Consolidation and Competition in Healthcare Markets

Luca Maini*

April 26, 2021

See most recent version here

CONTENTS

*These notes are a work in progress. Please check back frequently for an updated version, and email lmaini@email.unc.edu for corrections. The overview section is partially based on Zhenyu Lai's study notes from Harvard's second-year IO course.

In this chapter I examine the theoretical and empirical evidence behind various types of consolidation activity in healthcare markets. In doing so, we will also discuss various models of competition and their implications for anti-trust analysis. The chapter is broken down in five sections.

- In section 1, I present a basic introduction to the theory of mergers and antitrust.
- In section 2, I examine hospital mergers. The theory of hospital mergers is quite similar to standard theory of mergers in oligopoly markets, though with some differences driven by the presence of bilateral bargaining in many of these markets. The structure of the section loosely follows section 3.2 of Gaynor and Town's chapter in Vol. 2 of the Handbook of Health Economics (Gaynor and Town, 2012).
- In section 3, I cover insurers mergers. The theory behind insurer mergers is more complicated because insurers act as middle-men between patients and hospitals. While standard theory suggests that mergers increase markups, a larger insurer can also extract lower prices from providers, so the impact of insurance mergers is theoretically ambiguous.
- In section 4, I cover vertical integration of hospitals and physicians. Even though physicians and hospitals are both providers (and therefore offer substitute services on the healthcare market), there is a vertical aspect to this type of consolidation activity because oftentimes hospitals rely on referrals from physicians in order to receive patients.
- Finally, in section 5, I cover some results on consolidation across various healthcare markets, including results on consolidation of dialysis clinics, medical devices, etc.

1 BACKGROUND ON MERGERS AND ANTITRUST

We study mergers because they affect both efficiency and consumer welfare. Firms will merge for a variety of reasons, but the effect of mergers will almost always be to reduce competition, which usually leads to higher prices. For this reason, the government investigates any mergers that are likely to reduce competition in significant ways. It's important to note that while economists dislike higher prices because they reduce efficiency in the market, government agencies usually focus on consumer welfare. A simple way to express this difference graphically is through Figure 1.

If the price grows from the efficient $p^\star$ to a post-merger p^M (with a resulting drop in total quantity from $q^\star$ to q^M), the total welfare loss—which an economist would care about—is the total of the area of $A + B$. However, government agencies will generally focus on $A + C$ (i.e. the loss in consumer welfare). Note that even though the impact appears ambiguous, $A + B < A + C$ holds in virtually every conceivable application (the inequality would only revert if p^M grows above monopoly levels).

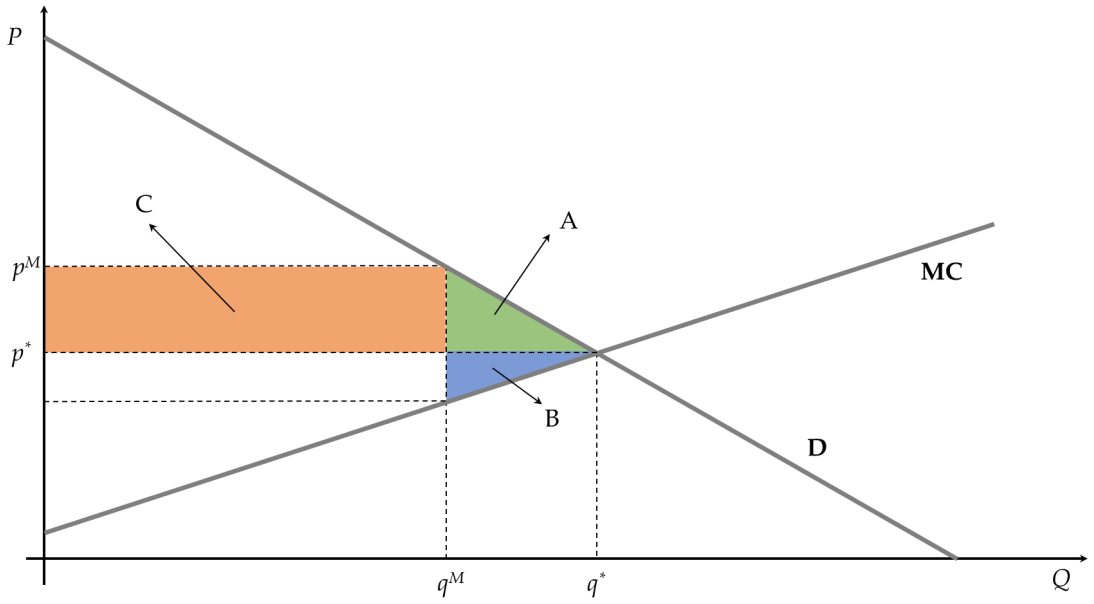

Figure 1: Change in welfare after a merger

1.1 Merger challenge guidelines

The two entities that regulate mergers in the US are the Federal Trade Commission (FTC), and the Department of Justice (DOJ). Firms have to report mergers/acquisitions if the assets involve surpass a certain thresholds (these thresholds are updated on a yearly basis). The vast majority of mergers and acquisitions goes through unchallenged. Occasionally, either the FTC or the DOJ can issue a request for more information or an an injunction to block the merger. If an injunction is filed, the firm has to decide whether to drop the merger, or fight the injunction in court. Firms can also generally negotiate a deal to avoid the injunction, for example by offering to divest certain assets or selling off specific product lines.

HHI

Because firms issue thousands of merger and acquisition reports every year, the FTC and the DOJ need to have quick and "cheap" ways of determining whether a merger should be challenged. The earliest such screens were developed with a very simple model of competition in mind: the Cournot quantity-setting model of competition. In that model, firms that produce homogeneous goods compete by choosing quantities. Prices are then set in a market-clearing equilibrium. Under the Cournot model the firm's first-order condition can be rearranged to look something like

$$\frac{p_i - c_i}{p_i} = \frac{s_i}{\varepsilon}$$

where the left-hand side is the markup, ε is the price elasticity of demand and s_i is the market share of product i. So we can write down the share-weighted markup (which is a good market

measure of how much higher the average market price is relative to marginal cost) as

$$\sum_i s_i \frac{p_i - c_i}{p_i} = \frac{\sum s_i^2}{\varepsilon}$$

The numerator of the right-hand side of the equation is known as the Hirschman-Herfindahl Index (or HHI for short), and is a common measure of market concentration. Before more sophisticated models came along, the DOJ and FTC used the HHI as a rule of thumb to figure out which mergers needed to be challenged. The rough thresholds for the HHI rules are

- HHI $< 1,000$: unconcentrated market, mergers rarely challenged.
- HHI $\in (1,000; 1,800)$: moderately concentrated market, mergers challenged if they produce an increase in HHI by more than 100 points.
- HHI $> 1,800$: concentrated market, mergers challenged if they product an increase in HHI by more than 50 points

You should notice immediately that market shares depend on market definition, and so it should not surprise you that this is where most of the contention happens. One way to do so (that is becoming progressively outdate) is geography. In general any "market definition" is going to be structurally inadequate because a firm is either "in" or "out" of a market. In reality, products are imperfect substitutes of one another, with varying degrees of substitutability.

UPP

More recent approaches use more nuanced concepts such as Upward Pricing Pressure, which is a measure of how likely the price is to go up and by how much. This is measured as

$$UPP_1 = -\underbrace{\frac{\partial D_2(P_2, P_1, \cdot)/\partial P_1}{\partial D_1(P_2, P_1, \cdot)/\partial P_1}}_{\text{diversion ratio}} \cdot \overline{M}_2 \frac{\overline{P}_2}{\overline{P}_1} - E_1\left(1 - \overline{M}_1\right)$$

where D_i is demand for good i, $\overline{M}_i$ are pre-merger markups, and E_1 is the efficiency gain from the merger. To compute this formula we take the difference in prices from a model where firms 1 and 2 first set prices separately, and then act as a joint venture.

There are some problems with this formula. First, it assumes away some second-order effects (notice we are only taking the first derivative of the demand function). Second, it disregards the responses of other firms, who are assumed to maintain the equilibrium from the pre-merger world (notice that this is a violation of Nash equilibrium). The argument that people try to give is that if other firms adjusted that would result in even more upward pricing pressure, because the natural response would be to raise prices as well, which then frees the new merged entity to raise them even more. In the vast majority of cases this is a reasonable argument, though you should keep in mind that (i) if we think prices will rise even more, we should probably try to figure out by how

much, or we risk letting too many mergers go through, and (ii) sometimes prices are not strategic complements.

1.2 Non-price considerations

Regardless of what screening tools are used, all economic models predict that all horizontal mergers lead to higher prices. So something else has to be happening at the same time to counteract that effect and justify letting the merger go through. Usually, firms will try to argue that by merging they can achieve some form of efficiency gain, such as economies of scale, integration of production processes, optimization of logistics and operations, etc. In health care, these sort of efficiency gains generally have to do with coordination of patient care, implementation of better clinical practices, and the ability to negotiate better deals with suppliers of medical devices and other inputs. This last effect is a bit different from the others, and I discuss it in the next section ("Horizontal mergers in vertical markets").

There are some issues you want to be aware of around these efficiency gains. First, the gains have to be significant in order to overcome the increase in markup that the merger will no doubt generate. Second, even if you believe these gains exist and are large enough, you should probably ask whether the only way to achieve them is through a merger. While the merit of these gains has to be evaluated on a merger-by-merger basis, so far there has been almost no evidence of systematic efficiency gains from mergers. That does not mean that these gains don't exist, though it probably means that they are not large enough to matter (or we would have been able to detect them).

There are other ways to argue in favor of a merger. Sometimes firms merge to avoid failure and exit of assets. This requires that the firm could not otherwise reorganize under chapter 11, and that efforts have been made to look for alternative solutions. In pharmaceutical markets, acquisitions of small biotech companies can be an alternative way to finance R&D (instead of financing primary research, a large pharmaceutical company can purchase a small biotech company with a pipeline of candidate drugs and carry those to market). It's important to note that money spent to acquire companies that have promising new technologies *can* be a form of R&D (the promise of a buyout is what motivates many of these companies to work on their innovations). It can also be an anticompetitive tool. Recent evidence shows that at least some fraction of recent pharmaceutical acquisitions resulted in the termination of R&D projects. We will cover this topic more in detail when we discuss innovation in the context of the pharmaceutical market.

Other considerations

There are also other issues that anti-trust policy can think about, though these generally are not considered in the first-pass screenings:

1. Product repositioning: occasionally, mergers can lead to a change in the products offered to the public. This is an issue in markets where changing the characteristics of the product is

easy (e.g. technology/apps).

2. Dynamic effects of mergers: issues to consider here are

   - **Entry**: mergers may increase or decrease the incentives to enter and exit in a market. These effects are very hard to evaluate in advance because entry costs are hard to assess.
   - **Incentive to invest**: the relationship between market structure (e.g. number of firms, and market shares) on the incentive to innovate has been a long-standing point of debate in economics. Theory has provided models that support a positive relationship (more competition $\implies$ more innovation), negative (the reverse), and U-shaped (more competition $\implies$ more innovation, but only up to a point). Until very recently, empirical evidence on the topic was pretty unreliable (it basically consisted of cross-industry comparisons, which are not very informative, because the cost and the benefit of innovation is vastly different across industries). The latest evidence on the topic (Igami and Uetake, 2020) suggests that competition is unequivocally better for innovation, though with two caveats: (i) diminishing returns set in pretty quickly (at around 4-5 firms), and (ii) the evidence comes from a pretty niche model and market (cost-cutting innovation in the the hard-drive disk industry).[1] The fact that we have very little empirical evidence on this topic is actually pretty worrying, since R&D effects can easily dominate any concern about prices in terms of welfare.

3. Impact on coordinated interactions: this is simply the idea that if a market has fewer firms it's easier to collude.

1.3 Horizontal mergers in vertical markets

The guidelines listed above are most appropriate for consumer-good markets. Most firms however, operate in markets where their customers are other firms. If a horizontal merger occurs among firms who are at the highest level of a production chain, most of the same intuitions apply. You have to do some extra work to track how an increase in the price of inputs for other firms' production functions trickles down to the price of the final product, but in terms of economic incentives, the intuition is largely unchanged.

Things are different if you are considering a horizontal merger between firms that act as intermediate links between other firms and final consumers. Firms that act as intermediaries have a dual role. Facing upstream, they are negotiating with their suppliers over the cost of inputs in their production function. Facing downstream they are either doing the same, but as sellers to other firms, or selling to consumers directly. Under standard economic models of negotiations, a merger of two firms results in a better bargaining position. This has two effects for an intermediary. On

[1] There are two issues here that matter. One is general, and it's that cost-reducing innovation is only one part of what innovation is (product innovation is the other). The second is more specific to the HDD market, and it's the fact that over the period being studied that market was slowly disappearing, as HDDs were progressively being replaced by solid-state drives.

the one hand, the standard merger intuition on the price of the final product still applies: because competition falls, markups will rise. However, the merged firm can also use their improved bargaining position to extract lower rates on the inputs they receive from other firms. So the marginal cost of the firm also goes down. As a result, the overall impact on prices is ambiguous.

When we discuss these sort of mergers in the context of healthcare markets I will focus on insurers as intermediaries. However, you can think of most firms in healthcare markets as intermediaries. For example, hospitals also negotiate prices with suppliers of medical devices. The cost of these devices in turn affects the rates that hospitals will try to extract from insurers.

2 Hospital Mergers

We begin our analysis of mergers in healthcare markets by examining hospital mergers. Studies of hospital mergers fall under three categories: (i) reduced-form, (ii) case-studies, and (iii) structural approaches. I cover categories (i) and (iii) here, but if you are interested in reading about some case studies, you can take a look at the references listed in Table 9.9 of the Handbook chapter by Gaynor and Town (2012).

The tradeoffs between approaches (i) and (iii) reflect the well-known tradeoffs of reduced-form vs. structural work. Reduced-form analysis is more agnostic about the underlying economic theory, but cannot easily account for more complicated relationships between market structure and price. In recent years, as the latter factor has become more prominent, the bulk of this research has moved away from reduced-form studies towards more sophisticated structural analysis. Nonetheless, good reduced-form work can still provide important insights on the mechanisms at play.

2.1 Reduced-form analyses of hospital mergers

Reduced-form analysis of hospital mergers generally follow a similar blueprint. These studies calculate a measure of hospital price and regress it on a measure of market concentration, occasionally using mergers as a source of variation. In terms of data construction, the building of both price and concentration measures presents some challenges.

Price measures: the main challenge here is that hospitals do not have a single "price", because they provide many services. Usually studies will try to come up with a sort of aggregate "price index". In doing so, one has to be mindful of various possible sources of error:

- *Service/patient composition*: different hospitals will perform a different mix of services on a different group of patients.
- *Contract structure:* contracts between insurers and hospitals can contain traditional fee-for-service components, but also elements of capitation, as well as adjustments and carve-outs for specific sets of services.

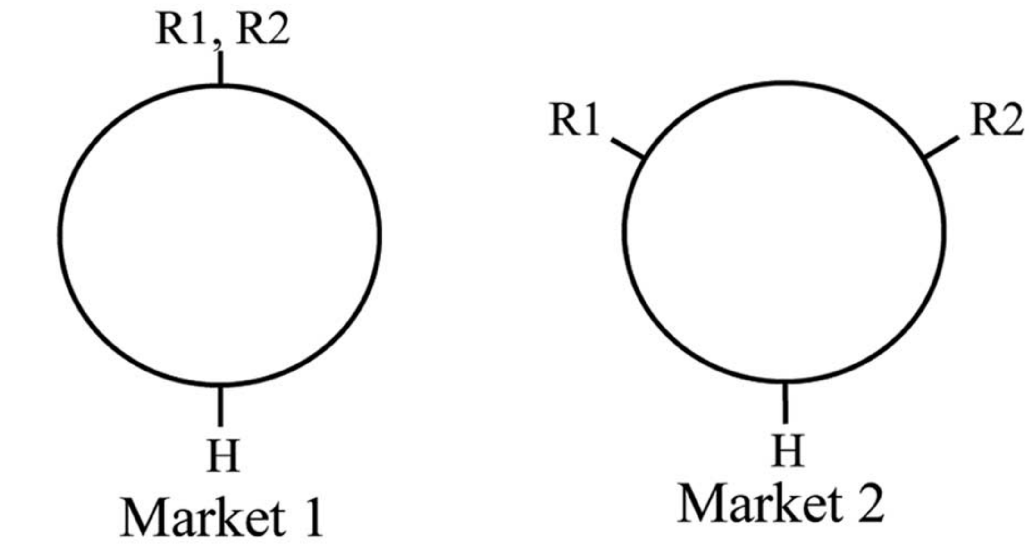

Figure 2: Examples of circular Hotelling markets

- *Non-market compensation:* hospitals also treat patients with government-sponsored insurance (e.g. Medicaid and Medicare enrollees). The reimbursement for the care of these patients does not necessarily react to market structure. For example, reimbursement for Medicare patients is statutorily determined through a formula.

Concentration measure: The most common of these measures (and the easiest to calculate) is HHI. However, HHI glosses over several considerations. Hospitals do not provide homogenous services and draw patients from different areas. Defining a market is therefore not a straightforward exercise here.

The impact of mergers on prices: Dafny (2009)

The main other issue that plagues reduced-form studies is the difficulty in cleanly identifying the impact of merger events. Keep in mind that mergers are not good natural experiments. The firms who are merging are selected endogenously, and even their rival hospitals may be different from other hospitals in ways that are hard to quantify.

Dafny (2009) proposes an instrumental variable strategy to deal with selection in mergers. She argues that hospitals that are closer to one another are more likely to merge, and therefore, hospitals who have rivals that are closely located to one another are more likely to be exposed to mergers.

The theoretical underpinnings of the paper are based on a version of the Hotelling model of competition with differentiated products. nstead of having a straight line, this version of the model consists of a circle. Firms are located along the circle at exogenously determined locations,

and offer otherwise identical products. Consumers have to travel to each firm and so prefer firms that are located closer to them. IFigure 2 shows two examples of these markets. In the first one, two hospitals are located in the same spot, opposite another hospital. In the second, three hospitals are located at equal distance from one another.

The intuition of the model is that hospitals will be able to charge higher prices if they are farther away from other hospitals. In the first market, hospitals R1 and R2 are forced to compete aggressively because they are essentially offering identical products. Hospital H will be able to charge a premium. In Market 2 all hospitals have some market power over consumers that are located closer to them than to their rivals and will be able to charge a markup.

In this model, hospitals that are closely located have a higher incentive to merge because doing so results in a bigger increase in price. Intuitively, if R1 and R2 merged in Market 1 they would be able to go from zero profits, to some positive amount of profits. In Market 2, R1 and R2 are already making some profits, and so their incentive to merge is lower.

The theoretical model provides some justification for using colocation as an instrument for mergers. The paper finds that empirically, being co-located is associated with a 6-7% increase in the likelihood of merging, while having a pair of colocated rivals increases the chance of seeing a rival merger increase by 11-12%.

Another key aspect of the model (that is perhaps obvious), is that prices are strategic complements. In other words, when a rival raises their price, the best response of other firms is to also increase their price. This provides a theoretical justification for running an analysis that looks at price growth for firms that did not merge themselves, but whose rivals did.

In the analysis, the paper ultimately finds that hospitals that are exposed to a merger increased their price by around 35% more than hospitals that were not exposed to a merger. This increase reflects changes between 1997 and 1988, so it translates roughly to an extra 3% price increase on a yearly basis.

Effect of banding together in hospital systems: Melnick and Keeler (2007)

We conclude this section by looking at another paper: Melnick and Keeler (2007). This paper examines price growth of hospitals between 1999 and 2003, and asks whether hospitals that are part of larger hospital systems were able to increase prices by a larger amount. This question is interesting because it provides some insight into why mergers can help firms. Not only do mergers (presumably) lead to higher markups, but they may also put hospitals in a better position going forward.

To put this paper in perspective we need a different model than the simple Hotelling price-setting circle. Instead, we consider a model of bargaining where hospitals set prices in negotiations with insurance companies. The main takeaway from this model is that hospitals that negotiate jointly will be able to extract higher prices from insurers than hospitals that negotiate alone. (If you don't remember why that is, go back to the notes on bargaining from the second week of class, where I cover this model).

The paper evaluates the following statistical model:

$$\log(P_{ht}) = \alpha_t + \beta X_{ht} + \gamma C_{ht} + \delta S_{ht} + \eta C_{ht} S_{ht} + \varepsilon_{ht}$$

where X_{ht} is a vector of control variables, C_{ht} is a measure of market concentration, and S_{ht} is a measure of the size of the system that hospital h is part of. The paper find that prices of hospitals in large system increase by about 30% relative to the prices of hospitals that are not part of any systems. Prices of hospitals that are part of smaller systems increase by about 16% in relative terms.

There are some methodological concerns with this paper. They do not control for hospital fixed effects in the main specification. When they do, the effects essentially disappear, falling to about 8-10% for both small and large systems. Neither estimate is statistically significant. They argue that fixed effects are not appropriate in this case because variables do not change much over time. That's fair (fixed effects can soak up some relevant variation), but it also suggests that they are not really doing a panel analysis as much as they are comparing prices in the cross-section. That's not how the results are interpreted however (the spin is specifically about price growth).

Fixed effects are important because in the absence of a true design for causal identification (which this paper does not have), they can control for potential unobservables that are otherwise impossible to observe. The practical impact of introducing a fixed effect is to limit the identifying variation of the parameter of interest to changes in the variable *within* each unit of observation. In other words, we estimate the impact of hospital systems using only variation in the price of hospitals that are part of such systems, instead of, say, the difference between the prices of hospitals that are part of hospitals systems and hospitals that aren't. The latter comparison is likely to be affected by many unobservables. Hospitals that are part of large systems may be higher quality, they may treat different patients, and they could be located in different areas. Looking at price changes within hospital is more rigorous because we only have to worry about unobservable changes in the characteristics of those hospitals (e.g. maybe quality or patient mix is changing over time). This is usually a much smaller concern (though it can be important in many cases).

2.2 Structural Approach: Gowrisankaran et al. (2015)

The previous two papers we covered are emblematic of the challenges of engaging in reduced-form analysis of mergers. In Dafny (2009), even with the use of IV, we can only look at the impact of a merger on rival firms, and not on the main entity of interest: the merged firm. In Melnick and Keeler (2007), an attempt to examine the implications of the size of hospital systems on prices was severely limited by the lack of an empirical design for causal identification.

In this section we cover a structural approach to the analysis of hospital mergers and competition. I have already discussed the paper we examine (Gowrisankaran et al., 2015) in the bargaining lecture notes. You can refer back to those for more details. Here, I go over the setup very briefly, and then focus on their counterfactual exercises.

Setup

They use a two-stage model that covers (i) hospital-insurer price bargaining, and (ii) hospital choice. Patients are automatically assigned to an insurer (so there is no real adverse selection), and the networks are already fixed.

Hospital choice is based on the Ho (2006) approach that I describe in the narrow networks notes. The main difference is that the negotiated price can affect choices through a coinsurance. Utility is specified as:

$$u_{ij} = \beta x_{ij} - \alpha c_i p_{m(i)j} + \varepsilon_{ij}$$

where c_i is the coinsurance that patient i faces, and $p_{m(i)j}$ indicates the hospital-MCO price (all patients of the same MCO face the same price, though not necessarily the same coinsurance). Note that the coinsurance rate here is fixed, so it cannot react to a change in $p_{m(i)j}$.

Hospital-insurer bargaining follows the standard Nash-in-Nash model adapted to vertical markets, following the theory in Horn and Wolinsky (1988), and the original empirical application in Crawford and Yurukoglu (2012):

$$\max_{\mathbf{p}_m} \left[\sum_{j \in \mathcal{J}_s} q_{mj}\left(\mathcal{N}_m, \mathbf{p}_m\right) \times \left[p_{mj} - mc_{mj}\right] \right]^{b_{s(m)}} \times \left[V_m\left(\mathcal{N}_m, \mathbf{p}_m\right) - V_m\left(\mathcal{N}_m \backslash \mathcal{J}_s, \mathbf{p}_m\right)\right]^{1-b_{s(m)}}$$

Counterfactuals

We skip the estimation and head straight to the counterfactuals. The motivation of the paper comes from a proposed merger that was challenged by the Federal Trade Commission and subsequently withdrawn. In 2006, a system of hospitals in Northern Virginia (Inova Health Systems), proposed the acquisition of an independent hospital, called Prince William Hospital (PWH). Figure 3 shows the geographic location of all hospitals in the market. Inova hospitals are marked in fuchsia, and PWH is circled in red.

Table 1 shows the results for the various counterfactuals that they run. We focus on the first two.

In the first counterfactual, the paper looks at what would have happened if the merger was allowed to go ahead. They find an increase in prices of 3.1%, which results in a fall in quantity by 0.5%, and an increase in profits of 9.3%. All numbers refer to the merged system. Notice that the fall in quantity reflects the (upward) adjustment in prices of rival hospitals. If one looks at the map in Figure 3, it is plausible to conjecture that a large part of this increase in profits comes from the area between PWH, and the Inova Fair Oaks and Fairfax hospitals. Patients in this area do not have a good outside option and therefore MCOs covering those patients will be willing to pay a high price to keep the merged Inova-PWH system in their network.

The most interesting counterfactual is number two. In this counterfactuals Inova and PWH are allowed to merge, but must negotiate separately. This is a remedy that has been used in the past for other hospital mergers. The authors find that this remedy is very inefficient at counteracting

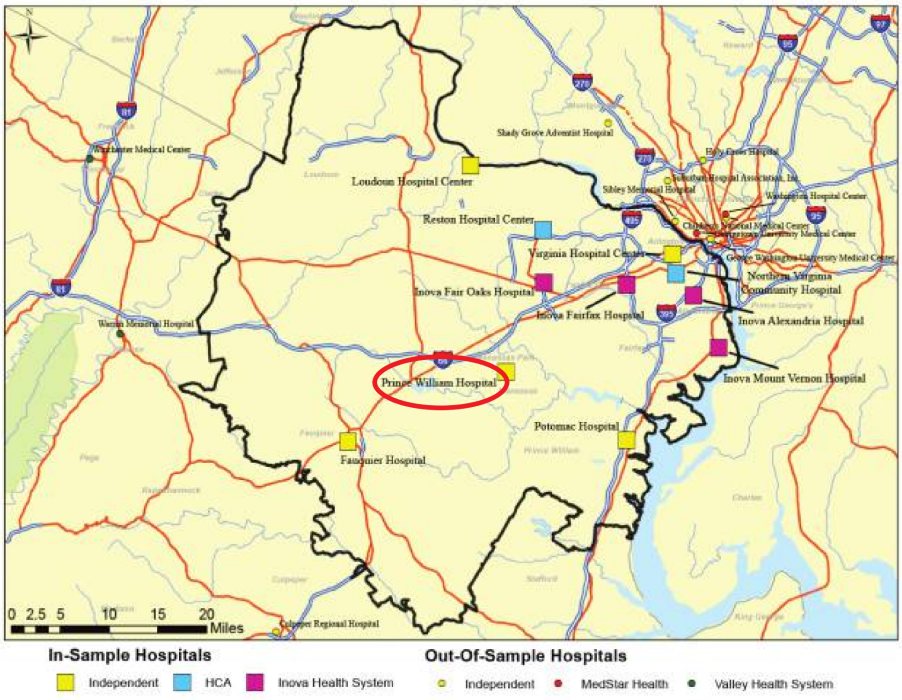

Figure 3: Northern Virginia hospital market (Figure 1 in Gowrisankaran et al., 2015)

the impact of a merger. The counterfactual analysis finds almost exactly the same change in prices and profits for the merged entity. The intuition here is that even though the two parties have to negotiate separately, they are able to internalize the substitution effect (i.e. setting a high price, or even disagreeing to terms results in many patient going to another Inova-PWH hospital, which is still beneficial). As a result, the Nash bargaining equilibrium will be higher.

The paper does not explicitly specify what Nash bargaining function they use in this second counterfactual (it's possible they have something in an appendix), and it's not immediately obvious what they should do. For example, should the profit function include revenues to Inova for PWH (and vice versa)? If it does, then there must be some assumption that the two hospitals have some form of revenue-sharing in place. However, if that's the case, then the better approach would be to include that revenue-sharing function specifically. Excluding revenues to the other party might make more sense if the hospitals are managed by administrators that earn money based on how their own hospital performs. Once again, the exact function may depend on the incentive structure faced by these administrators. Another aspect of this negotiation that you worry about is that there almost certainly is some degree of information sharing between hospital systems that does not occur with hospitals outside the system. So the traditional Nash-in-Nash assumptions (e.g. hold prices fixed in case of disagreement) may no longer apply.

Counterfactual	System	Percent Δ Price	Quantity	Profits
1. Inova/PWH merger	Inova & PWH rival hospitals change at	3.1	−0.5	9.3
	Inova+PW relative to PW base	3.6	1.2	12.0
		30.5	−4.9	91.5
2. Inova/PWH merger with separate bargaining	Inova & PWH rival hospitals	3.3	−0.5	8.8
		3.5	1.2	11.2
3. Loudoun demerger	Inova & Loudoun rival hospitals change at	−1.8	0.1	−4.7
	Inova relative to Loudoun base	−1.6	−0.2	−4.7
		−14.7	0.8	−38.5
4. Breaking up Inova	All hospitals	−6.8	0.05	−18.9

Table 1: Counterfactual results (Table 7 in Gowrisankaran et al., 2015)

2.3 Other aspects of hospital mergers

Evidence of cost-savings: Craig et al. (2020) This paper uses detailed hospital-level transaction data on the cost of medical supplies to estimate the impact of mergers on acquisition costs. The hypothesis is that merging could give hospitals a better bargaining position in negotiating the prices of medical supplies. They use a diff-in-diff approach that compares the change in transaction costs of merging hospitals before and after the merger to the transaction costs of a control group of hospitals that did not merge. They find basically no evidence of an economically significant change. Acquired hospitals tend to do a little better, and acquiring hospitals tend to do a little worse. The overall impact depends on the relative sizes of the two parties, but the average merger reduces acquisition costs by less than 0.5%.

Mergers across markets: Dafny et al. (2019) This paper uses a bargaining model similar to Gowrisankaran et al. (2015) to argue that hospital mergers that involve hospitals operating in different markets could still have an impact on price. Anti-trust has generally ignored these mergers because under simple models of competition, these mergers do not have an impact on price (and therefore do not affect patient welfare). For example, in the circular Hotelling model from Dafny (2009), two hospitals that merged in two different circular "cities" would leave the final equilibrium unchanged. Even in more sophisticated bargaining models, hospitals in different markets do not necessarily do better after merging. That's because there are no patients who are choosing between those two hospitals, which means that the substitution effect from Gowrisankaran et al. (2015) does not apply. Mergers across markets however, may matter if the entity negotiating with the hospital system is a large payer with members across all markets involved. For example, imagine a large employer with offices in all major cities across North Carolina. Each city is a separate market with three hospitals. If a hospital system owned a hospital in every single city, they would be in a very strong position when negotiating with this employer: if negotiations fail, the employer is left with at most two hospitals in every single market. Conversely, if negotiations

succeed, the employer is guaranteed at least one hospital in every market.

This example should clarify two crucial assumptions needed for this result:

1. The hospital system must be able to negotiate an "all-or-nothing" deal. Forcing every hospital in the system to negotiate separately would break the result.
2. There must be at least one payer with members scattered across the markets affected by the merger.

The paper uses a diff-in-diff approach to estimate the impact of these cross-market mergers on prices. They find that cross-market, within-state mergers lead to an increase in price of about 7-8%.

3 INSURER MERGERS

Unlike with hospitals, the effect of insurer concentration and mergers on prices (i.e. premiums) and consumer welfare is theoretically ambiguous. On the one hand, more market power leads to higher markups on premiums. On the other hand, the same market power can lower the rates of services negotiated with healthcare providers, which lowers marginal costs.

Empirical research on the impact of market concentration and mergers in insurance markets has lagged behind relative to work on hospital mergers. In part, this is due to the fact that data on insurance choice is much harder to come by than data on hospital discharges (which can be used to build market shares). The existing work can be split along the same lines as the work on hospital mergers: reduced-form studies and structural analyses. Leemore Dafny's testimony to Congress (Dafny, 2015) provides a good overview of what we have learned from this literature.

3.1 Reduced-form analyses of insurer mergers

We start by looking at reduced-form work. The general empirical strategy of these papers is to regress a measure of premium on market concentration. Just like with hospitals, how to define these measures is not obvious ex-ante. In particular, you want to be aware of a few issues.

Premium: Premium depends on generosity as well as negotiated rates. Insurer concentration may affect plan offering (this is something that virtually no paper considers, as questions of product offering are much more complex than questions of pricing). Beyond that, premiums can vary depending on age, gender and location (e.g. in the ACA Health Exchanges), as well as number of dependents, so you want to make sure you are appropriately accounting for these. For example, it would probably be wrong to use the average premium (since that depends on who signs up, which could change over time). A better option is to use the premium conditional on a set of demographics.

Market concentration: the insurance landscape is a bit more complex than the hospital landscape. Insurers can offer plans directly to patients (e.g. on the ACA Health Exchanges, or through

Medicare Advantage), or to their employers. They can also offer simple administration services to large employers who are self-insured. Mergers of insurers will affect these markets differently. In order to take into account these effects one most likely needs a structural model, though one of the papers I discuss below (Trish and Herring, 2015) is a good example of how one might go about setting up a reduced-form study that takes these different types of relationships into account.

Using mergers as instruments: Dafny et al. (2012)

Dafny et al. (2012) uses a merger of two large insurers to estimate the impact of market concentration on insurance premiums. To do so, they use a large proprietary dataset reporting premiums and enrollment data for a set of several hundred employers across many US markets. The data follows employers over ten years.

The empirical methodology combines the structure of the data and the events of the merger to obtain a plausible path to causal identification. The argument for causal identification follows two steps. First, they focus on the *growth* of employer-specific premiums. This is a way to implicitly control for a variety of factors that could affect premium *levels*, such as employee population, location-specific characteristics, etc. Second, they exploit the sharp change in market concentration following the merger between Aetna and Prudential Healthcare in 1999. Specifically, they exploit the fact that market concentration levels changed more in markets where Aetna and Prudential had a larger market share prior to the merger. I will discuss this strategy more in detail in a second. For now, the point to remember is that simply regressing premium growth on HHI changes is unlikely to yield useful estimates. That's because there are many factors that likely affect both outcomes, such as patient preferences, product offerings, marketing strategies, etc. For example, markets that are hit by adverse economic conditions will likely see both insurer exit, and increased prominence of low-priced (and less generous) options. These two effects will result in a higher market concentration, but also lower premiums. This does not imply that, all else equal, concentration results in lower premiums.

The measure of market concentration they use is HHI. However, as discussed above, simply regressing premium growth on HHI does not yield informative estimates (in the paper, they find basically no relation between the two). Instead, the paper looks at how the Aetna-Prudential merger of 1999 affected the HHI in each market. To do so, they calculate how HHI in 1999 would have been different if Aetna and Prudential had joined forces (1999 is the last year that Aetna and Prudential market shares can be distinguished, since they are reported as a single company starting in 2000). The difference is given by

$$\begin{aligned} sim\Delta HHI_m &= \left(\sum_{j \neq Aetna,Prud} s_{j,1999}^2 + \left(s_{Aetna,1999} + s_{Prud,1999} \right)^2 \right) - \left(\sum_{j \neq Aetna,Prud} s_{j,1999}^2 + s_{Aetna,1999}^2 + s_{Prud,1999}^2 \right) \\ &= \left(s_{Aetna,1999} + s_{Prud,1999} \right)^2 - s_{Aetna,1999}^2 - s_{Prud,1999}^2 \\ &= 2 \times s_{Aetna,1999} \times s_{Prud,1999} \end{aligned}$$

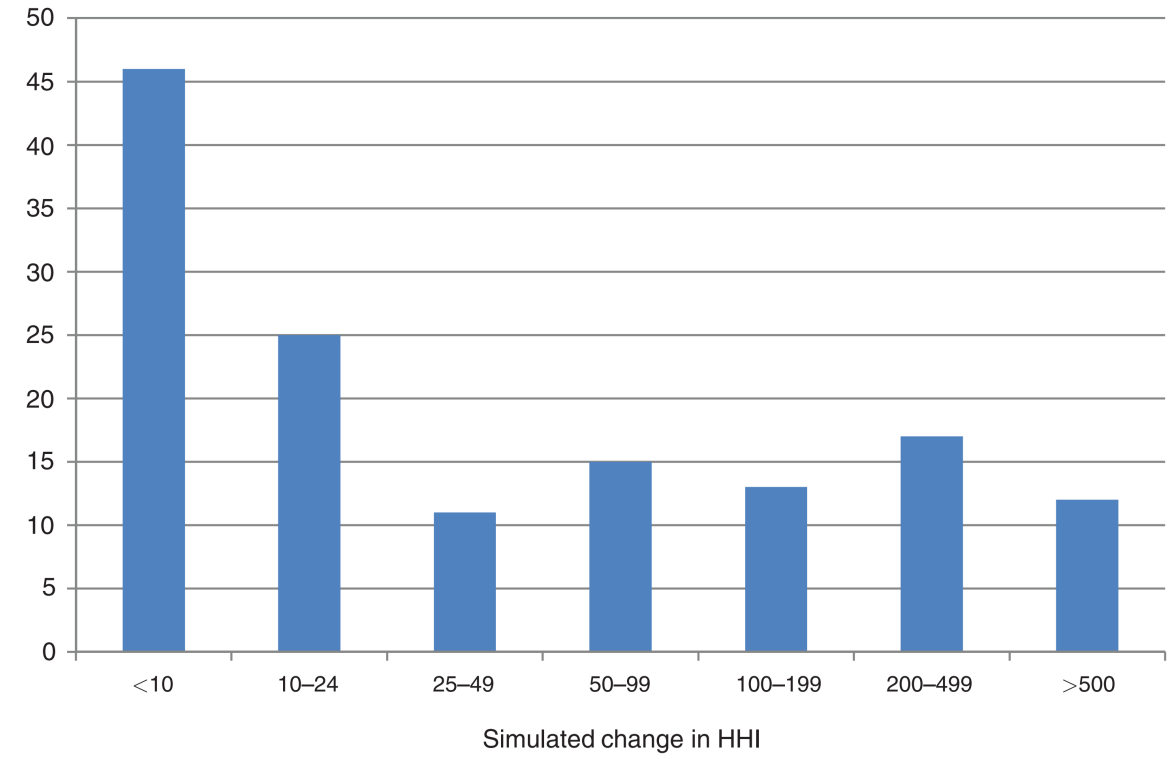

Figure 4: Distribution of $sim\Delta HHI_m$ (Figure 2 in Dafny et al., 2012)

Figure 4 plots the distribution of $sim\Delta HHI_m$ across the 139 markets in the data. You can see that there is a lot of variation in "exposure" to the merger. In particular, both insurance companies were very popular in Texas, so the Department of Justice required the two companies to sell off their Texas business.

In their main specification, the authors use $sim\Delta HHI_m \times \mathbb{I}\{y > 1999\}$ as an instrument for HHI. The idea here is that the merger was motivated by national-level strategic needs, and therefore is plausibly exogenous to local market dynamics. If that assumption holds, then the estimates of the IV regression can be interpreted as causal. The IV approach yields a positive and significant coefficient on HHI (when instrumented using $sim\Delta HHI_m \times \mathbb{I}\{y > 1999\}$). Backing out the effect on premiums is somewhat convoluted (because it matters *when* growth occurs—e.g. growth towards the end of the sample matters less), but they estimate that the overall impact of the consolidation they see over the 10 years of their sample data was responsible for an increase of about 7% in premiums (HHI increased by ~700 points on average during that span).

Some other results:

- They also find a positive impact on prices of firms that were not part of the merger. This mimics the results and strategy of Dafny (2009) on hospital mergers.
- They also look at impact of the merger on salary and employment of healthcare personnel.

They find noisy effects on wages, but large significant negative effects on employment levels. Presumably, the impact was driven by a fall in the negotiated rates with hospitals, which lead to downsizing as a cost-cutting measure. Depending on your point of view this could be a good thing (trimming inefficient spending, etc.), or a bad thing (fall in quality, similar considerations). Interestingly these effects are smaller for nurses than for physicians, which suggests that some of the changes resulted in the substitution of physicians with nurses.

- Dafny et al. (2015) uses a similar strategy to examine the impact of market concentration on premiums in the ACA Health Exchanges.

Taking into account different types of concentration:Trish and Herring (2015)

This paper maintains a reduced-form approach, but tries to think more carefully about how concentration in the insurer market affects both employers and hospitals. The main downside is that it does not necessarily have a strong claim to causal identification.

The key detail you want to keep in mind in terms of background overview is that insurers sell two types of plans to employers on two markets: fully-insured and and self-insured managed care plans. The first types of plans are for (generally smaller) employers who do not want to bear the risk of insurance themselves. The second types of plans are for large employers that bear the full risk of healthcare costs and are simply looking for a manager/administrator. You can generally think of these two markets as separate (and the authors of the paper do), though obviously there is probably some degree of substitution at the margin.

Insurance companies do not charge a premium for self-insured plans. However, since employers still use that insurer's network, the number of patients enrolled in self-insured plans will still matter when negotiating with hospitals. This is the key source of variation that the paper uses. They also control for concentration on the hospital side. The three measures of concentration they have are:

Insurer:Employer HHI This measure of concentration only considers fully-insured plans and it's supposed to measure the market power of insurance companies when setting premiums for employers. All else equal, higher Insurer:Employer HHI should result in *higher* premiums.

Insurer:Hospital HHI This measure includes all types of plans and it's supposed to measure the market power of insurance companies when negotiating rates with providers. All else equal, higher Insurer:Hospital HHI should result in *lower* premiums.

Hospital HHI This is a measure of concentration on the hospital side, which is used as a control variable. All else equal, higher hospital HHI should mean higher cost of insurance, and therefore *higher* premiums.

One thing you might be concerned about here is that the two measures of insurer HHI are quite correlated with one another. That's true, as Figure 5 shows, but there seems to still be plenty of variation across markets.

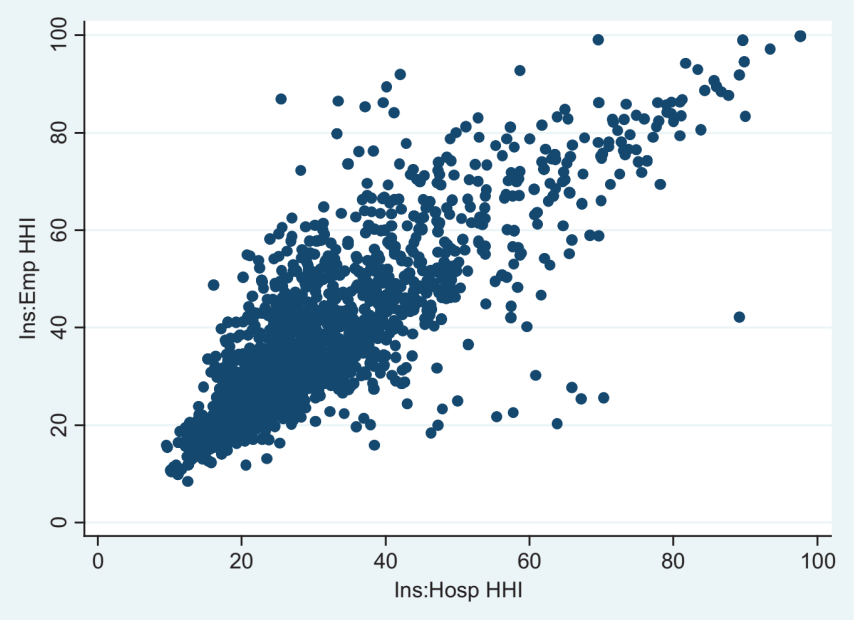

Figure 5: Correlation between Insurer:Employer HHI and Insurer:Hospital HHI (Figure 1, Panel A in Trish and Herring, 2015)

Their analysis consists of simply running an OLS regression that includes the three variables above, as well as a series of other control variables. The results seem to confirm the initial hypothesis. An increase of 100 points in Insurer:Employer HHI, Insurer:Hospital HHI, and Hospital HHI are associated respectively with a premium increase of 0.2%, a decrease by 0.2%, and an increase by 0.2%. As a secondary analysis, they look at the interaction between these various types of market concentrations by performing a stratified analysis. They find suggestive evidence supporting the idea that there are some interactions effects (if you are interested, check out Tables 3 and 4 in the paper, as well as the associated discussion).

One potential concern with the analysis is that the results are quite small. If you recall the results from Dafny et al. (2012), OLS tends to mask the relationship between concentration and premiums. You might suspect that something like this is going in this paper as well. So the most valuable thing that we take from this exercise is perhaps not the estimates themselves (which we suspect are biased downward), but the idea that different types of concentration matter in different ways. This ultimately suggests that a structural approach is the way to go, so that's where we move next.

3.2 A structural approach: Ho and Lee (2017)

Ho and Lee (2017) represents a significant step forward in our understanding and analysis of healthcare markets (it won the Frisch medal as the best applied paper published in Econometrica between 2016 and 2020). The main contribution of the paper is in providing a generalized framework for the analysis the interactions between providers, hospitals, and employers in the

market for health insurance plans. In doing so, they can theoretically separate the various incentives at play (most of which we have already discussed in this section), and then empirically assess their relative importance under various scenarios such as entry/exit of insurers, mergers, and changes in employer plan offerings. The most similar paper among the ones we have covered is Gowrisankaran et al. (2015). Relative to that paper, Ho and Lee (2017) introduces two stages: bargaining over premiums, and insurance choice.

Setup

The model considers the market for private health insurance, which most individuals obtain through an employer or benefits manager. Patients pay a monthly premium that is negotiated by the employer with the insurance company. The insurance company in turn negotiates payment rates to hospitals. It is very similar to many papers we have already covered. In particular, it builds on initial work by Ho (2006), and is almost identical to Ho and Lee (2019), with the exception of the network-formation stage (which is missing here). Like Ho and Lee (2019) (and unlike Gowrisankaran et al., 2015), they do not allow demand to vary with negotiated prices (which is similar to imposing no coinsurance, though not quite as extreme).

The timing of decisions is as follows:

1. Bargaining over premiums, and hospital prices occurs simultaneously (networks are taken as given). This is slightly different than the generalized setting I discussed in the previous set of notes, where premiums are set only after prices are negotiated. What's the difference? Basically, if the negotiation occurs simultaneously, the premium cannot react to a negotiation failure between a hospital and an insurer. Almost certainly this is a computationally convenient assumption, because it simplifies the calculation of the outside option from price negotiations (if you let premiums adjust, the outside option of the insurer is a plan with a smaller network, but a lower premium, and it's probably hard to predict how the premium would adjust).
2. Given networks and premiums, households choose insurance plans
3. Given insurance plan choices (and networks), patients who get sick choose hospitals.

They define the full network of hospitals and insurers as $\mathcal{G}$. This is just a matrix of 0s and 1s that indicates which pairs of hospital and insurers have a link. The set of links of a given insurer j is $\mathcal{G}_j^M$, while the set of links of a given hospital h is $\mathcal{G}_h^H$. The set of all insurers offered by an employer is given by $\mathcal{M}$.

Theoretical model

Stages 2 and 3 are pretty much the same as Ho (2006), and yield

- plan j-specific demand functions $D_{hj}^H(\mathcal{G}, P)$ for hospital h (as a function of the network $\mathcal{G}$ and premiums P)

- plan demand functions $D_j(\mathcal{G}, P)$

From stage 2, they can also derive welfare functions $W(\mathcal{M}, P)$ for each employer (these are simply the welfare implied by the demand functions minus premium payments).

Given those functions, they write profits for plan j as

$$\pi_j^M(\mathcal{G}, p, P) = D_j(\cdot) \times (P_j - \eta_j) - \sum_{h \in \mathcal{G}_j^M} D_{hj}^H(\cdot)\, p_{hj}$$

where p is the matrix of all negotiated hospital prices p_{hj}, and η_j are non-inpatient costs, and profits for hospital h as

$$\pi_h^H(\mathcal{G}, p, P) = \sum_{j \in \mathcal{G}_h^H} D_{hj}^H(\cdot)\,(p_{hj} - c_h)$$

and c_h is the marginal cost of the hospital.

Then, the two bargaining problems can be written as follows.

Employer-Insurer bargaining over premiums

$$P_j = \arg\max_P \left[\pi_j^M(\mathcal{G}, p, \{P, \boldsymbol{P}_{-j}\})\right]^{\tau^P} \times \left[W(\mathcal{M}, \{P, \boldsymbol{P}_{-j}\}) - W(\mathcal{M}\backslash j, \boldsymbol{P}_{-j})\right]^{1-\tau^P}$$

Here the assumption is that no employee will switch employer just because a given plan/insurer is no longer offered. Hence the threat point of the insurer is simply 0 (there is no way to get to those patients other than agreeing to terms with the employer). The threat point of the employer is welfare if that particular insurer is excluded. Premiums of other insurers are denoted in bold because they are held fixed.

Insurer-Hospital bargaining over prices

$$p_{hj} = \arg\max_p \left[\pi_j^M\left(\mathcal{G}, \left\{p_{hj}, \boldsymbol{p}_{-hj}\right\}, \boldsymbol{P}\right) - \pi_j^M\left(\mathcal{G}\backslash hj, \boldsymbol{p}_{-hj}, \boldsymbol{P}\right)\right]^{\tau_j} \times \left[\pi_h^H\left(\mathcal{G}, \left\{p_{hj}, \boldsymbol{p}_{-hj}\right\}, P\right) - \pi_h^H\left(\mathcal{G}\backslash hj, \boldsymbol{p}_{-hj}, \boldsymbol{P}\right)\right]^{1-\tau_j}$$

Here both the insurer and the hospital have an outside option. Severing the link could induce some patients to switch insurer (this is captured in the threat point of the hospital), and it also leads to substitution patterns for members who do not switch (this is captured in the threat point of the insurer). Notice that premiums stay the same however.

The moment conditions for estimation come from the first order conditions of these two bargaining problems:

$$\frac{\partial \pi_j^M(\cdot)}{\partial P_j} = \frac{1-\tau^P}{\tau^P} \times \frac{\pi_j^M(\cdot) \times \left(-\frac{\partial \Delta W(\cdot)}{\partial P_j}\right)}{\Delta W(\cdot)} \tag{1}$$

where $\Delta W(\cdot)$ is the change in welfare of the employer from removing insurer j, and

$$\underbrace{p_{hj}^{\star} D_{hj}^{H}(\cdot)}_{\substack{\text{total received} \\ \text{by hospital}}} = (1-\tau_j)\left[\underbrace{\Delta_{hj} D_j(\cdot) \times (P_j - \eta_j)}_{\substack{\text{1. change in revenue net} \\ \text{of non-inpatient costs}}} - \underbrace{\left(\sum_{\ell \in \mathcal{G}_j^M \setminus hj} p_{\ell j}^{\star} \Delta_{hj} D_j(\cdot)\right)}_{\substack{\text{2. change in payments} \\ \text{to other hospitals}}}\right] +$$

$$\tau_j \times \left[\underbrace{c_h D_{hj}^{H}(\cdot)}_{\text{3. hospital costs}} - \underbrace{\sum_{k \in \mathcal{G}_h^H \setminus hj} \Delta_{hj} D_{hk}(\cdot) \times (p_{hk}^{\star} - c_h)}_{\substack{\text{4. change in profits from patients} \\ \text{switching to other insurers}}}\right] \quad (2)$$

Implications of changes in insurance market structure

We can discuss how a change in the market structure of the insurance market would affect the equilibrium by looking at Equations 1 and 2. For example consider a change that removes an insurer from the market (e.g. through a merger).

Insurer prices

- Less competition among insurers means that it's harder for employers to exclude insurers (i.e. lower elasticity of demand with respect to premiums). This is the classic merger effect, and shows up in Equation 1 in the $\frac{\partial \pi_j^M(\cdot)}{\partial P_j}$ term. This effect would exist even in a world without bargaining (e.g. where $\tau^P = 1$).

- Since $\tau^P < 1$, we also have a change on the right-hand side of Equation 1. In particular, the gains from including any one insurer (i.e. $\Delta W(\cdot)$) generally go up when there is less competition. This has a negative overall effect on the LHS and leads to an even higher price (because usually $\pi_j^M(\cdot)$ is concave).[2]

Hospital prices

[2]This is not a given. In some strange cases, excluding an insurer can lower ΔW for other insurers. For example, if k is a high-cost option, including a low-cost j can be highly beneficial, because it can move some patients into lower-cost plans. However if k is excluded, having j is less important.

- With one less insurer around, other insurers will experience less substitution when excluding hospitals. This reduces term 1. in Equation 2, and puts downward pressure on prices.
- At the same time, less competition leads to higher premiums, so *on the margin*, each patient substituting to a different insurer causes a higher loss in profits. This increases term 1. in Equation 2, and puts upward pressure on prices.
- Changes in other terms are more complicated to sign intuitively because they essentially depend on the agreed upon prices between other hospitals and insurers.

Empirical setting, counterfactuals and implications

The data for the paper comes from the California Public Employees' Retirement System (CalPERS). CalPERS operates a market where three large insurers (Anthem Blue Cross, Blue Shield, and Kaiser Permanente) offer health insurance plans to over one million state employees and retirees (and their families).

Within this context, the paper estimates two main counterfactuals: the removal of Kaiser Permanente, and the removal of Anthem Blue Cross. The two exercises are different because Kaiser is a very popular option (39% overall market share), while Anthem BC is a smaller insurer, with generally higher premiums.

Removing Kaiser The effect of removing Kaiser on premiums is unequivocally one-sided: premiums go up by 14-17%. The effect on hospital prices is also positive though to a different degree for the two remaining insurers. Blue Shield prices remain constant on average (with market-by-market fluctuations). Anthem BC prices go up by 20% on average. What is causing this? There are many moving parts to hospital prices, but two important effects are (i) the reduced substitution effect, which should lower prices, and (ii) the effect of higher premiums, which should increase prices. The (empirical) implication of the model is that (i) and (ii) are roughly equivalent for Blue Shield, but (i) dominates for Anthem BC.

They check that this is the case, by running the simulation while holding premiums fixed. In this scenario, Anthem BC sees roughly no change in prices (remaining variation is also caused by the other three elements of the pricing equation). Their intuition for what is driving this heterogeneity in the response of each insurer is that Kaiser is a much closer substitute to Blue Shield than to Anthem BC (that's driven by the location of their enrollees—apparently Anthem BC's enrollees tend to live far away from Kaiser Hospitals). If that's the case, then the substitution effect for Anthem BC will not change much whether Kaiser is included or not. As a result, the premium effect dominates.

Removing Anthem BC Anthem BC is a smaller, higher cost insurer, and removing it has some counterintuitive implications. First, they find in simulations that premiums fall. The intuition behind this result is the following. Because Anthem BC has higher costs, having both Kaiser and

Blue Shield (as well as Anthem BC) is important because CalPERS wants to steer consumers away from this high-priced plan. If Anthem BC is excluded, then having both insurers may actually be less important, since the need to steer consumers away does not exist anymore. In "math", what happens is that the removal of Anthem BC decreases $\Delta W(\cdot)$, which leads to lower premiums (see Equation 1).[3]

Since premiums fall, so do hospital prices (because now both the premium and the substitution effect point in the same direction).

4 Hospital-Physician Integration

Hospitals employ a large number of physicians. Increasingly, hospitals have also been buying independent physicians practices. The physicians in these practices often keep operating pretty much in the same way as they had been prior to acquisition, but may occasionally be able to bill differently for the same services provided.

Economic research has only recently started to seriously examine the implications of integration of hospitals and physicians. The first forays on this topic were two papers in a 2006 issue of the Journal of Health Economics (Cuellar and Gertler, 2006; Ciliberto and Dranove, 2006).[4] Over the past few years the papers that have been written on the topic have focused on two aspects of integration: (i) the ability to redirect flow of patients through the use of referrals, and (ii) the presence of potentially suboptimal incentives in certain parts of Medicare that may have played a role in fostering greater integration.

In terms of antitrust policy, the initial attitude towards these types of acquisitions was to allow them without additional scrutiny. The initial evidence from Cuellar and Gertler (2006) and Ciliberto and Dranove (2006) pointed to a more mixed impact. Recent work has been less optimistic about the impact of hospital and physician integration, though the evidence is noisy, and not always consistent.[5] Capps et al. (2018) finds that integrated physicians charge higher prices, and may lead to increased patient spending. Baker et al. (2016) finds that integration of hospitals and physicians increases the within-system referrals in a way that prompts patients to go to more expensive and lower-quality hospitals. Heebsh (2020) finds that, after integrating, physicians are more likely to recommend more intensive procedures (which need to be administered in a hospital setting).

Historically, it's important to note the role that managed care had in encouraging integration. Payment reforms such as capitation that were introduced as part of the managed care wave gave a

[3] This weird premium effect disappears if premiums are set in a more traditional Nash-Bertrand equilibrium.

[4] Interestingly, a third paper in that same issue (Huckman, 2006), studied how horizontal integration between hospitals affected referral patterns. Referral patterns are a distinct "vertical" aspect (you can think of the referring agent as the upstream firm), and would later became a major part of studies of physician-hospital integration.

[5] Usually papers distinguish between two types of hospital-physician groups. Physician-hospital organizations (PHOs) are independent contracting entities that represent groups of otherwise independent physicians and hospitals. Integrated salary model (ISM) hospitals represent fully integrated groups of hospitals and physician groups. Different papers have found different effects for these types of organizations.

financial incentives for hospital and physicians to work together and achieve economies of scope through care coordination and information sharing. In many cases, the easiest way to do so was simply to join forces.

4.1 A theoretical model of hospital-physician integration

Some of the discussion in this section is based on Gaynor (2006). Vertical integration between hospitals and physicians can be efficiency enhancing by lowering transaction costs, ensuring supply of an input, or improving monitoring. In addition, there are gains from pricing coordination because doctors and hospitals sell complementary products. However, since the firms are independent, they do not coordinate pricing. This means there can be an efficiency gain from integration due to a single firm internalizing the effects that physician and hospital pricing decisions have on each other.

Anticompetitive concerns can arise due to foreclosure (i.e. redirecting patients to other providers within the system even when it may be suboptimal to do so), and increased bargaining power with insurers.

Since we have already discussed bargaining in our overview of hospital mergers and insurer mergers, here we focus on the incentives that may arise because of foreclosure.

To present the relevant incentives I borrow a model from a recent job market paper (Brotgoldberg and Vaan, 2018), which in turn borrows ideas from the literature on joint patient and physician decision-making started by Ellis and McGuire (1986). It's a partial equilibrium model, where prices do not adjust after vertical integration. That means that it may miss some effects. However, it does describe really well the foreclosure incentives we are interested in.

Choice function

We start from an assumption that observed choices in the data are generated from a decision-making process that takes into account both the financial windfall to the physician as well as the well-being of the patient. This can be a model of joint decision-making between the physician and the patient, or an altruistic model of physician choice.

The paper examines the role of PCPs in the market for surgical joint replacement procedures, but its insights are broadly applicable. The utility of patient i, in the care of PCP j, of receiving the procedure from orthopedist k is given by

$$u_{ijk} = \Psi_j f\left(X_i, r_i \mathbb{E}\left[Y_{ijk}\right], Z_k\right) + V_{jk} T_{ijk} \tag{3}$$

where

- $f(\cdot)$ represents benefit to the patient (as a function of patient characteristics X_i, expected cost-sharing $r_i \mathbb{E}\left[Y_{ijk}\right]$, and surgeon characteristics Z_k.
- Ψ_j measures the weight of patient benefit (relative to financial benefits)

- V_{jk} is an indicator for vertical integration between PCP j and surgeon k
- T_{ijk} is the financial compensation for an internal referral.

To this equation, one also needs to add a model of patient cost, e.g.

$$Y_{ijk} = g\left(X_i, Z_k, V_{jk}\right)$$

Notice that PCPs only receive compensation if they are integrated. This comes from the Stark Law, which prohibits kickbacks between unaffiliated hospitals and physicians (the law also regulates referral payments in integrated systems, but there are usually way to get around those limits by making the payment ties non explicit). One way to introduce some financial aspect for unaffiliated PCPs would be to put them under some form of capitation contract. In this case, you could add a term

$$B_{ij} - b_{ij}\mathbb{E}\left[Y_{ijk}\right]$$

where B_{ij} is the "target" and $b_{ij}\mathbb{E}\left[Y_{ijk}\right]$ represents the fraction of expected spending that would count against the target. Notice that in this case, capitation is essentially a supply-side instrument to deal with a demand-side issue (overutilization drive by moral hazard, which occurs whenever $r_i < 1$).

Adding a logit error to Equation 3 lets you go from utility to analytic choice probabilities.

Implications

Focusing on the implications of vertical integration on choices, there are two avenues for change:

1. The cost of treating a patient may change. The difference is given by

   $$\Delta Y_{ijk} = g\left(X_i, Z_k, 1\right) - g\left(X_i, Z_k, 0\right)$$

   We generally think (hope) that $\Delta Y_{ijk} < 0$, so this should incentivize PCPs to increase within-system referrals.

2. The PCP receives an additional financial incentive T_{ijk} that makes them even more likely to engage in within-system referrals.

Since both incentives move in the same direction we cannot point to increased within-system referrals as evidence of either one alone. (A decrease in within-system referrals could only be attributable to perverse cost effects, and also incredibly weird).

4.2 Empirical evidence on the impact of hospital-physician integration

Table 2 provides an overview of several recent papers on the topic of hospital-physician integration. The findings tend to paint a pretty negative picture of hospital-physician integration.

Paper	Empirical Setting	Methodology	Results
Cuellar and Gertler (2006)	Arizona, Florida and Wisconsin (1994–1998)	Panel analysis of price, cost, and volume for a sample if differentially integrated physician groups	Integration results in higher prices without any apparent change to cost or quality.
Ciliberto and Dranove (2006)	California (1994–2001)	Panel analysis of price, for a sample if differentially integrated physician groups	Integration results in a small but noisy decline in prices
Baker et al. (2016)	20% randomized Medicare sample (2009)	Random utility choice model of referrals and hospital choice	Integration associated to 30% higher probability to referral to the owning hospital
Capps et al. (2018)	Administrative claims from several states (2007–2013)	DID regression around integration.	Physicians that integrate with hospitals raise prices by 14.1%. Half of the increase comes from exploiting payment rules
Heebsh (2020)	20% randomized Medicare sample (2008–2012)	IV design using Medicare payment rules as instrument for VI	Integrated cardiologists are more likely to recommend intensive treatment, with worse patient outcomes

Table 2: Overview of recent evidence on hospital-physician integration

However, there are issues with most of these papers that could compromise the interpretation of the results. The biggest challenge is that integration is not random, and so finding a good natural experiment, or a good instrument is hard. This is the same problem we saw in the previous two sections. Unlike with hospital and insurer mergers however, to the best of my knowledge no one has published a fully-formed structural model of the integration process and its consequences.[6] This is a big gap in the literature.

A promising avenue to possibly obtain causal identification is a change in Medicare reimbursement rules, which Dranove and Ody (2019) show can explain part of the trend of hospital-

[6]The paper whose model I reference (Brot-goldberg and Vaan, 2018), is perhaps the lone exception. However, at the time of writing, the paper has been removed from the author's website, with an update coming soon. As a result, I refrain from discussing the results of the paper for the time being.

physician integration.[7]

The part of Medicare's reimbursement rules that affects vertical integration incentives is the following (this is lifted almost verbatim from Dranove and Ody). Physicians can perform procedures in numerous sites, such as an office, a hospital outpatient department, or an emergency room. Regardless of the site, Medicare reimburses physicians using the physician fee schedule (PFS). When an independent physician (i.e., a physician who is not a hospital employee) performs a procedure in a non-facility, the PFS covers all costs. If an independent physician performs a procedure in a facility, or anytime a VI physician (i.e., a hospital-employed physician) performs a procedure, Medicare reduces the PFS price and the facility submits a separate bill under Medicare's Outpatient Prospective Payment System (OPPS). The combined price is almost always higher than the single PFS price.

Why does this represent an incentive for vertical innovation? Because VI physicians don't need to actually perform the procedure in a facility setting. Hence, simply by integrating, the hospital and the physician can get a larger reimbursement from Medicare. Dranove and Ody (2019) use a 2010 shock to Medicare's PFS that made the discrepancy between physician-only and facility reimbursement even larger. The shock affected various procedures differently, so they estimate the impact of the Medicare site-of-service differential on vertical integration by comparing the likelihood of integrating of physicians whose service "profile" made them particularly good candidates for integration post-reform (this is almost identical to the identification strategy in Clemens and Gottlieb, 2014).

A recent job market paper (Heebsh, 2020) uses the same Medicare reimbursement reform as an instrument for vertical integration. As far as I can tell, this is one of the cleanest identification strategies for the impact of vertical integration. Moreover, given that his setting is relatively limited (treatment of arterial blockages), I think there is ample margin to repeat the same identification design in a different setting.

5 Other topics in healthcare and antitrust

5.1 Price discrimination and transparency in bargaining

One important aspect of the contractual relationships between providers, insurers, and employers is the negotiated nature of prices and (often) premiums. These negotiations result in contractual terms that are almost always undisclosed. In turn, the "shrouded" nature of contractual terms can result in different "buyers" paying substantially different prices for the same product or service.

This sort of discrimination often raises the issue of whether there could be some intervention that would make prices more uniform, and in so doing, lower the acquisition costs of buyers that are getting particularly bad deals. This question is tackled in a series of papers on the market for medical devices by Grennan (2013) and Grennan and Swanson (2020).

[7] This reimbursement rule is what Capps et al. (2018) refer to when they say "exploiting payment rules".

Grennan (2013) uses a large detailed dataset of transacted prices and quantities between hospitals and manufacturers of coronary stents to estimate a simple (relatively speaking) model of firm-to-firm bargaining. From the bargaining model he extracts estimates of marginal cost (per stent), and hospital-manufacturer-specific bargaining weights. He then examines three possible ways of uniforming prices:

1. Switching to a Nash-Bertrand equilibrium where manufacturers simply post prices;
2. Centrally negotiated prices for all hospitals;
3. Negotiations at the level of merged hospital systems.

Perhaps unsurprisingly, the implications of uniform prices depend on *how* one achieves such uniformity. Switching to Nash Bertrand yields overall price increases because hospitals tend to have pretty heterogeneous preferences over stents. The optimal response of manufacturers in this case is to price to the "inelastic" portion of the demand curve, e.g. try to extract as much surplus from the hospitals that have high preferences for their brand and pretty much abandon the rest of the market. Notice that this was not a foregone conclusion. In markets where preferences are more "vertical" (i.e. all hospitals have similar preferences over which stent is best), theory predicts that Nash Bertrand would lead to lower prices. Grennan argues that this equilibrium may be what happens if a government agency imposed complete price transparency, making it all but impossible to price-discriminate. I actually don't think that connection is quite as clean (and you'll see why in the next paper).

The outcome under central negotiations would depend on bargaining weights. A counterfactual analysis under the assumption that the bargaining weight of the entire set of hospitals is equal to the average estimated bargaining weight actually yields a slight price increase. More generous assumptions tend to lead to large savings. Here you probably worry a bit about what really is behind bargaining weights. There is no microfoundation (in this paper or anywhere else) that really matches to intuitive notions that apply to firms.[8] Finally, the counterfactual that simulates bargaining after various (hypothetical) mergers take place has much of the same intuition as the market-level bargaining counterfactual: the results depend on assumptions on the bargaining weights.

Grennan and Swanson (2020) tackles the "transparency" counterfactual in a more direct way. They take advantage of an industry database that collects purchasing costs and transactions by a set of hospitals. Participating hospitals can use the database to check how much other hospitals are paying, but must provide their own data, including a 12-month panel of historical data. This pre-panel allows the authors to examine the impact of information on rivals' costs on a hospital's bargaining outcomes.

Their main result is that hospitals with relatively high prices manage to obtain lower prices after accessing the database. They propose two candidate explanations. One is a model of asymmet-

[8] Rubinstein (1982) uses discount rates to microfound bargaining weights, but that does not seem like something that would really apply here.

ric information, where contracting prices of other hospitals reveal additional information about the marginal cost of the supplier. The second is an agency model, where information about other hospital's prices is essentially a monitoring technology that allows the hospital to check the effort of the agent negotiating the prices. To test which hypothesis fits the data best, they exploit the fact that if agency is the leading option then it is not necessary to have immediate access to information on other hospital prices, as long as this information will be revealed at some point in the future. They test this hypothesis using data on new brands (the idea is that information about prices of other hospitals for new brands will be revealed *after* the contract has been signed, since everyone is negotiating simultaneously, and for the first time). The results point towards asymmetry being the main mechanism.

A parting thought on this paper is that, compared to the previous one we discussed, it treats transparency with a lot more nuance. First, they have a story about heterogeneity in bargaining outcomes that is not simply about bargaining weights (the asymmetric information theory). Second, they propose what I would consider a pretty realistic (potential) microfoundation for bargaining weights through the agency model. That story ends up not panning out, but my personal instinct here is that the test is not too powerful, because it does require pretty strict rational expectation assumptions that a myopic or otherwise behavioral agent may not satisfy (plus their results are noisy, which does not mean that an effect does not exist, just that it's hard to tell).

5.2 Consolidation and governance

One aspect of consolidation that few papers focus on (and none of the ones we have previously discussed in these notes), is how mergers and acquisitions may impact the inner workings of the firm. This is largely for two reasons (i) simpler models of profit-maximization or bargaining are easier to solve and understand, and (ii) data on the "inner workings" hardly ever exists.

There are however some exceptions, and we cover them in this section. In particular we look at two papers that examine three different instances of consolidation; one in the dialysis industry (Eliason et al., 2020b), and one on nursing homes (Gupta et al., 2021). Note that these are very recent papers.The second one hadn't even come out when I started writing this chapter of the notes.

The common theme of these two papers is that not only do they study how the acquisition of firms affects market outcomes (like price and quality), but they can also observe (or infer) the operational changes that are causing these changes.

The first paper (Eliason et al., 2020b), examines how acquisitions of independently operating dialysis clinics by large national chains affects the care received by patients at these clinics. Dialysis is a medical procedure used to clean the blood of patients with end-stage renal disease (ESRD). Patients in this stage of the disease need regular treatment (or a kidney transplant) in order to survive.

The identification strategy of this paper is similar to that of other reduced-form papers we have seen previously. They use facility fixed effects and a diff-in-diff setup that implicitly compares the

change in outcomes within an acquired facility to a "control" group of other facilities that were not acquired. This strategy is not bulletproof, because acquisitions are not random. However, this paper can leverage a very detailed patient-level dataset to control for most potential sources of selection and endogeneity.

They find a series of results, all pointing towards two facts: facilities that are acquired by a large national chain alter treatment to (i) increase reimbursement, and (ii) decrease costs. In particular, they find that acquisition is linked to:

1. Dramatic increase in the use of prescription drug Epogen, which is reimbursed separately by Medicare, and carries large profit margins (especially for large chains, which can negotiate deeper discounts with the manufacturer);

2. A substitution in the prescription of an iron-deficiency drug (Ferrlecit) with a near-identical substitute that offers higher reimbursement;

3. An increase in the patient load per employee, and per dialysis station;

Taken together, these changes lower the quality of care received by patients. In particular, integration is linked to:

1. An increased likelihood of hospitalization;

2. An overall decrease in the survival rate;

3. A decrease in the likelihood of receiving a kidney transplant;

4. A decrease in a range of other metrics linked to quality of care and quality of life.

All the while, spending actually grows.

The paper goes a bit into the details of why these changes are implemented (or, more appropriately, why were they not implemented before by what we were assuming are profit-maximizing firms). Market power turns out to not matter all that much. Most of these acquisitions have a negligible impact on concentration, and there does not seem to be a differential effect for clinics in very concentrated markets. Rather, what seems to be happening is that the characteristics and behavior of acquired firms converges to that of the other clinics of their parent company. This suggests a different channel for the impact of acquisitions, one that has nothing to do with competition, but that can still have a very negative impact on patient welfare.

You may find, hidden somewhere in here, a sort of philosophical question about how we should regulate these sort of acquisitions. Is antitrust the right tool? Maybe not, because these acquisitions don't seem to necessarily have an anticompetitive aspect to them (then again, the large chains in this market control over 60% of the market, so maybe there is something to be said about antitrust). Better tools probably have more to do with (i) removing certain "easy" margins for profit, such as drug reimbursements (the government did do this, and the authors study the change in a follow-up paper, Eliason et al., 2020a) and (ii) implementing some quality control.

Note that the government does this, to an extent, through the Five Star rating system. Maybe the answer is to impose fines for facilities that consistently perform below a certain minimum threshold.

The second paper (Gupta et al., 2021), focuses on nursing homes (or skilled nursing facilities, or SNFs) that areacquired by private equity (PE) investment companies. By way of background, private equity firms are entities focused on the quick turnaround of unprofitable businesses.[9] Acquisitions occur primarily through debt financing, and the goal of an acquisition is to raise the profit margin and re-sell the business after a relatively short time (a few years or so). You can already see a few problems with this model. First, it is incredibly short-sighted, so it will sacrifice long-term viability for short-term profits. Second, it has very large incentives to cut costs (and therefore reduce quality). In a market where quality is hard to gauge, these sort of very high-powered incentives can produce suboptimal outcomes (for patients).

The identification strategy is once again very similar. They use facility fixed effects, and look at changes before and after acquisition. They don't do much (as far as I can tell) to address endogeneity/selection in the acquisition decisions. They do however, control for patient mix (which is one component of this sort of endogeneity and selection). Patient mix matters more here than in other contexts because SNFs can screen patients (and do: see Gandhi, 2020). Their instrument here is a distance-based instrument, which measures the difference in the zip-code distance between the patient's address and the closest PE-owned and non-PE-owned nursing homes.

The results are qualitatively similar to what Eliason et al. (2020b) find in the context of dialysis clinics: patient outcomes and measurable quality metrics get worse, and spending goes up. In particular:

1. Patient mortality goes up by 1.7pp (a 10% increase). This effect appears to be loaded on older, but healthier patients, and looks like it could be linked to a massive increase in the prescription of antipsychotics (by 3pp, or 50% of the mean). Use of antipsychotics is usually discouraged in older patients due to a higher chance of fatal side effects.

2. Total amount billed per patient stay goes up by almost 20% (11% if we include the 90 days following a discharge). The channel for this increase was not clear to me as I read the paper. As far as I can tell, this is not due to an increase in the use of antipsychotics (though it may be), because nursing homes are reimbursed on a per-diem basis.

Taken together, the results of these two papers raise some very challenging questions for regulators of healthcare markets. In both cases, firms that seem to operate with a profit-maximizing motive end up extracting higher margins through reduction in quality, and ultimately, loss of patient welfare. The papers do not explicitly propose an explanation for this phenomenon, but the main friction that could explain these results is that quality is very hard to assess in healthcare markets, not only before making provider choices, but sometimes even after provider choices have been made (after all, dialysis patients go back to these clinics 3 times a week). Given this

[9] Disclaimer: I know very little about private equity. Most of what I write here I learned from the paper itself.

evidence, it would not be a stretch to suggest that existing frictions in the healthcare market may make a traditional market equilibrium a poor fit for welfare maximization. At the very least, it is discouraging to see that what should be, by all metrics, "better" firms (e.g. firms with more information, more know-how, better-aligned incentives, etc.) succeed not by making a better product, but by shirking on hard-to-measure dimensions, and by gaming reimbursement rules.

5.3 The role of intermediators

Finally, one last topic I want to mention is the increased presence of negotiation intermediaries in healthcare markets. Examples of such intermediaries include group purchasing organizations (GPOs), who negotiate contracts for the supply of medicad devices and other inputs on behalf of hospitals, and pharmacy benefit managers (PBMs), who negotiate prices of pharmaceutical products on behalf of insurance companies.

There has been almost no specific research directed at the role of these organizations. The second counterfactual in Grennan (2013) (see Section 5.1) could be interpreted as simulating the introduction of a large intermediary. However, the role of such an intermediary is somewhat limited. They are basically a large negotiator intervenining on behalf of a bunch of different clients. Some of the research on physician-hospital consolidation could also speak to this question because one of the organizational units that links hospitals and physicians is a group that only links the two for negotiation purposes (much like a GPO).

All in all, this research does not really address the questions of (i) what do these intermediaries really do, and (ii) do we even need them around. The only paper I am aware of who tries to tackle these two questions is Feng (2019), which specifies a model of PBMs and tries to understand their role a little bit better. PBMs are of particular public policy interest because they have recently come under fire, together with drug manufacturers, for their role in the growing burden of out-of-pocket costs of prescription drugs.

6 BIBLIOGRAPHY

Baker, Laurence C, M Kate Bundorf, and Daniel P Kessler, "The effect of hospital/physician integration on hospital choice," *Journal of Health Economics*, 2016, *50*, 1–8.

Brot-goldberg, Zarek C and Mathijs De Vaan, "Intermediation and Vertical Integration in the Market for Surgeons," 2018.

Capps, Cory, David Dranove, and Christopher Ody, "The effect of hospital acquisitions of physician practices on prices and spending," *Journal of Health Economics*, 2018, *59*, 139–152.

Ciliberto, Federico and David Dranove, "The effect of physician-hospital affiliations on hospital prices in California," *Journal of Health Economics*, 2006, *25* (1), 29–38.

Clemens, Jeffrey and Joshua D. Gottlieb, "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?," *American Economic Review*, 2014, *104* (4), 1320–1349.

Craig, Stuart, Matthew Grennan, and Ashley Swanson, "Mergers and Marginal Costs: New Evidence on Hospital Buyer Power," 2020.

Crawford, Gregory S. and Ali Yurukoglu, "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, apr 2012, *102* (2), 643–685.

Cuellar, Alison Evans and Paul J. Gertler, "Strategic integration of hospitals and physicians," *Journal of Health Economics*, 2006, *25* (1), 1–28.

Dafny, Leemore, "Estimation and Identification of Merger Effects: An Application to Hospital Mergers," *Journal of Law and Economics*, 2009, *52* (August), 523–550.

_ , Jonathan Gruber, and Christopher Ody, "More Insurers Lower Premiums: Evidence from Initial Pricing in the Health Insurance Marketplaces," *American Journal of Health Economics*, jan 2015, *1* (1), 53–81.

_ , Kate Ho, and Robin S. Lee, "The price effects of cross-market mergers: theory and evidence from the hospital industry," *RAND Journal of Economics*, 2019, *50* (2), 286–325.

_ , Mark Duggan, and Subramaniam Ramanarayanan, "Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry," *American Economic Review*, apr 2012, *102* (2), 1161–1185.

Dafny, Leemore S., "Health Insurance Industry Consolidation: What Do We Know From the Past, Is It Relevant in Light of the ACA, and What Should We Ask?," 2015.

Dranove, David and Christopher Ody, "Employed for Higher Pay? How Medicare Payment Rules Affect Hospital Employment of Physicians," *American Economic Journal: Economic Policy*, 2019, *11* (4), 249–271.

Eliason, Paul J, Benjamin Heebsh, Riley J League, Ryan C Mcdevitt, and James W Roberts, "The Effect of Bundled Payments on Provider Behavior and Patient Outcomes," 2020.

Eliason, Paul J., Benjamin Heebsh, Ryan C. McDevitt, and James W. Roberts, "How Acquisitions Affect Firm Behavior and Performance: Evidence from the Dialysis Industry," *Quarterly Journal of Economics*, 2020, *135* (1), 221–267.

Ellis, Randall P. and Thomas G. McGuire, "Provider behavior under prospective reimbursement. Cost sharing and supply," *Journal of Health Economics*, 1986, *5* (2), 129–151.

Feng, Josh, "Pricing Intermediaries in Prescription Drug Markets: Impact and Policy Implications," 2019.

Gandhi, Ashvin, "Picking Your Patients: Selective Admissions in the Nursing Home Industry," *SSRN Electronic Journal*, 2020.

Gaynor, Martin, "Is vertical integration anticompetitive? Definitely maybe (but that's not final)," *Journal of Health Economics*, 2006, *25* (1), 175–180.

_ **and Robert J. Town**, *Competition in Health Care Markets*, Vol. 2, Elsevier B.V., 2012.

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

Grennan, Matthew, "Bargaining ability and price discrimination: Emprical evidence from medical devices," *American Economic Review*, 2013, *103* (1), 145–177.

_ **and Ashley Swanson**, "Transparency and negotiated prices: The value of information in hospital-supplier bargaining," *Journal of Political Economy*, 2020, *128* (4), 1234–1268.

Gupta, Atul, Sabrina T Howell, Constantine Yannelis, and Abhinav Gupta, "Does Private Equity Investment in Healthcare Benefit Patients? Evidence from Nursing Homes," *SSRN Electronic Journal*, 2021.

Heebsh, Benjamin, "Vertical Integration and Treatment Choices: Evidence from Cardiologists," 2020.

Ho, Kate and Robin S. Lee, "Insurer Competition in Health Care Markets," *Econometrica*, 2017, *85* (2), 379–417.

_ **and** _ , "Equilibrium provider networks: Bargaining and exclusion in health care markets," *American Economic Review*, 2019, *109* (2), 473–522.

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

Horn, Henrick and Asher Wolinsky, "Bilateral Monopolies and Incentives for Merger," *The RAND Journal of Economics*, 1988, *19* (3), 408–419.

Huckman, Robert S., "Hospital integration and vertical consolidation: An analysis of acquisitions in New York State," *Journal of Health Economics*, 2006, *25* (1), 58–80.

Igami, Mitsuru and Kosuke Uetake, "Mergers, Innovation, and Entry-Exit Dynamics: Consolidation of the Hard Disk Drive Industry, 1996-2016," *Review of Economic Studies*, 2020, *87* (6), 2671–2702.

Melnick, Glenn and Emmett Keeler, "The effects of multi-hospital systems on hospital prices," *Journal of Health Economics*, 2007, *26*, 400–413.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 1982, *50* (1), 97–109.

Trish, Erin E. and Bradley J. Herring, "How do health insurer market concentration and bargaining power with hospitals affect health insurance premiums?," *Journal of Health Economics*, 2015, *42*, 104–114.