Economics at the FTC: Estimating Harm from Deception and Analyzing Mergers

Michael Vita
Keith Brand
Miriam Larson-Koester
Nathan Petek
Charles Taragin
William Violette
Daniel H. Wood

Abstract: Economists in the U.S. Federal Trade Commission's Bureau of Economics perform economic analysis in support of the Commission's dual missions to protect consumers and competition by preventing anticompetitive, deceptive, and unfair business practices through law enforcement, advocacy, and education. This article first presents summaries of analyses that FTC economists performed to estimate the consumer harm from two different types of deception that involved misleading information about lease terms and suppression of negative product reviews. The essay next turns to economic analyses of mergers: We first consider the vertical issues that arose in a semiconductor merger; and then we provide a discussion of how complementarity between hospitals may affect the analysis of hospital mergers.
Keywords: Antitrust, Consumer Protection, Deception, FTC, Mergers

I.   Introduction

This article summarizes work done by economists in the Federal Trade Commission's (FTC's) Bureau of Economics (BE) in support of the commission's consumer protection and antitrust missions. In particular, it highlights the analyses that were performed by consumer protection economists to assess the harm to consumers who were provided bad information in two different settings: In the first, consumers were misled about the true terms of rent-to-own payment plans; and in the second, consumers were shown only a censored set of consumer reviews of apparel products. On the competition side, we first discuss the analysis of a proposed semiconductor merger that raised vertical concerns, and then we explore how the evaluation of hospital mergers can assess and account for complementarity between hospitals.

The staff of the Federal Trade Commission's Bureau of Economics is composed of about 74 Ph.D. economists, nine research assistants/statisticians, eight financial analysts, and five administrative professionals. BE mainly contributes to the FTC's missions by performing economic analysis in extensive law enforcement investigations and rulemakings that are conducted along with teams of FTC attorneys from the Bureau of Competition and the Bureau of Consumer Protection. The information that is generated by the economists and attorneys serves as an important input into the decisions that are made by the FTC Commissioners. When the Commission pursues a legal challenge in court, BE economists help to develop economic evidence and, in some cases, may testify in court.

BE economists also conduct economic research to address topics of economic importance outside the context of law enforcement investigations. Our economists write working papers and journal articles that are based on original research that they conduct in addition to their other professional duties. BE economists also conduct commission-wide studies and reports, often in collaboration with FTC attorneys.[1]

---

[1] For instance, the FTC announced in January of 2021 a study of physician group and healthcare facility acquisitions. The Commission voted to use compulsory process to acquire data from six major health insurers to study the effects of physician group and healthcare facility consolidation that occurred from 2015 through 2020; see https://www.ftc.gov/news-events/news/press-releases/2021/01/ftc-study-impact-physician-group-healthcare-facility-mergers.

BE also participates in the larger economics community via participation in seminars and conferences. Notably, on November 4 and 5, 2021, BE -- along with co-sponsor, the Tobin Center for Economic Policy at Yale – held the fourteenth FTC Microeconomics Conference.[2] The topics that were covered in keynote addresses and paper presentations included: gig work and digital platforms; data governance; pharmaceutical rebates; misinformation in advertising; product market concentration; and online privacy. On November 3-4, 2022, BE and the Tobin Center will again hold the FTC Microeconomics Conference in Washington, DC.[3]

Section II of this article describes the economic analysis of harm to consumers from allegedly misleading information about rent-to-own leasing terms that involved the firm Progressive Leasing. Section III presents a methodology for exploring the consumer harm from the censorship of negative reviews by Fashion Nova, a marketer of apparel. Section IV presents some of the economic issues that were considered by BE in the NVIDIA/Arm semiconductor merger investigation. Finally, Section V explores the conditions that may lead payers to view merging hospitals as complements and how that complementarity would affect the analysis of the merger's effects.

II.   Estimating How Many Consumers were Misled in the Progressive Leasing Investigation Using Store-Level Complaint Rates

       A.   Background

Progressive Leasing is a leading virtual rent-to-own business in the United States. In April 2020 the Federal Trade Commission reached a settlement with Progressive. The FTC alleged in its complaint that Progressive had represented, "expressly or by implication, that consumers [would] pay the retail or cash price to purchase merchandise," when in fact purchasing a good through a

---

[2] Copies of the papers that were presented along with a video of the conference are available at https://www.ftc.gov/news-events/events/2021/11/fourteenth-annual-federal-trade-commission-microeconomics-conference.

[3] Details are available at: https://www.ftc.gov/news-events/events/2022/11/fifteenth-annual-federal-trade-commission-microeconomics-conference. Should public health conditions dictate, this conference may need to be held virtually.

Progressive lease was always more costly than paying the retail price. The settlement provided $175 million in monetary relief in addition to putting Progressive under order not to misrepresent the cost of its leases in the future.[4]

One of BE's roles in the Progressive investigation -- which is a typical role for BE for FTC consumer protection investigations -- was to estimate consumer injury that was caused by the alleged practices. This section describes how we estimated a portion of the injury in this case.

Progressive operates through partnerships with retail companies such as Big Lots, Best Buy, and smaller retailers that offer Progressive leases in their stores as a form of financing: mostly for customers who do not qualify for other credit options that are available in the retailers' stores. If a customer purchases an item with the use of a Progressive lease, Progressive pays the retailer for that item and leases it back to the customer. Once the customer pays off the lease -- which are almost always 12 months in length -- the customer takes ownership of the item.

Paying a lease off over its full term is relatively expensive: In 2019, obtaining ownership of a good through paying off a Progressive lease over its full term involved a cost that was approximately twice the retail price. Progressive also offered a 90-day payment option if customers paid off their lease within 90 days, which was less expensive. In addition to payment of the retail price, the 90-day option required payment of a fee of $49 or $79. Last, Progressive offered an "early buyout option" (EBO) for consumers who paid off the lease after more than three months but less than a year. To exercise the EBO customers typically needed to pay an additional fee of 65% of their remaining balance, which meant that the lease cost under EBO payoff was somewhere between 1.5 and 2 times the retail price. Substantial fractions of customers exercised either the 90-day payoff option or the EBO.

At the time of the investigation, consumer complaints and other evidence indicated that some retailer partner stores or their sales associates sometimes misrepresented the cost of these leases. For instance, sales associates often mischaracterized Progressive's 90-day payment option as a "90-day same as cash offer" -- despite there being a fee that was involved in originating the lease. Sales associates also sometimes provided incorrect or misleading information about the full-

---

[4] See https://www.ftc.gov/legal-library/browse/cases-proceedings/182-3127-progressive-leasing for details on the case.

term lease cost, such as by claiming that it included "no interest," which was likely to mislead those consumers who interpreted it to mean that they would not owe more than the retail price in total.

An alternative way that customers might have learned the full information about lease costs was during the lease application process, which was conducted in-store using Progressive's web-based "Approve.me" platform. This multistep process guided consumers through selecting the goods that they were leasing and had them e-sign rent-to-own contracts. Unfortunately, Progressive made important lease cost information not visible by default on this platform. While the retail price was prominently displayed, both the total lease cost and the 90-day fee were placed behind an "Additional lease terms" clickable button. Omitting this key information likely misled many consumers, who might have misunderstood the cost of the lease for a variety of reasons, including but not limited to consumers who were affirmatively misled by sales associates' claims.

B.   Theory of Injury

In cases where a company provides misleading information to consumers, consumer injury can be measured as the difference between the consumer welfare that consumers actually obtained and the welfare that they would have obtained if the company had not misled the consumers. Therefore, we needed to understand what each consumer would have done in the counterfactual world in which Progressive disclosed the complete lease terms to him or her. A consumer's counterfactual behavior depends on that consumer's next best alternative to taking a Progressive lease and using the observed payoff option.

Many consumers who paid off Progressive leases paid them off full term or with EBO. But consumers had a strong incentive to pay off a Progressive lease within 90 days, as the cost increased discontinuously at the 90-day mark by potentially hundreds of dollars. In addition, Progressive leases were often consumers' financing option of last resort. Thus, it seems likely that most consumers who paid off their leases more slowly than 90 days were simply unable to pay them off more quickly. However, many other consumers did pay off their leases within 90 days.

Based on this evidence of differing abilities to pay off leases quickly, we divided consumers conceptually into two categories. The first category was financially constrained consumers who would not have been able to take advantage of the 90-day option. The second category was financially unconstrained consumers, who we defined as able to pay off the lease within 90 days.

Despite the solvency of these consumers, we inferred from the high take up of leases subsequently paid off within 90 days that the short-term flexibility offered by the leases was valuable to these consumers. While we did not observe the type of consumer that took a lease directly, nor did we possess consumer credit scores or other measures of potential liquidity, we were able to use consumer behavior to estimate the share of consumers of each type.

In this article we focus on how we estimated injury for the financially unconstrained group.[5] Because of the incentives to do so, financially unconstrained consumers who were not deceived would almost always pay off leases within 90 days; consequently, the share (of all leases) that had 90-day payoffs is a reasonable estimate for the share of financially unconstrained consumers who were not injured. Financially unconstrained consumers who were deceived would potentially pay off the lease more slowly, and thereby incur higher costs: the primary form of injury for this group. We estimated the total injury for this group by determining both the approximate injury from slower-than-optimal payoffs and the number of unconstrained consumers in this group that were deceived to their detriment. Total injury is the product of per-consumer injury and the number of injured unconstrained consumers.

C.   Per-Consumer Injury as Differences in Lease Costs

Members of the financially unconstrained group who were deceived to their detriment missed the opportunity to pay off the lease at the 90-day cost. It seems likely that these consumers eventually realized they had missed the 90-day option when payments on the lease continued to be collected after 90 days. Most of them upon realizing this would likely have exercised the early buyout option, which increased the lease cost by hundreds of dollars, but was significantly less costly than a full-term payoff. An additional injury would occur if these consumers would have chosen not to lease the item, but instead just bought it outright; in this case, paying the fee for the

---

[5] Financially constrained customers necessarily faced a lease cost that was significantly greater than the retail price. Because of that, it is likely that their largest injury stems from the decision to take the lease because they did not understand the lease cost, when they would not have leased from Progressive if they had known the true lease cost. Consumers who would not have taken a Progressive lease if they had been aware of the true cost could have forgone purchasing the good altogether or used an alternative financing option if a less expensive one was available to them. We estimated that injury in a different manner which we do not describe here.

90-day payoff is also an injury that they suffer. In summary, for a consumer in this group, he or she will exercise the EBO, and we estimated his or her injury as the difference in cost between his or her EBO lease cost and the lease cost under the 90-day option.[6]

Progressive provided the FTC with several datasets on its business operations, which we used to estimate the potential consumer injury to Progressive customers. One dataset was a lease-level dataset, which allowed us to calculate the average injury to a member of this group from being misled.[7]

D.   The Number of Injured Consumers

The more difficult question was how many financially unconstrained Progressive customers were misled to their detriment. To determine how many consumers were misled by Progressive's practices, we supplemented the lease data with store-level data to estimate how the 90-day payoff rate varied across individual retail stores with varying consumer complaint rates. We used an internal Progressive complaints dataset and counted the number of complaints that were from categories that could indicate customer confusion about lease costs: e.g., complaints about lease terms' not being explained. We then divided the number of these complaints by yearly counts of leases that were originated at each store to calculate the complaint rates at each retail store between April 2016 and March 2017 and also between April 2017 and March 2018.

This relationship between a store's 90-day payoff rates and complaint rates is useful because a store's complaint rate serves as a proxy for the quality of information that was provided to consumers visiting that store. Stores with no complaints were more likely to have not misinformed consumers about the leasing option, while stores with high complaint rates were more likely to have misled consumers through sales associates' omitting important cost information or

---

[6] We could not determine whether one of these consumers would have used the 90-day option or would have forgone the lease financing altogether. Because both of these alternatives were significantly cheaper than an EBO, in practice the difference between the EBO and 90-day cost captured the majority of injury. We also calculated and reported to the Commission an injury estimate that included the 90-day fee.

[7] A few consumers took multiple Progressive leases, but these consumers were less likely to be injured – they may have been satisfied by the first Progressive experience, hence were less likely to have been misled – so in analyzing injury we restricted the data to one-time consumers of Progressive.

even making affirmative misrepresentations, and thereby exploiting Approve.me's disclosure inadequacies. Financially unconstrained customers who were well informed knew they had a strong financial incentive to pay off the lease in 90 days, and would do so at a higher rate than customers with bad information about lease costs. Using this variation, it would be possible to estimate what the 90-day payoff rate would be in the counterfactual scenario in which Progressive or its retailer partners' salespeople did not mislead any consumers.

1. Econometric Model

We assumed that each store that offered Progressive leases was endowed with a propensity to mislead consumers who use Progressive Leasing to purchase an item at that store. For example, a particular retail store might mislead 5% of its Progressive consumers about Progressive lease costs, but another store might mislead none of its Progressive consumers. Complaint rates at stores are a proxy variable for the stores' underlying propensity to mislead consumers about Progressive lease costs.

A major complication in estimating the relationship between store-level complaint rates and take-up of the 90-day option is that complaint rates are an imperfect measure of the underlying propensity of the store experience to mislead consumers. Many stores had only a relatively small number of customers, and different consumers have idiosyncratic differences in the likelihood of their complaining. For both of these reasons, complaint rates are a noisy signal of the propensity of a store to mislead consumers.

Because of this noise, estimates of the relationship between complaint rates and 90-day payoff rates with the use of ordinary least squares (OLS) would suffer from attenuation bias, which would likely cause the estimates to be an underestimate of the relationship between bad information and 90-day payoff rates. We dealt with attenuation bias by using instrumental variables (IV) estimation, with stores' 2016 complaint rates instrumenting for their 2017 complaint rates. Because the 2017 random noise is avoided, the IV estimate is not biased towards finding a smaller than true effect (Hausmann 2001).

Our econometric model is a lease-level outcome linear probability model which explains customer i's decision to pay off the lease he or she received from store j in 90 days -- $ninety_{ij}$ -- as

a function of the complaint rate at store j and the retail price of the good that i leased in logs: $lnprice_i$:[8]

$$ninety_{ij} = \alpha + \beta\ complaintrate_j + \gamma\ lnprice_i + \varepsilon_{ij}.$$

Standard errors were clustered at the store level. We control for the retail price of the good because, ceteris paribus, more expensive goods are more difficult to pay off in 90 days.

Let $complaintrate_{j,t}$ be store j's complaint rate in year t. Progressive coded complaints from April 2016 to March 2017 with the use of several, but not all, of the complaint categories that it used starting in April 2017. Therefore, we used the complaint rate at store j from April 2016 to March 2017 to instrument for the 2017 to 2018 complaint rate. In other words, $complaintrate_{j,2017}$ is predicted using $complaintrate_{j,2016}$ along with $lnprice_i$ and we estimated β from the predicted $complaintrate_{j,2017}$ variable.

Both β and γ were negative, economically significant, and statistically significant at the 1% level. The first-stage F statistic indicated that $complaintrate_{j,2016}$ was not a weak instrument.

2. Robustness Checks

The estimates were a valid measure of the effect of store propensities to mislead with respect to 90-day payoff rates if we assume that consumers at stores that provided worse information were not systematically different from consumers at stores that provided better information -- controlling for the retail price of the leased good. We possessed limited demographic information on consumers: age; monthly gross income; approval amount (the maximum lease amount that Progressive would allow); and consumer zip code.

We tested the assumption that consumer characteristics are uncorrelated with store propensities in several ways. First, comparing consumers from stores that had the lowest third of complaint rates to consumers from stores that had the highest third of complaint rates did not reveal any meaningful differences in customer age, monthly gross income, approval amount, or retail price of the leased good. Second, neither including fixed effects for customer zip code nor including customer income as a control significantly affected our parameter estimates compared to the base model that was described above.

---

[8] A probit model yielded similar results.

3. Calculating the Number of Injured Consumers

We found that stores with higher customer complaint rates also on average had lower 90-day payoff rates. That allowed us to determine how much higher the 90-payoff rate would have been in the counterfactual world in which there were no consumer complaints, by predicting $ninety_{ij}$ if $complaintrate_j = 0$. Then we calculated total injury to the financially unconstrained consumers by calculating the number of those consumers and multiplying that by the average per-consumer injury described above. Our total injury estimate for this group was substantial.

We were also able to consider other counterfactuals. Taking zero complaints as the counterfactual outcome in which Progressive had not misled consumers assumes that all consumer confusion stemmed from Progressive's deception, rather than from misunderstandings that were attributable to the consumer's failure to understand a complex contract or other reasons for confusion. Therefore, the calculations that were described above with the zero-complaint counterfactual represented an upper bound on injury.

E. Discussion

Often in consumer protection investigations, the FTC begins with, or is able to obtain, consumer complaint data. Because most consumers do not complain, however, one challenge BE often faces is extrapolating from complaint data to the total number of consumers that were injured by the practices about which we have complaints.[9] In Progressive, we were able to leverage rich store-level variation in the degree to which consumers were misled -- combined with behavioral predictions as to how misled consumers would act -- in order to translate from complaints to injured consumers.

---

[9] Raval (2020) analyzed both complaint data and victim lists from fraud cases, and found considerable variation in the ratio of complaints per victim across cases.

III.   Does Hiding Negative Consumer Reviews Boost Sales? Evidence from the Fashion Nova, LLC Case

A.   Background

Over the past few years, the FTC has taken an active role against firms that manipulate the product reviews that are shown to consumers.[10] While some firms have turned to review manipulation in an effort to improve their brand image and artificially boost sales, it remains an open question as to whether this strategy is profitable. For example, hiding or distorting negative product reviews may reduce the product information that is available to consumers, which could lead consumers to make suboptimal purchases and increasing product returns to firms. Existing economic research has provided little insight into the consequences of review manipulation -- possibly because information about these practices is rarely available to academic researchers.[11]

The FTC's case against the fast-fashion company Fashion Nova, LLC provided BE economists with a unique opportunity to estimate the effect of review manipulation on sales and prices. We performed analysis to explore whether Fashion Nova's practice of hiding negative consumer reviews allowed Fashion Nova to maintain prices on low-rated products without suffering lower consumer demand. Like other deceptive or misleading advertising, the theory is that review suppressions would shift the consumer demand curve upward.[12] In March 2022, the

---

[10] See FTC press releases: "FTC Approves Final Consent Agreement with Sunday Riley Modern Skincare, LLC," November 2020; "FTC Puts Hundreds of Businesses on Notice about Fake Reviews and Other Misleading Endorsements," October 2021; "Fashion Nova will Pay $4.2 Million as part of Settlement of FTC Allegations it Blocked Negative Reviews of Products," January 2022.

[11] Existing research largely relies on partnerships with online platforms as well as tedious manual data collection. Luca and Zervas (2016) partnered with Yelp to study fake reviews that were identified by Yelp's machine learning algorithm. Hollenbeck et al. (2022) collected data manually from Facebook groups that facilitate sales of fake Amazon reviews.

[12] See Wood and Stone (2018).

FTC announced a $4.2 million settlement with Fashion Nova, LLC that addressed allegations that the company suppressed negative consumer reviews of its products from its website.[13]

B. Suppressing and Revealing Negative Reviews

Fashion Nova sells clothing products almost exclusively online through its website. The FTC learned that Fashion Nova had been suppressing all consumer reviews that were posted to its website with ratings three stars or below since late 2015. In response to FTC inquiries, Fashion Nova decided to post to its website all previously-suppressed reviews.

Several features of Fashion Nova's negative review posting make it an ideal setting for an economic analysis of review manipulation. First, Fashion Nova posted all reviews quickly over the course of a few days in November 2019. Second, as far as we are aware, the review posting was not accompanied by any substantial new product launches, advertising campaigns, or other design changes on the website. Finally, the review posting affected only products with any reviews that were three stars and below. Products with reviews that were only four stars and above appeared unchanged, providing a useful way to control for other trends in sales and prices over the review-posting period.

C. Measuring Reviews, Sales, and Prices

To measure the effect of the negative review posting, we received data on the daily total sales in units and revenue for each product from 2017 to 2020. Products include mostly women's clothing priced around $20. The data also provided the average star rating and total number of reviews that had been submitted as of November 2019 for each product. This average star rating included all reviews: the previously suppressed reviews and also the always-posted reviews. We divided products into those with average ratings above 3.5 stars – high-rated products -- and those with average ratings below 3.5 stars: low-rated products.[14] This approach leveraged the intuition

---

[13] See FTC press release "FTC Finalizes Order with Fashion Nova Over Allegations it Blocked Negative Reviews" March 2022.

[14] We also tried different cut offs for high- and low-rated products, and the results remained similar.

that low-rated products were more likely than were high-rated products to be affected by the review posting.[15]

Our empirical strategy attributed differential changes in sales and prices for low-rated versus high-rated products to the effect of the review posting.[16] Since the review posting occurred just before "Black Friday",[17] this approach allowed us to use trends for high-rated products around Black Friday to understand what might have happened to low-rated products in the absence of the review posting.

To support our empirical approach, we explored whether the low- and high-rated products followed similar trends in sales and prices leading up to posting all negative reviews. Parallel patterns would provide support for the assumption that low- and high-rated products would have continued to follow similar trends in the absence of the review posting.

D.   Estimating Difference-in-Differences

To isolate the effect of posting negative reviews, we estimated the following difference-in-differences equation:

$$y_{it} = \beta_0 + \beta_1 post_t + \beta_2 low\ rated_i + \beta_3 post_t \times low\ rated_i + X_t + \varepsilon_{it}$$

Where: Outcome y_{it} includes prices and unit sales for each product i and week t; $post_t$ is an indicator for weeks after the review posting; $low\ rated_i$ is an indicator for products with relatively low average star ratings; and X_t includes additional time controls to account for month-specific seasonality and the weeks that precede and follow Black Friday.

Our coefficient of interest -- β_3 -- captures the differential change in outcomes for low-rated products relative to high-rated products after the review posting. We interpret this coefficient as capturing the causal effect of the review posting under the assumption that outcomes for low-rated

---

[15] Unfortunately, by including average star ratings instead of individual reviews, the data prevented us from identifying products with only four stars or above reviews, which would have been relatively unaffected by the review posting.

[16] Implicit in this exposition of the analysis is an assumption that the release of reviews did not materially impact the number of consumers considering purchases at this retailer.

[17] This phrase refers to the Friday immediately after Thanksgiving in November, when retailers have often accumulated enough operating surpluses so as to have covered their annual fixed costs and are thenceforth operating "in the black". It is also a traditional day for special low-price limited-time "sales".

products would have evolved the same way as those for high-rated products in the absence of review posting: Without posting all reviews, all products would appear to be highly rated. This assumption would not hold if high-rated products are also affected by review posting: either directly because high-rated products may have some negative reviews, or indirectly because consumers substitute high-rated for low-rated products as a result of review posting.

Our analysis may underestimate the effect of the review posting because high-rated products may still have had some reviews that were three stars or below and therefore were at least partly affected by the review posting. At the same time, our analysis may overestimate the effect of the review posting if consumers responded to negative reviews on low-rated products by purchasing more high-rated products instead. We are unable to determine whether our results are biased upwards or downwards overall.

E.   Interpreting Estimated Effects

To maintain the confidentiality of Fashion Nova's data, we do not report the estimated effects, but we can discuss the implications that may be drawn from various combinations of estimated effects. The effect of the suppression of negative reviews on consumer welfare can be viewed as a question of whether it artificially inflated demand, which caused consumers to be willing to pay more than they would have with more complete information.

The important question is whether the analysis of the company's data from before and after the review posting indicates that the disclosure of negative reviews shifted demand for those products. A finding of estimated effects -- β_3 -- that have opposite signs for price and quantity sold would be inconclusive, because it would be consistent both with a shift of the demand curve and with movement along an unchanged demand curve.

However, if both estimated effects have the same sign, this would indicate a shift of the demand curve. For instance, negative price and quantity effects would imply that demand shifted inward when the reviews were posted, which would be indicative of consumers' having been harmed by not having the low ratings disclosed to them initially. Contrarily, positive price and quantity effects would imply that demand shifted outwards in response to the review posting.

Therefore, this analysis allows us to explore whether the data may be informative about the impact of the review suppression on consumers.

IV.   Vertical Arithmetic and Efficiencies in the NVIDIA/Arm Merger

A.   Background

On December 2, 2021, the FTC sued to block the acquisition of Arm Ltd. By the NVIDIA Corp. The vertical merger would have combined NVIDIA -- one of the largest semiconductor firms[18] -- with Arm: the largest semiconductor intellectual property (IP) supplier.[19] This would have been the largest semiconductor merger in history.[20]

The FTC complaint alleged that post-merger, NVIDIA would have both the ability and incentive to degrade its (downstream) semiconductor rivals' access to Arm technology in three product markets: Data Processing Unit (DPU) SmartNICs; High-Level Advanced Driver Assistance System (ADAS) Central Compute SoCs; and Arm-based Datacenter CPUs for Cloud Computing Service Providers. The complaint alleged various mechanisms by which Arm's IP could be foreclosed, including: price; customer support; timing of access; and technological development.[21] BE economists analyzed the qualitative and quantitative evidence that was relevant to the theories of harm and the relevant markets that were presented in the FTC complaint.

NVIDIA abandoned the transaction in February 2022 following the issuance of the FTC complaint, and while it was still under review by multiple other jurisdictions, including the EC and

---

[18] https://www.mirrorreview.com/top-semiconductor-companies/

[19] https://s3.i-micronews.com/uploads/2020/10/IR20186-Design-IP_Extract.pdf

[20] https://www.ftc.gov/news-events/news/press-releases/2021/12/ftc-sues-block-40-billion-semiconductor-chip-merger

[21] The complaint, http://www.ftc.gov/system/files/documents/cases/d09404_part_3_complaint_public_version.pdf, also alleged a reduction in innovation through two channels: the sharing of competitively sensitive information between NVIDIA and its rivals; and a change in the future path of innovation of Arm to favor technologies that do not compete with NVIDIA.

the UK.[22] NVIDIA announced that it plans to continue to "partner closely with Arm," while Arm's parent company, Softbank, announced plans to take Arm public rather than seek another buyer.[23]

A common framework for assessing the impact of a vertical merger is to weigh two effects: the incentive by the merged entity to raise its rivals' costs (RRC); and the incentive of the merged entity to expand its output owing to the elimination of double marginalization (EDM); see Sheu and Taragin (2021). Below, we discuss how the weighing of RRC and EDM effects differed in the context of NVIDIA/Arm, including nonstandard aspects of both effects.

First, we note below public evidence that Arm's per-unit profits from licensing its IP are relatively low compared to NVIDIA's per-unit profits from processors. This fact means that one standard method for assessing the incentive to raise rivals' costs -- "vertical arithmetic" (which will be explained below) -- showed a clear incentive to do so.[24] However, the parties claimed both that intense upstream competition meant that Arm had no ability to foreclose its customers, and that NVIDIA had no incentive to do so due to the risk of undermining the Arm "ecosystem."[25] Below we discuss what we can infer from upstream and downstream margins for foreclosure incentive and ability in this matter.

Second, we note that the parties did not claim EDM as a pro-competitive justification for the merger.[26] Perhaps one reason that they did not claim that the merger would reduce the per-unit royalty that NVIDIA currently pays Arm for its IP is that NVIDIA would retain its 20-year license

---

[22] https://www.ftc.gov/news-events/news/press-releases/2022/02/statement-regarding-termination-nvidia-corps-attempted-acquisition-arm-ltd

[23] https://nvidianews.nvidia.com/news/nvidia-and-softbank-group-announce-termination-of-nvidias-acquisition-of-arm-limited

[24] See e.g. Pittman (2021).

[25] Answer and Defenses of Respondents to FTC Complaint
https://www.ftc.gov/system/files/documents/cases/d09404_-_answer_and_defenses_of_respondentsv_nvidia_corporationsoftbank_group_corp._and_arm_ltd.pdf

[26] Answer and Defenses of Respondents to FTC Complaint
https://www.ftc.gov/system/files/documents/cases/d09404_-_answer_and_defenses_of_respondentsv_nvidia_corporationsoftbank_group_corp._and_arm_ltd.pdf

with Arm regardless of whether the merger was consummated.[27] Instead, the parties claimed that the merger would increase investment in Arm IP -- perhaps by aligning the investment incentives of NVIDIA and Arm. In Section C below, we discuss two mechanisms that were proposed by the parties for this claimed investment boost: complementarity between NVIDIA's products and products that use Arm as an input; and complementarity between Arm's investment and NVIDIA's investment.

### B.	Inferences from Relative Upstream and Downstream Margins

Because the parties' answer to the complaint did not claim that the merger would result in EDM, this article will not focus on the equilibrium effects of both RRC and EDM.[28] Moreover, in the interest of brevity, we focus our analysis here on one concern that was articulated in the Complaint: that Nvidia might, post-merger, withhold Arm IP completely from its non-integrated rivals. In the case of non-price or total foreclosure, it is not necessary to weigh the additional profit that would be received from raising the input price. Vertical arithmetic considers foreclosure tactics that do not result in input price increases and is therefore well-suited to the concerns in the FTC Complaint.

### 1.	Vertical Arithmetic in NVIDIA/Arm

When a vertically integrated firm chooses not to sell a marginal unit of input to an unintegrated rival, it forgoes the profit that it would have earned on that sale. It can nonetheless increase its total profits if this foreclosure increases its downstream profits by a sufficient amount. Vertical arithmetic is a simple calculation to compare the forgone upstream profits to the profits that are gained downstream. The calculation requires only upstream and downstream price-cost margins, and a measure of downstream diversion from input foreclosure (Moresi and Salop, 2013).

If a vertically integrated firm withholds one unit of input from a downstream competitor, it loses the upstream margin in dollars on that input: M_U. However, to the extent that withholding that unit of input from one competitor increases input sales to other competitors, the upstream loss may

---

[27] https://www.tomshardware.com/news/nvidia-big-cpu-plans-with-20-year-arm-license

[28] Merger simulation could shed light on the equilibrium interaction of both effects (Das Varma and De Stefano, 2020).

be less than the full upstream margin. Let F be the percentage of foreclosed upstream sales that are recaptured. Then the full loss from foreclosure is given by $(1 - F) * M_U$.

When an upstream firm forecloses an input from a downstream competitor, some sales will be diverted to the downstream merging firm. This diversion is captured by DR_{UD}.[29] If we assume that one unit of upstream input is required for one unit of downstream product, the firm earns the downstream margin in dollars -- M_D -- on those sales: $DR_{UD} * M_D$. Furthermore, the merging downstream firm can now obtain inputs at cost from the upstream firm, which means that its margin is now inflated from its pre-merger margin by the amount of the upstream margin: $M_D = M_D^{Pre} + M_U$.

Vertical math subtracts from the incremental profit that is gained the incremental profit that is lost, so as to assess the overall profitability of foreclosing a marginal unit of input:

$$DR_{UD} * (M_D^{Pre} + M_U) - (1 - F) * M_U$$

If this difference is positive, then foreclosure is on net profitable. If no other firms increase their use of the input following foreclosure $(F = 0)$, the upstream firm loses its whole margin on the units that are foreclosed. For example, if the upstream firm forecloses all firms that use the input at once (other than the merging downstream firm), then no other firms will increase their use of the input. Although the downstream merging firm will increase its use of the input, the impact of this increase on the profit of the integrated firm is through increased downstream sales, not upstream sales, as the integrated firm does not pay itself an input price. Then the formula can be re-written as follows:

$$DR_{UD} - M_U / (M_D^{Pre} + M_U)$$

The merged firm will have an incentive to foreclose if diversion to the downstream firm exceeds the ratio of the upstream to post-merger downstream margin in dollars. Therefore, relative upstream and downstream margins are a key component of a vertical arithmetic analysis.

Importantly, it is the ratio of upstream (Arm) and downstream (NVIDIA) margins in dollars that matters, not percentage. Although Arm could have a relatively strong percent margin given the

---

[29] Note that DR_{UD} is not necessarily the same as downstream diversion. The foreclosure of the input could result in input substitution without diverting sales to the downstream firm. DR_{UD} can be written as diversion*departure rate (O'Brien, 2022).

likely low per-unit costs of IP licensing, Arm may recognize very little dollar value relative to the profits downstream. Although determination of the correct measure of margins can be a difficult task that relies an on analysis of proprietary information, the exposition of this analysis can be demonstrated using some publicly reported information: For instance, Arm's royalty rates have been reported to be around 1-2% of the downstream selling price of the products for which it licenses IP, with some variation depending on the IP.[30] By contrast, NVIDIA earns around a 65% margin on its products based on its public financials.[31] If we assume that the products that compete with NVIDIA are similarly priced, Arm is earning only $1-$2 for every $65 that NVIDIA is earning.[32,33]

Plugging these margins into the vertical arithmetic formula implies that any diversion to NVIDIA that is greater than 1.5% or 3% -- $DR_{UD} > 1/(65 + 1)$ or $2/(65 + 2)$ -- suggests that NVIDIA has an incentive to foreclose. For example, under an assumption of diversion according to share with no input substitution, the vertical arithmetic implies foreclosure would be profitable in any market in which NVIDIA had at least a 3% share -- although (as we noted above) input substitution would push this threshold higher. The three markets described in the complaint would have to satisfy this condition in order for vertical math to imply an incentive to foreclose.

2.   Interpretation of Relative Upstream and Downstream Margins

The low estimated threshold for foreclosure results from the fact that NVIDIA would give up very little Arm profit by foreclosing rivals, relative to the added profits that could be earned downstream. However, Arm's relatively low share of downstream profit means that under many commonly used models of price-setting, Arm may appear to lack the market power to drive downstream substitution substantially (O'Brien, 2022).

---

[30] https://www.anandtech.com/show/7112/the-arm-diaries-part-1-how-arms-business-model-works/2

[31] https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2022

[32] Note that this assumes that Arm incurs no marginal costs to license its IP. Otherwise, Arm's margin would be smaller than the 1-2% that it charged in royalties. This would make foreclosure more likely.

[33] According to the FTC complaint, "Arm's licensing model is based on upfront license fees and royalties." Therefore, one could argue that some of Arm's upfront license fees could be pro-rated to the marginal unit of IP. In this case, Arm's margin would be higher, which would make foreclosure less likely.

One possible explanation for low relative upstream margins is substantial premerger upstream competition. In their response to the FTC complaint, the parties argued that there was substantial upstream competition: "Arm faces significant competitive threats from other CPU architectures, including x86 (Intel, AMD) and RISC-V (open source), and that competition will further prevent NVIDIA from engaging in the types of conduct the Complaint imagines."

However, this explanation is contradicted by the fact that Arm had no close substitutes for many applications.[34] Competitors to Arm are thus limited in their ability to replace Arm. The first architecture -- x86 -- was not licensed as of the time of the Complaint. The second -- RISC-V -- is viewed by many industry participants as inferior to Arm's technology for various applications.[35] The FTC Complaint stated, for example, that "Virtually all major DPU SmartNICs suppliers, including NVIDIA and its direct competitors, incorporate Arm processor technology", and "… industry participants have acknowledged that the automotive industry is reliant on Arm for ADAS development."[36]

Furthermore, it has also been reported that various industry participants publicly opposed the deal over concerns over access to Arm's technology, which might be unnecessary if Arm had close substitutes. Thus, one potential explanation for low relative Arm margins is not consistent with the other information about the relevant markets affected by the transaction.

Furthermore, the apparent paradox of low relative margins and upstream market power is not limited to Arm. In *FTC vs. Qualcomm Inc.,*[37] a Ninth Circuit panel upheld that Qualcomm had monopoly power in several types of wireless chips.[38] Qualcomm charged Apple a $7.50 royalty per iPhone, which Apple considered to be too high.[39] If Apple earned a 40% margin on a $1000 iPhone,[40] this royalty rate would constitute 1.8% of the downstream margin – which was

[34] FTC Complaint http://www.ftc.gov/system/files/documents/cases/d09404_part_3_complaint_public_version.pdf

[35] See FTC Complaint at Section II.

[36] See FTC Complaint at ¶65 and ¶85.

[37] https://www.ftc.gov/legal-library/browse/cases-proceedings/141-0199-qualcomm-inc

[38] https://www.clearygottlieb.com/-/media/files/alert-memos-2020/20200827-our-analysis-of-the-ninth-circuit-panel-decision-reversing-ft-pdf.pdf

[39] https://www.fool.com/investing/2019/01/14/heres-how-much-apple-was-paying-qualcomm-in-royalt.aspx

[40] https://appleinsider.com/articles/21/04/28/apples-gross-margin-is-highest-its-been-in-9-years

comparable to or slightly lower than Arm's relative margin based on the public sources that were cited above.

Below we discuss three conceptual models that could produce low premerger relative margins without necessarily undermining Arm's ability to foreclose rivals profitably.

First, consider that Arm may license its IP prior to the development of downstream products by its licensees, so that the surplus that is being divided when setting licensing fees and royalty rates is the ex ante expected profit -- which takes into account R&D costs and the probability of failure. Therefore, at the time that license fees and royalties are negotiated, the expected surplus may be low due to substantial development risks. Therefore, even if Arm is able to claim a large portion of ex ante expected profits, Arm may have a relatively small portion of ex post profits.

Post-merger, the integrated firm, when considering withholding services and support for incremental innovations, stands to lose only the margin that it negotiated prior to the development of the product.[41] However, foreclosure could occur after the development of a successful product through degradation of services and support. With a sufficiently high diversion to its downstream product, the integrated firm stands to gain its relatively high ex post downstream margin on diverted sales. Therefore, the relevant margins to compare are Arm's pre-development margins, negotiated ex ante, and the downstream margins on the successful product, determined ex post, as reflected in the vertical arithmetic above.

Second, were Arm to try to charge high royalty rates to downstream firms, some of the downstream investments might prove unprofitable to pursue. Accordingly, Arm might accept what appears to be a small portion of ex post surplus of successful investments to encourage these potentially valuable downstream investments.

Third, the parties identified bolstering Arm's "ecosystem" as a goal for the transaction in their response to the complaint: "NVIDIA's goal with this Transaction is to build a stronger Arm

---

[41] Without vertical integration, it is in Arm's interest to provide continuing support to all firms, because, among other things, the relatively low margin is likely similar for all licensees.

ecosystem that will benefit NVIDIA and all of Arm's potential customers."[42] They identify Arm's ecosystem as including "datacenter-server and PC software, middleware, database, and peripheral suppliers."[43] The existence of an "ecosystem" around Arm could imply that the adoption of Arm IP by one customer increases demand for Arm IP from other customers. Such network effects may give Arm an incentive to set lower royalties than if these effects were not present. Arm may want to encourage investment in complementary products to spur demand.

However, it is not clear to what extent these alleged network effects in the Arm ecosystem would mitigate harm from the merger. The incentive to raise rivals' costs would be a change in Arm's pricing incentives that would not occur but-for the merger.

In this section we have discussed three conceptual models for Arm's low relative margins that do not necessarily mitigate harm from a vertical merger: negotiation of royalties prior to risky downstream product investment; desire to incentivize downstream investment; and network effects with complementary products within the Arm ecosystem. The existence of complementary products outside the Arm ecosystem does not implicate Arm's relative margins, but will be discussed in more detail in the next section in the context of merger-specific investment efficiencies.

C.   Investment Incentives as a Merger Efficiency

In their response to the FTC complaint, the parties stated that the merger would increase investment by both NVIDIA and Arm: "The Transaction will result in billions of dollars of additional investment…"[44] Similar to EDM, this appears to be an incentive change as an efficiency but applied to investment rather than to price. Below, we discuss two potential mechanisms for a

---

[42] Answer and Defenses of Respondents to FTC Complaint

https://www.ftc.gov/system/files/documents/cases/d09404_-_answer_and_defenses_of_respondentsv_nvidia_corporationsoftbank_group_corp._and_arm_ltd.pdf

[43] Answer and Defenses of Respondents to FTC Complaint

https://www.ftc.gov/system/files/documents/cases/d09404_-_answer_and_defenses_of_respondentsv_nvidia_corporationsoftbank_group_corp._and_arm_ltd.pdf

[44] Note that the parties did not articulate specific types of investment that would be undertaken in their response to the complaint. https://www.ftc.gov/system/files/documents/cases/d09404_-_answer_and_defenses_of_respondentsv_nvidia_corporationsoftbank_group_corp._and_arm_ltd.pdf

merger-specific change in investment incentives. First, we discuss possible complementarity between GPUs (graphics processing units) -- NVIDIA's core product -- and datacenter CPUs (central processing units): an emerging market for Arm.[45] Second, we discuss the implications for the merger of possibly incomplete contracting over investment.

1. Complementarity of CPUs and GPUs

In their response to the FTC complaint, the parties argued that post-merger NVIDIA would have the incentive to substantially increase investment in Arm to increase demand for NVIDIA's GPUs, which are not built on Arm architecture: "…drive down CPU prices, and create more opportunities to sell NVIDIA's GPUs and other accelerators."[46] In economic terms, this could imply that Arm was underinvesting in CPUs that are complementary to NVIDIA GPUs: specifically, datacenter CPUs. This underinvestment could have occurred because Arm did not internalize the impact of this investment on NVIDIA's profit. If this investment could not be contracted for between NVIDIA and Arm, a merger may be required to better align investment incentives.

The validity of this claim depends crucially on the extent to which CPUs and GPUs are strong economic complements. However, the relationship between CPUs and GPUs in a datacenter is complicated and workload-dependent. Consequently, it is unclear under which use-cases CPUs and GPUs are weak complements, strong complements, or substitutes. In the case of weak complements, a lower price of CPUs would have a minimal, but positive impact on sales of GPUs. In the case of strong complements, the CPU price would have a large positive impact, and in the case of substitutes the impact would in fact be negative.

While it is the case that GPUs cannot be used without at least one CPU, the ratio of GPUs to CPUs in datacenter acceleration can vary substantially. For example, on Amazon Web Services,

---

[45] https://www.arm.com/blogs/blueprint/optimizing-data-center

[46] Answer and Defenses of Respondents to FTC Complaint
https://www.ftc.gov/system/files/documents/cases/d09404_-_answer_and_defenses_of_respondentsv_nvidia_corporationsoftbank_group_corp._and_arm_ltd.pdf

accelerated computing instances are available with CPU to GPU ratios between 4:1 and 32:1.[47] How the price of CPUs affects substitution across these types of instances and many other possible datacenter configurations will affect the degree of economic complementarity.

Moreover, CPUs and GPUs may also be substitutes. Datacenter customers can choose between CPU and GPU, or CPU-only, computing instances in various settings. In 2017, NVIDIA CEO Jensen Huang said that GPUs will replace CPUs for AI workloads; he was likely referring to the rise of GPU-accelerated computing.[48] Substitution between GPUs and CPUs is also implied in the parties' response to the FTC complaint: "NVIDIA's products accelerate computing systems by offloading certain tasks from the CPU, thus allowing the CPU to focus on other core tasks."[49]

A key question here is whether the magnitude of investment increase that we would expect from complementarity post-merger -- across any or all use-cases -- could outweigh the potential harm from the merger. The 2010 Horizontal Merger Guidelines (U.S. Department of Justice and the Federal Trade Commission, 2010) set forth standards for recognizing efficiency claims such as CPU/GPU complementarity, stating in part: "…it is incumbent upon the merging firms to substantiate efficiency claims so that the Agencies can verify by reasonable means the likelihood and magnitude of each asserted efficiency…" Given that CPUs and GPUs in datacenters could plausibly be either weak complements, or even substitutes, the "likelihood and magnitude" of this investment efficiency is not clear.

2. Incomplete Contracting over Investment

Integral to the parties' investment efficiencies argument was the idea that joint ownership of NVIDIA and Arm could far better align investment incentives than could imperfect contracting between independent firms. We discuss three mechanisms by which a merger theoretically could

---

[47] https://aws.amazon.com/ec2/instance-types/

[48] https://segmentnext.com/huang-says-nvidia-gpus-will-replace-cpus-future-moores-law-dead/

[49] Answer and Defenses of Respondents to FTC Complaint
https://www.ftc.gov/system/files/documents/cases/d09404_-_answer_and_defenses_of_respondentsv_nvidia_corporationsoftbank_group_corp._and_arm_ltd.pdf

increase investment -- double moral hazard; hold-up; and bilateral bargaining externalities -- we and consider each in the context of NVIDIA/Arm.

The alleged underinvestment posited by the parties could be thought of as a non-price analog to EDM. EDM describes an incentive to decrease price in mergers between complementary firms, such as those that make an input and a downstream product. Prior to the merger, neither firm internalizes the effect that its price has on the other firm's profit, and therefore sets too high a price. Likewise, when negotiating over investments, neither firm may internalize the full impact of those investments on the other firm's profit, and therefore underinvest. In other words, firms do not undertake investments that would be mutually beneficial, but that do not justify the cost individually. In the context of NVIDIA/Arm, if we assume most downstream investments occur after the technology is licensed, and to the extent this investment is non-contractible, double moral hazard is plausible (O'Brien 2020).

Another underinvestment problem that could occur is hold-up: If prices are determined after investment, firms invest less because the other firm may later extract the surplus from that investment (Klein, Crawford, and Alchian, 1978). For example, an Arm customer may become more locked-in to Arm IP by designing complementary products around the IP. Once the complementary products are designed, Arm may then choose to increase prices to capture some of the value that has been created downstream. Anticipating this, Arm customers such as Nvidia may be less likely to undertake investments that are specifically complementary to Arm IP.

Israel and O'Brien (2021) describe merger investment effects in the specific context of bilateral bargaining, which they claim also results in underinvestment: "When bilateral contracting externalities are present, underinvestment would still occur even if each investing firm captured all available incremental profit from the investment." Intuitively, when an upstream firm negotiates with any individual downstream firm, it cannot credibly commit to not sell to every downstream firm, which would decrease downstream profit. Therefore, under bilateral bargaining, the upstream firm cannot raise enough revenue to invest optimally, and no single downstream firm is willing to fund investment that will benefit its rivals. If we assume that Arm negotiates with individual licensees, and licenses its IP widely, bilateral bargaining externalities would appear plausible in theory.

All three mechanisms for investment efficiencies -- double moral hazard; hold-up; and bilateral bargaining externalities -- are theoretically plausible in a vertical merger context, depending on the specific facts of the industry.[50] However, the 2010 Horizontal Merger Guidelines (U.S. Department of Justice and the Federal Trade Commission, 2010) state: "Efficiency claims will not be considered if they are vague or speculative or otherwise cannot be verified by reasonable means." The response to the complaint does not name specific upstream or downstream investments, nor does it in any way quantify the magnitude of the investments that might be made with the merger. The factual predicates for applying any of the three discussed mechanisms for investment efficiencies in the context of NVIDIA/Arm are also unclear. Therefore, based on the public record, the claimed investment incentive changes appear unverifiable.

D. Discussion

The use of vertical arithmetic and the weighing of incentives to foreclose with efficiencies are a common practice in the analysis of vertical mergers. The NVIDIA/Arm matter featured interesting complications in both areas.

With regard to the use of vertical arithmetic, we were confronted with trying to understand how seemingly low premerger relative upstream margins could be reconciled with strong evidence that Arm had few good substitutes in many possible applications. The less that Arm earned in profit relative to downstream firms, the greater was the incentive to foreclose, as was shown by vertical arithmetic. Simultaneously, under many models, relatively low margins could signal low upstream market power, which would undermine Arm's ability to foreclose. The analysis presented above suggested several possibilities for how this tension could be resolved.

We also evaluated several investment incentive claims that could have been merger-specific, including: the economic complementarity of CPUs and GPUs; double moral hazard; hold-

---

[50] The applicability of Israel and O'Brien's (2021) model to the NVIDIA/Arm merger depends, in part, on how well the details of that industry match the assumptions of their model.

up; and bilateral bargaining externalities. While theoretically plausible, these theories were not supported by specific evidence.[51]

V.   Complementarity in Bilateral Bargaining

A recent literature explores the possibility that some combinations of health care providers may be complements in bargaining with insurers. (See, for example, Easterbrook, et al. (2019). This topic is also addressed at least briefly in Vistnes and Sarafidis (2013), Peters (2014), Dafny, Ho, and Lee (2019), and Brand and Rosenbaum (2019).) If two providers are complements from the perspective of a given insurer, then it is possible, as a matter of theory, that a merger between the two providers would create *downward* pricing pressure on the prices that are negotiated with that insurer even absent any cost efficiencies:[52] A merger between complementary providers may decrease their combined relative bargaining leverage in negotiations with insurers.

A.   Theory and Background

In this section, we describe the basic intuition behind complementarity in a bargaining setting and how mergers between complementary providers could generate price decreases, even absent efficiencies.[53] We also discuss the similarities of complementarity in a bilateral bargaining setting with complementarity in the more familiar posted-price setting. To distinguish between

---

[51] Answer and Defenses of Respondents to FTC Complaint

https://www.ftc.gov/system/files/documents/cases/d09404_-_answer_and_defenses_of_respondentsv_nvidia_corporationsoftbank_group_corp._and_arm_ltd.pdf

[52] Substitution from the perspective of insurers and substitution from the perspectives of patients are related, but distinct, concepts in bargaining theory. As will be discussed below, the more that patients view two hospitals as substitutes at the point of service, the more likely it is that the hospitals will be substitutes from the perspective of insurers in bilateral bargaining. The distinction between substitution for insurers in bilateral bargaining and substitution for patients at the point of service may be important. For example, it could be the case two hospitals are complements from the perspective of an insurer in bilateral bargaining and yet the hospitals compete directly for patients at the point of service.

[53] Throughout this discussion, we focus on the prices that are negotiated between providers and insurers, as opposed to the prices that insurers charge to consumers and employers.

complementarity in bilateral bargaining and complementarity in posted-price games, we use the term "bargaining complements" here.

Finally, we discuss some principles that we view as important in evaluating arguments that the parties to a given proposed merger are bargaining complements. Among these, we describe important components of bargaining theory that indicate that mergers between complementary providers may nonetheless *increase* their combined relative bargaining leverage in negotiating with insurers. Throughout this discussion, we focus on the principles that are derived from Nash-in-Nash bargaining theory.

We begin by describing bargaining complementarity and how a merger between complementary providers may lead to decreases in the prices negotiated with insurers. As is discussed in Easterbrook et al. (2019), a key consideration in this analysis is the fact that insurers often assemble health care provider networks to market with their risk-bearing and claims processing functions to employers and consumers. Any two providers -- A and B -- within a given insurer's network may be bargaining substitutes or complements from the perspective of the insurer depending on how A and B interact in terms of their value-added to the profitability of the insurer's network.

For example, the value-added for the insurer of provider A may depend on whether provider B is already in the network. If the presence of B reduces the value-added of A, and vice versa, then A and B are bargaining substitutes from the perspective of the insurer. Conversely, if the presence of B increases the value-added of A, and vice versa, then A and B are bargaining complements from the perspective of the insurer. Hence, the insurer's profits are convex in A and B.[54]

An alternative way to consider bargaining complementarity is that the insurer's incremental profit when including both providers A and B is greater than: (i) the insurer's incremental profit when including A but not B; plus (ii) the insurer's incremental profit when including B but not A. (This is a sufficient condition for the convexity definition above.[55])

---

[54] The literature often refers to the 'concavity', in the case of substitutes, or 'convexity', in the case of complements, of the insurer's profits in the providers A and B to describe these conditions.

[55] This definition of complementarity and the convexity definition in the previous paragraph are equivalent if the insurer's profit if it excluded both A and B is zero.

We illustrate this with a simple example that is similar to the example in Easterbrook et al. (2019): Suppose that the insurer would earn 100 if it had both A and B in network. Further suppose that the insurer would earn 30 if it had A but not B in network, and 20 if it had B but not A. Finally, suppose that the insurer would earn 0 if it had neither A nor B in network.[56] Throughout this example, we assume that the set of providers in the insurer's network other than A and B is fixed. Under these assumptions, excluding either A or B, but including the other, would be very costly for the insurer, and excluding both A and B would be only modestly more costly for the insurer. Given the definition of bargaining complementarity above, A and B are complements from the perspective of the insurer since $100 > 30 + 20$.[57,58]

In bargaining theory, the *bargaining leverage* of a provider can be measured by the reduction in profit for the insurer if the insurer did not include the provider in its network. Under the assumptions in the simple example above, the insurer would lose:

i) $100 - 20 = 80$ if the insurer excluded A but included B;

ii) $100 - 30 = 70$ if the insurer excluded B but included A; and

iii) $100 - 0 = 100$ if the insurer excluded A and B.

Hence, the bargaining leverage of A is 80, the bargaining leverage of B is 70, and the bargaining leverage of the merged AB is 100.

---

[56] Note that neither A nor B is a 'must-have' provider in the literal sense since the insurer would earn a positive profit if it excluded either A or B while including the other. However, the merged provider AB would be a 'must-have' in that the insurer would earn zero profit if it excluded both A and B.

[57] To connect this example with the notion of convexity, note that the value-added of A is only 30 if B is not in network but 80 if B is in network. Similarly, the value-added of B is only 20 if A is not in network but 70 if A is in network. Hence, the presence of A or B in the network increases the value-added of the other.

[58] As a simple example of bargaining substitutes, suppose that the insurer would earn 100 if it had both A and B in network. Further suppose that the insurer would earn 90 if it had A but not B in network, and 80 if it had B but not A in network. In this situation, A and B are substitutes from the perspective of the insurer since $100 < 90 + 80$. To connect this with the notion of concavity, note that the value-added of A is 90 if B is not in network but only 20 if B is in network. Similarly, the value-added of B is 80 if A is not in network, but only 10 if A is in network. Hence, the presence of A or B in the network decreases the value-added of the other.

In this simple bargaining framework, the negotiated payment to the provider is assumed to be proportional to the reduction in profit that is incurred by the insurer if it did not have the provider in the insurer's network:[59] The more important is a provider to an insurer's profitability, the more that the provider can extract from the insurer in their negotiations. We denote this proportion of what the insurer would lose that is paid to the provider as α.[60] Under this assumption: The negotiated payment to A is $\alpha(100 - 20) = \alpha80$; the negotiated payment to B is $\alpha(100 - 30) = \alpha70$; and the negotiated payment to the merged provider AB is $\alpha(100 - 0) = \alpha100$. The result is that the merged provider AB could negotiate less than the total amount that A and B could negotiate as separate entities, $\alpha100 < \alpha70 + \alpha80 = \alpha150$.[61,62] Hence, a merger between A and B would reduce the amount that the two providers could negotiate from the insurer.

The intuition underlying this result is that the insurer would lose a high percentage of its profit if it excluded either A or B, and so the insurer must negotiate with two entities (A and B) that each could severely damage the insurer's business if an agreement is not reached. Under the merger, the insurer must negotiate with just one such entity and, for this reason, may be in a better bargaining position: The insurer may prefer to negotiate with one entity that could reduce its profits by 100% than to negotiate with two entities: one of which could reduce its profit by 80%, and

[59] Similar assumptions on how the joint surplus from an agreement is split between the parties are usually made in more complex Nash bargaining models. See, for example, Peters (2014), Gowrisankaran, Nevo, and Town (2015), and Balan and Brand (*forthcoming*).

[60] Commonly, the value of α is assumed to be one-half. Under this assumption, the provider and the insurer evenly (50-50) split the joint benefit of reaching an agreement.

[61] Note that the merged provider AB could negotiate more than either A or B could on its own. That is, $\alpha100 > \alpha70$ and $\alpha100 > \alpha80$. Hence, AB has more bargaining leverage than either A or B has separately. But that is not the relevant comparison. Rather, the relevant comparison is whether AB has more or less bargaining leverage than the bargaining leverage of A *plus* the bargaining leverage of B.

[62] The assumption that the proportion α is not altered by the merger is crucial to this result. Lewis and Pflum (2015, 2017) explore theoretical and empirical evidence that mergers increase the proportion α that determines the negotiated payment to the merging providers. Such an increase in the proportion α would make it less likely that a merger of bargaining complements would reduce prices.

another that could reduce its profit by 70%. Insurers' views on whether this is a good description of the real world can be highly informative in evaluating a bargaining complements defense.

This result has some similarities to the price effects of mergers of complements outside the bargaining setting: e.g., posted-price markets. If products X and Y are complements -- the cross-price elasticity between them is negative -- then a decrease in the price of X increases the profits of Y because consumers will buy more of Y as the sales of X increase (because consumers use X and Y together). Pre-merger, X would not have the profit incentive to consider the effect of an increase or decrease in its price on the profits of Y, but X would have the incentive do so if X and Y merged. Hence, a price decrease that was unprofitable prior to a merger may be profitable under a merger.

The analogous principle here is that an agreement between A and an insurer significantly increases the bargaining leverage of B in negotiating with that insurer, and vice versa. A and B would not have the incentive to take that into account prior to the merger but would have that incentive under the merger.

One key distinction between the posted-price setting and the health-care-bargaining setting is that while the price decrease that follows the merger of complements in a posted-price market is profitable for the merging firms, this is unlikely to be the case in the bargaining-complements setting absent cost savings. The reason is that in the posted-price setting, the price decrease is profitable because the decrease in the firm's margin is more than offset by an increase in volume.

This is unlikely to occur in a health-care-bargaining setting. In health care markets, in which consumers generally do not face the full marginal price of care, a decrease in the negotiated price between a provider and an insurer will increase the volume of patients treated by the provider mainly through a decrease in the insurer's premium in the downstream insurance market. Absent narrow networks or steering provisions that shift a sufficiently large number of patients, the increase in patient volume that is treated by merging providers would likely be small because it will only increase on net if the price decrease brings in previously uninsured consumers (or consumers currently on public insurance) into the commercial market.

Moreover, such a volume increase that would be caused by such a premium decrease would likely be spread over all providers in the insurer's network -- as opposed to being captured by the merging providers: The providers that reduced their prices would not fully capture the volume

increase that is caused by the price decrease.[63] For this reason, it is highly unlikely that a decrease in price that is caused by a merger of bargaining complements would be profitable -- absent significant cost efficiencies.

Under what circumstances does Nash-in-Nash bargaining theory suggest that two providers are more likely to be bargaining complements? Generally, bargaining complementarity is more likely to occur if the two providers overlap little in terms of services provided and if the differentiated services that they provide create substantial value for the insurer's provider network. For example, if two providers that operate in the same geographic area overlap very little in most service lines and offer distinct services that employers and consumers value highly, then they may be bargaining complements. A similar argument could be made for geographic differentiation: If two providers operate in distinct geographic areas, but employers strongly prefer a provider network that covers both areas such that a network without one of the two hospitals would have little value, then they may be bargaining complements.

Easterbrook et al. (2019) note that most research and the empirical tools that have been applied in policy settings cannot account for bargaining complementarity. For example, two of of the empirical methods that have been developed to predict the price effects of mergers between health care providers that are considered in Garmon (2017) and Balan and Brand (forthcoming) -- Upward Pricing Pressure (*UPP*), and Willingness-to-Pay (*WTP*) -- assume that the merging health care providers are at least weak substitutes: *UPP* and *WTP* cannot yield a prediction that the merging providers are bargaining complements.

Despite this limitation, our view is that *UPP* and *WTP* are nonetheless useful tools in situations in which it is uncertain if the merging providers are bargaining substitutes or bargaining complements. Garmon (2017) and Balan and Brand (forthcoming) show that *UPP* and *WTP* are useful approximations to the underlying bargaining theory in merger analysis and provide reasonably accurate predictions of the effects of mergers on bargaining leverage. These tools apply health care event-level data to assess the extent to which consumers view the merging providers as substitutes at the point of service. While substitution from the perspective of consumers is distinct

---

[63] In the simulations in Balan and Brand (forthcoming), mergers that lead to price decreases always reduce the combined profits of the merging hospital systems.

from substitution from the perspective of insurers, bargaining theory indicates that they are closely related:[64, 65] The more that consumers view two providers as substitutes, the more likely it is that the providers are bargaining substitutes, as opposed to bargaining complements. Hence, large changes in *WTP* or large diversion ratios (an input in a *UPP* analysis) due to a merger -- while not ruling out bargaining complementarity -- make such a result less likely.

In addition, the results of more sophisticated analyses that *can* yield a prediction that the merging providers are bargaining complements are often highly sensitive to the values of calibrated parameters or modeling assumptions. Moreover, such results may be highly inconsistent with the qualitative evidence. In such situations, the results of the simpler metrics *UPP* and *WTP* may be more reliable -- particularly if the qualitative evidence in a given merger investigation is not consistent with the predicted outcome under bargaining complementarity that the merger will create downward pricing pressure.

B.   Comments Based on Experience in Recent Cases

In this section, we discuss issues that we have found to be important for assessing the likely effects of a hospital merger in situations where the merging hospitals could be bargaining complements.

1.   Mergers of Bargaining Complements Can Cause Upward Pricing Pressure

A key difference between posted-price settings and bargaining models is that mergers between bargaining complements do not necessarily result in lower prices, while in posted-price settings mergers between complements do result in lower prices. The reason is that convexity in the

---

[64] See Vistnes and Sarafidis (2014), Dafny, Ho, and Lee (2019), and Balan and Brand (forthcoming) for discussions of the relevant economic theory.

[65] In discussing the difference between substitution from the perspective of patients and substitution from the perspective of insurers in *ProMedica*, the Sixth Circuit noted: "ProMedica also argues that MCOs, rather than patients, are the relevant consumers here, and that the Commission therefore erred by 'assess[ing] substitutability from the patients' perspective.' But this is an argument about semantics. MCOs assemble networks based primarily upon patients' preferences, not their own; and thus the extent to which an MCO regards ProMedica and St. Luke's as close substitutes depends upon the extent to which the MCO's members do."

merging hospitals in the insurer's profit function does not necessarily imply convexity in the Nash bargaining objective function (Peters 2014).

As a result, mergers of bargaining complements can still cause upward pricing pressure. We provide some intuition for why prices may increase following a merger of bargaining complements and show how to quantify the net effect on price.

To understand why mergers of bargaining complements can still cause upward pricing pressure, it is useful to focus on how hospital prices are set in the Nash bargaining model: Ho and Lee (2017) decompose the first-order condition for the equilibrium total payments that a hospital receives from an insurer into: (i) the effect of having the hospital in-network on the insurer's revenue net of non-hospital costs (premiums and enrollment effect); (ii) the amount that the insurer pays to other hospitals (price reinforcement effect); (iii) the cost to the hospital of treating the insurers' enrollees (hospital cost effect); and (iv) the effect of going out of network on the hospital's profits from other insurers' (recapture effect).

The first two effects are part of the insurer's disagreement profit, and the latter two are part of the hospital's disagreement profit. Approaches to predicting the extent to which mergers increase prices under joint all-or-nothing bargaining have focused on the impact of the merger on the premiums and enrollment (the first effect) and have used proxies such as *WTP*.[66] Similarly, arguments that the hospitals are bargaining complements and, therefore, a merger between them will cause their prices to fall also focus on the premium and enrollment effect (Easterbrook et al. 2019).

However, the first-order effect of the merger on prices also depends on how the merger affects price reinforcement, hospital costs, and recapture, not just whether the hospitals are bargaining complements or substitutes. Thus, even mergers of bargaining complements can cause prices to increase if these other forces offset the downward pressure on prices from a merger of bargaining complements.

The recapture effect is an important channel through which a merger could increase prices on net despite the downward pressure that a merger of true bargaining complements puts on prices.

---

[66] See, e.g., Capps et al. (2003), Farrell et al. (2011), and Garmon (2017). Some analyses also include the price reinforcement effect. See, for example, Gowrisankaran, Nevo, and Town (2015) and Balan and Brand (forthcoming).

A merger that increases the recapture rate could cause the Nash bargaining objective function to be concave, via its effect on the hospital's disagreement profits -- even if the insurer's profit function is convex in the merging hospitals. For example, Peters (2014) decomposes the effect of a merger on prices into the substitutes/complements effect, which could be either positive (in the case of substitutes) or negative (in the case of complements) and the recapture effect, which must be weakly positive. The net effect of these two forces determines whether bargaining theory predicts upward or downward pricing pressure as a result of a merger.

One way to conceptualize recapture is the share of patients who continue to visit a hospital after the hospital is excluded from their insurer's network either by switching insurers or visiting the hospital despite its being out of network. (The latter choice is outside Peters' model, but it can happen in practice). Post-merger increases in the recapture rate increase the hospitals' bargaining leverage because their profits will not fall by as much if the patients leave an insurer's network: Increases in recapture increase the hospitals' disagreement payoff, which can increase the negotiated price.

It is likely that within-market hospital mergers increase the recapture rates in that more consumers will likely switch to a competing insurer if the focal insurer excludes both merging hospitals, as opposed to just one or the other. This is particularly true if the merging hospitals are the two most preferred options for a given patient. Patients who pre-merger would have gone to their second-choice hospital if their first choice was out of network may switch insurers post-merger rather than lose access to their two most preferred hospitals. This increase in the recapture rate may be large if the parties have a high degree of bargaining leverage pre-merger: e.g., the hospitals are close substitutes or there are few other options for care.

Thus, mergers for which commonly used metrics -- such as diversion ratios or changes in WTP -- suggest are harmful may also result in large increases in the recapture rate.[67] Increases in the recapture rate are especially plausible in markets with few options pre-merger. For example, recapture rates are likely to increase in markets with two or three hospitals that are close substitutes

---

[67] In extreme cases, the recapture rates may be close to one pre-merger, so any merger-induced increase would be small.

for a large fraction of patients and for which the remaining alternatives (if any) are far less attractive than the patients' second-choice hospital.

In addition, increases in recapture rates that are due to a merger of providers may be significant even if the merging providers are not particularly close substitutes for consumers at the point of service. Some consumers who value both providers may not switch insurers if the insurer is missing one or the other in its network but would switch if the insurer is missing both providers. This point is explored in Peters (2014).[68]

The division of the joint surplus from an agreement between the hospital and the insurer also matters for the ultimate effect of a merger of true bargaining complements on prices (Peters 2014). As the insurer's share of the joint surplus goes to one, the weight on the premium and enrollment effect goes to zero, which results in a price that allows the insurer to extract the entirety of the joint surplus. This means that as an insurer's bargaining ability (as measured by its share of the joint surplus) increases, the weight on the downward pressure that the complementarity effect has on prices will decrease, and the weight on the upward pricing pressure from increased recapture will rise. All else equal, increases in the insurer's share of the joint surplus make it more likely that a merger of bargaining complements will result in higher prices.

This fact is also relevant to instances where the merging hospitals are substitutes and there is an insurer that is arguably a powerful buyer whose share of the joint surplus approaches one.[69] Increases in the recapture rate will also cause increases in post-merger prices in the presence of powerful buyers even if the premium and enrollment effect is small.

---

[68] Peters (2014) considers the example of a merger between a hospital and physician group. Joint contracting post-merger could increase the bargaining leverage of the physician groups and the hospital even though they are not substitutes at the point of service. This is because more consumers would switch to another insurer if both the hospital and the physician group are excluded from the insurer's provider network than if either the hospital or physician group (but not both) is excluded.

[69] For an example of a merger in which the defendants made a powerful buyer argument in a bargaining setting see, Federal Trade Commission and State of North Dakota v. Sanford Health, Sanford Bismarck and Mid Dakota Clinic, P.C., IN THE UNITED STATES DISTRICT COURT FOR THE DISTRICT OF NORTH DAKOTA, MEMORANDUM OF DECISION, FINDINGS OF FACT, CONCLUSIONS OF LAW, AND ORDER, 12/15/2017 (https://www.ftc.gov/system/files/documents/cases/1710019_sanfordpiorder.pdf).

To assess quantitatively whether a merger of bargaining complements will nonetheless result in increased prices, one could use an index that is similar to the gross upward pricing pressure index (*GUPPI*) that has been applied in recent healthcare merger matters, but instead measures the upward pricing pressure that is generated by the increase in the recapture rate. We refer to this index as the Recapture Upward Pricing Pressure Index (*RUPPI*).

For hospital *A* that bargains with insurer *M*, *RUPPI* is determined by: (i) the increase in the hospital's recapture rate because of increased consumer switching to other insurers; and (ii) the hospital's margin with those insurers. These two terms give the change in the disagreement profit of provider *A* that is due to the merger, which we express as a percent of its price negotiated with insurer *M*:

$$RUPPI_{AM} = \sum_{n \neq M} \frac{p_{An} - c_{An}}{p_{AM}} \times Increase\ in\ Recapture\ Rate_{An}.$$

Similar to the interpretation of *GUPPI*, *RUPPI* measures the increase in the providers' disagreement payoffs, and hence their relative bargaining leverage, because of the increase in recapture that is due to the merger. The intuition behind *RUPPI* is incorporated in a relatively complete structural empirical model such as Ho and Lee (2017). We offer it here as a standalone metric for circumstances in which the analysis at hand accounts for only bargaining complementarity or bargaining substitutability of the merging providers. Using *RUPPI*, one can assess the net change in pricing pressure that accounts for both changes in recapture and bargaining complementarity or substitutability.

While estimating the increase in the recapture rate may be difficult without a structural model of competition in the insurance market, this analysis illustrates that quantitative or qualitative evidence that the merging parties are complements from the perspective of insurers does not fully address the question of whether the merger will cause upward or downward pressure on prices. Models that are designed to assess the effects of a proposed merger that allow for hospitals to be bargaining complements should not focus only on the question of complementarity or substitutability in the insurer's profit function. They should include additional factors that might affect price -- particularly changes in the recapture rate.

2.   Pre-Merger Bargaining Leverage and Bargaining Complementarity

Arguments that the merging hospitals are bargaining complements are often based on the claim that the hospitals have substantial bargaining leverage pre-merger. For example, the parties may argue that they are 'must-have' providers for insurers even absent the merger. Since insurers need both providers to be in-network in order to earn any profit, the parties will argue that they must be bargaining complements.

While this intuition could be consistent with the description of bargaining complements that was provided above, it is important to note that substantial bargaining leverage pre-merger is usually a basis for concern that a merger will be harmful. This intuition is reflected in bargaining theory in that higher relative bargaining leverage for the provider typically implies higher prices and therefore higher gross margins for the provider; and, other things equal, the price effects of mergers are generally increasing in the pre-merger gross margins of providers. (See, for example, Capps, Dranove, and Sattherthwaite (2003), Farrell and Shapiro (2010), and Balan and Brand (forthcoming)). Additionally, if the merging providers are strong bargaining complements, then it is possible that Nash-in-Nash with joint bargaining is not the right approach to predicting the effect of the merger on prices. We address these points in turn.

First, bargaining complements arguments often suggest that the merging firms are extremely important to insurers' networks, and thus the merging hospitals should have high pre-merger prices. For example, suppose the merging providers offer some distinct services that consumers value highly. An insurer may lose most of its enrollees if either of the merging parties was out of network since consumers may place little value on a network that does not include the services offered by each provider. Then, as in the simple example above, it would be plausible that the sum of the values of adding either provider to the network that already contains the other would be greater than the value of adding them simultaneously, and thus they would be bargaining complements. In this situation, both providers would have a tremendous amount of bargaining leverage pre-merger and would be able to negotiate high prices.

However, the merging providers may also offer many of the same services; and for these services, consumers view the merging providers as alternatives at the point of service. The merging providers may be the dominant providers of these overlapping services as well. In this situation,

while the providers may both be extremely important to insurers, it is not clear whether they are substitutes or complements from the perspective of the insurer.

As was discussed above, bargaining theory suggests that a high degree of bargaining leverage pre-merger, and the accompanying high prices, is usually a basis for concern that a merger will cause upward pricing pressure. This principle is illustrated in some of the simple metrics that have been applied in merger analysis. High bargaining leverage tends to result in high pre-merger incremental margins, which will increase the merging providers' incentive to raise price post-merger as long as patients view them as substitutes to some degree. This intuition is captured in the *GUPPI* that is analyzed in Garmon (2017) and Balan and Brand (forthcoming).[70]

Hence, evidence that the merging providers are each highly important to insurers even absent the merger does not necessarily suggest that the merging providers are bargaining complements. Such evidence on the importance of the providers to insurers likely suggests that the merger would result in upward pricing pressure (absent efficiencies) unless additional evidence supports analysis with Nash bargaining with bargaining complementarity.

Relatedly, the *RUPPI* that was discussed above is also increasing in the hospitals' pre-merger margins. For example, for a given increase in the recapture rate, the effect of increased recapture on price will be larger when the firms have the high margins that often accompany substantial bargaining leverage. Thus, as was discussed above, while arguments that the merging hospitals are bargaining complements might suggest a negative premium and enrollment effect, they also suggest higher price effects via other channels, all else equal.

Second, if hospitals are sufficiently important to insurers networks that they are actually bargaining complements, the Nash-in-Nash bargaining model might not accurately predict the price effects of the merger. Collard-Wexler, Gowrisankaran, and Lee (2019) point out that when two firms are strong complements, there is no Nash equilibrium. This is not purely a technical concern: If the hospitals are strong complements, the insurer would not be willing to contract with the

---

[70] As defined in Garmon (2017), Balan and Brand (forthcoming), and elsewhere, $GUPPI_A = Diversion_{AB} \times Margin_B \times \frac{Price_B}{Price_A}$. The pre-merger bargaining leverage of B is measured in $Margin_B$.

hospital at the prices that are implied by Nash-in-Nash bargaining because the insurer would earn negative profits prior to the merger.

As a result, either the hospitals would be out of the insurer's network at the outset, or their negotiations are unlikely to be described well by Nash-in-Nash bargaining. Therefore, empirical results that are based on Nash-in-Nash in this situation could be unreliable -- particularly if the results are highly inconsistent with the qualitative information about the nature of bargaining in this market and the views of market participants about the likely effects of the merger.

In sum, evidence that the merging providers are important to insurers pre-merger -- the providers have a high degree of bargaining leverage pre-merger, but a Nash equilibrium exists -- is consistent with both downward and upward pricing pressure from the price and enrollment effect. Moreover, such evidence also suggests that the merger likely has greater upward pricing pressure via other channels such as recapture. In extreme cases, such evidence may call into question whether Nash-in-Nash is the correct theoretical model. These considerations highlight the importance of qualitative evidence in cases in which bargaining complementarity is a possibility.

3. Applying Structural Models That Allow for Bargaining Complementarity

Calibrated or estimated structural models in the spirt of Ho and Lee (2017) could be used to simulate mergers in a framework that allows for bargaining complementarities. In our experience, we have found that, depending on the specifics of the model, a key parameter in this type of merger simulation is the price elasticity of demand in the insurance market. In these cases, hospitals were more likely to appear to be bargaining complements when insurance demand was assumed to be extremely elastic. However, the empirical evidence suggests that insurance demand is not unusually elastic, and the sorts of models that generate unusually elastic insurance demand are inconsistent with the nature of competition in insurance markets.

Because health insurance margins tend to be relatively low, one could conclude that insurance demand is quite elastic on the basis of publicly available data on insurer margins and the Lerner Rule. The Lerner Rule indicates that in a posted price setting, the elasticity of demand is just the inverse of the insurer's margin. For example, an insurer whose margins are regulated under the

Affordable Care Act (ACA) may have a margin of about 15%.[71] The Lerner Rule implies that this insurer's demand is quite elastic: -6.67.

However, the literature suggests that the demand for insurance is less elastic than would be implied by the Lerner Rule: Ho and Lee (2017) estimate elasticities that range from -1.23 to -2.95. Other studies have found elasticities that range from about -1 to -3.5 (Abraham et al., 2017; Cutler and Reber 1998; Ho 2006; Royalty and Solomon 1999).[72] This suggests that the Lerner Rule is unlikely to be the right approach to calibrating the demand for insurance.

There are at least two reasons why the Lerner Rule might not work well in this setting:

Employer premiums are determined through a process that more closely resembles bargaining or an auction, instead of a process that involves posted prices. Under bargaining and auctions, margins can be low without demand being extremely elastic. For example, Ho and Lee (2017) find that demand is much less elastic than would be implied by the Lerner Rule and that backing out margins on the assumption that competition is Nash-Bertrand yields margins that are far too high. They are better able to match publicly available insurer margins by adding a parameter that approximates insurers and employer bargaining over premiums.

Auctions have also been used to model insurer-employer premium setting: e.g., in the Anthem-Cigna merger litigation (Miller and Sheu 2021). Auction models may also rationalize the moderate insurance demand elasticity estimates in the literature and the low observed insurer margins: For a given logit demand function, the margins in the types of auctions that used to model premium setting are smaller than the margins that would arise under Nash-Bertrand competition (Panhans and Taragin 2022). In addition, to the extent that the ACA insurer margin regulations bind, it may be difficult to use margins to infer the elasticity of demand.

Structural models such as Ho and Lee (2017) could be used to study complementarities in hospital mergers. Parameters such as the insurance demand elasticity could have a substantial effect on the predicted effects of the merger. We suggest using parameter values that are consistent with

---

[71] Fully insured health insurance plans are subject to margin regulation that requires they pay out at least 80-85% of their premiums as claims for clinical services and/or quality improvement. See, https://www.cms.gov/CCIIO/Programs-and-Initiatives/Health-Insurance-Market-Reforms/Medical-Loss-Ratio.

[72] Saltzman (2019) found higher elasticities on the Affordable Care Act (ACA) exchanges in two states.

the literature and models that are consistent with the nature of competition in this market if one uses a structural model to analyze hospital mergers.

VI.   Conclusion

This article provides an indication of the heterogeneity in the topics that FTC economists address and also in the modes of economic analysis that they utilize in order to provide decision-makers with the best possible information. Some assignments -- such as the two consumer protection cases that were described here -- require economists to develop empirical strategies to estimate specific numbers of interest. As was demonstrated by both the semiconductor merger and the discussion of hospital mergers, BE's economic analysis can be focused on the applicability of theories in specific real-world settings.

VII.   References

Abraham, J. M., Drake, C., McCullough, J. S., & Simon, K. (2017). What drives insurer participation and premiums in the Federally-Facilitated Marketplace? International Journal of Health Economics and Management, 17(4), 395-412.

Balan, D., & Brand, K. (forthcoming). Simulating Hospital Merger Simulations. Journal of Industrial Economics.

Brand, K., & Rosenbaum, T. (2019). A review of the economic literature on cross-market health care mergers. Antitrust Law Journal, 82(2), 533-550.

Capps, C., Dranove, D., and Satterthwaite, M. (2003). Competition and Market Power in Option Demand Markets. RAND Journal of Economics, 34(4):737-763.

Collard-Wexler, A., Gowrisankaran, G., and Lee, R. S. (2019). Nash-in-Nash Bargaining: A Microfoundation for Applied Work. Journal of Political Economy, 127(1):163-195.

Cutler, D. M., & Reber, S. J. (1998). Paying for health insurance: the trade-off between competition and adverse selection. The Quarterly Journal of Economics, 113(2), 433-466.

Dafny, L., Ho, K., and Lee, R. S. (2019). The Price Effects of Cross-Market Hospital Mergers: theory and evidence from the hospital industry. RAND Journal of Economics, 50(2):286-325.

Das Varma, G., & De Stefano, M. (2020). Equilibrium analysis of vertical mergers. The Antitrust Bulletin, 65(3), 445-458.

Easterbrook, K. F., Gowrisankaran, G., Aguilar, D. O., & Wu, Y. (2019). Accounting for complementarities in hospital mergers: Is a substitute needed for current approaches. Antitrust Law Journal, 82(2), 497-532.

Farrell, J., Balan, D. J., Brand, K., & Wendling, B. W. (2011). Economics at the FTC: Hospital mergers, authorized generic drugs, and consumer credit markets. Review of Industrial Organization, 39(4), 271-296.

Farrell, J. and Shapiro, C. (2010). Antitrust Evaluation of Horizontal Mergers:
An Economic Alternative to Market Definition. The B. E. Journal of Theoretical Economics, 10(1).

Garmon, C. (2017). The Accuracy of Hospital Merger Screening Methods. RAND Journal of Economics, 48(4):1068-1102.

Gowrisankaran, G., Nevo, A., and Town, R. J. (2015). Mergers When Prices are Negotiated: Evidence from the Hospital Industry. American Economic Review, 105(1):172-203.

Hausman, Jerry (2001). Mismeasured Variables in Econometric Analysis: Problems from the Right and Problems from the Left. Journal of Economic Perspectives, 15(4), 57-67.

He, Sherry, Brett Hollenbeck, and Davide Proserpio (2022). The market for fake reviews. Marketing Science.

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

Ho, K., & Lee, R. S. (2017). Insurer competition in health care markets. Econometrica, 85(2), 379-417.

Israel, M. A., & O'Brien, D. P. (2021). Vertical Mergers with Bilateral Contracting and Upstream and Downstream Investment. Available at https://ssrn.com/abstract=3886048.

Klein, B., Crawford, R. G., & Alchian, A. A. (1978). Vertical integration, appropriable rents, and the competitive contracting process. The Journal of Law and Economics, 21(2), 297-326.

Lewis, M. S., & Pflum, K. E. (2015). Diagnosing hospital system bargaining power in managed care networks. American Economic Journal: Economic Policy, 7(1), 243-74.

Lewis, M. S., & Pflum, K. E. (2017). Hospital systems and bargaining power: Evidence from out-of-market acquisitions. The RAND Journal of Economics, 48(3), 579-610.

Luca, Michael, and Georgios Zervas (2016). Fake it till you make it: Reputation, competition, and Yelp review fraud. Management Science, 62.12: 3412-3427.

Miller, N. H., & Sheu, G. (2021). Quantitative methods for evaluating the unilateral effects of mergers. Review of Industrial Organization, 58(1), 143-177.

Moresi, S., & Salop, S. C. (2013). vGUPPI: Scoring unilateral pricing incentives in vertical mergers. Antitrust Law Journal, 79, 185.

O'Brien, D. P. (2020). The Economics of Vertical Restraints in Digital Markets. The Global Antitrust Institute Report on the Digital Economy, 9.

O'Brien, D. P. (2022). Tethering Vertical Merger Analysis. Available at https://ssrn.com/abstract=4061410.

O'Brien, D. P. (2022). Technical Note. Available at https://www.microfoundations.com/_files/ugd/f829dc_d2af395de9044e128f4b107963c02378.pdf

Panhans, M., & Taragin, C. (2022). Consequences of Model Choice in Predicting Horizontal Merger Effects. FTC Bureau of Economics Working Paper. Available at https://www.ftc.gov/system/files/ftc_gov/pdf/working_paper_348.pdf.

Peters, C. T. (2014). Bargaining Power and the Effects of Joint Negotiation: The Recapture Effect. U. S. Department of Justice, Economic Analysis Group Discussion Paper, No. 14-3r. Available at https://www.justice.gov/sites/default/files/atr/legacy/2014/09/26/308877.pdf.

Pittman, R., (2021). Three Economist's Tools for Antitrust Analysis. U. S. Department of Justice, Economic Analysis Group Discussion Paper, No. 21-2. Available at https://www.justice.gov/atr/page/file/1404436/download.

Raval, D. (2020). Whose voice do we hear in the marketplace? Evidence from consumer complaining behavior. Marketing Science, 39(1), 168-187.

Royalty, A., & Solomon, N. (1999). Health Plan Choice. Price Elasticities in a Managed Competition Setting. The Journal of Human Resources, 34(1).1-41.

Sheu, G., & Taragin, C. (2021). Simulating mergers in a vertical supply chain with bargaining. The RAND Journal of Economics, 52(3), 596-632.

U.S. Department of Justice and the Federal Trade Commission (2010). Horizontal Merger Guidelines. Available at: https://www.ftc.gov/sites/default/files/attachments/merger-review/100819hmg.pdf

Vistnes, G. and Sarafidis, Y. (2013). Cross-Market Hospital Mergers: A Holistic Approach. Antitrust Law Journal, 79(1):253-293.

Wood, D. & Stone, D. (2018). You Can Fool Some of the People All of the Time: Heterogeneity in Consumer Deception. Working Paper. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3206620.