h_1\})$.
- $\{h_2\}$ Both patients visit h_2 , and h_2 is paid $2b_2(\{h_2\})$ for a surplus of $(4 + 5) - 2b_2(\{h_2\})$.

To replicate the two-part tariff, set $b_1(\{h_1, h_2\}) = 2b_1(\{h_1\})$.⁷ Since h_1 is treating half as many patients in the full network as it does in the h_1 -only network, h_1 's revenue (profit) is then the same for either network $\{h_1, h_2\}$ or $\{h_1\}$. Furthermore, by construction, if $b_1(\{h_1, h_2\}) = B_1$ and $b_2(\{h_1, h_2\}) = B_2$, then the insurer's choice problem with network specific prices is identical to the problem with lump-sum prices of B_1 and B_2 .

The network-specific prices then produce precisely the same outcomes as with lump-sum prices, and yet add the essential benefit that the supply side is insulated from unsustainably aggressive downstream competition.

I also claim that network-specific pricing can be thought of as fostering more aggressive bidding than simple per-patient pricing. To see this, consider the per-patient prices of $b_1 = b_2 = 2$, and recall that these are the minimum rationalizable bids. At price profile $(2, 2)$, consider each hospital's

⁷In the formal model, two-part tariffs can generally be replicated by network-specific bidding.

incentive to undercut. In order for hospital h_1 to incentivize the insurer to drop h_2 , he must lower his price to $b_1 = 1 - \epsilon$; this can be seen in Figure 2 above. Such an undercut would lower hospital profits from 1×2 to $2 \times (1 - \epsilon)$, and therefore – at prices of $b_1 = b_2 = 2$ – no hospital wants to lower its price.

Now introduce the obvious corresponding network-specific prices: $b_1(\{h_1, h_2\}) = b_1(\{h_1\}) = 2$ and $b_2(\{h_1\}) = b_2(\{h_2\}) = 2$. When each hospital sets the same price on every network, this is equivalent to simple per-patient pricing. Thus, at these prices, the insurer also selects the full network, $\{h_1, h_2\}$.

Reconsider hospital h_1 's incentive to undercut with network-specific prices. If h_1 wants to increase his quantity, it must incentivize the insurer to switch from the full network to the h_1 -only network. To do this, it lowers its price *only on* the $\{h_1\}$ network, $b_1(\{h_1\})$. This is key. If a hospital wants to change the insurer's decision, it should lower its price only on an option that is not currently being selected.

With simple per-patient prices, if hospital h_1 lowers its price b_1 , this impacts the insurers network choice decision in two ways: (1) it is now more valuable for the insurer to drop h_2 and select the $\{h_1\}$ network, but (2) it *also* more valuable for the insurer to simply keep the full network unchanged. That is, with simple per-patient prices, the benefit to the insurer from undercutting enters both options. Since the insurer retains the option to assemble its favorite network, there will be some bid profiles where a hospital would lower its price *only if* the insurer would commit to dropping the other hospital.

This is the power of network-specific prices: h_1 can now target a lower price precisely on the one option that benefits h_1 . This forces the insurer to commit to dropping the other hospital if it wants to accept h_1 's lower price. Decoupling the prices across the networks has therefore had an important, yet unmentioned, consequence: it shuts down part of the option value of the insurer, and this will translate to more aggressive bidding.

In particular, with simple linear prices, if $b_1 = 1.5 - \epsilon$ and $b_2 = 2$, then the insurer continues to select the full network. But if $b_1(\{h_1\}) = 1.5 - \epsilon$ while $b_1(\{h_1, h_2\}) = b_2(\{h_1\}) = b_2(\{h_2\}) = 2$, the insurer will now switch to the $\{h_1\}$ network. For concreteness, the relevant calculations are listed below.

Simple Per-patient

($b_1 = 1.5 - \epsilon$ and $b_2 = 2$)

$$\begin{array}{lll} \{h_1, h_2\} & 10 - (1.5 - \epsilon) - 2 & = 6.5 + \epsilon \\ \{h_1\} & 9 - 2(1.5 - \epsilon) & = 6 + 2\epsilon \\ \{h_2\} & 9 - 2(2) & = 5 \end{array}$$

Network-specific

($b_1(\{h_1\}) = 1.5 - \epsilon$ and others at 2)

$$\begin{array}{lll} \{h_1, h_2\} & 10 - 2 - 2 & = 6 \\ \{h_1\} & 9 - 2(1.5 - \epsilon) & = 6 + 2\epsilon \\ \{h_2\} & 9 - 2(2) & = 5 \end{array}$$

As the calculations clarify, with simple per-patient pricing, the price drop increases the insurers surplus on both the full network, $\{h_1, h_2\}$, as well as the h_1 -only network, $\{h_1\}$. With network-specific pricing, the price decrease only increases the surplus on the targeted network. The key point is that simple per-patient pricing requires hospital h_1 to offer an even larger price decrease to incentivize the insurer to drop the rival hospital h_2 than with network-specific pricing. This produces a range of prices where the hospital will undercut with network-specific pricing, but will not undercut with simple per-patient pricing. The particular prices chosen here fall in that range: $b_1(\{h_1\}) = 1.5 - \epsilon$ is a profitable deviation for h_1 : treating two patients at a price of $1.5 - \epsilon$ is better than treating one patient at a price of 2.

This is precisely the result: by allowing the hospital to target its prices to specific networks, it prices more aggressively. In equilibrium, the prices are $b_1(\{h_1, h_2\}) = b_2(\{h_1, h_2\}) = 1$ and $b_1(\{h_1\}) = b_2(\{h_2\}) = 0.5$, and the insurer makes the efficient selection: the full network.

The central message of this paper is that this example generalizes. Simple, bilateral, per-patient pricing creates a strong pricing complementarity and ultimately generates a pricing distortion that is economically significant. The essence of two-part tariffs solve this problem, but two-part tariffs are ill-suited for this setting due to the (unmodeled) downstream competition. The presented alternative, *network-specific pricing*, overcomes these difficulties and produces the desired outcomes.

Outline and Contributions

The current project is broken into two separate papers: theory (Part I) and empirics (Part II).

This paper is Part I and contains the theoretical framework that generalizes the opening example. After reviewing the literature, I provide a brief overview of some details of the healthcare market that motivate some of the modeling assumptions. I then present the formal models: one baseline model that represents the current market structure, and one counterfactual market design setting where I introduce pre-exchanges and network-specific pricing.

While the theoretical findings of Part I provide clean, qualitative insights, the most striking result is the magnitude of this pricing distortion. In the example above, with simple per-patient pricing, the minimum rationalizable bid is twice as high as the equilibrium network-specific prices. Part II contains the empirical work that quantifies this impact of pre-exchanges in a practical setting. I use recent health care data from Massachusetts, both hospital data as well insurance exchange data, to estimate consumer preferences and then simulate pricing outcomes under each market: the current baseline and the counterfactual pre-exchanges. This companion paper, Part II, can be found on my website listed on the cover page.

2 Related Literature

Below I describe the relation of my work to the fields of industrial organization, market design and health economics. Several papers that are both IO and health are located in the IO section.

Industrial Organization

The introduction of health insurance pre-exchanges can be thought of as a policy intervention that fosters upstream competition and, therefore, shares some similarities with recently proposed regulations of cable channel (un)bundling as studied by Crawford and Yurukoglu (2012). They estimate a baseline bargaining model, and then examine the counterfactual impact of requiring cable companies to sell channels à la carte.⁸ The advance made by Crawford and Yurukoglu (2012) over previous work is accounting for the regulatory effects on pricing incentives.⁹ In particular, allowing channel prices to adjust after regulation results in dramatically higher per-channel prices and a corresponding ambiguous impact on welfare.

The distinction between our frameworks is the specification of the **counterfactual**. They keep the bargaining bilateral and introduce à la carte channel sales. In contrast, I retain the sale of bundles (i.e. access to hospital networks), but I enrich the network formation process to allow for explicitly multilateral contracting.

An analogy helps clarify the distinction. Consider a simple market with only two channels, ESPN and Fox Sports. In their counterfactual, there are two prices total: ESPN sets an à la carte price π_E and Fox Sports sets an à la carte price π_F . The consumer then chooses from four options: (i) buy both channels at a total price of $\pi_E + \pi_F$ (ii) buy just ESPN at price π_E (iii) buy just Fox Sports at price π_F or (iv) buy nothing. Note that if ESPN increases its price π_E , then consumer surplus is decreased on both options (i) and (ii).

In my counterfactual, there are four prices total: ESPN sets two prices, $\pi_E(\{E\})$ and $\pi_E(\{E, F\})$, and Fox Sports similarly sets two prices, $\pi_F(\{F\})$ and $\pi_F(\{E, F\})$. Each channel is setting a bundle-specific price for each bundle that includes it. The consumer now chooses from four options: (i) buy both channels at a total price of $\pi_E(\{E, F\}) + \pi_F(\{E, F\})$ (ii) buy just ESPN at price $\pi_E(\{E\})$ (iii) buy just Fox Sports at price $\pi_F(\{F\})$ or (iv) buy nothing. A crucial distinction is that no single price enters two options. The prices are, by design, bundle-specific. As in the opening example, I argue that decoupling prices into bundle-specific components has a dramatic impact on pricing incentives by effectively unwinding the (problematic) strategic complementarity present in the upstream market; i.e. channels or hospitals.

⁸Recent legislation in Massachusetts similarly requires insurers to offer two networks: one standard and one “limited.” I will discuss this in greater detail below.

⁹See Crawford (2008), Byzalov (2008) and Rennhoff and Serfes (2008) for work on bundling in the cable industry.

A recent paper that also (implicitly) introduces this complementarity is Brand, Gowrisankaran, Nevo, and Town (2012). They study the impact of hospital mergers on post-merger prices and welfare, and their techniques build on the work of Town and Vistnes (2001) and Capps, Dranove, and Satterthwaite (2003) (discussed below). The key innovation (for my purposes) of Brand, Gowrisankaran, Nevo, and Town (2012) over the earlier literature is a seemingly simple observation: the incremental value to an insurer of adding a high-quality-high-price hospital to a network is less than the incremental value of a hospital that is high-quality-low-price. That observation drives them to include prices in a more explicit fashion than the previous literature, and precisely the complementarity I highlight in this paper is present in their work as well. Since Brand, Gowrisankaran, Nevo, and Town (2012) retain the bilateral pricing upstream, it would be interesting to see how merger incentives change in an environment with pre-exchanges or network-specific pricing.

An important recent contribution to the literature on network formation is Lee and Fong (2012). They introduce a dynamic model that captures the evolution of networks through time. Among their results, they show that networks can be inefficient, and I provide a similar result in the empirical work below. As in Brand, Gowrisankaran, Nevo, and Town (2012) and Crawford and Yurukoglu (2012), their dynamic framework utilizes linear bilateral pricing in each stage game.

The closest theoretical IO paper is recent work by Rennhoff and Serfes (2009).¹⁰ They study network formation and à la carte regulation in a simple model with two upstream firms and two downstream firms. Each upstream firm names one price that applies uniformly to the downstream firms. Aside from the absence of bundle-specific bidding, this uniform pricing assumption is an additional degree of separation from my work. My theoretical results suggest that the richer contracting may be (much) more important than à la carte sales, and it would be interesting to see how their results change under the more general pricing protocol outlined in this paper.

There are several other papers that study vertical contracting but the key point of differentiation is that the ultimate downstream consumer (patient) has preferences that depend inherently on the upstream network or upstream contracting. In that sense this is closer to a three-sided matching market than a typical vertical market. The consumer also wants access to a great number of upstream firms (i.e. hospitals), which implies that the health care market shares some structural properties with the cable television market, mentioned above.

In the empirical companion piece, I estimate patient demand for insurance, and there have been several advances in recent years. I share the general demand estimation concepts with Ho (2006, 2009) which built on work by Town and Vistnes (2001) and Capps, Dranove, and Satterthwaite (2003). My estimation benefits from micro patient choice data as well as rather unique network

¹⁰There is other work by Gal-Or (1997) that presents a similarly simple theoretical framework. The modeling assumptions are more relevant to the strict HMO regime witnessed in the 1990s, and her main result is somewhat inconsistent with modern network formation, namely that insurers' networks are differentiated.

and choice set variation. I view these potentially important differences and discuss them at length in the empirical section. I also contribute to the growing literature using **latent class models** for demand estimation. A latent class model is the discrete analog of the more commonly used (continuous) mixed logit, and the potential benefits have been recently noted by Shen (2009) and Hess, Ben-Akiva, Gopinath, and Walker (2011). In particular, the formulation allows for closed form solutions that render computationally intensive simulations unnecessary. Latent class models also allow for flexible interaction between observable characteristics and the distribution of the latent variables.

Market Design The theoretical model is formally a combinatorial (procurement) auction: the insurer will solicit bids from hospitals and select a subset of those hospitals for her network (Immorlica, Karger, Nikolova, and Sami (2005), Cramton, Shoham, and Steinberg (2006)). Importantly, the quantity of patients treated at any particular hospital h depends on the full network choice and not simply upon whether h is in or out of the network. This externality introduces a price-quantity tradeoff for the bidding hospitals that is implicitly defined through the insurer’s network choice. The analogous tradeoffs in other multi-unit auctions (Ausubel (2006)) are typically much more straightforward by design.¹¹ In practice (as in the baseline model), hospitals submit simple linear per-patient bid schedules to the insurer. These “simple” bids significantly complicate the analysis, and a key contribution of this paper is demonstrating *how* the restriction to simple, linear prices distorts incentives.

From a design perspective, typical solutions to these types of problems (e.g. two-part tariffs and non-linear pricing of Wilson (1997) and Armstrong and Vickers (2010)) are ill-suited to this environment due to downstream competition; i.e. the insurers will form networks and then compete downstream for enrollees.¹² In the present insurer-hospital setting, as with many oligopoly markets with vertical contracting, two-part tariffs render the downstream price competition between insurers too intense, as discussed above.

Standard auction theoretic analysis often maintains the assumption of private consumption values for both the auctioneer and the bidders, but this is violated in this industrial setting in two respects: (1) there are competing insurers and (2) each insurer’s value of a network depends upon the networks of the other insurers. In other words, there are competing auctioneers, and the auctioneers’ reservation values are interdependent and therefore endogenous. I know of no paper that studies both of these effects, but there has been recent work by Pai (2012) studying competition among

¹¹Ausubel and Cramton (2002) study the incentives for demand reduction in uniform price auctions. There is some flavor of demand reduction in the present paper, but a hospital’s bid controls its quantity implicitly as it manipulates the insurer’s network choice. That is, a hospital h can increase its quantity by lowering its bid and enticing the insurer to foreclose on some of h ’s rivals.

¹²An important example of unmodeled downstream competition is the series of FCC spectrum auctions (Bulow, Levin, and Milgrom (2009)). The value of a spectrum license to a telecommunications company, e.g. AT&T, importantly depends upon the licenses obtained by its competitors.

auctioneers.

The desire to introduce richer contracting while maintaining per-patient pricing leads to the innovation of this paper: network-specific pricing. This is closely related to package pricing (or package bidding) in the auction literature (Bikhchandani and Ostroy (2002)), which is often motivated by the presence of complementarities in the auctioned items (Ausubel and Milgrom (2002)).¹³ In the present model, there are no complementarities, but there are allocative externalities. As such, a fundamental portion of the current paper (i.e. the insurer’s network selection protocol) shares structural properties with common agency games and menu auctions. The seminal references are Bernheim and Whinston (1986a,b). In that literature, bids are often contingent upon the full outcome chosen by the common agent, e.g. the insurer and his network choice. In the baseline model (as in practice), hospital prices are too coarse to fully express some relevant tradeoffs. Maintaining this assumption retains important qualitative features of the market, but also greatly increases the difficulty of the analysis. One contribution of the current paper is providing a framework for analyzing complete information first-price bidding games that resist analytical characterization with standard tools.

Other recent work in market design has focused on bilateral contracting (Hatfield and Kominers (2010)) and competitive equilibrium pricing (Gul and Stacchetti (1999)) in the absence of strategic considerations. In the present setting, competitive equilibria (in the usual sense) exist and suggest the unrealistic outcome of marginal cost pricing; i.e. every hospital would make zero profits in equilibrium. I feel that incorporating strategic considerations is particularly important here since this (naive) adaptation of Gul and Stacchetti (1999) produces precisely the opposite result I obtain. In fact, the contrast is much stronger: linear pricing leads hospitals to earn excessively high profits in my setting, and excessively low profits (zero profits) in their setting.¹⁴ Given that the qualitative nature of the result pivots on the full consideration of incentives, I view this as an important contribution of my work.

Health Economics I have already mentioned several health economics papers on network formation and pricing, and there are many others (Wedig (2012), Van Horn and Wedig (2011), Bardey and Rochet (2010), Burgess, Carey, and Young (2005), Sorensen (2003)).

There is an additional line of research I would like highlight that is nearly unique to health economics: cost containment. If cost containment is the major unsolved problem in health care, there is a host of well-known contributing subproblems; e.g. over-utilization, inefficient practice variation, misaligned physician incentives, and various difficulties associated with having a large government payer.¹⁵ These are extraordinarily thorny problems that have been known and debated for many

¹³See Ausubel and Milgrom (2002) for a discussion of complementarities and the “exposure problem.”

¹⁴I define “excessively high” as hospitals earning more than their marginal contribution to society.

¹⁵Reforms include the (1) reducing consumer moral hazard by taxing overly generous insurance plans (2) reducing

years. While recent reforms have done a remarkable job at targeting these specific subproblems, most analysts agree that they collectively will not solve the more general problem of cost containment – just as the opening quote from Christie Romer suggests. The Kaiser Institute offers a comprehensive survey of the relevant work.

In this paper, I address network formation as a means of containing costs and this is a pro-market reform aimed at the supply side. An alternate reform was recently introduced in Massachusetts which resembles the cable unbundling mentioned earlier. In 2010, the state passed legislation requiring all insurers to (begin to) offer insurance plans with at least two networks: one standard network and one limited network. The goal is precisely the same as the goal of cable unbundling: increase upstream competition for downstream consumers. Given the ambiguous impact of unbundling in the cable market (Crawford and Yurukoglu (2012)), it is unclear if the introduction of limited networks alone will have the desired effects. In fact, there are other reasons why limited networks may be problematic in health insurance that have no counterpart in the cable industry. The most important is adverse selection. The patients that buy the limited network plans are likely to be healthier on average. I avoid the issue by focusing on a single insurer with a single network, but this is an important area for future work. In particular, network-specific bidding together with multiple network offerings is a promising avenue that deserves attention.

Before presenting the formal model, I introduce some aspects of the health care market and some intuition for the results.

3 Brief Overview of the Health Care Market

The details of the health care market will shape the modeling assumptions, and this section highlights the key features: network access, bilateral contracting and per-patient pricing.

A health insurance contract typically permits an enrollee access to a network of hospitals. That is, the enrollee pays an upfront premium to the insurer, and then – if sick – can seek treatment at any of a pre-specified set of hospitals at little-to-no additional cost. I refer to this as *network access*.

The hospital care is then being paid for by the insurer out of the upfront premium payments, and the enrollee – if sick – is visiting her preferred hospital in the network. Note that the person choosing the hospital, i.e. the enrollee (now patient), pays effectively the same amount across all hospitals.

If an insurer committed to include all hospitals in the network, then hospitals would have strong supply side moral hazard by aligning incentives through new payment schemes and Accountable Care Organizations (3) reviewing payments to providers by introducing the Independent Payment Advisory Board and (4) encouraging proper treatment choice by fostering Cost Effectiveness Research.

incentives to charge very high prices; i.e. once a network is fixed, a hospital's own-price elasticity is close to zero. Insurers counteract this by threatening to drop hospitals from the network that set prices too high. This effectively reverses the timing: hospitals set prices and then insurer networks are formed. This practice is deemed *selective contracting*.¹⁶ The hospitals are then competing to be included in the insurer's network, and this process fosters a type of downward price pressure.

In practice, the hospital price setting and insurer network decisions are the result of individual negotiations between hospitals and insurers. In particular, each hospital-insurer pair meets privately to discuss terms. Hospitals offer terms and rates, and insurers typically have departments devoted to assessing the value a hospital may add to their existing network. If the terms are agreeable to the insurer, the hospital is added to the network. An important descriptive detail of these contracts: they are not explicitly interdependent. Given the large number of illnesses and treatment options, the contracts themselves are lengthy and complex, but there is effectively zero cross-referencing across hospitals for a given insurer. I refer to this as *bilateral contracting*.

Reinhart (2006) contains a detailed account of this price-setting process, and he describes how the insurer-to-hospital payments are typically either: per procedure, per diem, or per episode,¹⁷. For my purposes, I abstract from the precise details and note that these are all linear per-patient price schedules. I refer to this as *per-patient pricing*.

Network access implies that an insurer will not be encouraging its enrollees to visit different in-network hospitals on the basis of price differential. Rather, the insurer will discipline hospital prices solely through the threat of exclusion from the network. *Bilateral contracting* and *per-patient pricing* imply that the hospitals will be offering linear supply schedules to the insurer as they compete to be included to be in the network.

¹⁶A similar model of competition is often in place when the end user does not pay on the margin, but rather pays an upfront fee for access. A helpful example is Netflix's streaming service: customers pay roughly \$8 per month for unlimited access.

¹⁷An episode can be thought of as group of procedures.

Theory

The formal model extends the previous example to a more general setting. To isolate the pricing incentives of the hospitals, I continue to focus on the network selection of a single insurer, but that insurer is now dealing with several hospitals.

The insurer’s network selection is modeled as a competitive bidding procurement auction: the hospitals compete to be included in the insurer’s networks by submitting per-patient bids, and the insurer will then select its preferred network. I consider two variants of the auction that vary along a single dimension: the bid space. In the first variant, hospitals will submit bilateral bids, and I think of this as representing the current market structure. In the second variant, I introduce the network-specific bids in what I call a pre-exchange. Framing issues aside, the only formal distinction is the bid space, and the main conjecture is that, in equilibrium:

bilateral bids \gg network-specific bids.

The equilibrium analysis of the bilateral bidding game, unfortunately, is not entirely straightforward. To overcome this, I introduce two new concepts: pairwise stable equilibrium and marginal product. I proceed to verify the main conjecture by ultimately proving:

bilateral bids $>$ pairwise stable eq’a $>$ marginal products $>$ network-specific bids.

There are two subsections: one model of the market market (bilateral bids) and one model of the pre-exchanges (network-specific bids).

4 Model: Current Market Structure (Bilateral Bidding)

There is a single insurer i and a set of hospitals H . The insurer wants to buy access to the hospitals (for enrollees) by selecting a network $N \subset H$. The insurer’s value of access to a network N is denoted $v(N)$, which I think of as the revenue he can extract from enrollees downstream or, relatedly, the total downstream enrollee welfare.

Hospital $h \in H$ has per-patient treatment cost c_h , and each hospital is assumed to have enough capacity to service the entire market. If the insurer has access to network N , it sends $q_h(N) > 0$ patients to hospital $h \in N$. If the insurer does not have access to a hospital, the insurer sends zero patients to that hospital. For instance, in the opening example:

$$q_1(\{h_1, h_2\}) = q_2(\{h_1, h_2\}) = 1 \quad \text{and} \quad q_1(\{h_1\}) = q_2(\{h_2\}) = 2.$$

The hospitals compete to be included in the insurer’s network by submitting per-patient bids. Given a profile of bids $b = (b_1, \dots, b_H)$, if the insurer selects network $N \subset H$ its payoff is:

$$V(N, b) = v(N) - \sum_{h \in N} b_h q_h(N).$$

Similarly, if network N is chosen, hospital $h \in N$ realizes profit:

$$\Pi_h(N, b) = (b_h - c_h)q_h(N).$$

The insurer will select the network N to solve the following program.

$$V(b) = \max_{N \subset H} V(N, b) \tag{OPT-B}$$

That is, given a profile of hospital bids, the insurer is selecting the network that yields him the highest surplus. The value of the program is denoted OPT-B to emphasize bilateral bidding.¹⁸ If the insurer is indifferent, it selects any network that maximizes the sum of hospital profits.¹⁹ In what follows, fix a selection rule consistent with OPT-B.

For a bid profile b , let $N^*(b)$ denote the insurer’s optimal network choice, and let $q_h^*(b) = q_h(N^*(b))$ denote hospital h ’s subsequent quantity of patients. Given a profile of bids b , hospital h ’s profit is then:

$$\Pi_h(b) = (b_h - c_h)q_h^*(b) \tag{h’s profit}$$

Note that the quantity of patients the insurer sends to a given hospital depends on the hospital bids only insofar as the bids determine the insurer’s network choice. This is meant to capture intuition of network access: for a given network, an insurer has a negligible effect on his own quantity.²⁰ An insurer’s bid *will*, however, effect the network chosen.

The next two assumptions capture the intuition that hospitals are substitutes in two respects: (i) if the insurer adds a new hospital to the network, the number of patients that the other hospitals treat decreases weakly and (ii) the insurer does not decide to drop a hospital after other rival hospitals increase their prices. If either of these is failing, then there is some sort of complementarity among the hospitals.

ASSUMPTION 1. The hospitals are network substitutes: *the number of patients treated by a*

¹⁸In the next section, we will see a similar program OPT-NS with network-specific bidding.

¹⁹This tie-breaking rule uses some value information, but as discussed in Day and Milgrom (2008) is fairly innocent. See Simon and Zame (1990) for the original discussion.

²⁰This has been estimated recently by Brand, Gowrisankaran, Nevo, and Town (2012): for a 1% increase in a hospital’s price, it sees 0.03% fewer patients. This estimates holds the network fixed and does not account for the increased likelihood of being dropped from the network.

hospital h is weakly decreasing in the size of the network:

$$h \in N \subset N' \quad \Rightarrow \quad q_h(N) \geq q_h(N')$$

ASSUMPTION 2. *The **hospitals are price substitutes** for the insurer: if hospital h is included at some bid profile $b = (b_h, b_{-h})$, then h continues to be included if the other hospitals (weakly) increase their bids to $b'_{-h} \geq b_{-h}$.*

$$\forall b'_{-h} \geq b_{-h} : h \in N^*(b_h, b_{-h}) \Rightarrow h \in N^*(b_h, b'_{-h})$$

Furthermore, I would like to capture the notion that the hospitals are (horizontally) differentiated and, relatedly, that insurer networks are large. Given that (i) hospitals are geographically dispersed and (ii) the most important factor determining a patient's hospital choice is driving distance, it should be unsurprising that insurers often include all hospitals in their networks. See Ho (2009) for a full discussion. This motivates our next assumption.

ASSUMPTION 3. *The **full network is efficient**: if each hospital bids its true marginal cost $b = c = (c_1, \dots, c_H)$, then the insurer includes all hospitals in the network:*

$$V(H, c) > V(N, c) \quad \forall N \neq H$$

Let B^{FN} denote the set of bid profiles that result in the full network:

$$B^{FN} = \{b : N^*(b) = H\}.$$

By our tie-breaking assumption, B^{FN} includes all bid profiles for which the full network is weakly optimal for the insurer:

$$V(H, b) \geq V(N, b) \quad \forall N \Rightarrow b \in B^{FN}.$$

Now consider the insurer's problem under two bid profiles: $c = (c_1, \dots, c_H)$ and $c' = (\infty, c_{-h})$. When all hospitals are bidding their true marginal costs c , the insurer is selecting the efficient full network. In c' , hospital h is bidding so high as to effectively remove himself from the the problem, $V(c') = V(H \setminus \{h\}, c)$. Therefore, $V(c) - V(c')$ is the amount by which h 's presence increases the total welfare. I record this in the following definition.

DEFINITION 1. *Hospital h 's **marginal product** MP_h is the amount by which h increases total welfare.*

$$MP_h = V(H, c) - V(H \setminus \{h\}, c).$$

A market system that gives each firm or agent its marginal product often has many desirable properties, including proper incentives for investment and entry. However, in most settings, it is

infeasible to give each agent its marginal product. A classic example is the bilateral bargaining work of Myerson and Satterthwaite (1983). If there is one buyer and one seller, then giving each agent its marginal product would entail giving the buyer the entire surplus and also giving the seller the entire surplus. This is clearly infeasible, and most market mechanisms must find some way to divide the surplus which yields each participant less than its marginal product. Foreshadowing the results below, in the current market with bilateral pricing, each hospital is earning more than its marginal product – often much more. These high prices are producing distortions throughout the broader economy. For details, see the 2009 work by the Council of Economic Advisors: “The Economic Case for Health Care Reform.”²¹

Solution Concept

I have now described most of the network formation game between the hospitals and the insurer, and I complete the model with a solution concept. There is a growing literature in industrial organization regarding network formation, and the two solution concepts most commonly used are Nash equilibrium (Ho (2009)) and pairwise Nash bargaining (Crawford and Yurukoglu (2012), Brand, Gowrisankaran, Nevo, and Town (2012), Horn and Wolinsky (1988)). I now describe each solution concept in this context.

DEFINITION 2. *A profile of hospital bids $b = (b_1, \dots, b_H)$ is a **Nash equilibrium** if*

$$\Pi_h(b_h, b_{-h}) \geq \Pi_h(b'_h, b_{-h}) \quad \forall h \forall b'_h$$

DEFINITION 3. *Given a profile of bargaining weights $\alpha = (\alpha_1, \dots, \alpha_H)$, the profile of hospital bids $b = (b_1, \dots, b_H)$ is a **pairwise Nash bargaining solution** if*

$$b_h \in \operatorname{argmax}_{b'_h} \left[(\Pi_h(b'_h, b_{-h}))^{\alpha_h} (V(b'_h, b_{-h}))^{1-\alpha_h} \right] \quad \forall h$$

where $V(b)$ is payoff to the insurer given bid profile b .

Note that the definitions are equivalent if the hospitals have all of the bargaining power, i.e. $\alpha_h = 1$.

Nash equilibrium is standard, but pairwise Nash bargaining is less familiar. The interpretation of the latter is as follows: each hospital h meets privately with the insurer to bargain over the bid b_h . Holding fixed the bids of the others at b_{-h} , hospital h and the insurer participate in standard Nash bargaining with bargaining weights α_h and $1 - \alpha_h$. Specifically, they choose b'_h to maximize their Nash product $((b'_h - c_h)q_h^*(b'_h, b_{-h}))^{\alpha_h} (V(b'_h, b_{-h}))^{1-\alpha_h}$, holding b_{-h} constant. The insurer

²¹http://www.whitehouse.gov/assets/documents/CEA_Health_Care_Report.pdf

similarly meets with the other hospitals; a pairwise Nash bargaining solution can then be thought of as a Nash equilibrium between the Nash bargains.

One major drawback with both these solution concepts is non-existence. The opening example shows that Nash equilibria may fail to exist, and therefore pairwise Nash bargaining solutions may also fail to exist when hospitals have the bargaining power.

REMARK 1. *Pure strategy Nash equilibria do not necessarily exist.*

As noted above, this issue stems from the types of sharp discontinuities that are consistently present in small examples but are likely to be “smoothed out” in practice. Given this, I propose three reasonable courses of action: (i) analyze the mixed strategy Nash equilibria of the game (ii) switch to a richer model in hopes of re-establishing equilibria or (iii) use a different solution concept. I discuss the merits of each below.

Mixed Strategy Nash Equilibria There are both technical and substantive reasons why mixed strategy Nash equilibria are unsuitable for this analysis. Technically, since the hospital bid space is a continuum, equilibria become difficult to characterize. But more substantively, randomization would imply a level of network instability that is simply not present in the data. In fact, many insurers have virtually no changes to their networks from year to year. Failure to match this significant qualitative feature of the data suggests mixed strategy Nash equilibria may be unsuited for the current analysis.

A Richer Model Although the model presented above is a static bidding game, the true underlying process we seek to model is dynamic. In practice, an insurer-hospital pair may meet every year or every other year to discuss terms. If the hospital terms are acceptable to the insurer, the pair sign a (one- or two-year) contract and the insurer includes that hospital in the network for the duration. An insurer similarly meets with other hospitals over time, and timeline of these contracts overlap. This suggests that adding a dynamic element may be helpful, and – in fact – this has been recently investigated by Lee and Fong (2012), as noted earlier. While the line of research seems quite promising, dynamic models are inherently more complex. Their model is rich enough to capture meaningful dynamics, but unfortunately resists analytical tractability. Furthermore, I want a solution concept that can be imported into the empirical framework (outlined below), and one common expense of dynamic models is a large computational burden.²²

Solution Concept: *Pairwise Stable Equilibrium* I would therefore like to retain the essence of both Nash equilibrium and pairwise stability in a solution concept that yields reasonable predictions in the simple static model. Furthermore, since I will argue that bilateral equilibrium prices are too high, I would like a *conservative* concept that likely underestimates those prices. In what follows, I

²²The framework of Lee and Fong (2012) would be computationally infeasible in my empirical setting.

illustrate that a concept based on a *subset of the necessary conditions*²³ of both Nash equilibrium and pairwise stability satisfies these criteria: ***pairwise stable equilibrium***.

To motivate the use of pairwise stable equilibrium, I have the following story in mind. Hospitals and insurers are meeting periodically over time, but during a meeting a hospital does not simply name a myopically profit-maximizing bid. Rather, it considers both its static profit as well as the incentives of the insurer to renegotiate with other hospitals. Re-examine the discontinuities in the example. Those prices were unstable because hospital h_1 would increase his bid up so far that hospital h_2 would undercut and therefore leave h_1 dropped from the network. If h_1 has even a moderate amount of sophistication or foresight, I would expect him not to trigger such an event.²⁴ I therefore think of each hospital as submitting its profit-maximizing bid *among the bids that do not trigger a renegotiation with other hospitals*. This captures the idea that the hospitals generally have most of the bargaining power, but each hospital also understands that the insurer retains the option to recontract with other hospitals.

Pairwise stable equilibria will coincide with the “hospital optimal” pairwise stable outcome. This is consistent with the general idea of Nash equilibrium (giving the bargaining power to the hospitals) as well as with pairwise Nash bargaining (the prices must be pairwise stable for each pair).²⁵

Just as Nash equilibria are fixed points of standard (Nash) best responses, pairwise stable equilibria are fixed points of stable best responses (defined below). Furthermore, pairwise stable equilibria are conservative in the following sense: stable best responses are everywhere “below” Nash best responses. This statement will be made precise in THEOREM 1. In the example outlined above, the pairwise stable equilibrium is weakly below all rationalizable bids.

DEFINITION 4. A bid profile $b = (b_h, b_{-h})$ is **pairwise stable for h** if there exists no other bid b'_h such that both the hospital h and the insurer prefer (b'_h, b_{-h}) to (b_h, b_{-h}) :

$$\Pi_h(b_h, b_{-h}) \geq \Pi_h(b'_h, b_{-h}) \quad \forall b'_h \leq b_h$$

Let B_h^{PS} denote the set of bid profiles that are pairwise stable for h .

Note that formal statement restricts attention to lower bids $b'_h \leq b_h$ since higher bids always make the insurer weakly worse off.

²³This type of selection is often used in empirical work to estimate parameters in strategic settings. Loosely speaking, if an agent is best responding, then her chosen action satisfies a set of payoff inequalities. The full set of inequalities often poses difficulties, and a subset of these (moment) inequalities is used to estimate parameters.

²⁴Note that this is very different than the nature of undercutting in standard Bertrand models. There, a firm undercuts to gain all or most of the market share. Here, an additional bid increment results in a small increase in profit.

²⁵See Jackson and Watts (2002) for a dynamic foundation of pairwise stability.

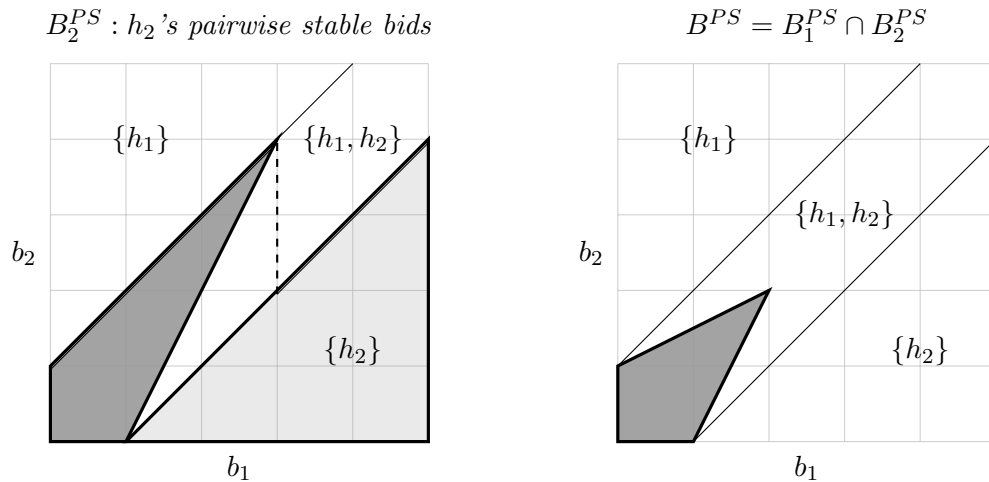
DEFINITION 5. A bid profile b is **pairwise stable** if it is pairwise stable for h for each h :

$$B^{PS} = \bigcap_h B_h^{PS}.$$

There are several ways to capture the idea of pairwise stability, but this statement – together with the tie-breaking rule – ensures that the set of pairwise stable bids is closed.

Figure 6 illustrates the idea in the context of the opening example. The left panel also depicts h_2 's best response, which is the upper envelope of the bid profiles that are pairwise stable for h_2 .

Figure 6: Pairwise Stable Bids



Examining h_2 's pairwise stable bids, note that it is comprised of two components: the darker region corresponds to full network bid profiles, and the lighter corresponds to profiles where only h_2 is in the network. Unsurprisingly, if a bid profile leads to the h_2 -only network, h_2 has no incentive to renegotiate a lower bid (since he is already servicing the entire market).

The intuition for unstable full network bids is slightly more subtle: if a bid profile $b \in B^{FN}$ and is unshaded, hospital h_2 must prefer to undercut. For instance, if $b_2 = b_1 - 1 + \epsilon$ so that the full network $\{h_1, h_2\}$ is selected, then h_2 strictly prefers to drop his bid slightly to $b_2 = b_1 - 1 - \epsilon$ to incentivize the insurer to select the h_2 -only network. A similar argument applies to the entire wedge of unshaded, full network bids.

This also rules out all bid profiles above h_2 's best response: if a lower bid by hospital h_2 would make h_2 more profit, then that lower bid is beneficial to both h_2 and the insurer. Thus, pairwise stable bids must be below profit-maximizing bids. This argument is imprecise since the best responses

are not technically well-defined (due to an “open set” problem). The correct statement – which captures the same idea – is recorded in the following theorem.

THEOREM 1. *If a bid profile b is pairwise stable for h , then any more profitable bid must be larger than b_h :*

$$\Pi_h(b'_h, b_{-h}) > \Pi_h(b) \quad \Rightarrow \quad b'_h > b_h$$

While the “proof” here is particularly straightforward, all proofs are in the appendix.

This result makes precise the notion that any particular pairwise stable bid profile is “more conservative” than Nash equilibrium. Furthermore, if a Nash equilibrium exists, it is pairwise stable; i.e. at b , if each agent is attaining its maximum profit given b_{-h} , then a fortiori there cannot be a bid b'_h that makes both the insurer and h better off. I record this fact in the following theorem.

THEOREM 2. *If b is a Nash equilibrium, then b is pairwise stable:*

$$\Pi_h(b_h, b_{-h}) \geq \Pi_h(b'_h, b_{-h}) \quad \forall h \forall b'_h \quad \Rightarrow \quad \Pi_h(b_h, b_{-h}) \geq \Pi_h(b'_h, b_{-h}) \quad \forall h \forall b'_h \leq b_h$$

The next few results help sketch out the structure of the set of pairwise stable bids B^{PS} , and these structural results will culminate in establishing B^{PS} as a lattice in the usual component-wise ordering.

LEMMA 1.

(i) *The set of pairwise stable bid profiles is non-empty:*

$$B^{PS} \neq \emptyset.$$

(ii) *If a bid is pairwise stable, then all hospitals are in the network:*

$$B^{PS} \subseteq B^{FN}.$$

The first part of the lemma is straightforward: if all hospitals bid cost, then the insurer cannot be made any better off without hospital profits becoming negative. This implies that the bid profile of marginal costs is pairwise stable and so the set of pairwise stable bids is nonempty.

The second part follows from two assumptions: if all hospitals are included in the network when all hospitals bid marginal cost (ASSUMPTION 3), then our assumption that hospitals are price substitutes (ASSUMPTION 2) guarantees that any excluded hospital can get himself back into the network by bidding arbitrarily close to marginal cost. Therefore, any bid profile that excludes a hospital cannot be pairwise stable.

Figure 6 also suggests a sense in which the set B^{PS} is “monotone,” and the following lemma formalizes that idea.

LEMMA 2.

- (i) Fix b_h, b_{-h} and $b'_{-h} > b_{-h}$ such that both (b_h, b_{-h}) and (b_h, b'_{-h}) lead to the full network. If (b_h, b'_{-h}) is pairwise stable for h , then (b_h, b_{-h}) is also pairwise stable for h .

$$(b_h, b'_{-h}) \in B_h^{PS} \Rightarrow (b_h, b_{-h}) \in B_h^{PS}$$

- (ii) Fix b_h, b_{-h} and $b'_h > b_h$ such that both (b_h, b_{-h}) and (b'_h, b_{-h}) lead to the full network. If (b_h, b_{-h}) is pairwise stable for h , then (b'_h, b_{-h}) is also pairwise stable for h .

$$(b_h, b_{-h}) \in B_h^{PS} \Rightarrow (b'_h, b_{-h}) \in B_h^{PS}$$

LEMMA 2 says that hospital h is more willing undercut (or renegotiate with the insurer) when its competitor’s prices are high or its own prices are low. The first stems from the idea that if your competitors are charging low prices, you would need to offer a deep discount to incentivize the insurer to drop a (cheap) rival. The second is different: you are more content with your quantity when you are making a higher margin than when you are making a lower margin.

These results together imply that the set of pairwise stable bids is constructed as the intersection of monotone lower bounds. As suggested in the right panel of Figure 6, these lower bounds trace out a “diamond shape”; i.e. a lattice.

THEOREM 3. *The set of pairwise stable bids is a complete lattice in the usual component-wise ordering:*

$$b, b' \in B^{PS} \Rightarrow b \vee b', b \wedge b' \in B^{PS}$$

where $b \vee b' = (\max\{b_1, b'_1\}, \dots, \max\{b_H, b'_H\})$ and $b \wedge b' = (\min\{b_1, b'_1\}, \dots, \min\{b_H, b'_H\})$.

The full solution concept, pairwise stable equilibrium, will be a fixed point of a well-defined best response operator. From the discussion above, I want to capture both the spirit of a Nash best response as well as the notion that firms are not myopically bidding to maximize profit. I call this a stable best response.


DEFINITION 6. *Given bids $b \in B^{PS}$, a bid $b_h^*(b)$ is a **stable best response** for h if it maximizes h ’s profit among the bids that are not renegotiated.*

$$b_h^*(b) \in \operatorname{argmax}_{b'_h} \{ \Pi_h(b'_h, b_{-h}) : (b'_h, b_{-h}) \in B^{PS} \}$$

DEFINITION 7. A bid profile b is a **pairwise stable equilibrium** if:

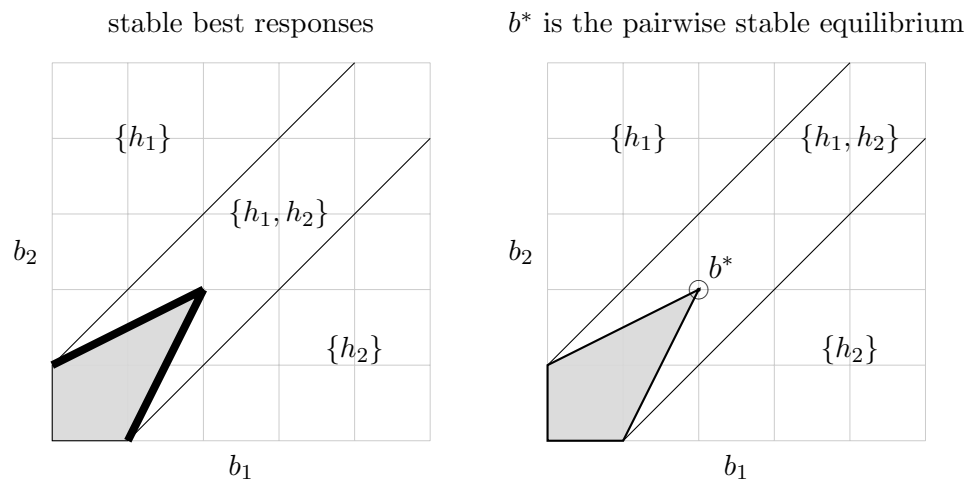
$$b_h = b_h^*(b) \quad \forall h.$$

Figure 7: Stable Best Response for h_1



Given the bids of the hospitals, hospital h knows that some of its bids will be rejected by the insurer, and some of its (accepted) bids will trigger renegotiation. A bid profile is an equilibrium if each hospital is choosing its most profitable (acceptable) bid among the bids that do not trigger a such a renegotiation. This captures the sophistication and stability outlined above.

Figure 8: Pairwise Stable Equilibrium



Since the set of pairwise stable bids is a lattice, the stable best response for h is increasing in the bids of the other hospitals $-h$. This is precisely the strategic complementarity we are concerned about. Note that the restriction to pairwise stable prices dampens some of the strategic complementarity: the stable best response is “steeper” than the Nash best response. This keeps with general the conservative approach taken here.

COROLLARY 1. *The stable best response for h is increasing in the bids of hospitals $-h$:*

$$b, b' \in B^{PS} \quad \& \quad b'_{-h} \geq b_{-h} \quad \Rightarrow \quad b_h^*(b') \geq b_h^*(b).$$

Given the structure of the set of pairwise stable bids, the stable best responses render this a game of strategic complements. Since the best stable responses are monotone increasing, Tarski’s theorem implies that the set of pairwise stable equilibria is a complete lattice.

THEOREM 4. *The set of pairwise stable equilibria is a complete lattice.*

The proof of THEOREM ?? presents an algorithm for computing equilibria: iterate the stable best response operator. A related, generalized algorithm is used in the empirical companion piece.

Effective Marginal Product Bids

Recall the definition of marginal product from earlier.

$$MP_h = V(H, c) - V(H \setminus \{h\}, c)$$

The next result shows that a hospital h ’s stable best response to the other hospital’s truthful bidding yields hospital h exactly his marginal product MP_h .

LEMMA 3. *Fix a hospital h . If hospitals $-h$ submit truthful bids c_{-h} , then hospital h receives his marginal product MP_h by bidding his stable best response, $b_h^*(c)$.*

This leads to the following definition.

DEFINITION 8. *The **effective marginal product bid** of hospital h is*

$$b_h^{MP} \equiv b_h^*(c) = \frac{V(H, c) - V(H \setminus \{h\}, c)}{q_h(H)} + c_h$$

and let $b^{MP} = (b_1^{MP}, \dots, b_H^{MP})$.

The intuition can be read from the figures: if a hospital h is submitting its stable best response, then the insurer is indifferent between keeping and dropping hospital h . On a more subtle note, this occurs precisely when all other hospitals bid marginal cost, and not more generally. For instance, when $b_2 > 0$, the profile $(b_1^*(b_2), b_2)$ does not lie on the boundary between the $\{h_1, h_2\}$ and $\{h_2\}$ regions, and therefore the insurer is not indifferent between keeping and dropping h_2 .

Figure 9: Effective Marginal Product Bids




Figure 9 suggests a finding which holds more generally.

THEOREM 5. *The effective marginal product bid profile is pairwise stable:*

$$b^{MP} = (b_1^{MP}, \dots, b_H^{MP}) \in B^{PS}.$$

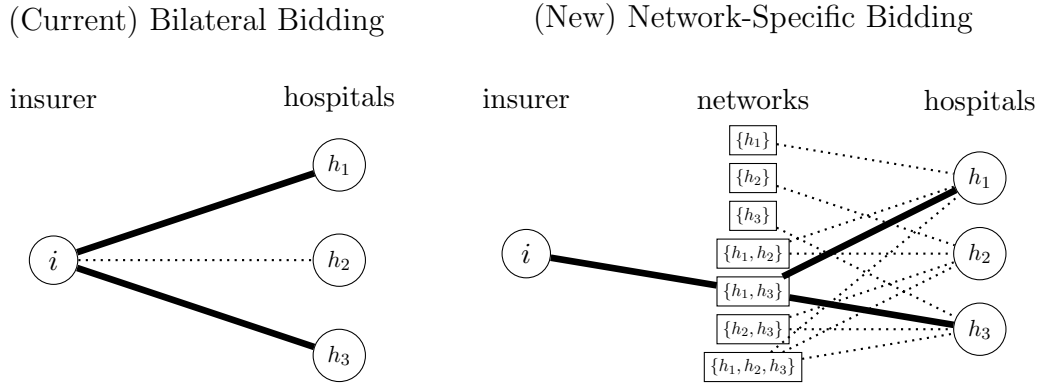
As I argued above, a natural upper bound on hospital profits is their marginal product, and yet – in this example – the hospitals are earning *twice* their marginal products; i.e. $b_h^* = 2b_h^{MP}$ for each h so that the markup the hospitals are charging in equilibrium is 100% higher than the desired upper bound. As mentioned in the introduction, the magnitude of this gap is one of the main motivations for the current project.

This brings me to the first main theoretical result: in equilibrium, each hospital bids higher than its effective marginal product.

THEOREM 6. *For any pairwise stable equilibrium bid profile b^* , $b^* \geq b^{MP}$.*

The preceding results were constructed to make the proof **THEOREM 6** straightforward: the effective marginal product bid is an interior point of the lattice of pairwise stable bids. If every hospital

Figure 10



bids marginal cost, then the profile of stable best responses is b^{MP} (LEMMA 3). Since the stable best responses are monotone increasing (COROLLARY 1), b_h^{MP} is the smallest stable best response for hospital h . Therefore, in any pairwise stable equilibrium, h must be bidding at least b_h^{MP} in any pairwise stable equilibrium.

5 Market Design: Pre-Exchanges (Network-Specific Bidding)

The previous section established that the current market structure of network access and bilateral bidding lead to each hospital's equilibrium profit being higher than its marginal product. The current section aims to address this issue through a better market design.

The key innovation will be a centralized *pre-exchange* where the insurer matches with a network of hospitals. Crucially, the hospital bidding will no longer be bilateral with the insurer. Rather, in the pre-exchange, hospitals will submit *network-specific* bids to the insurer.

On the left of Figure 10 is a depiction of current bilateral bidding framework. While all three hospitals are submitting bids to the insurer, the insurer only selects the $\{h_1, h_3\}$ network. On the right, a hospital h is now submitting one bid for each distinct network N such that h is in N . The insurer is again selecting the $\{h_1, h_3\}$ network.

More generally, each hospital is now submitting $\vec{b}_h = (b_h(N))_{N \ni h}$ which includes one per-patient bid for each network N such that $h \in N$.²⁶ The full profile is denoted $\vec{b} = (\vec{b}_1, \dots, \vec{b}_H)$. Given a

²⁶It is often rightly contested that submitting a bid profile of this length would be impossible for any given hospital. In practice, think of the hospitals as submitting bids to exclude their closest substitutes. For the formal results, I assume the full bid profile, but I investigate the impact of reduced bidding in the empirical work.

profile of network-specific bids \vec{b} , if the insurer chooses network N it realizes a surplus of:

$$V(N, \vec{b}) = v(N) - \sum_{h \in N} b_h(N) q_h(N)$$

and hospital h 's profit is:

$$\Pi_h(N, \vec{b}) = (b_h(N) - c_h) q_h(N).$$

If the insurer selects a network that does not include some hospital, that hospital's profit is zero. As in the baseline model, the insurer selects the network that maximizes its surplus:

$$V(\vec{b}) = \max_{N \subset H} V(N, \vec{b}). \quad (\text{OPT-NS})$$

Let $N^*(\vec{b})$ denote the insurer's optimal network choice. If the insurer is indifferent, it breaks ties in favor of maximizing total hospital profits. Let $q_h^*(\vec{b}) = q_h(N^*(\vec{b}))$ denote hospital h 's subsequent quantity of patients. The program is now tagged OPT-NS to note the presence of network-specific bids, and this contrasts with the earlier OPT-B and bilateral bidding.

Before demonstrating that network-specific bids produce desirable outcomes, I provide a quick discussion of dominated bids.

Dominated Bidding

Given that the hospitals are now submitting network-specific bids, any single hospital h in a network N can effectively preclude the insurer from selecting N by simply bidding $b_h(N) = \infty$. For instance, in the opening example, suppose $b_1(\{h_1\}) = b_2(\{h_2\}) = 0$ and $b_1(\{h_1, h_2\}) = b_2(\{h_1, h_2\}) = \infty$. In this case, the insurer is indifferent between either small network, but obviously will not select the full network. In fact, these bids form a Nash equilibrium (defined below), but setting such high prices on unselected networks is a never an essential component of any optimal bidding strategy by a hospital. This motivates the following definition.

DEFINITION 9. *A bid profile \vec{b} is **weakly dominated** if for some network $N \neq N^*(\vec{b})$ there exists a hospital $h \in N \cap N^*(\vec{b})$ that would get a strictly higher profit if the unchosen network N were selected:*

$$(b_h(N) - c_h) q_h(N) > (b_h(N) - c_h) q_h(N^*(\vec{b})).$$

*A bid profile \vec{b} is **undominated** if it is not weakly dominated.*

To see why these bids are unreasonable, fix a bid profile \vec{b} that is weakly dominated and the corresponding h , N , and $N^*(\vec{b})$. If $N^*(\vec{b})$ is selected given \vec{b} , then hospital h can be made no worse off by lowering his bid slightly on his preferred option N . If the slight move incentivizes the insurer

to switch to network N , h is strictly better off. If the move does not induce a switch, h 's profit is unchanged. Following this logic, h can be no worse off by moving his $b_h(N)$ bid all the way down until:

$$(b_h(N) - c_h)q_h(N) = (b_h(N^*(\vec{b})) - c_h)q_h(N^*(\vec{b})).$$

In this sense, weakly dominated bids are (weakly) inconsistent with profit-maximization, and this justifies the following restriction.

ASSUMPTION 4. *For the remainder of the paper, I limit attention to undominated bids.*

This is similar to (but weaker than) the typical assumptions of “truthful bidding” used in the menu auction games of Bernheim and Whinston (1986b) and the related “truncation reports” in the core-selecting auctions of Day and Milgrom (2008).²⁷ In the menu auction setting, a bid is truthful if the bidder (e.g. hospital) is simply bidding a lump-sum markup uniformly across all the options (e.g. networks) that might be chosen. This makes the bidders indifferent among the possible choices the auctioneer (e.g. insurer) might select. In my setting, restricting attention to undominated bids only precludes the bidders from bidding a strictly larger lump-sum markup on options that are not being selected.²⁸

If multiple hospitals are submitting dominated bids, then there may be equilibria which lead to unreasonable predictions for precisely the same reasons as in the menu auction games. Bernheim and Whinston (1986b) refer to these types of bids as not being “serious.”

Undominated bids have an important and useful characteristic: if a hospital is not included in the chosen network, it must be bidding marginal cost. To foreshadow the results, refer back to the network-specific bids in our opening example: $b_1(\{h_1, h_2\}) = 2b_1(\{h_1\})$ and $b_2(\{h_1, h_2\}) = 2b_2(\{h_2\})$. If *either* hospital is bidding greater than 1, then that bid profile is dominated. This is illustrated in Figure 11.

Results

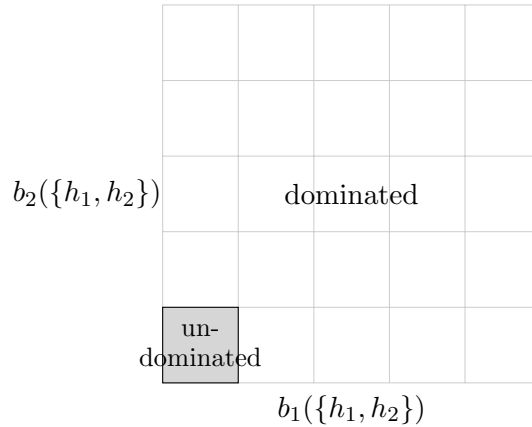
While limiting attention to undominated bid profiles is a reasonable restriction, another simple but useful observation is that all hospitals agree with the insurer's network selection.

LEMMA 4. *If a bid profile \vec{b} is undominated, then – holding \vec{b} fixed – the insurer selects the network*

²⁷This is also similar to the use of truncation strategies in two-sided matching markets. See Roth and Rothblum (1999).

²⁸These types of bid limitations have been utilized in other practical settings. See Ausubel and Milgrom (2002) for a discussion of “activity rules” in various implemented auctions.

Figure 11: Dominated Bidding



that maximizes each hospital's profit:

$$\Pi_h(N^*(\vec{b}), \vec{b}_h) \geq \Pi_h(N, \vec{b}_h) \quad \forall h \quad \forall N$$

This lemma is just a restatement of the definition, but will be useful in establishing the efficiency of network-specific bidding equilibrium outcomes.

THEOREM 7. *If a bid profile \vec{b} is undominated, then the efficient (full) network is selected by the insurer:*

$$N^*(\vec{b}) = H.$$

By LEMMA 4, if a bid profile is undominated, then – holding fixed the bid profile – all hospitals would choose the same network as the insurer. Since total welfare is the sum of all payoffs, and the transfers simply cancel in the welfare calculation, every hospital can agree with the insurer only if the welfare maximizing network is chosen.

Furthermore, as illustrated in Figure 11, the following result highlights the power of network-specific bidding.

THEOREM 8. *At any undominated bid profile \vec{b} , no hospital earns more profit than its marginal product.*

Most of the economic content of this subsection is contained in THEOREM 8: if hospitals are not submitting dominated bids, then no hospital earns more than its marginal product. Below I describe a particular bid profile, but this result is quite strong: in any reasonable network-specific bid profile, the cost of health care is contained.

Given that the marginal product is an upper bound for hospital profits, the remainder of the section constructs a bid profile \vec{b}^{MP} such that (i) each hospital earns exactly its marginal product and (ii) \vec{b}^{MP} is both a pairwise stable equilibrium and a Nash equilibrium. The equilibrium bid profile has a natural interpretation as a two-part tariff (as in the opening example), so it is presented before any definitions of equilibrium concepts.

Recall the effective marginal product bid profiles from bilateral bidding, b^{MP} where

$$b_h^{MP} = \frac{V(H, c) - V(H \setminus \{h\}, c)}{q_h(H)} + c_h.$$

This can be easily extended to a corresponding network-specific bid profile as follows:

$$b_h(N) = \frac{(V(H, c) - V(H \setminus \{h\}))}{q_h(N)} + c_h.$$

If the insurer selects network N that includes hospital h , the total payment from the insurer to the hospital is $c_h q_h(N) + V(H, c) - V(H \setminus \{h\})$. This resembles a two part tariff, just as in the opening example: the hospital agrees to treat the patients at cost, $c_h q_h(N)$, plus a lump-sum payment of $V(H, c) - V(H \setminus \{h\})$, which is independent of the network chosen.

Denote this bid profile as \vec{b}^{MP} . By construction, hospital h earns a profit of $V(H, c) - V(H \setminus \{h\})$ on each network so that \vec{b}^{MP} is undominated. To complete the picture, I verify below that \vec{b}^{MP} is both a Nash equilibrium as well as a pairwise stable equilibrium.

DEFINITION 10. *A bid profile \vec{b} is a **Nash equilibrium** if, for each h , \vec{b}_h maximizes h 's profit, holding fixed \vec{b}_{-h} :*

$$\Pi_h(\vec{b}) \geq \Pi_h(\vec{b}'_h, \vec{b}_{-h}) \quad \forall \vec{b}'_h \quad \forall h$$

Let \vec{B}^{NE} denote the set of Nash equilibrium bid profiles.

Given that no hospital earns more than its marginal product (THEOREM 8), the profile \vec{b}^{MP} is a Nash equilibrium provided that no hospitals are excluded from the network; i.e. $N^*(\vec{b}^{MP}) = H$.

This is true under the following mild condition originally introduced by Shapley (1962).

DEFINITION 11. *The hospitals are **Shapley substitutes** if:*

$$V(H, c) - V(N, c) \geq \sum_{h \in H \setminus N} V(H, c) - V(H \setminus \{h\}, c)$$

The condition poses no restriction in many settings, e.g. the assignment problem of Shapley (1962), but is needed here since the insurer's value function was left fully general. As shown by

Bikhchandani and Ostroy (2006), this is precisely the right condition to ensure that each hospital can receive its marginal product. See De Vries, Schummer, and Vohra (2007) for a longer discussion.

THEOREM 9. *Given bid profile \vec{b}^{MP} , the insurer selects the full network if hospitals are Shapley substitutes.*

COROLLARY 2. *The bid profile \vec{b}^{MP} is a Nash equilibrium if hospitals are Shapley substitutes.*

The following (stability-related) definitions are similarly the network-specific bidding counterparts of the earlier bilateral bidding versions.

DEFINITION 12. *A bid profile $\vec{b} = (\vec{b}_h, \vec{b}_{-h})$ is **pairwise stable for h** if there exists no other bid \vec{b}'_h such that both the hospital h and the insurer prefer $(\vec{b}'_h, \vec{b}_{-h})$ to $(\vec{b}_h, \vec{b}_{-h})$:*

$$\Pi_h(\vec{b}_h) \geq \Pi_h(\vec{b}'_h, \vec{b}_{-h}) \quad \forall \vec{b}'_h \leq \vec{b}_h$$

Let \vec{B}_h^{PS} denote the set of bid profiles that are pairwise stable for h .

DEFINITION 13. *A bid profile \vec{b} is **pairwise stable** if it is pairwise stable for h for each h :*

$$\vec{B}^{PS} = \bigcap_h \vec{B}_h^{PS}.$$

DEFINITION 14. *A bid profile $\vec{b} \in \vec{B}^{PS}$ is a **pairwise stable equilibrium** if, for each h , \vec{b}_h maximizes h 's profit among all pairwise stable bid profiles, holding fixed \vec{b}_{-h} :*

$$\Pi_h(\vec{b}) \geq \Pi_h(\vec{b}'_h, \vec{b}_{-h}) \quad \forall (\vec{b}'_h, \vec{b}_{-h}) \in \vec{B}^{PS} \quad \forall h$$

Let \vec{B}^{PSE} denote the set of pairwise stable equilibrium bid profiles.

As with bilateral bids, Nash equilibria are also pairwise stable equilibria, which yields the next result.

COROLLARY 3. *The bid profile \vec{b}^{MP} is a pairwise stable equilibrium.*

This concludes the formal results.

Discussion

The comparative picture is now complete: bilateral bidding produces equilibrium prices which are systematically higher than the equilibrium prices of network-specific bidding. This systematic bias stemmed from an additional degree of strategic complementarity that is composed of two ingredients: (1) the insurer’s network choice decision and (2) the resulting pricing incentives for the hospitals.

With bilateral bidding, the insurer’s network selection reflects a degree of relative cost-savings: if some hospital h increases its price, the insurer is more eager to accept a different hospital h' because this will channel some patients away from the now-more-expensive h . Equivalently, after h increases its price, the insurer is more reluctant to drop h' because some of those patients will flow to the now-more-expensive h . This reasoning produces the long $\{h_1, h_2\}$ tranche along the diagonal in the opening example.

Since each hospital wants to extract both its value added to the insurer as well the cost-savings it provides, hospital h' demands a strictly higher price after the rival h increases its price.

In other words, this price increase by h has made the insurer worse off in two ways:

1. The insurer is directly worse off as it is now paying more to that hospital h for every patient it treats.
2. The insurer is indirectly worse off since it is now in a worse bargaining position with the other hospitals – h in the network.

This is precisely why network-specific pricing is so powerful in this setting: the indirect effect is completely absent.²⁹ If the insurer is selecting some network N and hospital h demands a slightly higher price *on network* N , this has exactly zero effect on the insurer’s reluctance to drop some other hospital h' : the prices for network $N \setminus \{h'\}$ have not changed. The direct effect is present with both bilateral bidding and network-specific bidding, but it is this indirect effect, this complementarity, that encourages prices to reach excessive levels. This is precisely the motivation for decoupling the prices, and the previous results illustrate the impact.

²⁹This also suggests why à la carte pricing may not have the desired effect: the strategic complementarity is still present.

6 Conclusion

The current network formation process is a highly evolved system that has settled on an approximate solution: bilateral contracting. The goal of this paper is to establish that this approximate solution can be improved substantially through a better market design. In particular, the theoretical results above suggest a systematic upward bias in hospital prices, which inflates the overall cost of health care.

Much of the economic intuition can be found in the opening toy example, where hospital profits are 100% higher than their marginal contribution to society. Furthermore, this two-hospital example is the best case scenario in the following sense: if one hospital incentivizes the insurer to drop its rival, the remaining hospital gains all of the dropped hospital's market share. In a more practical setting, if a hospital is dropped, those patients are typically dispersed among all of the remaining hospitals, not just to a hospital that drops its price. This further weakens the incentives to undercut, and this threat of undercutting is the key to restraining prices.

Along these lines, investigations into more general (numerical) problems suggest a degree of excess hospital markups that is even more alarming than the simple example. Unfortunately, this is a type of quantitative result that is difficult to capture in an analytically tractable model. This motivates the empirical work that can be found in the companion paper, Part II: Empirics. The current paper, Part I: Theory, illustrates the general direction of the bias as well as the underlying reasoning. The empirical work demonstrates that this is a large, economically meaningful effect.

While this paper uses tools of industrial organization and market design, it is fundamentally a health economics paper. My proposal of network-specific pricing will require coordination among both the insurers and the hospitals. I believe in the qualitative nature of these results, and hope that many states seriously consider centralizing the network formation process along the lines of the pre-exchanges outlined in the current work.

References

- ARMSTRONG, M., AND J. VICKERS (2010): “Competitive Non-linear Pricing and Bundling,” *The Review of Economic Studies*, 77(1), 30–60.
- AUSUBEL, L. (2006): “An Efficient Dynamic Auction for Heterogeneous Commodities,” *The American Economic Review*, 96(3), 602–629.
- AUSUBEL, L., AND P. CRAMTON (2002): “Demand Reduction and Inefficiency in Multi-Unit Auctions,” *mimeo*.
- AUSUBEL, L., AND P. MILGROM (2002): “Ascending Auctions with Package Bidding,” *Frontiers of theoretical economics*, 1(1), 1–42.
- BARDEY, D., AND J. ROCHET (2010): “Competition Among Health Plans: A Two-Sided Market Approach,” *Journal of Economics & Management Strategy*, 19(2), 435–451.
- BERNHEIM, B., AND M. WHINSTON (1986a): “Common Agency,” *Econometrica*, 54(4), 923–942.
- (1986b): “Menu Auctions, Resource Allocation, and Economic Influence,” *The Quarterly Journal of Economics*, 101(1), 1–31.
- BIKHCHANDANI, S., AND J. OSTROY (2002): “The Package Assignment Model,” *Journal of Economic theory*, 107(2), 377–406.
- (2006): “Ascending Price Vickrey Auctions,” *Games and Economic Behavior*, 55(2), 215–241.
- BRAND, K., G. GOWRISANKARAN, A. NEVO, AND R. TOWN (2012): “Mergers When Prices Are Negotiated: Evidence from the Hospital Industry,” Discussion paper, University of Arizona.
- BULOW, J., J. LEVIN, AND P. MILGROM (2009): “Winning Play in Spectrum Auctions,” Discussion paper, National Bureau of Economic Research.
- BURGESS, J., K. CAREY, AND G. YOUNG (2005): “The Effect of Network Arrangements on Hospital Pricing Behavior,” *Journal of health economics*, 24(2), 391–405.
- BYZALOV, D. (2008): “Unbundling Cable Television: An Empirical Analysis,” Discussion paper, Harvard.
- CAPPS, C., D. DRANOVE, AND M. SATTERTHWAITE (2003): “Competition and Market Power in Option Demand Markets,” *RAND Journal of Economics*, 34(4), 737–763.
- CRAMTON, P., Y. SHOHAM, AND R. STEINBERG (2006): *Combinatorial Auctions*. MIT Press.

- CRAWFORD, G. (2008): “The Discriminatory Incentives to Bundle in the Cable Television Industry,” *Quantitative Marketing and Economics*, 6(1), 41–78.
- CRAWFORD, G., AND A. YURUKOGLU (2012): “The Welfare Effects of Bundling in Multichannel Television Markets,” *American Economic Review*, 102(2), 643–685.
- DAY, R., AND P. MILGROM (2008): “Core-selecting Package Auctions,” *International Journal of game Theory*, 36(3), 393–407.
- DE VRIES, S., J. SCHUMMER, AND R. VOHRA (2007): “On ascending Vickrey auctions for heterogeneous objects,” *Journal of Economic Theory*, 132(1), 95–118.
- GAL-OR, E. (1997): “Exclusionary Equilibria in Health-Care Markets,” *Journal of Economics & Management Strategy*, 6(1), 5–43.
- GUL, F., AND E. STACCHETTI (1999): “Walrasian Equilibrium with Gross Substitutes,” *Journal of Economic Theory*, 87(1), 95–124.
- HATFIELD, J., AND S. KOMINERS (2010): “Matching in Networks with Bilateral Contracts,” in *Proceedings of the 11th ACM conference on Electronic commerce*, pp. 119–120. ACM.
- HESS, S., M. BEN-AKIVA, D. GOPINATH, AND J. WALKER (2011): “Advantages of Latent Class Choice Models over Continuous Mixed Logit Models,” Discussion paper, University Press, Harrisburg.
- HO, K. (2006): “The welfare effects of restricted hospital choice in the US medical care market,” *Journal of Applied Econometrics*, 21(7), 1039–1079.
- (2009): “Insurer-Provider Networks in the Medical Care Market,” *American Economic Review*, 99(1), 393–430.
- HORN, H., AND A. WOLINSKY (1988): “Bilateral Monopolies and Incentives for Merger,” *RAND Journal of Economics*, 19(3), 408–419.
- IMMORLICA, N., D. KARGER, E. NIKOLOVA, AND R. SAMI (2005): “First-Price Procurement Auctions,” Discussion paper, Northwestern University.
- LEE, R., AND K. FONG (2012): “Markov-Perfect Network Formation An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks,” Discussion paper, New York University.
- MILLIOU, C., E. PETRAKIS, AND N. VETTAS (2004): “Efficient Trading Forms in Competing Vertical Chains,” Discussion paper, University of Athens.

- MYERSON, R., AND M. SATTERTHWAITE (1983): “Efficient mechanisms for bilateral trading,” *Journal of economic theory*, 29(2), 265–281.
- PAI, M. (2012): “Competing Auctioneers,” Discussion paper, University of Pennsylvania.
- RENNHOFF, A., AND K. SERFES (2008): “Estimating the Effects of a la Carte Pricing: The Case of Cable Television,” *Available at SSRN 1085392*.
- RENNHOFF, A., AND K. SERFES (2009): “The Role of Upstream-Downstream Competition on Bundling Decisions: Should Regulators Force Firms to Unbundle?,” *Journal of Economics & Management Strategy*, 18(2), 547–588.
- ROTH, A., AND U. ROTHBLUM (1999): “Truncation Strategies in Matching Markets—In Search of Advice for Participants,” *Econometrica*, 67(1), 21–44.
- SHAPLEY, L. (1962): “Complements and Substitutes in the Optimal Assignment Problem,” *Naval Research Logistics Quarterly*, 9, 45.
- SHEN, J. (2009): “Latent Class Model or Mixed Logit Model? A Comparison by Transport Mode Choice Data,” *Applied Economics*, 41(22), 2915–2924.
- SIMON, L., AND W. ZAME (1990): “Discontinuous Games and Endogenous Sharing Rules,” *Econometrica*, 58(4), 861–72.
- SORENSEN, A. (2003): “Insurer-Hospital Bargaining: Negotiated Discounts in Post-Deregulation Connecticut,” *The Journal of Industrial Economics*, 51(4), 469–490.
- TOPKIS, D. (2001): *Supermodularity and Complementarity*. Princeton University Press.
- TOWN, R., AND G. VISTNES (2001): “Hospital competition in HMO networks,” *Journal of Health Economics*, 20(5), 733–753.
- WEDIG, G. (2012): “The Value of Consumer Choice and the Decline in HMO Enrollments,” *Economic Inquiry*.
- WILSON, R. (1997): *Nonlinear Pricing*. Oxford University Press, USA.

Appendix – Theory

Omitted Proofs

Proof of THEOREM 1. This follows directly from the definition of pairwise stability. \square

Proof of THEOREM 2. The proof is obvious from the statement. Pairwise stability must satisfy only downward deviations, while Nash equilibria satisfy both upward and downward deviations. Therefore, Nash equilibria satisfy the constraints of pairwise stability. \square

Proof of LEMMA 1.

- (i) If all hospitals bid cost $b_h = c_h$, then the insurer cannot be made better off, and therefore $c \in B^{PS}$ and so $B^{PS} \neq \emptyset$.
- (ii) By ASSUMPTIONS 2 and 3, any hospital will be included in the network if they bid marginal cost. Therefore, any excluded hospital can make itself no worse off by submitting a bid of cost, and this makes the insurer strictly better off. By continuity, there exists a bid that would make both the excluded hospital and the insurer better off. Therefore, if a hospital is not included at a bid profile b , then b is not pairwise stable. \square

I present the proof of LEMMA 2(ii) before the proof of LEMMA 2(i).

Proof of LEMMA 2(ii). Suppose (b_h, b_{-h}) is pairwise stable for h so that for any $\hat{b}_h < b_h$,

$$(b_h - c_h)q_h(b_h, b_{-h}) \geq (\hat{b}_h - c_h)q_h(\hat{b}_h, b'_{-h}).$$

Both initial bid profiles lead to the same network

$$q_h(b_h, b_{-h}) = q_h(b'_h, b_{-h})$$

and therefore

$$(b'_h - c_h)q_h(b'_h, b_{-h}) \geq (b_h - c_h)q_h(b_h, b_{-h}) \geq (\hat{b}_h - c_h)q_h(\hat{b}_h, b_{-h})$$

for all $\hat{b}_h < b_h$. This completes the proof. \square

The following result is useful for proving both LEMMA 2(i) and THEOREM 3.

LEMMA 5. Consider b, b' such that $b'_1 > b_1$ and $b'_2 > b_2$. If $(b'_1, b_{-1}), (b'_2, b_{-2}) \in B^{PS}$ then $(b'_1, b'_2, b_{-12}) \in B^{PS}$.

Proof of LEMMA 5. Fix b, b' such that $b'_1 > b_1$ and $b'_2 > b_2$. If $(b'_1, b_{-1}), (b'_2, b_{-2}) \in B^{PS}$, then the price substitutes condition (ASSUMPTION 2) neither h_1 nor h_2 can profitably undercut at $(b'_1, b'_2, b_{-12}) \in B^{PS}$. That is, if h_1 wants to undercut at (b'_1, b'_2, b_{-12}) then h_1 also wants to undercut at (b'_2, b_{-2}) by LEMMA 2(ii). Similarly for h_2 .

Now suppose h_3 wants to undercut at (b'_1, b'_2, b_{-12}) from b_3 to $\hat{b}_3 < b_3$. Since $(b'_1, b_{-1}), (b'_2, b_{-2}) \in B^{PS}$, h_3 must prefer not to undercut at either of those profiles. By (ASSUMPTION 2, h_3 's quantity is the same at each of $(b'_1, b_{-1}), (b'_2, b_{-2}), (b'_1, b'_2, b_{-12})$ since they are all full network. This implies that h_3 's quantity must be strictly higher at $(b'_1, b'_2, \hat{b}_3, b_{-123})$ than at $(b'_1, \hat{b}_3, b_{-13})$ or $(b'_2, \hat{b}_3, b_{-23})$. However, if h_1 is dropped at $(b'_1, b'_2, \hat{b}_3, b_{-123})$, then h_1 must also be dropped at $(b'_1, \hat{b}_3, b_{-13})$. Similarly for h_2 . This implies that h_3 's quantity is the same across $(b'_1, b'_2, \hat{b}_3, b_{-123}), (b'_1, \hat{b}_3, b_{-13})$, and $(b'_2, \hat{b}_3, b_{-23})$, a contradiction. \square

Proof of LEMMA 2(i). This follows from successively applying LEMMA 5 componentwise. \square

Proof of THEOREM 3. The lattice structure follows from LEMMA 2. Completeness follows from the compactness of B^{PS} , which is closed and bounded by construction. \square

Proof of COROLLARY 1. This is a direct consequence of the lattice structure. \square

Proof of THEOREM 4. Define the stable best response mapping $F : B^{PS} \rightarrow B^{PS}$ defined as

$$F(b) = (b_1^*(b), b_2^*(b), \dots, b_H^*(b)).$$

Since this mapping is monotone (COROLLARY 1), apply Tarksi's fixed point theorem to complete the proof. For reference, see Topkis (2001). \square

Proof of LEMMA 3. By the definition of the stable best reply, the insurer must be exactly indifferent between having hospital h in or out of the network: $V(b_h^*(c), c_{-h}) = V(\infty, c_{-h})$. However, by assumption, all hospitals are in network under both bid profiles (c_h, c_{-h}) and $(b_h^*(c), c_{-h})$:

$$q_h(c) = q_h(b_h^*(c), c_{-h}).$$

Collecting and canceling terms yields the desired equality:

$$\begin{aligned} V(c) - V(\infty, c_{-h}) &= V(c) - V(b_h^*(c), c_{-h}) \\ &= (b_h^*(c) - c_h)q_h(c). \end{aligned}$$

□

Proof of THEOREM 5. Given the substitutes assumption, the proof is fairly straightforward: at the bid profile b^{MP} , each hospital is making a weakly positive profit. Any given hospital would only find a lower bid more profitable if it increases its quantity. Suppose hospital h lowers its bid all the way to marginal cost. Since each other hospital h' is included in the network (by construction) at the bid profile $(c_h, b_{h'}^{MP}, c_{-hh'})$, hospital h' must continue to be included in the network at the higher bid profile $(c_h, b_{h'}^{MP}, b_{-hh'}^{MP})$. Therefore, even if h lowers its bid to c_h , each hospital is still included in the network and h has not increased his own quantity. This cannot increase h 's profit, and I conclude that $b^{MP} \in B^{PS}$. □

Proof of THEOREM 6. The stable best responses to the marginal cost profile are the effective marginal products:

$$(b_1^*(c), b_2^*(c), \dots, b_H^*(c)) = (b_1^{MP}, \dots, b_H^{MP}).$$

Since the stable best responses are monotone increasing, the effective marginal product bid profile is the smallest stable best response. This completes the proof. □

Proof of LEMMA 4. By definition, if $N^*(\vec{b})$ is chosen at undominated bid profile \vec{b} , then – by definition – for all other N :

$$(b_h(N^*(\vec{b})) - c_h)q_h(N^*(\vec{b})) \geq (b_h(N) - c_h)q_h(N) \quad \forall h \in H$$

which is the desired result. □

Proof of THEOREM 7. Fix an undominated bid profile b . If the insurer selects network $N \neq H$, then

$$v(N) - \sum_{h \in N} b_h(N)q_h(N) > v(H) - \sum_{h \in N} b_h(H)q_h(H) - \sum_{h \in H \setminus N} c_h q_h(H)$$

Rearranging and collecting terms yields:

$$V(N, c) - \sum_{h \in N} (b_h(N) - c_h)q_h(N) > V(H, c) - \sum_{h \in H} (b_h(H) - c_h)q_h(H)$$

or

$$V(N, c) + \sum_{h \in H} (b_h(H) - c_h)q_h(H) > V(H, c) + \sum_{h \in N} (b_h(N) - c_h)q_h(N).$$

Since the full network is efficient and the bid profile is undominated, both addends on the right are larger than their counterparts on the left, a contradiction. \square

Proof of THEOREM 8. Hospital h 's marginal product is:

$$\begin{aligned} MP_h &= \left[v(H) - c_h q_h(H) - \sum_{h' \neq h} c_{h'} q_{h'}(H) \right] - \left[v(H \setminus \{h\}) - \sum_{h' \neq h} c_{h'} q_{h'}(H \setminus \{h\}) \right] \\ &= [v(H) - v(H \setminus \{h\}) - c_h q_h(H)] + \left[\sum_{h' \neq h} c_{h'} q_{h'}(H \setminus \{h\}) - \sum_{h' \neq h} c_{h'} q_{h'}(H) \right]. \end{aligned}$$

At an undominated bid profile \vec{b} , the insurer prefers the full network H network to the smaller network $H \setminus \{h\}$:

$$v(H) - b_h(H)q_h(H) - \sum_{h' \neq h} b_{h'}(H)q_{h'}(H) \geq v(H \setminus \{h\}) - \sum_{h' \neq h} b_{h'}(H \setminus \{h\})q_{h'}(H \setminus \{h\}).$$

The profit of h is:

$$\begin{aligned} & b_h(H)q_h(H) - c_h q_h(H) \\ & \leq \left[v(H) - \sum_{h' \neq h} b_{h'}(H)q_{h'}(H) \right] - \left[v(H \setminus \{h\}) - \sum_{h' \neq h} b_{h'}(H \setminus \{h\})q_{h'}(H \setminus \{h\}) \right] - c_h q_h(H) \\ & = [v(H) - v(H \setminus \{h\}) - c_h q_h(H)] + \left[\sum_{h' \neq h} b_{h'}(H \setminus \{h\})q_{h'}(H \setminus \{h\}) - \sum_{h' \neq h} b_{h'}(H)q_{h'}(H) \right]. \end{aligned}$$

Combining the first equality with the second inequality:

$$\begin{aligned}
& MP_h - \overbrace{(b_h(H)q_h(H) - c_hq_h(H))}^{h\text{'s profit}} \\
& \geq \left[\sum_{h' \neq h} c_{h'}q_{h'}(H \setminus \{h\}) - c_{h'}q_{h'}(H) \right] - \left[\sum_{h' \neq h} b_{h'}(H \setminus \{h\})q_{h'}(H \setminus \{h\}) - b_{h'}(H)q_{h'}(H) \right] \\
& = \sum_{h' \neq h} \left(\underbrace{(b_{h'}(H) - c_{h'})q_{h'}(H) - (b_{h'}(H \setminus \{h\}) - c_{h'})q_{h'}(H \setminus \{h\})}_{\geq 0 \text{ since bids are undominated}} \right) \\
& \geq 0.
\end{aligned}$$

□

Proof of THEOREM 9. The surplus to the insurer of selecting the full network is:

$$V(H, c) - \sum_{h \in H} (V(H, c) - V(H \setminus \{h\}, c)).$$

The surplus from selecting network N is:

$$V(N, c) - \sum_{h \in N} (V(H, c) - V(H \setminus \{h\}, c)).$$

Therefore, the insurer selects the full network exactly when:

$$V(H, c) - V(N, c) \geq \sum_{h \in H \setminus N} (V(H, c) - V(H \setminus \{h\}, c)).$$

□

Proof of COROLLARY 2. This follows directly from THEOREM 8 and THEOREM 9. □

Proof of COROLLARY 3. To verify that \vec{b}^{MP} is a pairwise stable equilibrium (also defined below), I verify that any hospital h cannot incentivize the insurer to drop another hospital, even if h bids down to marginal cost. But – by design – bidding down to marginal cost only saves the insurer one lump-sum payment of $V(H, c) - V(H \setminus \{h\}, c)$, and does not change the marginal incentives to include or exclude other hospitals. Therefore no hospitals are dropped after the price decrease, and \vec{b}^{MP} is a pairwise stable equilibrium. □