Simulating Hospital Merger Simulations*

David J. Balan
Keith J. Brand
Bureau of Economics
Federal Trade Commission
dbalan@ftc.gov
kbrand@ftc.gov

November 24, 2009

[VERY PRELIMINARY DRAFT. PLEASE DO NOT CITE, CIRCULATE, OR QUOTE]

Abstract: In recent years, researchers have developed a number of new methods for predicting the price effects of hospital mergers. Though there are several variants, the basic steps are the same. First estimate a discrete choice model of hospital choices; then use these estimates to generate a hospital-level measure of market power (a large part of the innovation was in creating new market power measures that have certain attractive properties); and then use the market power measure as an independent variable in a hospital price regression. Finally, use the estimated relationship between the market power measure and price to simulate the effects of mergers. In this paper, we seek to test the accuracy of these simulation methods. To do this, we set up a simple model of hospital competition which can, for any given values of the parameters of the model, generate the "true" effects of a merger between any two hospitals. These "true" effects are then compared to the effects predicted by the simulation methods described above. We repeat this exercise 32,400 times and, using each of several market power measures, derive results regarding the conditions under which the simulation method does or does not generate predicted effects that are close to the "truth." Our preliminary results suggest that the simulation methods slightly under-predict merger effects on average, and that this under-prediction becomes more pronounced as the diversion between the merging hospitals increases.

* We are grateful to Chris Garmon, Robert McMillan, David Schmidt, and Robert Town; to seminar participants at George Washington University, the University of Illinois, the University of Oklahoma, and the 2008 International Industrial Organization Conference; and to brown bag participants at the Department of Justice and at the Federal Trade Commission for their helpful comments. The views expressed in this paper are those of the authors and do not represent the views of the Federal Trade Commission or of any individual Commissioner.

I. Introduction

In recent years, researchers have developed a number of new methods for predicting the price effects of hospital mergers.[1] These methods are mostly variations on the same theme. First they use patient discharge data to estimate a discrete choice model of hospital choices. Second, they use these estimates to construct a measure of each hospital's market power. (As discussed in detail below, these measures are cleverly designed and have some attractive properties). Third, the market power measures are included on the right-hand side of a hospital price regression. Fourth, the estimates of the effect of market power on price are used to generate predictions regarding the price effects of mergers. The purpose of this paper is to make a contribution to evaluating the accuracy of these methods.

One obvious way to test these methods would be to gather data on a large number of actual hospital mergers, apply the methods to pre-merger data to generate predictions of the merger effects, use post-merger data to measure the actual effects of each merger, and then compare. The problem with such an exercise is its scope: gathering good data on a sufficient number of mergers that such a test would have statistical power would be a very daunting task. This problem is made worse by the fact that the best available price data, which is data on actual transaction prices, is generally proprietary and unavailable to researchers.[2]

In this paper, we take a different approach. Instead of gathering data on a large number of hospital mergers, we randomly generate our own "data" on the attributes of hospitals (geographical location and hospital fixed effects) and on patient attributes and preferences (location, the degree of their distaste for travel, and idiosyncratic preferences for each hospital). We then set up

[1] See Town & Vistnes (2001); Capps, Dranove & Satterthwaite (2003); Gaynor & Vogt (2003); Capps & Dranove (2004); and Melnick & Keeler (2007).
[2] The alternative, which is used by most academic researchers, is to use publicly available "billed charges" data and then to estimate price by multiplying these charges by the publicly available "cost-to-charge" ratios.

a Nash bargaining game between each hospital and a monopoly Managed Care Organization (MCO). We solve the game for the equilibrium prices of each hospital and for the MCO's profit-maximizing premium for its insurance product, given the pre-merger market structure. We then change the market structure by merging two of the hospitals and re-solve the game for the new equilibrium prices. The difference between the two sets of prices represents what we call the "true" merger effect.

We then take the "data" that would be available to a hypothetical econometrician trying to prospectively predict the effect of a hospital merger (pre-merger prices and patient flows) and apply each of the merger simulation methods. This generates predicted merger effects, which can then be compared to the "true" effects. The closer the match, the more accurate the methods can be said to be.

The remainder of the paper is organized as follows. Section II describes the simulation methodologies that have been proposed by other researchers, and that this paper's primary purpose is to test. Section III lays out the bargaining model that we use to generate the "true" merger effects. Section IV contains the numerical exercise in which both the "true" and the predicted merger effects are calculated for each of 32,400 simulated mergers and are compared to see how close they are to each other. Section V discussed future extensions of the model. Section VI concludes.

II. The Simulation Methodologies:

Traditional price-concentration studies involved delineating geographic markets, calculating concentration in each market using the Herfindahl-Hirschman Index (*HHI*), and then regressing average price in a market on the market's *HHI*. This is problematic, both because of the familiar problems with the whole idea of market definition,[3] and because the relationship between average concentration in a market and average price may be a poor approximation of the relationship of interest, which is the relationship between the market power of a given *hospital* and its price.

The market power measures developed in the more recent literature avoid these problems. They are designed so that the unit of analysis is a hospital rather than a market. That is, the analysis involves regressing hospital prices on a hospital-level measure of market power. The measures also have the attractive property that no *a priori* market definition is necessary: the scope of geographic competition is determined by the data themselves.[4] Finally, the measures are able to capture the differences in the degree to which different hospitals compete with each other, which is important because it means that the amount by which the measures will change as a result of a (simulated) merger will depend on the degree of pre-merger substitutability between the merging hospitals.

We consider two market power measures. The first is Willingness-to-Pay (*WTP*), and the second is Hospital-Specific *HHI*. In the remainder of this section we discuss these measures and the simulation methods that incorporate them.

A. Willingness-to-Pay (*WTP*).

The first market power measure we consider is "Willingness-to-Pay" (*WTP*). Although Capps, Dranove, & Satterthwaite (2003) were the first to use that term, the measure developed by Town

[3] Economists have generally preferred to avoid geographic market exercises when alternatives are available, as these exercises have the unappealing feature that all firms that are found to be "in" the market are treated as equally competitively significant (not allowing for differences in substitution patterns across different products), and firms that are found to be "out" of the market are treated as if they do not exist.

[4] The new market power measures are designed so that any disagreement about which geographic areas or other hospitals serve to constrain the market power of a particular hospital of interest can be resolved directly by simply including the relevant data in the calculation of that hospital's market power measure. The disputed hospital or geographic area is competitively significant to the hospital of interest if and only if its inclusion causes the measure to change to a meaningful extent.

and Vistnes (2001) is very similar, and the differences that do exist are irrelevant for our study. In the remainder of this section, we will use the Capps, Dranove, & Satterthwaite (CDS) formulation of the market power measure. (We briefly discuss the Town & Vogt (TV) measure, and why for our purposes the two measures do not meaningfully differ, in the appendix). Though the two *measures* are the same for our purposes, there is an important difference in how the two sets of researchers actually perform their simulations. This difference is discussed below, and we will include both the CDS and the TV approaches in our analysis.

WTP is a way of capturing the value-added of a hospital or hospital system to the provider network of a Managed Care Organization (MCO). To see how it does this, begin by considering a standard discrete choice problem in which consumer *i*'s preferences over each hospital *j* in a choice set *G* are given by:

(1) $$U_{ij} = V_{ij} + \varepsilon_{ij}, \quad \forall j \in G$$

where V_{ij} is a linear-in-parameters index of hospital characteristics and interactions of hospital and consumer characteristics, and ε_{ij} is an independently and identically distributed Extreme Value error term. Consumer *i*'s actual hospital choice will depend on the realization of the ε_{ij} terms. But it is also possible to calculate how much consumer *i* values having access to *all* the hospitals in *G* *before* the realization of the ε_{ij} terms. This will be the expected value of the maximum utility: each hospital has some probability of providing more utility than any other (and so of being chosen), and each has an expected utility conditional on being chosen, and from this it is possible to calculate the utility that consumer *i* expects to receive from whichever hospital turns out to provide the highest utility. Because the ε_{ij} terms are distributed extreme value, there is a simple and familiar expression for the expected value of the maximum:

(2) $$E_{\varepsilon}[MAX_{j\in G}\{V_{ij}+\varepsilon_{ij}\}]=\gamma+\ln\sum_{j\in G}e^{V_{ij}}$$

where $\gamma \equiv E[\varepsilon_{ij}]$, which is known as Euler's constant (approximately 0.5772) and is the mean of the unconditional Extreme Value distribution.

The expected value of the maximum in (2) can be calculated for any choice set. So a measure of the value that consumer *i* places on a particular hospital *k* would be the difference between the expected value of the maximum for choice sets *G* and *G\k*. This turns out to have the following neat closed-form solution:

(3) $$WTP_{ik}=\left(\gamma+\ln\sum_{j\in G}e^{V_{ij}}\right)-\left(\gamma+\ln\sum_{j\in G\backslash k}e^{V_{ij}}\right)=\ln\left(\frac{\sum_{j\in G}e^{V_{ij}}}{\sum_{j\in G\backslash k}e^{V_{ij}}}\right)=\ln\left(\frac{1}{1-prob_{ik}}\right)$$

where

(4) $$prob_{ik}^{G}=\Pr[U_{ik}>U_{ij},\forall j\neq k]=\frac{e^{V_{ik}}}{\sum_{j}e^{V_{ij}}}$$

The *WTP* of consumer *i* for hospital *k* is a straightforward function of $prob_{ik}$, which comes directly out of the choice model in (1) above. The total *WTP* for Hospital *k* (WTP_k) is:

(5) $$WTP_{k}=\sum_{i}WTP_{ik}=\sum_{i}\ln\left(\frac{1}{1-prob_{ik}}\right)$$

The expression in (5) is what CDS put on the right-hand side of their price regressions. They hypothesize that WTP_k should be positively related to prices, ceteris paribus.

As described above, WTP_k is a hospital-level measure since it is determined solely by the valuation of the choice set that excluded only Hospital *k*. As such, it is unaffected by the industry market structure, and so cannot be the basis for simulating a merger. However, *WTP* can also be defined at the level of a hospital system *s* as follows:

(6) $$WTP_s = \sum_i \ln\left(\frac{1}{1-\sum_{j\in s} prob_{is}}\right) = \sum_i \ln\left(\frac{1}{1-prob_{is}}\right)$$

While for our purposes there is no meaningful difference between TV and CDS in the construction of their market power measures, there is an important difference in how they generate predicted merger effects. To see this, consider a merger between hospitals *k* and *l*. In TV, the effect of the merger on the market power of Hospital *k* is equivalent to:

(7) $$\sum_i \ln\left(\frac{1}{1-prob_{ik}-prob_{il}}\right) - \sum_i \ln\left(\frac{1}{1-prob_{ik}}\right) = \sum_i \ln\left(\frac{1-prob_{ik}}{1-prob_{ik}-prob_{il}}\right)$$

In TV, the change in the market power of Hospital *k* as a result of the merger with Hospital *l* is the simple change in the relevant right-hand side variable. The change for Hospital *l* is calculated similarly, and the change will not be the same for both hospitals. In general, the increase in bargaining power is inversely related to relative pre-merger bargaining power. This is true because the bargaining power measure is constant across all members of the post-merger system. Hence, the change in bargaining power due to a merger will be larger for the hospital (or system) with the lower level of pre-merger bargaining power.

In contrast, CDS derive their simulating merger effects by taking the difference between the *WTP* of the merged system and the sum the pre-merger *WTP*. Hence:

(8) $$\begin{aligned} \Delta WTP_k &= \Delta WTP_l = WTP_{kl} - WTP_k - WTP_l \\ &= \sum_i \ln\left(\frac{1}{1-prob_{ik}-prob_{il}}\right) - \sum_i \ln\left(\frac{1}{1-prob_{ik}}\right) - \sum_i \ln\left(\frac{1}{1-prob_{il}}\right) \\ &= \sum_i \ln\left(\frac{(1-prob_{ik})(1-prob_{il})}{1-prob_{ik}-prob_{il}}\right) \end{aligned}$$

In CDS, the change in the market power variable is defined at the *post-merger* system level and there appears to be no well-defined way to back out different price effects for *k* and *l*. However, the CDS definition has the intuitive property that the change in bargaining power is close to zero if consumers do not view *k* and *l* as substitutes, i.e., $prob_{ik} < d$ or $prob_{il} < d$ from some positive *d* arbitrarily close to zero for all *i*. The TV definition does not have this intuitive property.

In this study, we consider both approaches. That is, we regress prices on the CDS measure *WTP* and evaluate the predicted price changes based on the TV approach in (7) and on the CDS approach in (8). In the CDS approach, we simply assume that the predicted price increase will be the same for both *k* and *l.*

Two final notes on both the TV and CDS approaches should be made. First, both have the unattractive property that the bargaining power measures do not change for non-merging hospitals. Hence, the models predict no price effect for these hospitals even if they are close competitors of the merging firms. Second, in their frameworks, the bargaining power of a hospital system is based on its threat to exclude the entire system. Hence, "all-or-nothing bargaining" while not explicitly considered in TV or CDS, is nonetheless built into the derivation of both market power measures. As our model will show below, however, mergers can have price effects even if the merged entities continue to bargain separately. In that case, the TV and CDS measures can still be used as measures of market power, as they are still based on the choice model in (1), but they cannot be interpreted exactly the way that TV and CDS describe them.

The *WTP* measure has the desirable property that hospitals and hospital systems have higher *WTP* if they face fewer competitors, or if they have some attribute that makes them more desirable.[5] It also has the unfortunate property that a hospital system will have higher *WTP* just by virtue of "pure" size. For example, a large hospital system comprised of non-competing hospitals will have a very high *WTP*. This is problematic because there is little reason to think that a hospital system that is large by virtue of a pure size effect will command a particularly high price.[6] This represents a problem for the estimation: a large hospital system with a very high *WTP* but a not-particularly-high price will cause the regression of price on *WTP* to have a poor fit. This is not much of an issue in the current version of the paper, but will likely become an issue in future revisions that will allow for large pre-merger systems.

An alternative specification of willingness-to-pay that addresses this problem is to define a "per-person" *WTP* (denoted by *WTP_PP*) which is equal to:

(9)
$$WTP\_PP_s = \frac{WTP_s}{\sum_i prob_{is}}$$

That is, the *WTP_PP* is simply the *WTP* divided by the predicted total consumers who choose system *s*. This has the effect of assigning willingness to pay on the basis of consumer valuation of the hospital's characteristics, rather than by the sheer size of the system.

Having calculated the WTP, the next step is to estimate the price-concentration equation. The simplest version of this equation is:

(10)
$$price_j = \varphi_0 + \varphi_1 WTP_j + \nu_j$$

B. Hospital-Specific HHI.

An alternative measure of concentration is known as Hospital-Specific *HHI*, and is used by Capps & Dranove (2004) and Melnick & Keeler (2007). As with the *WTP* measures, the first step is to estimate a choice model like (1) and then derive, for each consumer, the probability of

---

[5] This means that it is possible for small independent hospitals to have higher *WTP*, and so to command higher prices, than do larger hospitals if they face few competitors or have some attractive attribute such as a desirable location.

[6] The mechanism by which pure size can influence price is if one party to the negotiation has a payoff function that is more concave than the other. It seems intuitive that the magnitude of this effect should be small, and there is empirical evidence of this in Sorensen (2003).

choosing each hospital. With this in hand, it is possible to regard *each consumer* as a "market" with "shares" that correspond to the choice probabilities.[7] It is then possible to calculate an *HHI* for each consumer using the familiar formula:

$$HHHI_i = \sum_j prob_{ij}^2 \tag{11}$$

We refer to the measure in (11) as $HHHI_i$ because it is calculated using "hospital probabilities," meaning that it ignores any joint ownership of hospitals. Below this will be contrasted with $SHHI_i$ which is calculated using "system probabilities."

Having calculated the $HHHI_i$ for each consumer i, the next step is to calculate a Hospital-Specific HHI_j for each hospital. This is done by calculating a weighted sum of the $HHHI_i$, where the weights are defined as follows: the weight given to consumer i in calculating the Hospital-Specific HHI_j for hospital j is the total contribution of consumer i to hospital j's total expected patients. That is:

$$Hospital\ Specific\ HHI_j = \sum_i \frac{prob_{ij}}{\sum_i prob_{ij}} HHHI_i \tag{12}$$

The intuitive interpretation of a Hospital-Specific HHI_j is that a hospital for which a large proportion of its patients are drawn from consumers with high $HHHI_i$ can be thought of as operating in a generally concentrated environment, and vice-versa.

The definition described above is applicable if all hospitals are independent, and must be modified if some hospitals are part of multi-hospital systems. That is, a "System-Specific" HHI_j must be calculated instead. This can be calculated in one of two ways. Both begin in a manner

[7] An alternative to treating each patient as a "market" is to divide the patient population into discrete "bins" (for example a bin could be a zip-code/diagnostic code combination), and then to apply (11) and (12) to these bins. This is the actual approach taken by Capps & Dranove (2004) and by Melnick & Keeler (2007).

similar to (11), but these use the system shares $prob_{is}$ and are denoted $SHHI_i$ instead of $HHHI_i$. That is:

(13) $$SHHI_i = \sum_s prob_{is}^2$$

Like the Hospital-Specific HHI_j, the System-Specific HHI_j are calculated by taking a weighted average of the individual $SHHI_j$. The weights can either be the system weights or the individual hospital weights. If hospital weights are used, then the System-Specific HHI_s for hospital system s is calculated as:

(14) $$System\ Specific\ HHI_s = \sum_i \frac{prob_{ij}}{\sum_i prob_{ij}} SHHI_i$$

If system weights are used, then the System-Specific HHI_s are calculated as:

(15) $$System\ Specific\ HHI_s = \sum_i \frac{prob_{is}}{\sum_i prob_{is}} SHHI_i$$

If system weights are used, then the System-Specific HHI_s will be the same for all of the hospitals in a system, whereas this is not the case if the hospital weights are used.

Having calculated the System-Specific *HHI*, the next step is to estimate the price-concentration equation. The simplest version of this equation is:

(16) $$price_j = \phi_0 + \phi_1 System\ Specific\ HHI_j + \upsilon_j$$

One advantage of the System-Specific *HHI* approach is that, unlike in the case of *WTP* above, a merger affects the predicted prices of the non-merging hospitals.

III. The Bargaining Model:

A. Setup.

There is a set of hospitals G with a cardinality of g. A monopoly managed care organization (MCO) enters into separate and simultaneous bargaining with each hospital in G. A deal between the MCO and a hospital, if a deal is reached, consists of a linear per-patient price. With its network and its negotiated prices in place, the MCO sets the profit-maximizing premium for its insurance product. Given this premium, consumers choose whether to purchase insurance from the MCO. Consumers who purchase insurance become sick with some probability, and seek care at their most preferred hospital in the MCO's network (people with no insurance do not use any hospital). For convenience, we set this probability equal to one, in order to abstract from the issue of risk-aversion. The MCO, as we model it, is an assembler of a network and not an insurer. The out-of-pocket costs faced by consumers are the same for all hospitals in the network, regardless of the price negotiated with each hospital by the MCO. In other words, the MCO cannot "steer" consumers by giving them incentives to use hospitals with which the MCO has a lower contracted price.

B. Insurance Premiums.

The valuation of patient i for the MCO's insurance product is determined as follows. Following TV and CDS, we assume that consumers do not know the realization of their idiosyncratic preference shocks ε_{ij} when they are deciding whether to buy insurance. Rather, they know the distribution of those shocks, and their valuation for insurance is equal to the expected value of the maximum of the utilities from using each hospital, which is equal to $\gamma + \ln \sum_j e^{V_{ij}}$. Consumer i also has an idiosyncratic valuation for having insurance coverage ζ_i, which is also assumed to be distributed Type 1 Extreme Value and is assumed to be unknown to both the MCO and the hospitals. This valuation is still in utils, and must be converted to dollars, which we do as follows:

$$Valuation_i = \lambda_1 + \lambda_2 \ln \sum_j e^{V_{ij}} + \zeta_i \tag{17}$$

The MCO calculates the g+1 optimal premiums for a candidate vector of hospital prices ***price***, taking expectations over the distribution of both idiosyncratic components, ζ_i and ε_{ij}. This seems to be a reasonable approach in that the MCO will not know exactly who will buy insurance and who will go to which hospital when bargaining with hospitals. In addition, the idiosyncratic component ζ_i also ensures that the MCO's objective function is differentiable in the premiums in that it avoids the use of step functions. This simplifies solving the MCO's optimization problem significantly without sacrificing any basic intuition.

The MCO's optimal premium will depend on the probabilities of patient insurance take-up and hospital choices of the marginal patient. Specifically, the MCO solves the maximization problem:

$$\max_{prem^G} \left\{ \sum_i \left[\left(prem^G - \sum_j price_j prob_{ij}^G \right) \left(1 + \exp\left\{ prem^G - \lambda_1 - \lambda_2 \ln \sum_j e^{V_{ij}} \right\} \right)^{-1} \right] \right\} \tag{18}$$

The interpretation of (18) is as follows. For a candidate premium $prem^G$, the first term in the parentheses represents the expected margin that the MCO will receive for patient i. This margin will only be realized conditional on patient i actually purchasing insurance (i.e., having a valuation greater than the premium). Since the MCO does not know the idiosyncratic component ζ_i, the MCO takes expectation over its distribution and multiplies the expected margin by the probability of insurance take-up (the term in the second parentheses). The expected number of patients who use hospital j is equal to:

(19) $$n_j^G = \sum_i prob_{ij}^G \left(1 + \exp\left\{prem^G - \lambda_1 - \lambda_2 \ln \sum_{k \in G} e^{V_{ij}}\right\}\right)^{-1}$$

Similarly, for a given hospital *k,* the MCO finds the premium that maximizes its profit if hospital *k* is excluded by solving:

(20) $$\max_{prem^{G\backslash k}} \left\{ \sum_i \left[\left(prem^{G\backslash k} - \sum_j price_j prob_{ij}^{G\backslash k} \right) \left(1 + \exp\left\{ prem^{G\backslash k} - \lambda_1 - \lambda_2 \ln \sum_{j \in G\backslash k} e^{V_{ij}} \right\} \right)^{-1} \right] \right\}$$

where $prob_{ij}^{G\backslash k}$ denotes the probability that patient *i* would choose hospital *j* given that hospital *k* was not available. Note that in this case, the marginal valuations are taken over the choice that excludes hospital *k*:

(21) $$Valuation_i^{G\backslash k} = \lambda_1 + \lambda_2 \ln \sum_{j \in G\backslash k} e^{V_{ij}} + \zeta_j$$

The expected number of patient admitted to hospital *j* under the exclusion of hospital *k* is:

(22) $$n_j^{G\backslash k} = \sum_i prob_{ij}^{G\backslash k} \left(1 + \exp\left\{prem^{G\backslash k} - \lambda_1 - \lambda_2 \ln \sum_{l \in G\backslash k} e^{V_{il}}\right\}\right)^{-1}$$

C. Bargaining.

Prices are set by *G* separate Nash bargains: each of the *g* hospitals in *G* has a separate negotiation with a representative of the MCO. Negotiation proceeds under the standard Nash assumptions that: (i) all negotiations happen simultaneously; (ii) no party to any negotiation observes or is in any way affected by what happens in any of the other negotiations; (iii) both parties to each negotiation believe that all the other negotiations will be successful (i.e., that all other hospitals will be included in the MCO's network); and (iv) both parties to each negotiation have beliefs (which turn out to be correct in equilibrium) about the prices agreed to in the other negotiations. The bargaining equation between the MCO and an independent hospital *k* is as follows:

$$(23)\quad \max_{price_k}\left[(n_k^G(price_k - c) - 0)^{\alpha}\left(\sum_{j\in G} n_j^G(prem^G - price_j) - \sum_{j\in G\setminus k} n_j^{G\setminus k}(prem^{G\setminus k} - price_j)\right)^{1-\alpha}\right]$$

As is standard in Nash bargaining, the equilibrium negotiated price for hospital k ($price_k$) is the product of the increase in hospital k's payoff if a deal is reached times the increase in the MCO's payoff if a deal is reached.[8] The payoff to hospital k if it reaches a deal with the MCO is $n_k^G(price_k - c)$, where G represents the complete set of hospitals or the "network of the whole," and n_k^G is the expected number of consumers who buy insurance and choose hospital k when the network offered by the MCO is the network of the whole, and $price_k$ is the price that hospital k negotiates with the MCO. The payoff to hospital k if it fails to reach a deal with the MCO is zero. The parameter α denotes the division of joint surplus between the hospital and MCO. For example, it could capture the relative skill of the negotiators involved in the bargaining.

The payoff to the MCO if it reaches a deal with hospital k is as follows. Since, by assumption, both parties to the negotiation believe that all the other hospitals will be in the MCO's network, the premium if a deal is reached between the MCO and hospital k will be $prem^G$, which is the premium that maximizes the MCO's profits given that its network is the network of the whole. For each hospital $j\in G$ in the network of the whole, there are n_j^G consumers who buy insurance and choose hospital j, and each of these consumers will generate a profit for the MCO of $prem^G$ minus hospital j's price $price_j$. Similarly, the payoff to the MCO if it fails to reach a deal with hospital k is as follows. $G\setminus k$ represents the set of all hospitals except k. For each hospital $j\in G\setminus k$ in the network, there are $n_j^{G\setminus k}$ consumers who buy insurance and choose hospital j. For each of these

[8] An equivalent way to describe the solution to a Nash Bargaining game is to identify the total joint surplus accruing to the parties if a deal is reached, and then splitting that deal-specific joint surplus between the two parties according to the bargaining power of each side (i.e., the hospital would get a share α and the MCO would get a share $1-\alpha$).

consumers, the MCO receives a margin equal to $prem^{G\setminus k}$ (which will be lower than $prem^{G}$) minus the negotiated price $price_j$.

D. Equilibrium.

From the above discussion, we see that the equilibrium Nash bargain price for hospital *k* is a function of all the other *g*-1 prices, both directly and via the premiums $prem^{G}$ and $prem^{G\setminus k}$, and of consumer choice, given those premiums, of whether to buy insurance and of which hospital to use conditional on buying insurance. Solving the system is complicated by the fact that the *g*+1 profit-maximizing premiums $prem^{G}$ and $prem^{G\setminus j}$ each depend not only on the price vector, but also on n_j^{G} and $n_j^{G\setminus j}$ for each set of candidate premiums. The profit-maximizing premium, in turn, depends not only on the individual hospital prices, but also on the hospital choice of the marginal consumer. To see this, suppose that the MCO, given a vector of hospital prices, is trying to decide whether it is profitable to cut the premium by enough to attract one more consumer to buy insurance. Whether or not doing so is profitable will depend on which hospital that marginal consumer would choose if he/she bought insurance, which in turn depends on the negotiated price with that hospital.

A proposed vector ***price*** is an equilibrium if the following is true: (i) the set of *g*+1 premiums $prem^{G}$ and $prem^{G/j}$ are each profit-maximizing for the MCO given ***price***; (ii) ***price*** is a solution to the system of Nash bargaining equations given those g+1 premiums and given consumer behavior (i.e., given the choices that consumers make about whether to buy insurance given the premiums and which hospital to use conditional on buying insurance).

The joint surplus from reaching a deal is equal to: (hospital's payoff if a deal is reached − hospital's payoff if no deal is reached) + (MCO's payoff if a deal is reached − MCO's payoff if no deal is reached).

D. Mergers.

Now suppose that formerly independent hospital *k* merges with formerly independent hospital *l* into a joint hospital system. The Nash bargaining equations of the remaining $g-2$ hospitals will remain unchanged. The bargaining equations for the merged hospitals will depend on whether the merged hospitals continue to bargain independently or they bargain on an "all or nothing" basis.[9] We consider each of these possibilities in turn.

i. The Merged Hospitals Continue to Negotiate Independently.

We first consider the case where the merged hospitals continue to negotiate independently. That is, in this sub-section we assume that negotiators for hospitals *k* and *l* continue to engage in separate negotiations with the MCO, but that in their negotiations they each internalize the fact that they have a profit stake in the other. An intuitive way to interpret the post-merger payoff to hospitals *k* and *l* is to imagine that the pre-merger owner of each hospital remains in charge of negotiating that hospital's price, but now gets half of the merged entity's profits.[10] If a deal is reached with hospital *k*, then the payoff is half of the profits of a two-hospital network. The bargaining equations for this case are as follows.

$$(24a)\ \max_{price_k}\left[\left(\frac{n_k^G(price_k-c)+n_l^G(price_l-c)}{2}-\frac{n_l^{G\backslash k}(price_l-c)}{2}\right)^{\alpha}\left(\sum_{j\in G}n_j^G(prem^G-price_j)-\sum_{j\in G\backslash k}n_j^{G\backslash k}(prem^{G\backslash k}-price_j)\right)^{1-\alpha}\right]$$

$$(24b)\ \max_{price_l}\left[\left(\frac{n_l^G(price_l-c)+n_k^G(price_k-c)}{2}-\frac{n_k^{G\backslash l}(price_k-c)}{2}\right)^{\alpha}\left(\sum_{j\in G}n_j^G(prem^G-price_j)-\sum_{j\in G\backslash l}n_j^{G\backslash l}(prem^{G\backslash l}-price_j)\right)^{1-\alpha}\right]$$

---

[9] The decision of which form the negotiations should take may itself be reflective of the relative bargaining power of the two sides. We ignore the means by which the form of bargaining was chosen and simply lay out the implications of each of the two possibilities.

[10] An alternative, equivalent assumption is that the merged entity assigns a negotiator to each negotiation, and then provides those negotiators with incentives such that they act as if they get half of the profits of the merged entity.

Note that the terms in the right-hand sets of parentheses are the same as in (23) above. What changes after the merger is the payoff to each hospital if no deal is reached. In the pre-merger situation, the payoff to hospital *k* if no deal was reached was zero. In the post-merger situation, the payoff is half of the profits hospital *l*. To see the effect of the merger on the equilibrium price, we begin by considering the case where there was no pre-merger competition between *k* and *l*. That is, we consider the case where nobody who would choose *l* if *k* was not in the network would choose *k* if it were in the network, and vice-versa (i.e., $n_l^G = n_l^{G\backslash k}$ and $n_k^G = n_k^{G\backslash l}$). In this case, the terms in the left-hand set of parentheses in each bargaining equation will be the same as in (23) and the merger will have no effect.[11]

If there is pre-merger competition between *k* and *l*, then the merger will have an effect. The reason is that if $n_l^G \neq n_l^{G\backslash k}$ and/or $n_k^G \neq n_k^{G\backslash l}$, then each hospital's negotiator internalizes the fact that if they fail to reach a deal, some of the patients that will be lost as a result will use the other merged hospital instead, and so some of those lost profits will be recaptured. This causes each side to bargain more aggressively, and to get a higher equilibrium price. The magnitude of the merger effect is determined by the magnitude of this diversion. This, of course, is a variation on the standard intuition in which diversion is the source of price increases resulting from mergers between substitutes.

As discussed above, the derivation of the *WTP* measures depends on the assumption that following the merger the negotiations will proceed on an "all or nothing" basis. This may call into question the appropriateness of using *WTP* as a measure of market power for analyses in which the merging firms continue to negotiate separately even post-merger. Nevertheless, we do include *WTP* in our independent-bargaining analyses. We do this because *WTP*, regardless of how

[11] The term will be multiplied by a constant of ½, but this will have no effect on the bargaining outcome.

it is derived, can also be thought of simply as a measure of market power (since it is based on choice probabilities), and as such can be tested for how well it predicts merger effects.

ii. The Merged Hospitals Negotiate on an "All or Nothing" Basis.

The merged entity might instead merge their negotiations and offer their hospitals on an all-or-nothing basis. In this case, the post-merger bargaining equation for the MCO and the merged entity will be:

$$(25)\quad \max_{price_k, price_l}\left[\left(n_k^G(price_k - c) + n_l^G(price_l - c) - 0\right)^{\alpha}\left(\sum_{j\in G} n_j^G(prem^G - price_j) - \sum_{j\in G\backslash kl} n_j^{G\backslash kl}(prem^{G\backslash kl} - price_j)\right)^{1-\alpha}\right]$$

Under all-or-nothing bargaining, the MCO and the merged entity are bargaining over two prices ($price_k$ and $price_l$), but there is only one bargaining equation because both parties care only about the total payment from the MCO to the merged entity if a deal is reached. There are infinitely many combinations of $price_k$ and $price_l$ to reach any given total payment.

The difference between (25) and (23) above is that in (25), both the merged entity and the MCO are bargaining over two hospitals instead of one. This means that if the negotiations fail, the merged entity will lose two hospital contracts instead of one, and the MCO will suffer a two-hospital hole in its network instead of a one-hospital hole. This doubling of the stakes will not necessarily result in a price increase: if the payoff functions of both the hospitals and the MCO are linear, then the doubling of what the hospitals stand to lose will be exactly offset by the doubling of what the MCO stands to lose, and there will be no price effect. However, if the two merged hospitals are substitutes for each other, then the MCO's payoff function will be concave and not linear: the loss to the MCO from losing two hospitals will be more than twice as large as the loss from losing one hospital, and this effect will be greater the substitutability between the two hospitals, and the less substitutable are the other hospitals.[12] As shown by Chipty & Snyder (1999), a negotiator with a concave payoff function is made worse when its opponent is larger.[13] The simulations currently included in the paper do not yet include any all-or-nothing bargaining cases. These will be added in the next revision.

IV. Numerical Analysis and Merger Simulation:

A. Generating the Data.

The first step in the numerical analysis is to generate our "data." In each simulation, we generate data on ten hospitals. For each hospital j, we take three random draws on U[0,1] and call them ρ_{jF}, ρ_{jX}, and ρ_{jY}. We draw a fixed effect η_j for each hospital j, which is the $100\rho_{jF}$ percentile of N(0,1.2). We also choose two locations for each hospital. For the first location, the x-coordinate of hospital j's location is the $100\rho_{jX}$ percentile, and the y-coordinate is the $100\rho_{jY}$ percentile, of U[-6,6]. For the second location, the x-coordinate of hospital j's location is the $100\rho_{jX}$ percentile, and the y-coordinate is the $100\rho_{jY}$ percentile, of N(0,16). That is, we draw one set of hospital locations from a distribution in which hospitals are equally likely to be located in any location in the support, and another set from a distribution in which hospitals are more likely to be located closer to the center.

[12] The reason is that the reduction in consumer valuation of the MCO's insurance product if one hospital is excluded is mitigated by the fact that the other, competing hospital is included and is a close substitute for at least some consumers. Post-merger if both hospitals are excluded from the network, then consumers who regard the two merging hospitals as roughly equally satisfactory, but who regard the other hospitals as much less satisfactory, will have a larger reduction in their valuation of insurance, and so the MCO will have a larger reduction in its payoff.

[13] There are other factors that could cause the payoff functions of the hospitals or of the MCO to be concave rather than linear, such as risk aversion. To the extent that this is true, the merger may result in a price increase or a price decrease for reasons that have nothing to do with competition, but simply are a result of the fact that the stakes of the negotiation have increased. However, Sorensen (2003) shows that this effect is modest: hospitals can command slightly higher prices simply by virtue of being larger, but this benefit is much smaller than the benefit from facing less competition.

In each simulation, we generate data on 100,000 patients. For each patient *i*, we take two random draws on U[0,1] and call them μ_{iX}, and μ_{iY}. We choose two locations for each patient. For the first location, the *x*-coordinate of patient *i*'s location is the $100\mu_{iX}$ percentile, and the *y*-coordinate is the $100\mu_{iY}$ percentile, of U[-10,10]. For the second location, the *x*-coordinate of patient *i*'s location is the $100\mu_{iX}$ percentile, and the *y*-coordinate is the $100\mu_{iY}$ percentile, of N(0,36). Note that the parameters of both the uniform and the normal distributions were chosen so that hospitals are more likely to have relatively central locations compared to the patient population.[14] The distributions of these hospital and patient characteristics are summarized in table 1 below.

Table 1: Distributions of Simulated Hospital and Patient Characteristics

Characteristics	**Distribution**
Hospital Location	U[-6,6] x U[-6,6], N(0,16) x N(0,16)
Hospital Fixed Effect	N(0,1.2)
Patient Location	U[-10,10] x U[-10,10], N(0,36) x N(0,36)

We define consumer preferences as:

$$U_{ij} = -\gamma_1 dist_{ij} - \gamma_2 dist_{ij}^{\ 2} + \eta_j + \varepsilon_{ij} \qquad (26)$$

where $dist_{ij}$ denotes the distance from consumer *i* to hospital *j*, η_j denotes a hospital-specific fixed effect which captures hospital quality, and ε_{ij} denotes the idiosyncratic component and is assumed to be distributed Type 1 Extreme Value. We use three sets of values for (γ_1, γ_2) to reflect low, medium, and high travel costs. The values we choose are (0.3, 0.003), (0.5, 0.005), and (0.7, 0.007).

[14] We allowed the location of the hospitals to be determined randomly, rather than as the result of profit-maximizing location, quality investment, and entry/exit decisions. This is partly for simplicity, but it is the case that many hospi-

We calibrate the model parameters as follows. We set the bargaining power α of each hospital to ½.[15] The parameters λ_1, λ_2, c, the travel cost parameters, and the parameters of the location distributions and hospital fixed-effect distribution are chosen based on the follow criteria.

1) Generate pseudo-R^2 values that we typically observed in real-world data (0.40 – 0.60).

2) The variation in the component of hospital valuation that varies across consumers (locations) explains about twice as much of the observed hospital choices (conditional on insurance take-up) as the variation in hospital valuation that does not vary across consumers (the hospital fixed-effect). This is the rough proportion that we have observed in real-world data.

3) Generate reasonable insurance take-up rates (0.85 - 0.95).

4) Generate reasonable hospital price-cost margins (0.12 - 0.17).

To generate data with the above characteristics, in addition to the aforementioned location and fixed-effect distributions and travel cost parameters, we choose the following parameter values: $\lambda_1 = 22$, $\lambda_2 = 0.4$, and $c = 4$.

B. Numerical Solution of the Model.

The next step is to solve the model using these generated data. The search algorithm used to solve for the equilibrium vector ***price**** employs a straightforward Newton-based approach. The only complication in this application is that for a candidate vector ***price***, a series of searches must be carried out to find a set of optimal premiums for the MCO. The solution algorithm proceeds as follows:

---

tals were built many years ago, and so their original locations might no longer be optimal. Moreover, in many states entry decisions are complicated by the existence of "certificate of public need" regulations.

1) Choose a starting price vector $\boldsymbol{price}^{\mathbf{0}}$.

2) Given $\boldsymbol{price}^{\mathbf{0}}$, solve the optimal g+1 MCO premiums. Since this is a single variable problem, we use a straightforward bisection method.

3) Given the premiums from step (2), evaluate the derivatives of the vector of Nash bargaining problems with respect to own-prices. Update guesses of the price vector using a Newton-Raphson method. Convergence occurs when the Euclidean norms of the vector of derivatives and the price update are within a tolerance of zero. Here, we use 10^{-12}.

4) Repeat (2) and (3) to convergence. Convergence occurs when, for a given price vector $\boldsymbol{price}^{*}$ and a given premium vector ($prem^{G}$, $prem^{G\backslash 1}$,…, $prem^{G\backslash J}$), the Nash bargaining problems in step (3) and the MCO optimization problems in step (2) are simultaneously solved.

Note that we do not re-solve for the optimal premiums for each updated guess of the price vector. We re-solve for the optimal premiums only after the derivatives are solved for prices given a set of premiums. This saves a great deal of computation time. A reasonable and natural alternative would be to simply code the derivatives of the optimal premiums with respect to price and incorporate that into the Newton search. While we may ultimately use this approach, our current approach avoids more the complicated coding associated with this alternative and the bisection method approach to solving the MCO's problem is very inexpensive and reliable.

As noted above, the inclusion of ζ_i in the consumer's insurance valuation ensures that the MCO's objective function is differentiable in its premiums. By extension, this ensures that the Nash bargaining problems are differentiable in prices. At this time, we do not have a formal

---

[15] We also ran the same set of simulations for α = .6 and for α = .4. As expected, merger effects are larger when hos-

proof that the Nash bargaining problems are globally concave in prices, or more generally, that the joint bargaining problem is convex. However, we have tested for convexity using 2000 randomly chosen starting values for the price vector. In this test, the minimums and maximums at the solution for each hospital's price deviated by no more than 10^{-8}. Hence, at the present time, we are reasonably sure that the joint Nash bargaining problem has a unique solution.

It is important to note the solving the model does *not* require that the uncertainty represented by ζ_i and ε_{ij} be resolved. This is true both before the merger and after, which means that measuring the "true" merger effect does not depend on the uncertainty being resolved. As will be made clear below, the same is not true for the simulated merger effect.

C. Merger Simulations.

The "true" effect of a merger can be simulated by re-solving the model after changing the bargaining according to (24*a*) and (24*b*) above. We do 120 merger simulations. As discussed above, each simulation contains six configurations: two sets of hospital/patient locations (one from the uniform distributions and one from the normal distributions), times three sets of travel cost coefficients. Each simulation includes every possible merger (Hospital 1 merging with Hospital 2, Hospital 1 merging with Hospital 3, etc.). Since there are ten hospitals, the total number of possible mergers is 45. For each merger, there are two price effects, one for each of the merging firms.[16] So altogether, we simulate 120 x 2 x 3 x 45 = 32,400 mergers. Since each merger involves two firms, this gives us a total of 64,800 separate sets of prices (with each set containing a pre-merger price, a simulated post-merger price, and a "true" post-merger price).

---

pitals have more bargaining power. Furthermore, the tendency of the simulation methodologies to under-predict the true merger effects is greater the larger is α.

D. Estimating the Price-Market Power Relationships.

Given the solution to the Nash bargaining problems in the pre-merger world, $\boldsymbol{price}^*$, the next step is to estimate the relationships between these "true" prices (which correspond to the pre-merger prices than a merger analyst would observe) and the market power variables according to (10) and (16) above. Since the market power measures are based on actual patient choices, this requires that the uncertainty represented by ζ_i and ε_{ij} be resolved. To do this, we randomly draw values of ζ_i and ε_{ij} from the Type 1 Extreme Value distribution.[17] This allows us to determine the set of consumers who would buy insurance from the MCO given the optimal premium that emerges from the solution of the Nash bargaining problem, and which patients will choose which hospitals. We then estimate the conditional logit model, construct the competition measures, and run the price regressions using only the consumers who buy insurance. In this way, we replicate discharge data that a researcher would use in this type of analysis in that real-world discharge data are drawn after the consumers' insurance decisions have already been made.

E. Simulation Results.

The results are summarized in Figures 1-4. The *x*-axis in the figures represents the "diversion" between each hospital and its merger partner. That is, it represents the fraction of those insured patients whose first choice would have been Hospital *k* (if Hospital *k* had been in the MCO's network), who would choose its merger partner Hospital *l* if Hospital *k* were omitted.[18] The *y*-axis

[16] The *WTP* measures (unrealistically) constrain rival effects to be zero. The *Hospital-Specific HHI* measure allows for rival effects to be positive, but there is no economical way to present these given our large number of simulations.

[17] Implicitly, there is an Extreme Value draw associated with not buying insurance as well. So, more precisely, ζ_i is drawn as a difference of independent Extreme Value random variables.

[18] We could have done this calculation in such a way as to take account of the fact that the number of people who buy insurance will be different if a hospital is actually omitted from the network, but instead we do it in the way that it would have to be done in a real merger investigation. Specifically, we take the patients in the dataset as the uni-

represents the quantity (predicted post-merger price – "true" post-merger price)/pre-merger price. That is, the *y*-axis represents the *percentage point difference* between the estimated merger effect and the true one.[19]

As can be seen in Figure 1 (which compares the true effect to the effect predicted using the "willingness-to-pay per person" simulation method), the predicted effects tend to be below the true ones, and this effect is greater the larger the diversion.[20] An alternative graphical representation of these same results can be found in Figure 5, which shows a kernel density estimate of the results for all mergers, for those in which diversion ≥ .2, and for those in which diversion ≥ .4. Yet another representation of the same result can be found in Table 2, which shows that the predicted effect is on average 0.01 percentage points lower than the true effect, but is 1.1% lower for mergers in which diversion ≥ .2, and is 2.1% lower for mergers in which diversion ≥ .4. The probability that the simulation will *over*-predict the true merger effect by at least five percentage points is correspondingly small: 2.2% for the full sample of mergers, 1.4% for mergers in which diversion ≥ .2, and 0.9% for mergers in which diversion ≥ .4. In other words, our results suggest that the probability of a substantial Type II error is very small.

Table 2 and the remaining figures contain results for the other three concentration measures: "Willingness-to-Pay" (*WTP*) as used by Town & Vogt (TV), *WTP* as used by Capps, Dranove, & Satterthwaite (CDS), and Hospital-Specific *HHI*. The results are broadly similar across the different concentration measures. They all under-predict the true merger effects on average, and in all of them this tendency is more pronounced the greater the diversion.[21] The CDS version of

---

verse, estimate the conditional logit model on that universe of patients, and then see how many patients have Hospital *k* as their first choice and Hospital *l* as their second.

[19] For example, if the predicted merger effect was 10% and the true effect was 15%, the *y*-axis value would be -.05.

[20] Figure 1 should have 64,800 "dots" on it. However, the graphing software that we used only permits 32,000 dots. For this reason, Figure 1 contains 32,000 dots randomly drawn from the 64,800. The same is true for Figures 2-4.

[21] Note that in Figures 3 and 4, the difference between the predicted effect and the true effect is very close to zero for low levels of diversion, but this is not the case in Figures 1 and 2. The reason is as follows. When diversion is zero,

WTP has the greatest tendency to under-predict the true effect (it under-predicts by an average of 4.7 percentage points for mergers in which diversion ≥ .4), and *WTP_PP* has the least tendency (2.1%).

Table 3 contains the same information as Table 2, but includes only those merger simulations in which the estimated relationship between price and the concentration measures is positive and has a p-value < .10. It is noteworthy that this involves throwing out quite a lot of observations (for example, in the case of Hospital-Specific *HHI* it causes the number of observations to fall from 64,800 to 36,270. However, it should be kept in mind that these relationships are each estimated with only ten data points (one for each hospital), so it is perhaps not too surprising that a positive and significant relationship is not always estimated. As one would expect, throwing out those observations in which the estimated relationship between price and concentration is negative or slightly positive increases the average prediction of the merger effect, and so reduces the degree to which the predicted effects are smaller than the actual effects.

At this stage, we are not certain whether this tendency to under-predict the true merger effects is a robust result, or whether it is an artifact of the particular simulations that we have run so far. We expect that this will be clarified as we run more simulations (see below). One factor that may be contributing to this result is the fact that, at present, there are no hospitals that are already jointly owned before the mergers that we study. This means that for mergers where diversion is high, the post-merger values of the concentration measures are out-of-sample: no hospital in the pre-merger data had values of *WTP* or Hospital-Specific *HHI* as high as some post-merger hospi-

---

the bargaining model will always return a "true" merger effect of zero. The CDS version of *WTP* and Hospital-Specific *HHI* have the property that when diversion is zero, a merger does not change the concentration measure at all, and so the predicted merger effect will be zero, and hence the percentage point difference between the two will also be zero. The TV version of WTP (and the *WTP_PP* measure that is based on it) does not have this property. For this reason, the percentage point difference between the true effect and the predicted effect need not be zero.

tals did. Since the estimation of the relationship between price and the concentration measures is done on pre-merger data, this might bias the estimate, and lead to inaccurate predictions.

V. Extensions

The primary goal of this paper is to test the accuracy of the merger simulation methodologies based on Willingness-to-Pay and Hospital-Specific *HHI*. However, it should also be possible to use our bargaining model to simulate specific mergers directly. That is, it should be possible to set up the model in such a way that the locations of the relevant hospitals, populations, geographic barriers, and so on reflect the actual realities of the case in question. The model parameters could then be calibrated to match the observed pre-merger prices, and the merger could be simulated directly. This approach has advantages and disadvantages relative to the approach outlined in this paper, and we leave it as a subject for future research.

VI. Conclusion

In recent years evidence has mounted that some hospital mergers can be expected to result in substantial price increases. As a consequence, researchers have become interested in developing new methods of predicting the effects of hospital mergers. The purpose of this paper is evaluate two of those methods by comparing their predicted post-merger price increases with the "true" price increases for a hypothetical merger using simulated data.

This research can be extended by collecting more results from numerous variations in the environment. The obvious adjustments to our methodology would allow us to test the performance of these merger simulation methodologies allowing some of the non-merging hospitals to be jointly owned, allowing multi-hospital system acquisitions of independent hospitals, and allowing mergers between multi-hospital systems, varying travel costs, the locations of the hospitals and the distribution of the populations, the relative bargaining power of the hospital and the MCOs, and so on. The ultimate goal is to arrive at a situation where it will be possible to identify which version of the price-concentration analysis is most appropriate for a given fact pattern of a merger under consideration. That is, we intend for this work to be of direct use to practitioners of hospital merger analysis (like us).

Appendix

In this appendix, we discuss the differences between the TV and CDS market power measures and why they are irrelevant for our study. The way that CDS construct their *WTP* measure is described in the text. TV construct their market power measure slightly differently (and do not use the term *WTP*). Specifically, TV construct their measure as follows. Like CDS, they begin by estimating a discrete choice model as in (1). Also like CDS, the valuation of a hospital for consumer *i* is equal to the expected value of the maximum as in (2). Both methods then aggregate these individual valuations to generate a measure of the market power of the hospital. For TV, the market power of hospital *k* is defined as the difference between the mean valuation for the network *G*, and the mean valuation for the network $G \backslash k$.

$$(A1) \qquad \Delta W = W^{G} - W^{G \backslash k} = \frac{1}{N}\sum_{i}\ln\sum_{j \in G} e^{V_{ij}} - \frac{1}{N}\sum_{i}\ln\sum_{j \in G \backslash k} e^{V_{ij}}$$

In their price regressions, TV include only $W^{G \backslash k}$ (the second term in *A*1) as a right-hand side variable since the first term (the average valuation of the whole choice set *G*) will be the same for all observations and so will just be swept into the intercept. In the framework that we use in this paper, this will produce exactly the same estimate of the effect of the relationship between price and the market power variable as we get using the CDS measure (though the sign will be reversed), since the two are linear transformations of one another. The two methods will produce different intercepts, but that does not matter for our purposes.

Note that the two measures would not produce the same estimates if we did our analysis using logs of the market power variables instead of levels. And TV do use logs, whereas CDS use levels. This is not an issue for us, however, as we follow CDS and run our regressions in levels.

There several other differences between TV and CDS. First, TV define $W^{G \backslash k}$ under two circumstances: (i) the mean expected welfare of hospital *k* is defined on the set $G \backslash k$; and (ii) the mean expected welfare of hospital *k* is defined on the set {*G**k*,*k*'}, where *k*' denotes the best substitute for *k* of the currently excluded hospitals. TV then use a switching regression framework to incorporate the uncertainly (from the perspective of the researcher) about which state is generating the data, (i.e., which network is the most relevant in constraining the price of hospital *k*). CDS do not observe network exclusions, and so this is irrelevant for their study. In our study, the equilibrium network is always the network of the whole. Hence, the switching regression framework employed by TV is irrelevant for us as well.

The second difference is that TV use DRG weights from Medicare's Prospective Payment System to weight utilities defining $W^{G\backslash k}$ over various clinical conditions. The rationale is that more serious conditions should be given more weight in defining consumer valuation of a network. CDS do not use this approach and state in their appendix that TV do not apply the weights correctly. In our study, there is, in effect, only one clinical condition, so the weights are irrelevant.

The third difference is that TV use price as a dependent variable while CDS use profits. Since we are primarily interested in price effects, we use prices as our dependent variable.

References:

Capps, Cory., and Dranove, David. “Hospital Consolidation and Negotiated PPO Prices.” *Health Affairs* 23 (2004): 175-81.

Capps, Cory., Dranove, David., and Satterthwaite, Mark. “Competition and Market Power in Option Demand Markets.” *Rand Journal of Economics* 34 (2003): 737-63.

Dor, Avi., Grossman, Michael., and Koroukian, Siran M. “Hospital Transaction Prices and Managed-Care Discounting for Selected Medical Technologies.” *American Economic Review* 94 (2004): 352-6.

Gaynor, Martin., and Vogt, William B. “Competition Among Hospitals.” *Rand Journal of Economics* 34 (2003): 764-85.

Melnick, Glenn., and Keeler, Emmett. “The Effects of Multi-Hospital Systems on Hospital Prices.” *Journal of Health Economics* 26 (2007): 400-13.

Simpson, John., and Shin, Richard. “Do Nonprofit Hospitals Exercise Market Power.” *International Journal of the Economics of Business* 5 (1998): 141-57.

Sorensen, Alan T. “Insurer-Hospital Bargaining: Negotiated Discounts in Post-Deregulation Connecticut.” *Journal of Industrial Economics* 51 (2003): 469-90.

Town, Robert., and Vistnes, Gregory. “Hospital Competition in HMO Networks.” *Journal of Health Economics* 49 (2001): 63-84.

Vita, Michael G., and Sacher, Seth. “The Competitive Effects of Not-for-Profit Hospital Mergers: a Case Study.” *Journal of Industrial Economics* 49 (2001): 733-53.

Vogt, William B., and Town, Robert. “How has Hospital Consolidation Affected the Price and Quality of Hospital Care?” *Robert Wood Johnson Foundation Research Synthesis Report No. 9* (2006).

Table 2: Summary Statistics

Deviation from Truth as a percent of Pre-Merger Price

		Mean	St Dev	F(0.02) – F(-0.02)	F(0.05) – F(-0.05)	1- F(0.05)
WTP_PP	All	-0.001	0.018	0.858	0.971	0.022
	Div > 0.2	-0.011	0.022	0.671	0.956	0.014
	Div > 0.4	-0.021	0.024	0.406	0.907	0.009
WTP (TV)	All	0.003	0.021	0.833	0.957	0.033
	Div > 0.2	-0.010	0.028	0.593	0.916	0.030
	Div > 0.4	-0.025	0.032	0.326	0.814	0.023
WTP (CDS)	All	-0.009	0.013	0.884	0.982	0
	Div > 0.2	-0.029	0.018	0.371	0.902	0
	Div > 0.4	-0.047	0.021	0.021	0.666	0
HSHHI	All	-0.006	0.012	0.921	0.987	0.000
	Div > 0.2	-0.019	0.021	0.577	0.928	0.002
	Div > 0.4	0.034	0.029	0.300	0.765	0.005

Table 3: Summary Statistics

Deviation from Truth as a percent of Pre-Merger Price

Coefficient Estimate > 0 and p-value < 0.10 Only

		Mean	St Dev	F(0.02) – F(-0.02)	F(0.05) – F(-0.05)	1- F(0.05)
WTP_PP N=57,960	All	-0.000	0.019	0.850	0.968	0.024
	Div > 0.2	-0.009	0.020	0.677	0.958	0.016
	Div > 0.4	-0.019	0.023	0.434	0.920	0.010
WTP (TV) N=32,760	All	0.008	0.024	0.789	0.933	0.062
	Div > 0.2	-0.000	0.028	0.647	0.920	0.057
	Div > 0.4	-0.013	0.031	0.447	0.875	0.044
WTP (CDS) N=32,760	All	-0.008	0.012	0.889	0.984	0
	Div > 0.2	-0.023	0.017	0.388	0.912	0
	Div > 0.4	-0.045	0.020	0.038	0.690	0
HSHHI N=36,270	All	-0.004	0.010	0.946	0.993	0.000
	Div > 0.2	-0.012	0.019	0.711	0.962	0.003
	Div > 0.4	-0.022	0.028	0.382	0.864	0.010

Figure 1: (WTP_PP Estimate - Truth)/Pre-Merger Price on Diversion

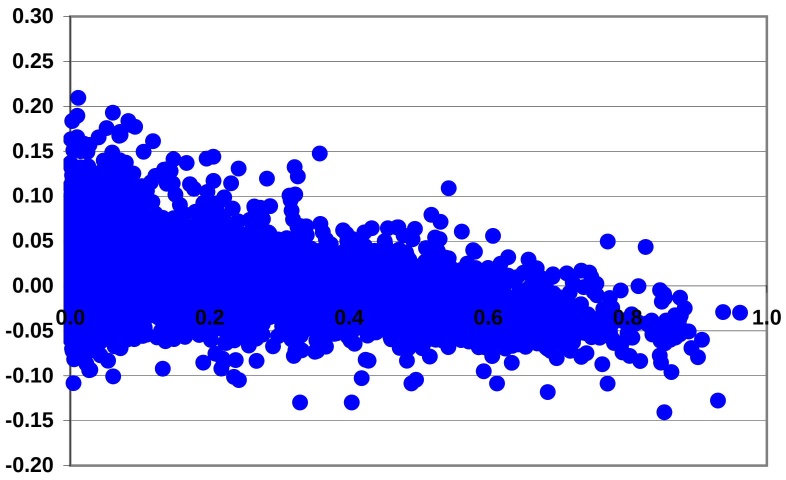

Figure 2: (WTP(TV)- Estimate - Truth)/Pre-Merger Price on Diversion

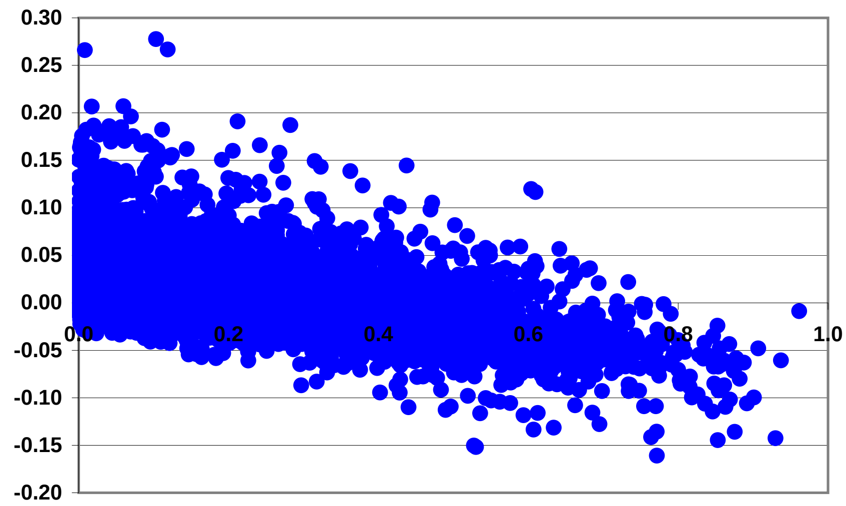

Figure 3: (WTP(CDS) Estimate - Truth)/Pre-Merger Price on Diversion

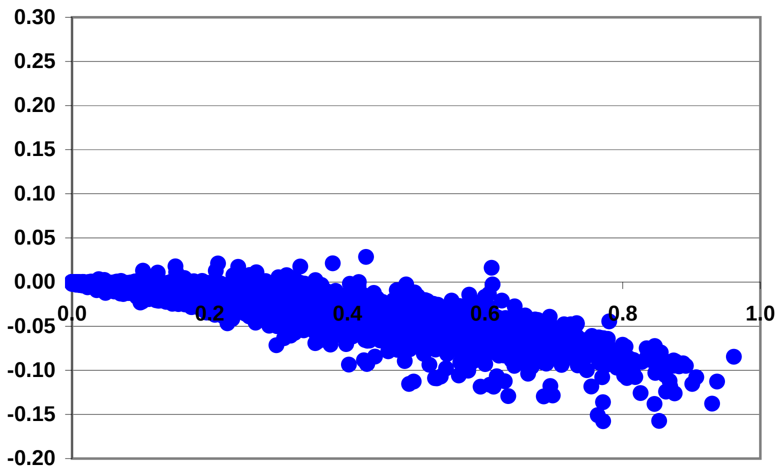

Figure 4: (HSHHI Estimate - Truth)/Pre-Merger Price on Diversion

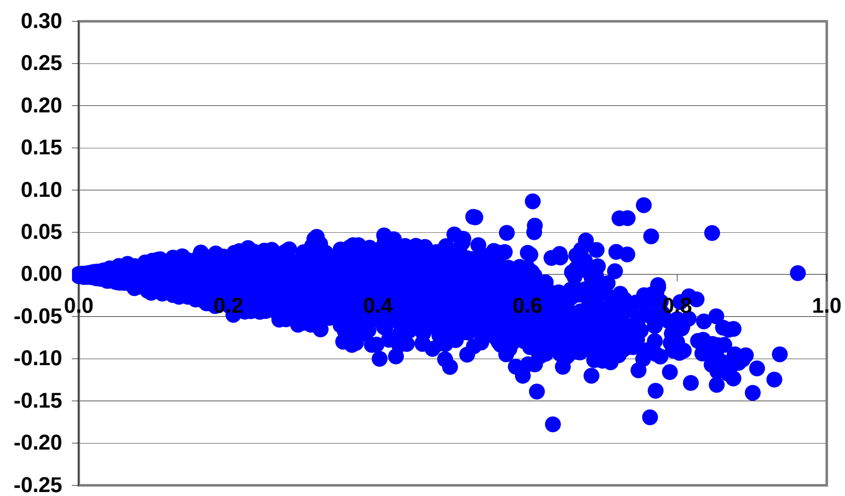

Figure 5: Kernal Densities of (WTP_PP - Truth)/Pre-Merger Price

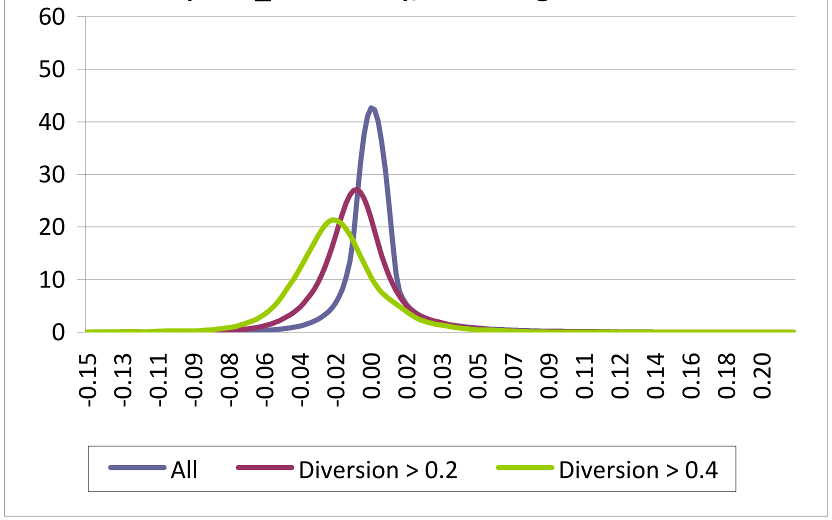

Figure 6: Kernal Densities of (WTP(TV) - Truth)/Pre-Merger Price

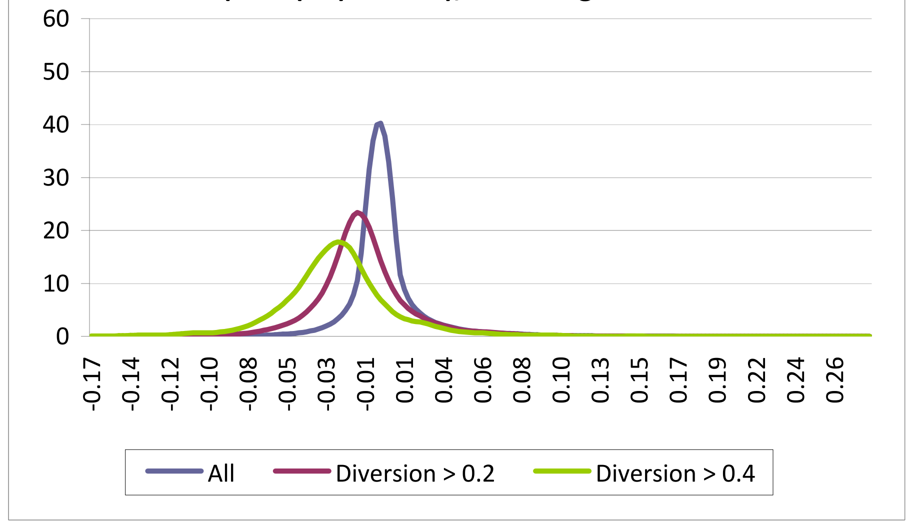

Figure 7: Kernal Densities of (WTP(CDS) - Truth)/Pre-Merger Price

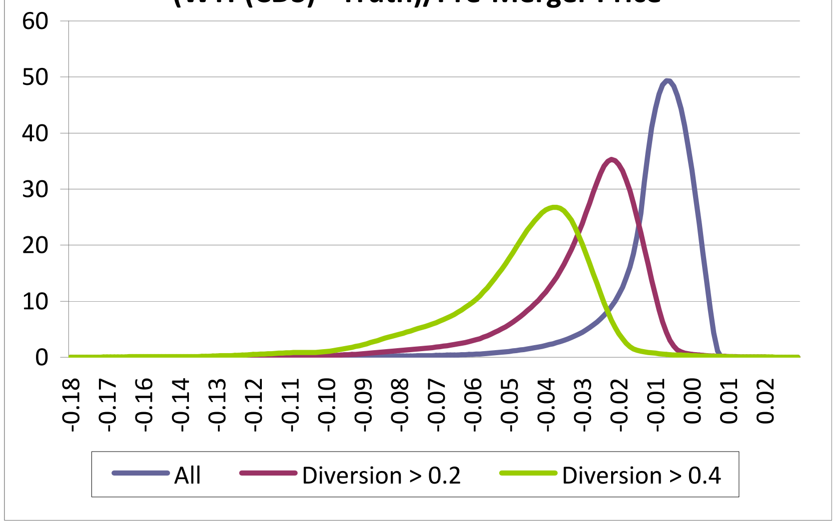

Figure 8: Kernal Densities of (HSHHI - Truth)/Pre-Merger Price