

NBER WORKING PAPER SERIES

HOW TO GET AWAY WITH MERGER:
STEALTH CONSOLIDATION AND ITS EFFECTS ON US HEALTHCARE

Thomas G. Wollmann

Working Paper 27274
<http://www.nber.org/papers/w27274>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020, Revised July 2021

I thank Jimmy Roberts for suggesting the setting, Paul Eliason for helping navigate the data, and Fiona Scott Morton (NBER discussant), Judy Chevalier (ASSA discussant), John Asker, Dave Balan, Dennis Carlton, Thom Covert, Austan Goolsbee, Neale Mahoney, Ariel Pakes, Nancy Rose, Chad Syverson, Haris Tabakovic, and Mike Whinston, and seminar participants at Harvard University, Yale University, the Federal Trade Commission, Duke University, KU Leuven, MIT, Northwestern University, Stanford University, Washington University in St. Louis, Columbia University, University of Maryland, ITAM, Boston University, Georgetown, Chicago Booth, the NBER Winter IO Meeting, the ASSA Annual Meeting, Microsoft Research, Charles River Associates, the Hal White Antitrust Conference, the ASHEcon Annual Conference, and the Midwest IO Conference for thoughtful comments. Credit for the title belongs with Mike Sinkinson. (Note that the title was suggested and adopted before the paper contemplated mortality and is not intended to make light of that outcome.) Excellent research assistance was provided by James Kiselik, Paulo Henrique Ramos, Divya Vallabhanen, and Jason Yang. Financial support from the Becker Friedman Institute's Health Economics Initiative is gratefully acknowledged. The data reported here have been supplied by the United States Renal Data System (USRDS). The interpretation and reporting of these data are the responsibility of the author and in no way should be seen as an official policy or interpretation of the US government. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Thomas G. Wollmann. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How to Get Away with Merger: Stealth Consolidation and Its Effects on US Healthcare
Thomas G. Wollmann
NBER Working Paper No. 27274
May 2020, Revised July 2021
JEL No. D4,D43,I11,K21,L0,L1,L11,L13,L4,L40

ABSTRACT

Only large mergers must be reported to the US government, so many anticompetitive deals may effectively avoid antitrust scrutiny. This paper uses data from the dialysis industry to show that premerger notifications are essential to antitrust enforcement and that ensuing enforcement actions preserve competition and protect consumers, evidenced by lower hospitalization and mortality rates. It then studies the effect of eliminating premerger notification exemptions. To do so, I estimate a structural model of quality competition. I find the counterfactual policy saves thousands of lives, yielding benefits that far exceed reasonable estimates of the additional costs of enforcement.

Thomas G. Wollmann
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
Thomas.Wollmann@chicagobooth.edu

1 Introduction

Governments around the world require firms that are interested in merging to notify them in advance, which provides antitrust authorities time to evaluate the competitive effects of deals prior to their completion. Yet, most transactions are *exempt* from premerger notifications on the basis of their size. In the US, these exemptions depend mainly on statutes established in the 1950s to 1970s, which in turn rely on an assumption that only large deals have enough impact on competition to merit prospective merger reviews [Wollmann, 2019].¹ However, the claim has never been empirically evaluated. This is despite the fact that merger activity has grown to transfer trillions of dollars of ownership rights annually in the US alone, providing terrific scope for potential harm, and despite a potential flaw in reasoning, which is highlighted by decades of work in industrial organization economics.

The issue with the claim is that in segmented industries, where products are differentiated to serve heterogeneous consumers, even minor deals can result in major changes in market structure and firm behavior. To illustrate, consider an industry where, in every city in the country, exactly two independently run establishments compete with one another. Further, suppose that each is valued at around \$80 million and that legislators set the reporting threshold at \$90 million, meaning mergers valued below that amount need not be reported to the government. Then, assuming premerger notifications are critical to enforcement and firms attempt to acquire close competitors, entire swathes of an economically important industry could merge to monopoly—all without any notice to the nation’s antitrust authorities. I call this outcome *stealth consolidation*. However, this eventuality is not guaranteed. Enforcement officials can learn about deals many ways, and they’re typically free to investigate any of them, regardless of whether they are reportable or even completed.

This paper studies premerger notifications in the US dialysis industry. I provide evidence that premerger notifications are essential to antitrust enforcement, that the resulting enforcement actions prevent consolidation, and that doing so protects consumers from harm. (In the dialysis industry, where prices are largely set by Medicare, firms compete on quality, so harm equates to higher hospitalization and mortality rates.) These findings raise the question of whether society would benefit, on net, from eliminating premerger notification exemptions in the dialysis industry. To answer this, "reduced form" estimates will not suffice: concentration indices are one-dimensional approximations of complicated substitution and ownership patterns, to which quality responds non-linearly in equilibrium. Thus, I

¹Hearings before the Senate Subcommittee on Antitrust and Monopoly in 1975 contain many examples. For one, “[lowering] the pre-merger notification requirement to include companies with a hundred million dollars in sales or assets would cover literally thousands of transactions which have no legally significant anticompetitive impact” (statement of A. G. W. Biddle).

estimate a structural model of competition in quality. Then, I predict how market structure would evolve under the counterfactual policy, use estimates of the structural model to calculate the benefits (in terms of lives saved), and compare those to the costs of additional enforcement.

The industry I study treats patients with a loss of kidney function. This setting provides an almost ideal laboratory to examine stealth consolidation for several reasons. First, it allows me to link ownership changes, premerger notifications, enforcement actions, and welfare-relevant outcomes at an establishment-year (i.e., facility-year) level. Second, dialysis markets have consolidated rapidly over the last two decades following thousands of ownership changes [Cutler et al., 2016, Eliason et al., 2019, Eliason, 2018]. Third, since patients require several treatments per week, they attend facilities near home, so while the industry accounts for tens of billions in annual healthcare expenditures, individual locations are valued at only about \$4 million each. This means firms can acquire rival chains operating several competing facilities without reaching the transaction value threshold set by Congress (currently \$92 million). Fourth, enforcement actions are common. The Federal Trade Commission (FTC), which has de facto responsibility for policing dialysis mergers, has effectively blocked hundreds of ownership changes. Fifth, entry into established dialysis markets is slow and expensive, so profits derived from acquiring close competitors are not easily dissipated.

To describe the setting more precisely, some definitions are helpful at the outset of the paper. I define "facility acquisitions" as changes in ownership of individual facilities, and I define "mergers" as simultaneous facility acquisitions between the same target/acquiror. Mergers whose transactions values meet or exceed the threshold set forth in the Hart-Scott-Rodino Antitrust Improvements Act (HSR Act) require a premerger notification. I call these mergers "reportable," refer to the remaining mergers as "exempt," and apply these same terms to the facility acquisitions with which they are associated. Loosely speaking, when the FTC concludes that part of a merger will reduce competition in a particular market, it requires the target to sell off the offending facilities as a precondition of closing the rest of the transaction. I call these sales "divestitures," use the term interchangeably with "enforcement actions," and say these events "block" a facility acquisition. Finally, at early stages of the enforcement process, it is common practice to summarize consolidation using predicted concentration changes. I take this approach as well when presenting descriptive evidence. To be precise, I construct "simulated HHI changes" at the merger-market level, which equal the amount HHI would increase were the transaction completed without any divestitures (based on market shares just prior to the merger).

My data consist of annual observations spanning 1996 to 2017. For each facility, I observe a unique Medicare identifier in addition to the the facility's name, location, owner, and other key characteristics

including quality. (I measure quality using standardized hospitalization and survival rates, described below.) An average of around 4,500 facilities each year blanket the United States. Ownership changes reflect facility acquisitions, which map to mergers. I observe and/or can accurately infer whether each merger was reportable and whether each facility acquisition required a divestiture. There are more than 4,300 facility acquisitions over the 22 year panel. Only about half are reportable. That is, the FTC never received any notice of over 2,000 of them. Nonetheless, the agency required almost 300 divestitures. For every dialysis patient, I observe histories of diagnosis, treatment, and hospitalization, dates of birth and death, payment methods, demographic information, baseline health measures, and residence ZIP code. Based on the treatment history, patients are then assigned to their modal treatment facilities. On average, 70% are hospitalized each year and 15% die, underscoring the "at-risk" nature of persons receiving dialysis.

I begin by establishing three important reduced form facts. First, premerger notification exemptions severely reduce antitrust enforcement. Among reportable facility acquisitions, divestiture rates track closely with simulated HHI changes. For example, while divestiture rates are almost zero when the acquired facility does not compete with any acquiring facilities prior to the transaction, divestiture rates near 100% when the facility acquisition would result in a monopoly or duopoly. In extreme contrast, divestiture rates are close to zero among exempt facility acquisitions *regardless of their predicted effect on market structure*. Many of these involve simulated HHI changes over 2,000 points—more than an order of magnitude higher than what agency guidelines stipulate is "presumed to be likely to enhance market power" [US Department of Justice and Federal Trade Commission, 2010].

Second, enforcement actions facilitated by premerger notifications have long-lasting effects on market structure. Note that this outcome is not guaranteed. If the acquirors of divested facilities struggle to operate those locations so that the facilities eventually fail, then the remedies required by the FTC are ultimately ineffective. Yet, in the dialysis industry, divestitures induced by premerger notifications significantly reduce concentration, even if one looks several years into the future.

Third, consolidation that results from premerger notification exemptions hurt patients. Among exempt facility acquisitions, those involving relatively large simulated HHI changes witness a post-merger increase in hospitalization rates and decrease in survival rates, which is consistent with the internalization of business-stealing effects following the purchase of a close rival. Among reportable facility acquisitions, however, there is no relationship between simulated HHI changes and within-facility changes in healthcare quality. Intuitively, facility acquisitions that would otherwise lead to poor patient health outcomes are effectively blocked by the agency (*vis-a-vis* a divestiture). In other words,

FTC enforcement actions save lives.

To address these issues more precisely, I present a structural model that captures equilibrium responses to market structure changes. Patients choose a facility (or the outside option, home hemodialysis) to maximize utility, which depends on facility characteristics such as quality and close proximity to their homes. Facilities choose quality to maximize profits. By taking the first order condition of profits with respect to quality and then changing ownership structure, one can clearly see why consolidation reduces quality: prior to acquiring a close rival, firms have incentives to raise quality and steal patients from the competing location, but post-acquisition those incentives disappear (since the firm would be stealing those patients from itself). This thought exercise also highlights that the incentives of a firm setting quality subject to fixed prices are very closely analogous to those of a firm setting prices subject to fixed quality. I then estimate the model, recovering demand and marginal costs estimates, which enable me to simulate equilibrium quality under any hypothetical market structure.

Finally, I calculate and compare the costs and benefits of a counterfactual policy that eliminates premerger notification exemptions in the dialysis industry. The calculations involve three steps. First, I estimate enforcement rates on reportable mergers and retroactively apply them to exempt mergers, which simulates how merger activity and market structure would evolve (stochastically) under the counterfactual policy. In the process, I allow for a degree of deterrence. (Briefly, facility acquisitions that are very likely to face enforcement actions are unlikely to be attempted in the first place.) Second, I forecast benefits by simulating equilibrium under the various possible market structures, computing the expected number of life years saved, and multiplying the result by the statistical value of a life year. Third, I forecast the additional costs imposed on the FTC by the counterfactual policy. Although a comprehensive study of the agency's production function is far beyond the scope of this paper, one can at least bound the added burden using historical agency data (e.g., the number of mergers evaluated, staffing and salary levels, etc.). The main finding is that *the benefits of the counterfactual policy are nearly two orders of magnitude higher than the costs.*

This paper contributes to several areas of research. Alongside recent work, it introduces the notion of stealth consolidation. Wollmann [2019] uses economy-wide merger and enforcement data to study an abrupt increase in US thresholds in 2001, which lead to a sharp decline in notifications and enforcement actions as well as a rise in mergers among rivals. Its broad scope, though, precludes realistic measures of competition: firms are assumed to compete if they occupy the same 4-digit SIC industry group, which is a much wider definition than courts would use in practice. It also precludes studying the effects of those mergers. The present work varies substantially in both regards. Together they suggest

premerger notification program changes contributed, at least in part, to secular trends in economy-wide consolidation [Furman and Orszag, 2015]. Motivated partly as an explanation of the fall in the otherwise historically stable labor share of output [Elsby et al., 2013, Karabarbounis and Neiman, 2013], various papers have related rising concentration to that decline [Autor et al., 2017] as well as to falling private investment [Gutiérrez and Philippon, 2016], rising price-cost markups [De Loecker and Eeckhout, 2017], and rising "profit shares" [Barkai, 2016], though few have attempted to explain rising concentration itself.²

This paper also relates closely to prior work on the industrial organization of the healthcare sector and dialysis industry in particular. Cutler et al. [2016] carefully document the sharp rise in dialysis industry concentration and study the effect of mergers among large providers. (This paper is described in more detail in Section 4, as are the following two papers.) Eliason et al. [2019] study how large firms transfer their operational strategies to the independent facilities they acquire. They find that national chains are able to boost reimbursements while cutting quality, effectively raising healthcare costs while jeopardizing patients. The authors thereby establish that industry consolidation presents concerns beyond its effects on competition—a potentially serious problem given the secular trends described in the preceding paragraph. Eliason [2018] estimates a structural model of spatial competition among dialysis providers and provides compelling evidence that local market power harms consumers—even to the point that subsidizing patients' travel meaningfully improves survival rates. Rapid consolidation has occurred outside dialysis as well.³ Moreover, concern extends beyond static considerations to effects on innovation, with Cunningham et al. [2017] finding that acquisitions of rival drug developers enable incumbents to "buy and kill" innovation that might later compete with them. Notably, many of the transactions they study fall below the HSR thresholds.

The paper proceeds out as follows. Section 2 provides institutional details. Section 3 describes and summarizes the data. Section 4 produces evidence that premerger notifications effect enforcement, market structure, and ultimately healthcare quality. Section 5 presents the model and reports the result of estimation. Section 7 forecasts the costs and benefits of the counterfactual policy. Section 8 concludes.

²An exception is Autor et al. [2017], who document that productive "superstar" firms have grown their share of output over past decades, which they argue may reflect the effect of increasing price sensitivity on firms with low production costs. Another exception is Peltzman [2014], who proposes the shift towards more permissive merger policy that followed the 1982 *Merger Guidelines* and work of Robert Bork.

³Physicians groups are an example Capps et al. [2017]. Koch et al. [2020] show that movements of physicians between provider groups can harm patients. Acute care hospitals are another example. Prior work shows concentration is broadly associated with higher prices. See Gaynor et al. [2015] for a review.

2 Institutional details and data

2.1 Merger reviews and premerger notifications

Changes in ownership that affect domestic commercial activity are subject to US competition law. Most critically, they must comply with Section 7 of the 1914 Clayton Act, which prohibits transactions where the effect "may be substantially to lessen competition, or to tend to create a monopoly."⁴ Though it is regarded today as the cornerstone of merger-related enforcement [Baer, 2014], the Clayton Act proved flawed as originally written for reasons that are central to this paper. Even after most substantive loopholes were closed in 1950,⁵ many direct competitors found that they could still successfully merge. To avoid scrutiny, these firms negotiated quietly and combined their operations covertly, telling as few parties as possible that they were doing so. Then, when the agencies eventually learned of these deals, it was too late to unwind them: since information was already shared and assets were already commingled, doing so would prove as hard as "unscrambling eggs" [Baer, 1996]. Thus, while the 1914 legislation provided the means to arrest anticompetitive deals prior to their completion, it did not provide the means to learn about these deals in their incipiency.

This practice abated in the late 1970s when Congress provided bipartisan support for the HSR Act, thereby establishing the US premerger notification program. The act requires parties to all proposed ownership transfers—unless explicitly exempt—to notify the agencies of their intentions, provide relevant details, and wait up to thirty days.⁶ During this time, filings are initially reviewed and, in the event they raise concerns, one (but not both) of the agencies would be prompted to seek "clearance" to investigate. Notably, the act bars firms from structuring transactions in ways that intentionally avoid reporting requirements and stipulates significant penalties for non-compliance.⁷

Exemptions are based mainly on size. For most mergers, the criteria are straightforward and can be summarized as follows. Under the act as it was originally written, transactions require notification only when the target has \$10 million or more in domestic assets (or has \$10 million or more in domestic

⁴ Other laws provide standing to sue but either rely on similar standards or are utilized far less frequently. As an example of the former, the FTC typically also sues under Section 5 of the FTC Act, which bars "unfair methods of competition in or affecting commerce." As an example of the latter, the DOJ, state attorneys general (AGs), and private plaintiffs can also sue under Section 1 of the Sherman Act, which bars arrangements "in restraint of interstate or foreign trade or commerce," or Section 2 of the Sherman Act, which bars monopolization of, or attempts to monopolize, "any part of interstate or foreign commerce." In some cases, other federal agencies will weigh in or even have jurisdiction. Examples include mergers involving banks, airlines, motor carriers, or entities for whom a change in control might jeopardize national security. Provided the transactions affect residents or resident businesses, state AGs can also sue under state statutes (e.g., California's Unfair Competition Law).

⁵ Most notable are the Celler-Kefauver amendments. Prior to them, for example, Section 7 did not address asset sales.

⁶ Waiting periods depend on attributes of the deal (e.g., the consideration offered) and have been amended since the original legislation, but the details are unimportant for this paper's purposes.

⁷ Penalties for willful avoidance can exceed \$42,000 per day.

revenue, in the event the target is engaged in manufacturing).⁸ Under the act as it was amended effective 2001, transactions require notification only if the aforementioned asset test is met *and* the deal involves consideration of \$50 million or more. Further, all thresholds have nominally increased with gross national product since 2004. As a result, at present, the act exempts most transactions valued below \$92 million. For many segmented industries, these thresholds absolve all but the largest transactions from reporting responsibilities.

To be clear, the act does not provide any party safe harbor from prosecution or affect the legality of any transaction. The congressional report on the act explicitly states, "[The] bill in no way alters the substantive legal standard of Section 7." This echoes earlier remarks of one of its sponsors, who said, "Let me emphasize one thing . . . this bill . . . merely provides for an effective *procedural* mechanism" (emphasis added). In other words, the clear intention of legislators is that anticompetitive deals should be challenged irrespective of their exemption or completion status. Moreover, if a nonreportable deal is discovered and raises initial concerns, then the review and enforcement procedures that are followed very closely resemble the ones applied to reportable transactions.⁹

Which mergers are challenged in practice, however, is a separate issue. As one widely referenced antitrust manual states, "If a merger clearly violates the guidelines, it is reasonably likely to be challenged *if* the government learns of the acquisition *and* discovers the potential violation" (emphasis in original) [Hills, 1985]. Agency officials appear to disagree—or at least obfuscate—about the degree to which reportability matters. For example, in 2014, an FTC Commissioner stated, "We don't have control over the mergers that are brought to us. They walk in the door. We're generally challenging HSR-reportable mergers" [Ohlhausen, 2014]. Yet in the same year, a Justice Department official stated, "The HSR reporting thresholds . . . are not synonymous with the contours of antitrust enforcement" [Overton, 2014]. Recent work provides at least preliminary evidence that the prior view holds [Wollmann, 2019]. To more completely resolve the issue as well as study its implications, I turn to merger activity and enforcement actions within the dialysis industry, which is described immediately below.

2.2 Kidney disease and treatment providers

The dialysis industry treats people suffering from a loss of kidney function. The kidneys are organs in the human abdomen that remove toxins and excess fluids from the body, but they gradually deteriorate

⁸This pared-down description is far from comprehensive but nonetheless covers most transfers of control. Supporting this claim are tight links between mergers, notifications, and enforcement actions, which are reported in this paper.

⁹For example, Hills [1985] writes that "non-HSR merger investigations will be very similar to HSR investigations" and "the scope of investigation is also identical." However, some procedural differences do exist. See Scher and Martin [2015] for examples.

with prolonged exposure to conditions such as diabetes and high blood pressure.¹⁰ When their filtration capacity—known as renal function—reaches as low as 10-15% of normal levels, patients are said to have end stage renal disease (ESRD). At this point, survival requires immediate care, with treatment regimens supervised by specialized physicians called nephrologists. One option is to receive a healthy kidney via a transplant, though these are suitable for only a small portion of the population and often involve long waiting lists, so they treat less than 3% of persons with ESRD.¹¹ The other option is dialysis, which involves diffusing a solution in the blood, passing the mixture across a semipermeable membrane, and filtering out unwanted substances.

Hemodialysis accomplishes this process outside the human body, with a machine serving as an external kidney.¹² Patients typically receive 3 4-hour treatments each week, during which time they may engage in sedentary activities (e.g., reading or watching television). Hemodialysis is the most common form of dialysis worldwide and by far the most common in the US, serving 90% of the nation's half-million dialysis patients. (Given its prevalence, the remainder of this paper refers to hemodialysis simply as "dialysis.") The number of these patients have grown over time alongside the prevalence of ESRD. While the proportion of the population receiving treatments, adjusted for age and comorbidities, has remained mostly flat since the mid-1990s, the unadjusted proportion has risen by one-quarter, reflecting older and more obese Americans. Mortality rates among patients fell 30% by 2012, with little consensus around the cause(s), have not moved much since then, and are still quite high. Persons on dialysis face 130 deaths per 1,000 person patient years [USRDS, 2017].

Since patients require such frequent treatments, they prefer locations that are very close to home. As a result, facilities are spread across the country, and patients rarely travel more than 20 miles for care.¹³ Markets are narrowly defined geographic units, and competition is a local phenomenon. (In the present paper, markets roughly correspond to US counties, but the conclusions are robust to alternative definitions where providers compete with nearby rivals for patients. See the Online Appendix for detail.) Ownership over these providers has consolidated rapidly over the past two decades. An industry once populated by independent owners is now dominated by two multinational corporations and 3 smaller national chains.

Although geographic coverage has expanded apace with demographics, entry into markets with

¹⁰They serve other functions, too, such as synthesizing erythropoietin, a hormone that stimulates red blood cell production.

¹¹This figure reflects the efficiency of the exchange mechanism. Economists are working to improve their design. See, e.g., Agarwal et al. [2019] and the literature cited therein.

¹²Peritoneal dialysis accomplishes this within the human abdomen, often while the patient is at home, asleep. Hemodialysis is ten to twenty times more prevalent over most of the panel, so I ignore peritoneal dialysis for the remainder of the paper. The reasons for its unpopularity are debated but beyond the scope of my analysis.

¹³See the Online Appendix for the empirical distribution of travel distances.

an incumbent provider is often expensive if not impossible. Many areas of the country require that entrants prove that the community "needs" additional capacity, creating large entry barriers. Kaiser's attempt to enter the Hawaiian island of Maui, where Liberty Dialysis had a monopoly, provides an illustration. Kaiser filed a Certification of Need application in September 2009, which was granted in May 2010 (subject to several preconditions including, among other things, a guarantee to provide service to particular communities for up to ten years). Nonetheless, Liberty requested and was granted reconsideration. After four pre-hearing conferences, one public hearing, and one private hearing, the state granted Kaiser's application in February 2011. Even so, Liberty successfully filed an appeal. After briefs, other submissions, and multiple rounds of oral arguments, the state again granted Kaiser's application in December 2011. Remarkably, Liberty successfully filed a secondary appeal. After additional litigation, Kaiser was eventually granted the certificate in June 2013—nearly four years later.¹⁴

2.3 Antitrust enforcement amid dialysis consolidation

The industry witnessed thousands of facility acquisitions over the past two decades, and preliminary investigations into many of these raised serious competitive concerns (as the subsequent section shows). When this happens, almost irrespective of the industry, the US merger review process proceeds in two steps: the agency cleared to investigate the transaction requests additional information to assess the competitive effects and, in the event concerns persist, formally challenges. The merging parties and agency typically find mutually agreeable terms under which the deal can proceed, but when they cannot, the case goes to trial.

Precisely how this all plays out—the venue, timing, tactics employed, etc.—depends on a host of procedural factors, although these idiosyncrasies can often be ignored when the analysis is confined to a single industry. This is true in dialysis, where merger enforcement is highly uniform. First, all cases are brought by the federal government and in particular the FTC (with a sole exception being a very small transaction unwound by the Minnesota Attorney General). Second, substantially every transaction involves a buyer whose principal business is to operate dialysis facilities and whose size exceeds that of its target.¹⁵ Relatedly, vertical considerations are limited. Third, as a "local" service industry, imports and exports play little role in determining market structure and behavior. Fourth, all economic analysis

¹⁴*Liberty Dialysis-Haw., LLC v. Rainbow Dialysis LLC*, SCAP-12-0000018 (June 27, 2013).

¹⁵There are a small number of potential exemptions related to entry by private equity, depending on how one views interactions between their portfolio businesses. The acquisition of Innovation Dialysis Systems by Ambulatory Services of America presents one potential example.

occurs at the market level, evidenced by each FTC complaint outlining the precise geographic areas where the associated merger will substantially lessen competition. Fifth, all enforcement actions equate to market-specific facility divestitures. In other words, in markets where the FTC and the merging parties agree that competition will be significantly reduced, target facilities must be sold to a third party that is approved by the agency. Informally, I refer to this outcome as the FTC "blocking" the acquisition of the facilities. However, to be precise, the agencies cannot unilaterally block deals, only challenge them. If the merging parties elect to try the case in court, then a judge or jury will decide their fate.

Though mergers among dialysis providers are straightforward to study, they are not especially unique in this regard. Federal agencies bring nearly all merger challenges. State AGs may join suits but often rely on national resources.¹⁶ Also, procedural differences exist between investigations originating at the DOJ and FTC, but these are second order considerations in nearly all cases. Further, agency actions mostly address horizontal combinations, and the legal standard of Section 7 of the Clayton Act applies. Finally, enforcement actions are typically market level rather than merger level decisions, with divestitures comprising the vast majority of remedies.¹⁷

3 Data

3.1 Sources

The Centers for Medicare & Medicaid Services (CMS) combines data from their regional offices to form Provider of Service (POS) files. They mainly reflect form CMS-1539, which surveys providers at events such as certifications, re-certifications, on-site visits, and changes of ownership. Annual files contain an individual record for each Medicare-approved facility. Since 1973, Medicare covers almost all persons with ESRD, so the dataset amounts to full coverage of the industry. For each facility I extract the unique Medicare provider number, name, address, opening and closing dates, parent company and recorded dates of any ownership changes. Together these provide a panel of facility-year observations spanning 1996-2017. (The Online Appendix provides more detail.)

United States Renal Data System (USRDS) collects and analyzes detailed operational and clinical data related to chronic kidney disease and ESRD, and it provides access to interested investigators. For each facility I observe basic attributes (e.g., the Medicare provider number) as well as the number

¹⁶"The number of mergers investigated by state attorneys general is small in comparison to the number investigated by the federal agencies" (although the author also mentions that the number is growing) [Gotts, 2001].

¹⁷For instance, a retrospective study by the FTC in 2015 of enforcement from 2006 to 2012 revealed that divestiture-only remedies accounted for 85% percent of relief. Remedies with some structural (i.e., divestiture-related) component accounted for 91%. Behavioral remedies accounted for less than 7% [FTC, 2015].

of machines and staff available to treat patients. For each ESRD patient, I observe an extensive list of demographic and baseline health variables: age, race, gender, body mass index (BMI), and renal function at the onset of treatments; dates of diagnosis, first dialysis session, and Medicare coverage; facility, modality, and dates of treatment; dates of hospitalizations; and date of death [USRDS, 2017]. To construct the dataset, I first standardize hospitalization and survival rates.¹⁸ Then, following Eliason [2018], I assign patients to their modal dialysis provider in each calendar year.

CMS also combines data from its Healthcare Cost Reporting Information System (HCRIS), which rely on annual surveys that are mandatory for all certified institutions. This source serves to supplement the prior two. From facility-year level records I extract the Medicare provider number, which provides a link to the aforementioned data, as well as chain affiliation and the numbers of machines and staff available to treat patients.¹⁹ To be specific, HCRIS provides timely information about ownership changes in years when facilities did not complete form CMS-1539 and hence did not update chain affiliation details. It also provides employment information in the first half of the panel, during which time USRDS was not collecting this data.

The FTC summarizes each enforcement action and publishes the associated legal documents. Consent orders, in particular, list the names and addresses of each facility that the parties and the agencies agree must be divested for the deal to close. Often, Medicare provider numbers are also provided. Complaints are also informative. Beyond merely identifying the parties, they describe their operations, enumerate the relevant geographic markets, and discuss the entry conditions. FTC data help identify facility acquisitions that would have been completed but for an enforcement action.

3.2 Summary

Table I summarizes the data at the facility-year level. In the average year, about 4,500 facilities are active.²⁰ Each houses 18 machines and serves 39 patients. Highlighting the at-risk nature of those individuals, patients face a roughly 70% chance of being hospitalized and a 15% of dying in a given year. Mean HHI, an unweighted average of the concentration index across markets in a given year, averages over 5,000—dialysis markets are typically very concentrated. As a point of comparison, the current *US Horizontal Merger Guidelines* consider a market with an HHI above 2,500 "highly concentrated" [US

¹⁸Control variables used to standardize rates are race, age, sex, BMI, and baseline kidney function, the latter of which is measured by glomerular filtration rate (GFR). The NIH states that the two most widely-used equations for estimating GFR are the CKD-EPI equation and the MDRD Study equation, so I include both measures.

¹⁹In about 10% of observations, though, the cost reports suffer from aggregation across affiliated, co-located providers (e.g., outpatient dialysis facilities associated with and located adjacent to hospitals).

²⁰Reassuringly, facility counts at the firm-year level match those that appear in FTC complaints and orders. See the Online Appendix for more information.

Department of Justice and Federal Trade Commission, 2010]. Moreover, there's considerable underlying heterogeneity. While some markets are highly fragmented, many others are monopolies. Separately, chain owners characterize about three-quarters of observations.

[Table I about here.]

Before summarizing merger and enforcement activity, I should note that while the confidentiality of HSR filings preclude the FTC from posting a comprehensive database of notifications it receives, exemptions are still relatively easily inferred for dialysis mergers. Reporting requirements depend on transaction values, which one can accurately predict using information on target firm operations. In particular, real transaction values are very closely tied to the number of facilities operated by the target firm. One can then verify the accuracy of this approach using "Early Termination" notices, which disclose basic transaction information for the vast majority of reportable mergers,²¹ and SEC filings. From the combined sources, I locate an HSR filing for 98.97% of the facility acquisitions that I predict are reportable,²² and I locate an HSR filing for only 0.48% of the facility acquisitions I predict are exempt. Thus, I conclude that exemption status is measured without meaningful error. The Online Appendix describes the process of inferring and verifying reportability in detail.

Table II summarizes the data at the facility acquisition level. About half of all facility acquisitions are exempt—over 2,000 establishment-level ownership changes were never reported to the government. The FTC requires about 270 divestitures, effectively blocking 6.2% of all facility acquisitions (as they are originally planned). The average simulated HHI change is about 500 points, although, as above, this figure masks important heterogeneity. (Figures H.2 and H.3 in the Online Appendix plot the full distribution.) In many cases, target and acquiror facilities are located in distinct markets prior to the transactions, while in other cases the facility acquisitions constitute mergers to duopoly and even monopoly. Finally, while reportable and exempt acquisitions are characterized by very similar *simulated* HHI changes, they're characterized by very different *realized* HHI changes. The root cause of this difference is explored in detail below.

[Table II about here.]

²¹The identities of parties requesting expedited HSR reviews are published in the *Federal Register*.

²²Even the small number of discrepancies probably do not reflect incorrect predictions about exemption status. All occur in the earliest years of the panel, during which time the FTC was not recording Early Terminations on its website. For this period, I am forced to search ad hoc through the *Federal Register*.

4 Descriptive evidence

4.1 Effects on enforcement

This subsection examines the relationship between premerger notifications and antitrust enforcement. US agencies are free to investigate and challenge any transaction, and their staff can learn about harmful deals many ways (e.g., through trade press, statements by brokers, SEC filings, etc.). However, most mergers are never publicized, and for those that are, details rarely surface before the transaction closes. Thus, the notification-enforcement relationship is an empirical question.²³

Enforcement actions equate to divestitures in the dialysis industry. They effectively block a portion of the total number of facility acquisitions associated with a merger, allowing the remainder to proceed. Thus, to study effects on enforcement, my outcome of interest is a divestiture. I index facility acquisitions by a and construct an indicator, $DIVEST_a$, which equals one when the merging parties divest any facility in the market where a occurs.²⁴

The estimation strategy depends on two sources of variation. The first source involves comparisons between the merger transaction values and exemption thresholds, which determine whether facility acquisitions must be reported to the FTC.²⁵ To illustrate, suppose that firm A buys firm B, facilities sell for \$4 million each, and the threshold stands at \$75 million in the year of the merger. Then, if B operates 18 facilities, none of the facility acquisitions associated with the merger are reported to the government, but if B operates 20, all of them are. I construct an indicator, $REPORTABLE_a$, which equals one if the facility acquisition is reportable (and zero if it is exempt). The second source involves differences in premerger market structure, which determine simulated HHI changes, denoted ΔHHI_a . Following theory, agency guidelines hold that transactions involving large changes in HHI are more likely to enhance market power, so these deals will face higher enforcement rates [US Department of Justice and Federal Trade Commission, 2010]. Hence, if premerger notification exemptions facilitate antitrust enforcement, then the effect of simulated HHI changes on divestiture rates may be stronger

²³See Section 1 for why prior work is inconclusive.

²⁴This approach accounts for a minority of cases where acquiror rather than target facilities are divested. Alternative approaches that avoid undercounting (e.g., proportionally allocating divestitures of acquiror facilities to target facilities in the same market) produces very similar results.

²⁵Given steep fines that can exceed \$42,000 per day for intentionally structuring a deal to avoid an HSR filing, I assume transaction values are not manipulated. Consistent with the assertion, I don't observe bunching of transaction values at the threshold. Moreover, only about 10% of proposed facility acquisitions are partial sales of the targets, so if acquirors are buying rivals in chunks to thwart detection, the practice is at least unlikely to drive my result. However, as Fiona Scott Morton points out, my results suggest intentional avoidance is tough for the FTC to spot but profitable for the firms when it goes undetected. Interestingly, the roughly 20 cases of single facility sales from one large chain to another involve disproportionately large simulated HHI changes, which is consistent with her remarks. In any case, this will not affect the interpretation of results once I condition on, or interact with, predicted market structure changes.

for reportable facility acquisitions.

To fix ideas, consider an estimating equation given by

$$\text{DIVEST}_a = \theta_0 + \theta_1 \text{REPORTABLE}_a + \theta_2 \Delta \text{HHI}_a + \theta_3 \text{REPORTABLE}_a \Delta \text{HHI}_a + \eta_a. \quad (1)$$

The interaction term $\text{REPORTABLE}_a \Delta \text{HHI}_a$ is designed to test whether the relationship between simulated HHI changes and enforcement actions is stronger for reportable facility acquisitions than exempt ones. Given the stark nature of the forthcoming results, I omit further discussion of the formal test in lieu of simply presenting the data and discussing its interpretation.

Figure I reports enforcement rates by deal type. The solid line reflects reportable facility acquisitions. For simulated HHI changes at or near zero, which indicate little or no competition between the acquired facility and any ones owned by the acquiror prior to the merger, there are almost no divestitures. In other words, when there are no competitive concerns, there are no enforcement actions. Yet, as simulated HHI changes rise, enforcement rates rise as well. For example, reportable facility acquisitions involving $\Delta \text{HHI} \approx 2,000$ face a nearly 75% chance of divestiture, which rises even further as ΔHHI nears 5,000. In short, the FTC blocks nearly all reportable facility acquisitions resulting duopoly and monopoly. In sharp contrast, the dashed line reflects exempt facility acquisitions. These ownership changes witness *effectively no enforcement actions, regardless of simulated HHI change*. This includes dozens of facility acquisitions involving $\Delta \text{HHI} > 2,000$, several of which involve ΔHHI near 5,000.

[Figure I about here.]

The lack of enforcement among nonreportable deals is even more surprising when juxtaposed with the *Horizontal Merger Guidelines*. The *Guidelines* state that "moderately concentrated markets that involve an increase in the HHI of more than 100 points" as well as "highly concentrated markets that involve an increase in the HHI of between 100 points and 200 points . . . potentially raise significant competitive concerns and often warrant scrutiny." Moreover, the *Guidelines* state that "mergers resulting in highly concentrated markets that involve an increase in the HHI of more than 200 points will be presumed to be likely to enhance market power" [US Department of Justice and Federal Trade Commission, 2010]. In other words, many exempt facility acquisitions yield concentration changes that are an order of magnitude higher than what the agencies presume to be anticompetitive. Importantly, however, this does not indicate the agencies are in any way at fault. Much more likely, it reflects the obvious difficulty in detecting privately negotiated contracts in their incipency.²⁶

²⁶This challenge was, after all, precisely what motivated Congress to establish the premerger notification program. As one of

In the Online Appendix, I explore robustness to alternative specifications. To assess whether the notification-enforcement relationship is sensitive to the way markets are defined, I redraw market boundaries entirely, using a 20 mile radius around each target facility. Separately, I recompute simulated HHI changes based on the number of machines rather than patients at each facility. Figure H.5 reports the results, which very closely resemble the relationship reported in Figure I. I then estimate equation 1, adding various controls including chain and year fixed effects. I also include simulated HHI in *levels* and its interaction with REPORTABLE. Table H.3 reports that estimates of θ_0 , θ_1 , θ_2 , and θ_3 are stable across the specifications, and that the predicted level of post-merger concentration is not an important determinant of divestitures.

I also investigate whether these findings reflect prosecutorial preferences over merger size. For instance, if big deals garner media attention, and if FTC staff are drawn towards "headline" transactions, then large deals will naturally face higher enforcement rates, which biases my results. To address this, I limit the sample to reportable facility acquisitions—ones I am sure the agency was aware of—and plot the relationship between divestiture rates and merger size.²⁷ Panel A of Figure H.6 reports the result, which indicates no such relationship.²⁸

4.2 Effects on market structure

Up until this point, nothing ensures that antitrust enforcement actually prevents consolidation. If acquirors of divested facilities struggle to operate those locations, then the facilities will shrink or even close, rendering the remedies ineffective. To evaluate this concern, I split the sample into exempt and reportable facility acquisitions and plot simulated HHI changes (on the x-axis) against observed ones (on the y-axis). The interpretation is straightforward. If facility acquisitions do not involve divestitures or if those divestitures are unsuccessful, then plotted points will lie near the 45 degree line; however, if divested facilities survive and ultimately prevent subsequent consolidation, then plotted points will diverge from the 45 degree line.²⁹

the bill's sponsors stated, "The only method the Division and FTC had to be aware of pending mergers and acquisitions was to read the general and trade press. In other words, if the *Wall Street Journal* missed one, so well may the FTC and the Division." ("Division" here means the Antitrust Division of the US Department of Justice.)

²⁷Real transaction values are very closely associated with the number of target facilities involved in the deal, so I use the latter measure to proxy for merger size. See Figure H.1 in the Online Appendix for a plot of real transaction values against the number of facilities owned by targets just prior to their acquisition.)

²⁸If "big business" draws negative sentiment, FTC staff may also be drawn towards large acquirors. Panel B of Figure H.6 rejects this idea as well.

²⁹Another way to understand this relationship is to imagine how the plot would look if the (a) agency blocks mergers if and only if involve simulated HHI changes above a threshold, (b) market shares were constant over time, and (c) my data was free of measurement error. Then, actual HHI changes would lie along the 45 degree line until the simulated HHI change threshold, at which point they would form a horizontal line.

Figure II reports the results, which show enforcement actions have a long lasting effect on market structure. In each panel, the base year for calculating observed HHI changes is the year prior to the acquisitions. In Panel A, I compute observed HHI changes one year forward from the base year. One year, of course, provides little time for facilities to fail. Thus, in Panels B and C, I repeat the exercise but compute the changes 3 and 5 years forward from the base year, respectively. Regardless of the horizon, exempt facility acquisitions lie near the 45 degree line, reflecting the absence of enforcement, while reportable facility acquisitions level off at around 900-1,100 points, reflecting the success of divestitures in preventing consolidation.

[Figure II about here.]

4.3 Effects on patients

The preceding subsections show premerger notifications are essential to effective antitrust enforcement. In their absence, providers acquire close competitors, eliminating erstwhile competition. Since dialysis providers compete on quality, which tends to decline when competition is eliminated,³⁰ these facts beg the question of whether reportability affects patients' health. This subsection answers it. It also motivates the structural model and counterfactual exercise presented later in the paper, which explicitly studies the provision of quality in equilibrium.

To ascertain effects on patients, my outcomes of interest are hospitalization and survival rates. I denote these measures y_{at} , where t indexes the year in calendar time. Both are natural choices to summarize quality in the industry. First, hospital admissions and deaths are unambiguously welfare relevant. Second, they're easily understood by policymakers and researchers—even those unfamiliar with the industry. Third, they're measured accurately in the data.³¹

The estimation strategy depends on three sources of variation. As in the prior subsection, I rely on variation in merger size, which determines exemption status. In this subsection, it's more natural to think in terms of exempt rather than reportable facility acquisitions, so I construct an indicator, $EXEMPT_a$, which equals $1 - REPORTABLE_a$. Again, I also rely on variation in premerger market structure, which determines predicted concentration changes (as well as realized ones, as evidenced by Figure II). To ease initial exposition, I discretize the measure of predicted consolidation, separating the

³⁰Firms in competitive markets provide high quality care to steal patients from competing facilities. Acquiring competing facilities reduce or even eliminate those incentives. See Section 5 for detail.

³¹By contrast, changes in inputs are tough to measure and interpret. Staff reductions, for example, provide no sense of whether quality has declined or the provider has improved efficiency. Moreover, they are mainly inferred from survey responses, which may be delayed or inaccurate.

top half of simulated HHI changes from the bottom half.³² Specifically, I construct an indicator, ΔHHI_a^* , which equals one if $\Delta\text{HHI}_a \geq 487$ points. I return to the continuous measure later in the section. The third source of variation relies on differences in the timing of mergers, which allow me to focus on within-facility changes in quality around acquisitions. I construct an indicator, POST_t , which equals one for post-acquisition observations and zero for pre-acquisition observations.

If acquisitions of rival facilities lead firms to reduce quality, then facility acquisitions involving large simulated HHI changes will lead to worse health outcomes for patients. However, premerger notifications facilitate antitrust enforcement, which block acquisitions of direct competitors, attenuating this relationship. Therefore, I test whether the effect of simulated HHI changes on premerger-to-postmerger changes in health outcomes is stronger for exempt facility acquisitions than for reportable ones. The panel consists of 10 years of observations immediately before and after each facility acquisition.³³ The approach requires quality to evolve similarly around different types of facility acquisitions, conditional on the controls.³⁴ The estimating equation is given by

$$y_{a\tau} = \lambda_1\text{POST}_\tau + \lambda_2\text{POST}_\tau\text{EXEMPT}_a + \lambda_3\text{POST}_\tau\Delta\text{HHI}_a^* + \lambda_4\text{POST}_a\text{EXEMPT}_a\Delta\text{HHI}_a^* + \mu_a + \kappa_t + \zeta_{a\tau}^c + \nu_{a\tau}. \quad (2)$$

τ indexes the year in event time. μ_a , κ_t , and $\zeta_{a\tau}^c$ represent fixed effects specific to facility acquisition a , calendar year t , and chain c .

Columns 1 and 2 of Table III report the results from estimating equation 2. For facility acquisitions involving large simulated HHI changes, premerger notification exemptions increase hospitalization rates by 3.6 percentage points and decrease survival rates by 1.8 percentage points. Both estimates are statistically (and clinically) significant. Notably, exempt facility acquisitions involving large simulated HHI changes are the only ones that, conditional on the control variables, affect quality.³⁵ In other words, the FTC is successful in blocking facility acquisitions that would otherwise reduce healthcare quality. More succinctly, antitrust enforcement saves lives.

³²This cutoff is merely chosen to balance the number of observations above and below—it yields only descriptive evidence and doesn't map to theory, but this is broadly true of concentration indices anyway. A natural alternative might be 200 points, which the *Horizontal Merger Guidelines* use to distinguish acquisitions presumed to increase market power from those that are not. This would not change the conclusions of this section.

³³Windows of 5, 7, and 10 years produce similar results, as shown below. A prior version of this paper considered a balanced but shorter window around each acquisitions. Again, similar results obtain.

³⁴The restriction is fundamentally untestable, but it can at least be challenged by examining the pre-acquisition evolution of the coefficient of main interest, which I do below.

³⁵This does not imply the acquisitions are benign. Chain dummies absorb the effect of large national providers taking over independently run facilities, which is the subject of work by Eliason et al. [2019], described below.

[Table III about here.]

In the Online Appendix, I investigate the robustness of these findings. Separately, I (a) redraw market boundaries using a 20 mile radius around each target facility, (b) recompute simulated HHI changes using the number of machines rather than patients at each facility, (c) reconstruct the panel using a 7 year window around each facility acquisition, and (d) reconstruct the panel using an even shorter 5 year window around each facility acquisition. I then re-estimate equation 2 under these alternatives and report the results in Table H.4. The coefficients are similar to the ones reported above. For example, when health outcomes are measured in survival rates, resulting estimates of λ_4 —the coefficient on the interaction between POST, EXEMPT, and ΔHHI^* —are -1.8, -1.8, -2.0, and -1.9 percentage points for (a)-(d), respectively. All are near the -2.0 percentage point change reported in column 2 of Table III. Next, I assess the possibility that preexisting trends in quality drive these results. To do so, I estimate a version of equation 2 where λ_1 , λ_2 , λ_3 , and λ_4 are allowed to vary by year in event time. I then plot the evolution of the main coefficient of interest and report the result in Figure H.7. There is no evidence of pre-acquisition trends for either outcome measure. Finally, I study facilities owned by acquiring firms, whose incentives to provide quality are similarly reduced by acquisitions of close competitors. I estimate a version of equation 2 that has been adapted to acquiring facilities, and I report the result in Table H.5. Consistent with the preceding findings, exempt facility acquisitions involving large simulated HHI changes increase hospitalization rates and decrease survival rates.

To ease exposition, preceding results relied on a binary measure of competition. Here, I relax that restriction, replacing ΔHHI^* with ΔHHI and re-estimating equation 2. Columns 3-4 of Table III report the results. Coefficients are similar in sign and significance to prior ones. For every 1,000 point change in simulated HHI, hospitalization rates increase 0.83 percentage points more for exempt facility acquisitions relative to reportable ones. When quality is instead measured by survival rates, the relative difference is a 0.50 percentage point decline.

While the negative relationship between concentration and quality in the dialysis industry squares with prior work by Eliason [2018], who relies on very different variation, it appears at odds with work by Cutler et al. [2016], at least at first glance. They study mergers between national chains, such as DaVita's \$3 billion purchase of Gambro's dialysis operations in 2005, and find that consolidation does not meaningfully impact quality. In fact, I reach exactly the same conclusion—for large mergers, which are all reportable to the FTC, predicted concentration changes do not affect quality, since the truly harmful elements of those deals are blocked by the agency. (This is evidenced by coefficients in rows

2 and 4 of Table III). In other words, Cutler et al. [2016] study precisely the sorts of deals that face aggressive effective enforcement, thereby introducing an empirical puzzle that I address.

Also related to these results, Eliason et al. [2019] study the transfer of management practices from large chains to independently run facilities that they acquire (i.e., they focus on changes in the identity of the owners, not on consolidation). Provocatively, they find that the chains transfer operational strategies that not only increase the amount of money collected for dialysis treatments but also reduce the quality of care (independent of any changes in market structure). As the authors point out, the effectiveness of these strategies implies competition does not perfectly discipline provider behavior—if it did, substitution away from low quality providers would drive them out of business. Though it is not apparent from Table III because chain fixed effects absorb the mean effect of chain ownership on quality, my estimates are consistent with their main result,³⁶ and I reach the same broad conclusion about imperfect competition.³⁷

5 Model

Descriptive evidence presented above suggests premerger notification exemptions cost lives in the dialysis industry. However, concentration indices are one dimensional approximations of complicated substitution and ownership patterns, to which quality responds non-linearly in equilibrium. Hence, judging the scale of the problem from the preceding evidence is hard, and using it to accurately predict how quality would respond to changes in market structure induced by an alternative notification policy is virtually impossible. Thus, in this section, I explicitly model that response.

³⁶The code that produces Table III designates the chain owner of non-chain facilities as "Independent" and enforces that these observations are the omitted category in the calculation of chain fixed effects. Hence, the approximate overall (negative) effect of chain ownership on healthcare outcomes can be quickly ascertained by computing the weighted mean of the reported coefficients on the chain dummy variables. Note this isn't a surprise finding—we rely on very similar data sources/samples.

³⁷Further along in the analysis, one of their specifications suggests competition has no effect on quality, but it defines competition as any target-acquiror overlap within very broadly defined geographic areas. Section G of the Online Appendix provides more detail. In any case, the point is not central to their paper, and amending their specification does not alter their main result, as indicated by the immediately preceding footnote.

5.1 Demand

Each patient i in geographic market g and year t chooses a facility j or the outside option, home dialysis, to maximize total utility.³⁸ If i chooses j , then he or she receives

$$u_{ijgt} = \beta_0 + \beta_x x_{jgt} + \beta_n n_{jgt} + \phi_t + \chi_{jgt}^c + \zeta_{jgt} + \alpha_1 dist_{ijgt} + \alpha_2 dist_{ijgt}^2 + \alpha_3 dense_{gt} dist_{ijgt} + \alpha_4 dense_{gt} dist_{ijgt}^2 + \epsilon_{ijgt}, \quad (3)$$

while if i chooses the outside option, u_{i0gt} is zero. Utility depends patients' distaste for travel, which in turn depends on the distance between i 's home and facility j , denoted $dist_{ijgt}$, and the population density of the surrounding area, denoted $dense_{gt}$. Utility also depends on the quality of care measured in survival rates, denoted x_{jgt} , and the number of machines, denoted n_{jgt} . ϕ_t and χ_{jgt}^c denote year and chain fixed effects, and ζ_{jgt} and ϵ_{ijgt} denote preference shocks. The former are i.i.d. over facilities and time, while the latter are i.i.d. over patients, facilities, and time. I assume ϵ_{ijgt} are distributed Type 2 extreme value.

This specification yields logit individual choice probabilities equal to

$$s_{ijgt} = \frac{e^{\alpha_{ijgt} + \beta_0 + \beta_x x_{jgt} + \beta_n n_{jgt} + \phi_t + \chi_{jgt}^c + \zeta_{jgt}}}{1 + \sum_k e^{\alpha_{ikgt} + \beta_0 + \beta_x x_{kgt} + \beta_n n_{kgt} + \phi_t + \chi_{kgt}^c + \zeta_{kgt}}} \quad (4)$$

where $\alpha_1 dist_{ijgt} + \alpha_2 dist_{ijgt}^2 + \alpha_3 dense_{gt} dist_{ijgt} + \alpha_4 dense_{gt} dist_{ijgt}^2$ is written as α_{ijgt} merely for notational convenience. The probability of choosing the outside option is given by the same expression except that the numerator of equation 4 is replaced with 1. Integrating the individual choice probabilities out over i yield market shares, s_{jgt} , and the outside share, s_{0gt} .

5.2 Quality

Each firm f chooses quality to maximize profits, which equal

$$\pi_{fgt} = \sum_{j \in J_{fgt}} (P_{gt} - mc_{jgt}) s_{jgt} M_{gt}. \quad (5)$$

³⁸This setup borrows heavily from Eliason [2018]. My specifications depart from his in several ways, though, because we study very different issues and counterfactual policies (e.g., he does not study antitrust enforcement). See also Footnote 39.

J_{fgt} denotes the set of facilities owned by f in g at t , P_{gt} denotes price, and M_{gt} denotes market size.³⁹ mc_{jgt} denotes marginal cost, which I parameterize as an additively separable function of quality, quality squared, the number of machines, and an unobserved shock, ω_{jgt} . ω_{jgt} is i.i.d. over facilities and time. γ denotes the vector of coefficients on the observable components.⁴⁰

Profit maximization dictates that firms set the first order condition of equation 5 equal to zero. For every $k \in J_{fgt}$,

$$\sum_{j \in J_{fgt}} \left((P_{gt} - mc_{jgt}) \frac{\partial s_{jgt}}{\partial x_{kgt}} \right) - \frac{\partial mc_{kgt}}{\partial x_{kgt}} s_{kgt} = 0, \quad (6)$$

where $mc_{jgt} \equiv \gamma_0 + \gamma_1 x_{jgt} + \gamma_2 x_{jgt}^2 + \gamma_3 n_{jgt} + \omega_{jgt}$ and $\frac{\partial mc_{kgt}}{\partial x_{kgt}} \equiv \gamma_1 + 2\gamma_2 x_{kgt}$. Notice that equation 6 bears a close resemblance to the conditions that determine Nash prices under exogenously determined product characteristics (with which many readers will be more familiar). This is not coincidental: firms setting x subject to fixed P face analogous tradeoffs to firms setting P subject to fixed x .

Equation 6 also clarifies why consolidation impacts quality. To see this, suppose f initially owns only facility j but then acquires a rival location j' . Prior to the acquisition, the provider chooses x_j such that $(P - mc_j) \frac{\partial s_j}{\partial x_j} = \frac{\partial mc_j}{\partial x_j} s_j$ (where g and t subscripts are omitted for ease of notation). Increasing quality at the margin attracts new patients on whom the firm will earn positive profits, but it also raises the cost of serving all current patients, and in equilibrium these forces exactly offset one another. After the acquisition, however, the provider appends $(P - mc_{j'}) \frac{\partial s_{j'}}{\partial x_j}$ to the left-hand side of the equation. The additional terms reflects how increasing x_j steals market share (i.e., diverts patients away) from j' , which the provider now owns. (Colloquially, it "cannibalizes" sales from the other location.) Thus, at least if one holds all other facility decisions fixed, x_j must fall to balance the left- and right-hand sides of the equation, meaning incentives to provide quality at j are reduced by the acquisition of j' .

³⁹This specification implicitly assumes machines are exogenously determined. Eliason [2018] endogenizes them, which is essential in his setting, since firms make entry decisions, thereby making long-run determinations of facility attributes.

⁴⁰Medicare sets reimbursement rates for a large majority of patients, so I implicitly assume here that firms in the same market face the same price. The substance of the assumption with respect to the counterfactual exercise is that consolidation doesn't affect price. In practice, a minority of treatments are privately insured, so modest price effects are possible. However, that bargaining process might be determined at a much higher level of market aggregation. In any case, I lack data on insurer-provider contracts, so price effects are beyond this paper's reach. Since firms rarely add or remove machines, I also assume these evolve exogenously. (This is relatively straightforward to relax.)

6 Estimation and results

6.1 Demand estimation

To estimate demand, I begin by rewriting equation 3 as

$$u_{ijgt} = \delta_{ijt} + \alpha_1 \text{dist}_{ijgt} + \alpha_2 \text{dist}_{ijgt}^2 + \alpha_3 \text{dense}_{igt} \text{dist}_{ijgt} + \alpha_4 \text{dense}_{igt} \text{dist}_{ijgt}^2 + \epsilon_{ijgt}, \quad (7)$$

substituting δ_{jgt} for the patient-invariant components of utility. The estimation process proceeds in two steps. The first step recovers estimates of $\alpha \equiv [\alpha_1 \ \alpha_2 \ \alpha_3 \ \alpha_4]$. For any guess $\tilde{\alpha}$, I (a) obtain $\delta(\tilde{\alpha})$, the vector of δ_{jgt} that exactly fit market shares under $\tilde{\alpha}$ [Berry, 1994], (b) compute the individual choice probabilities at $\tilde{\alpha}$ and $\delta(\tilde{\alpha})$, and (c) compute the log likelihood of observing the choices patients make in the data.⁴¹ I search over $\tilde{\alpha}$ to arrive at $\hat{\alpha}$, the vector that maximizes the log likelihood, and define $\hat{\delta} \equiv \delta(\hat{\alpha})$.

The second step recovers estimates of the coefficients on the patient-invariant observable components of utility. Estimating β_x presents the main challenge. Since firms know the vector of demand shocks when they choose quality, observed quality depends on ζ , so regressing $\hat{\delta}$ on x (and the other observables) produces biased estimates. The solution lies in the descriptive evidence reported in Section 4: exempt facility acquisitions involving large predicted concentration changes reduce healthcare quality but are plausibly uncorrelated with demand shocks. In other words, stealth consolidation shifts x independently of ζ .

To operationalize the idea, I construct an instrument such that

$$z_{jgt}^1 = \sum_a H_a \{j = j_a\} \{g = g_a\} \left(\{t \in (t_a, t_a + \rho]\} - \{t \in [t_a - \rho, t_a)\} \right), \quad (8)$$

where

$$H_a = \text{EXEMPT}_a \Delta \text{HHI}_a^{**} \left\{ \sum_{j,g,t} \left(\{j = j_a\} \{g = g_a\} \{|t - t_a| \leq \rho\} \right) \geq (2\rho + 1) \right\}. \quad (9)$$

j_a , g_a , and t_a index the facility, geographic market, and year associated with a , respectively. H_a selects facility acquisitions and equals the product of three terms. The first two terms select exempt facility acquisitions involving large predicted concentration changes. (Precise construction of HHI_a^{**} is described below.) The third term ensures there are at least ρ -number of pre- and post-acquisition observations.⁴²

⁴¹In practice, I obtain $\delta(\tilde{\alpha})$ using the SQUAREM procedure proposed by Reynaerts et al. [2012], which converges four times faster than the contraction mapping in my setting.

⁴²For example, if a facility enters the data in 2000 and is acquired in 2002, then it offers only 2 pre-acquisition observations. If

The remaining terms in z_{jgt}^1 construct a difference such that ρ pre-acquisition observations are subtracted from ρ post-acquisition observations. (In practice, I set ρ to 3, but slightly shorter/longer windows produce similar results.⁴³) Descriptive evidence presented in Section 4 indicates $\mathbb{E}[z^1 x | \phi, \chi] < 0$. Moreover, so long as facility-specific unobserved demand shocks do not induce acquisitions, $\mathbb{E}[z^1 \xi | \phi, \chi] = 0$.

Note that H_a depends on ΔHHI_a^{**} , not ΔHHI_a^* . At least at first glance, the latter appears a natural candidate to select facility acquisition involving large predicted concentration changes. It is not. ΔHHI_a^* depends on $s_{j,g,t-1}$, which in turn depends on ξ_{j,g,t_a-1} . Hence, if z^1 were constructed this way, then $\mathbb{E}[z^1 \xi | \phi, \chi] \neq 0$. Fortunately, one can predict concentration changes using other variables that are independent of the unobservable product-specific demand shocks.⁴⁴ I construct ΔHHI_a^{**} exactly as I would ΔHHI_a^* but compute market shares using numbers of machines rather than patients assigned to each facility.⁴⁵ Incidentally, this distinction turns out to be more of a technical than practical matter—I estimated the model both ways and arrived at similar estimates.

Finally, I construct a matrix Z_1 , whose columns comprise a constant, n , Z^1 , year dummies, and chain dummies. I assume $\mathbb{E}[\xi | Z_1] = 0$ and construct a GMM objective function given by $(Z_1' \hat{\xi})'(Z_1' \hat{\xi})$, which is based on moments implied by this restriction [Berry et al., 1995]. $\hat{\beta}_0$, $\hat{\beta}_x$, $\hat{\beta}_n$, $\hat{\phi}$, and $\hat{\chi}$ minimize the value of this function.

6.2 Marginal cost estimation

To recover marginal costs, I construct two sets of moments. To obtain the first set, I start with the first order condition of profits with respect to quality. Substituting mc_{jgt} in equation 6 with its parameterization yields that for all k in J_{fjgt} ,

$$\sum_{j \in J_{fjgt}} \left((P_{gt} - \gamma_0 - \gamma_1 x_{jgt} - \gamma_2 x_{jgt}^2 - \gamma_3 n_{jgt} - \omega_{jgt}) \frac{\partial s_{jgt}}{\partial x_{kgt}} \right) - (\gamma_1 + 2\gamma_2 x_{kgt}) s_{kgt} = 0. \quad (10)$$

I "pre-multiply" equation 10 (or more precisely, its equivalent in linear algebra notation) by the inverse of a matrix of derivatives, which provides a system of equations that are additively separately in the

$\rho = 3$, then $\{\sum_{j,g,t} \{j = j_a\} \{g = g_a\} \{|t - t_a| \leq \rho\}\} = 6$ which is less than $< (2\rho + 1) = 7$.

⁴³If ρ is much larger, H_a selects a very limited number of facility acquisitions, since its construction requires observations are balanced around the facility acquisitions. An alternative is to construct H_a using an unbalanced panel and weight pre- and post-acquisition observations; however, this complicates exposition, and Figure H.7 suggests little is gained by doing so.

⁴⁴I thank an anonymous referee for suggesting the robustness test that ultimately motivated this specification.

⁴⁵Other methods of construction are possible. One could simply count facilities (i.e., assign equal market shares to each j).

marginal cost disturbances. For every (g, t) -pair, this step yields

$$P_{gt} - \gamma_0 - x_{gt}\gamma_1 - x_{gt}^2\gamma_2 - n_{gt}\gamma_3 - \omega_{gt} - A_{gt}^{-1}s_{gt}\gamma_1 - 2A_{gt}^{-1}x_{gt}s_{gt}\gamma_2 = 0. \quad (11)$$

P_{gt} , x_{gt} , n_{gt} , ω_{gt} , and s_{gt} denote vectors of prices, quality, machines, marginal cost disturbances, and market shares, respectively. A_{gt} represents a matrix where each (j, k) element equals $\frac{\partial s_{jgt}}{\partial x_{kgt}}$ if j and k are jointly owned and zero otherwise. I construct matrix Z_2 , whose columns comprise a constant, n , an indicator for z^1 equal to one, an indicator for z^1 equal to -1, and P . I assume $\mathbb{E}[\omega|Z_2] = 0$ and form the first set of moments by interacting Z_2 with ω .

To obtain an additional moment, I exploit information on profit margins. Throughout the panel, the largest US dialysis provider has published firm-wide annual average per-treatment revenues and costs of their domestic facilities.⁴⁶ These figures imply stable but steadily rising profit margins over the panel. For example, the firm reports 29% and 31% in 2000 and 2001, respectively, and 33% and 34% exactly 10 years later, respectively. I assume constant marginal costs of treatments at the firm level and form a moment that matches the firm's annual average profit margins. Its construction is straightforward:

$$\sum_t \sum_g \sum_{j \in J_{fgt}} P_{gt}(1 - \overline{\text{MARGIN}}_{ft}) - \gamma_0 - x_{jgt}\gamma_1 - x_{jgt}^2\gamma_2 - n_{jgt}\gamma_3 - \omega_{jgt} = 0, \quad (12)$$

where $\overline{\text{MARGIN}}_{ft}$ represents the average profit margin reported by DaVita in t . I construct a GMM objective function based on the full set of moments. The vector $\hat{\gamma} \equiv [\hat{\gamma}_0 \hat{\gamma}_1 \hat{\gamma}_2 \hat{\gamma}_3]$ minimizes the value of this function.

6.3 Estimates

Panel A of Table IV reports demand-side estimates. Patients prefer facilities that are very close to home, consistent with them making several trips to their provider each week. The preference is especially pronounced in areas of the country with high population density, reflecting the time and hassle associated navigating congested roads. Patients also prefer facilities with more machines, where appointments are easier to schedule. Finally, patients value healthcare quality, measured in one-year risk-adjusted survival rates. Putting aside the constant, all demand-side estimates are significant at the 5% level. Distance-related estimates are especially precise, due at least in part to the large number of

⁴⁶See, e.g., the 2017 Annual Report of DaVita Inc., which is available at <https://investors.davita.com/static-files/ca97d703-5935-4708-9fcb-4d527b258651>.

patient-year observations in the panel.

[Table IV about here.]

The magnitudes of the coefficients imply sensible tradeoffs. On average, patients are indifferent between 2.96 miles of additional travel, a one percentage point increase in quality, and 3.9 more machines. Three miles may seem small at first, but recall that patients receive 156 treatments per year, equating to 312 one-way trips from/to their provider annually. Moreover, these trips directly precede/follow 4 hours of treatment, so patients are presumably very anxious to arrive at the facility and eager to return home. At an average speed over local/arterial roads of 22.5 miles per hour [Cardelino, 2012] and an average pretax time value of \$30 per hour,⁴⁷ the additional travel equates to \$1,248 per year. Notably, if I multiply this figure by 100, thereby translating a percentage point increase in survival rates into the statistical value of a life-year, then I arrive at \$124,800, which is close to the roughly \$150,000 commonly assumed and/or arrived at in the literature studying "health capital."⁴⁸

Panel B reports estimates of γ . Costs are increasing in the number of machines operated by the facility, consistent with trade press that indicates these machines require periodic maintenance, regardless of actual use, as well as substantial floor space, which translates to higher rent. Each raises per-treatment expenses by \$1.08. Costs are also increasing in quality (at an increasing rate). Calculated at the average observed quality level, a one percentage point increase adds \$4.12 per treatment.

Profit margins and price-cost markups average 31.3% and 50.8%, respectively. These measures improve with stealth consolidation. In the three year period immediately preceding exempt facility acquisitions involving $\Delta HHI_a^{**} = 1$, profit margins average 31.2%. In the years immediately following, they average 34.3%—an increase of 3.1 percentage points, or just over 10%.

7 Counterfactual policy analysis

Preceding sections show premerger notifications enable antitrust enforcement, which in turn promotes competition and saves lives in the US dialysis industry. These findings raise the question of whether society would be better off under stricter reporting requirements. Such a change, however, would generate additional work for the agency, which would need to hire staff to review and potentially

⁴⁷According to the BLS *Quarterly Census of Employment and Wages*, annual wages averaged \$55,390 in 2017. Adding 8% for retirement and marginal supplemental benefits and assuming a 2000 hour work-year implies a per-hour wage of \$29.91.

⁴⁸For example Cutler and Richardson [1998] assume \$100,000 in 1998, while Lee et al. [2009] arrive at \$129,090 in 2009. These figures translate to \$150,040 and \$147,800 in 2017 dollars, respectively.

challenge additional transactions. This section studies the costs and benefits of eliminating premerger notification exemptions in the dialysis industry.

7.1 Counterfactual merger activity and market structure

The first step of this analysis involves predicting how market structure would evolve were all dialysis mergers reported to the FTC in their incipiency. To make these predictions, I flexibly estimate the relationship between simulated HHI changes and divestiture rates among reportable facility acquisitions and then retroactively apply predicted enforcement rates to exempt facility acquisitions.⁴⁹ For example, if reportable facility acquisitions involving simulated HHI changes of 1,000 points face a 40% chance of divestiture, then exempt facility acquisitions involving similar simulated HHI changes will also be blocked 40% of the time under the counterfactual policy. Using this approach, many possible market structures may be realized, so ownership evolves stochastically, not deterministically.

To accurately characterize market structure under the counterfactual policy, I allow for deterrence.⁵⁰ Filing the premerger notification is among the very last steps in the process of acquiring another firm. Reaching that point requires incurring significant costs (e.g., the time managers spend evaluating the opportunity, attorneys spend drafting and negotiating the terms of sale, auditors spend conducting due diligence, etc). Thus, an acquiror facing a high risk that a relatively small merger will be blocked won't attempt it in the first place. This fact has important consequences for the design of premerger notification programs. So long as the agency commits to challenge anticompetitive deals of which it's made aware, the counterfactual policy deters many anticompetitive facility acquisitions, which means the agency never actually needs to challenge—let alone review—these transactions. In other words, deterrence provides "enforcement for free." To account for deterrence (and significantly reduce the computational burden), I assume that no firm proposes a merger where the average probability of divestiture across target facilities exceeds 50%.⁵²

The deterrence assumption combined with the empirical relationship between simulated HHI changes and divestiture rates among reportable facility acquisitions fully specify the range of possible market structures, indexed by r , and the probability each is realized, denoted $\psi(r)$. These objects are the basis for calculating the benefits of the counterfactual policy. Moreover, they fully specify the number

⁴⁹See Figure H.8 in the Online Appendix, which shows a second order polynomial provides a close approximation.

⁵⁰See also Wollmann [2019] and Clougherty and Seldeslachts [2013] for motivating discussions.

⁵¹Putting aside deterrence, my approach does not allow firms to adjust their acquisition plans in response to agency policy changes. See Nocke and Whinston [2010, 2013] for a compelling discussion of why and where dynamic strategies are important. (In my setting, incorporating them is unfortunately far too computationally burdensome.)

⁵²The true cutoff is probably below 10%. Fifty percent is conservative in the sense that lower cutoffs produce larger benefits under the counterfactual policy. See the Online Appendix for a detailed discussion.

and nature of transactions the FTC will have to review and potentially challenge, which form the basis for calculating the costs. Both sets of calculations are provided below.

7.2 Benefits of the counterfactual policy

The second step of the analysis involves forecasting the benefits of eliminating premerger notification exemptions, which depend on the number of lives saved. For each (g, t) -pair, I recompute equilibrium quality under each r as well as under the observed market structure (call it \bar{r}). The expected number of lives saved by the counterfactual policy in year t is then given by

$$\Lambda_t = \sum_g \sum_{r \in \mathcal{R}_{gt}} \sum_j R_{gt}^{-1} M_{gt} x_{jgt}^r s_{jgt}^r \psi(r) - \sum_g \sum_j M_{gt} x_{jgt}^{\bar{r}} s_{jgt}^{\bar{r}}, \quad (13)$$

where \mathcal{R}_{gt} denotes the set of r in (g, t) , R_{gt} counts the number of elements in \mathcal{R}_{gt} , $x_{jgt}^{\bar{r}}$ denotes the quality offered by j in g at t under \bar{r} , and $s_{jgt}^{\bar{r}}$, x_{jgt}^r , and s_{jgt}^r are defined analogously. Note that when I compute x_{jgt}^r , I do *not* change the value of the chain fixed effects (χ_{jgt}^c) assigned to the observations, so differences between x_{jgt}^r and $x_{jgt}^{\bar{r}}$ reflect only competitive effects. (If I were to have done so, the counterfactual policy would produce even larger benefits.) To obtain a dollar-based measure of benefits in real terms, I multiply that difference by the statistical value of a life-year, which equals \$100,000 in 1998 [Cutler and Richardson, 1998], and translate nominal figures into constant 2017 dollars.

Note that the benefits of the counterfactual policy do not end in 2017 simply because my data does. Rather, they recur until, for example, new technology changes the nature of renal replacement therapy or disrupts the way providers compete. To calculate the annuity value of lives saved, I let the differences in quality induced by the counterfactual policy persist for 7 to 15 years after the panel ends.⁵³ I grow the patient population at 1% and discount future benefits at 5% (in real terms). Both choices err on the side of producing conservative estimates.^{54,55}

Before turning to the costs of the counterfactual policy, it's useful to have a sense as to how quality responds to eliminating premerger notification exemptions. To provide it, I plot the distribution of equilibrium quality separately under the current and counterfactual policies. To focus on locations that are actually impacted by the counterfactual policy, I narrow the sample to markets affected by at

⁵³At the time of writing, nearly 3.5 years have passed since the data ended. To the best of my knowledge, there is no indication that the average market has fragmented or that market power has abated, so this seems like a natural starting point.

⁵⁴High growth rates predict higher benefits, and the ESRD population persistently grows at 3%. Low discount rates also predict higher benefits, and most authors rely on a 3% [Cutler and Richardson, 1998].

⁵⁵Since dialysis benefits are uncorrelated with stock market values, the *riskless* real rate of return is appropriate here.

least one exempt facility acquisition involving a large simulated HHI change.⁵⁶ Figure III reports the result. Under the counterfactual policy, the distribution of quality shifts right, consistent with the idea that less consolidation equates to more competition and ultimately better patient outcomes. Averaging over facility-year observations, the difference in mean quality is just over one percentage point higher. Nearly all of it is driven by facilities whose quality falls in the bottom tercile of the distribution under the current policy. A close inspection of the data reveals these observations reside in markets with either a single large simulated HHI change or multiple facility acquisitions.

[Figure III about here.]

7.3 Costs of the counterfactual policy

The third step of the analysis involves forecasting the costs of eliminating premerger notification exemptions. Although a comprehensive study of the agency's production function is beyond the scope of the current paper, one can nonetheless construct reasonable bounds on the added burden placed on the FTC by the counterfactual policy. To do so, I gather historical data on FTC operations that covers the same years as the panel of dialysis facilities, make assumptions about the relationship between merger activity and enforcement costs, and use the combination to predict the additional expenses borne by the agency.

To more precisely predict these costs, I make an important distinction between costs incurred by the FTC Premerger Notification Office (PNO) and those incurred in post-PNO activities. When providers submit premerger notifications, their filings are routed to the PNO first, which conducts initial antitrust reviews, prepares summaries of the transactions, and effectively screens out transactions that do not raise competitive concerns.⁵⁷ Most prospective merger reviews end here; few go on to involve formal investigations. In other words, all reportable mergers increase PNO costs, but only anticompetitive ones would increase post-PNO costs. The distinction is especially relevant for the counterfactual I consider, which deters most mergers that raise competitive concerns (were the firms to attempt them).

For a sense of resources consumed by the agency over the sample period, I turn to Figure IV. Panel A reports total enforcement costs specific to the PNO. To obtain these figures, I multiply office headcount, which is published in agency annual reports, by the average salary per FTC staff attorney, which is reported by the Office of Personnel Management, and inflate the product 46.2% to cover

⁵⁶To be precise, I narrow the sample to (g, t) -pairs such that an exempt facility acquisition involving a large simulated HHI change occurred in g at year $\tau \leq t$.

⁵⁷Put differently, "The initial antitrust assessment performed by the PNO is designed to separate nonproblematic transactions from transactions that appear to present competitive problems" Gotts [2001].

non-salary benefits, employer-paid taxes, and incidentals.⁵⁸ PNO costs range from \$3.5 to \$6 million annually—remarkably low when one considers the agency formally evaluates an average of 2,047 mergers per year, including several complicated multi-sector billion dollar transactions. To add further perspective, Panel C reports the dollar volume of reportable mergers, published in the HSR Annual Report. The office reviews ownership changes valued at \$1 and \$4 trillion annually, or about 5-10% of gross domestic product in most years. Panel B reports total Post-PNO enforcement costs, which range from \$31 to \$41 million annually. Post-PNO activities include investigations, negotiations with merging parties, and even litigation, so it's not surprising that they average almost an order of magnitude higher than figures reported in Panel A. These efforts are spread over fewer but larger mergers. The agency investigates an average of 190 mergers each year, which Panel D shows equate to about \$300 billion annually.⁵⁹

[Figure IV about here.]

To construct bounds, I make functional form assumptions about agency costs. Based on the notion that large deals are naturally more time consuming than small ones, I compute a lower bound by assuming agency costs are proportional to the value of transactions it evaluates. However, merely initiating a review and/or launching an investigation may consume agency resources, irrespective of deal size. Based on that idea (and with input from FTC staff), I compute an upper bound by assuming half of agency costs are fixed at the merger level (i.e., half of FTC expenses are proportional to the number of deals, while the remainder are proportional to the value). This approach bounds from above because none of the newly reportable mergers—ones that are reportable under the counterfactual policy but not the current one—will incur costs beyond the point of a "Second Request," which is the truly expensive portion of enforcement. (See Section H of the Online Appendix for details.)

7.4 Findings

Total benefits of eliminating premerger notification exemptions in the dialysis industry range from \$3.6 billion to \$4.7 billion. These figures consist of (a) an average of \$117 million in benefits each year for 21 years between 1997 and 2017 and (b) an "annuity" value of benefits that persist after the panel ends, which range from 5.7 to 10.5 times \$214 million. (The range of values depend on whether benefits

⁵⁸This approach implicitly ignores expenses that are essential to agency operations (e.g., staff that handle agency budgeting), although these costs are not likely to vary much with respect to premerger notification program changes, at least at the margin.

⁵⁹HSR Annual Reports tabulate the number of investigated mergers by transaction value. The average investigated merger is 4 times larger than the average reported merger, which presumably reflects more deterrence of small deals.

persist for as short as 7 years or as long as 15 years, and the \$214 million represents benefits realized in the terminal year of the panel, 2017.)

For a transparent look into how the policy generates gains on the order of \$100 million per year, consider the year 2007, which lies at the midpoint of the panel. By that time, the counterfactual policy thwarts—more precisely, it facilitates divestitures for or simply deters—an average of 259 facility acquisitions. These ownership transfers are located in 159 markets housing just under 80,000 patients. The mean difference in quality across facilities under the counterfactual policy relative to the current one is 1.2 percentage points. (See Figure III above for more detail.) If I multiply the patient count by the quality difference, I arrive at 920, which represents the expected number of life-years saved by the counterfactual policy. If I further multiply that product by the value of a life-year, I arrive at \$109 million—very near the average benefits over the panel of \$117 million.

Total costs are much smaller. The bounds on additional costs to the PNO created by the counterfactual policy are [\$45,842, \$2,163,122]. For post-PNO activities, the bounds are given by [\$235,223, \$47,068,592]. The sum is bounded between \$281,065 and \$49.2 million. Note that these figures cover the whole panel—not annual outlays—and do not assume any expenditures are offset by penalties or fees.

Why are costs so low? Partly they reflect the efficiency of the agency. For example, as shown above, the PNO evaluates trillions of dollars of merger activity each year on a budget of less than \$5 million. Considering that legal and advisory fees typically fall around 5-10% of a merger's transaction value, the fact that the FTC conducts effective preliminary investigations for only about 0.0002% of transaction values is striking. Partly they reflect deterrence. An average of three-quarters of mergers that would incur post-PNO costs (were they attempted) are deterred by the counterfactual policy.

7.5 Discussion

The expected benefits of eliminating premerger notifications in the US dialysis industry are two orders of magnitude higher than the expected costs (calculated at the upper bound of the costs). Nonetheless, this finding does not immediately advocate for broadly eliminating premerger notification exemptions, at least on its own. One reason is that expanding reporting requirements also requires expanding agency budgets. Enacting the first change without the second would either draw resources away from other areas of enforcement or force the additional notifications to go unexamined.⁶⁰ Notably, this expansion would come at little or no cost to taxpayers if HSR filing fees were extended to the newly

⁶⁰This assumes no slack in the production function, which is a safe assumption in the author's opinion.

reportable mergers. As evidence, if the current minimum fee \$45,000 was collected from all newly reportable mergers in the counterfactual exercise, then the proceeds cover the upper bound of costs.⁶¹

Another reason is that expanding reporting requirements will also burden firms, which could discourage mergers that improve consumer welfare. Even absent fees, filings are costly to complete.⁶² Beyond disclosing basic transaction attributes, the parties must detail ownership interests, enumerate subsidiaries, report revenues by 6-digit NAICS industry code (and 10-digit NAPCS code if engaged in manufacturing), and transmit "studies, surveys, analyses, and reports" that evaluate the transaction and prepared by/for any officer, board member, investment banker, consultant, or advisor. That said, there is no evidence in the present setting that mergers would actually improve consumer welfare. Irrespective of any change in market structure, acquisitions by large chains are generally associated with postmerger declines in quality [Eliason et al., 2019]. It's worth reiterating at this point that in order to focus on competitive effects, I did not alter the chain dummies in the counterfactual exercise; had I not done so, quality would presumably have improved further, meaning total benefits would have been higher.

Importantly, clever implementation might reduce the costs of expanding reporting requirements in the public as well as private sector. For instance, Scott Morton [2019] argues it may suffice to collect only basic transaction information for small mergers. Further, the agency could capture this information in a way that facilitates transactions screening.⁶³ To illustrate, once the process establishes that the target and acquiror operate dialysis facilities, it may suffice to know the locations of their facilities. Moreover, if this approach alerts the agencies to anticompetitive mergers at a high enough rate, then it will not only facilitate enforcement but also provide deterrence.

Then there is the broader question about how these benefits extend to other settings. Clearly more work is needed to say for sure, but the answer likely depends on several factors. Theory indicates that low entry barriers are likely to attenuate the gains. Mergers among close rivals aren't nearly as costly for consumers when they also attract entrants in a timely manner. That said, slow or expensive entry may characterize many economically important industries. Were it fast and cheap, entry would tend to dissipate corporate profits. Yet, most recent research indicates they are relatively high and increasing

⁶¹Note that extending fees to newly reportable mergers will not automatically increase agency resources. In practice, HSR filing fees are typically thought to merely "offset" appropriations. Thus, planners must explicitly request the funds needed to cover the additional workload posed by the alternative policy.

⁶²See the discussion in De Loecker et al. [2008] at 2.

⁶³See Scott Morton's testimony before the House Judiciary Committee and in particular her discussion of lower thresholds. She points out that most transaction details can be recorded in standardized formats—lines of business reported using "drop-down" menus, locations entered using numeric ZIP codes, etc. This not only simplifies the firms' reporting process but also facilitates the agencies' screening process, automating the earliest stages of the enforcement process.

[Barkai, 2016, De Loecker and Eeckhout, 2017].

Informed, sophisticated consumers are also likely to diminish the gains. If buyers are apprised of problematic ownership changes and aware that they should inform the authorities, then premerger notifications will not be as essential to enforcement. Indeed, in other areas of the healthcare sector the FTC has learned about—and ultimately prosecuted—several exempt mergers, and while the origin of the cases is rarely disclosed, it's reasonable to assume insurers are a frequent source. Moreover, judging by the sheer volume of large law firms' memoranda warning of customer complaints with respect to exempt mergers, the issue is clearly salient to practicing attorneys.⁶⁴ However, data on economy-wide enforcement actions and merger activity indicate that even within the healthcare sector, the dialysis industry does not stand out as particularly susceptible to stealth consolidation.⁶⁵ Furthermore, investigations prompted by customer complaints often occur after the merger has been consummated, making enforcement much more costly. Former FTC Commissioner J. Thomas Rosch estimates that "investigations of consummated mergers at the FTC take on average twice as long to complete as investigations of unconsummated mergers" and goes on to emphasize that the agency "need[s] to challenge and unwind anticompetitive transactions as quickly as possible to minimize consumer injury" [Rosch, 2012]. Thus, even if consumers can facilitate eventual remedies, premerger notifications may be much more efficient.

However, other factors may amplify the benefits. One in particular relates to the HSR Act's "size-of-persons" test. Currently, when neither the target or acquiror have \$18.4 million or more in assets (or in sales, in the event the party is engaged in manufacturing), the transaction value threshold is set at \$376 million, not \$92 million. In many industries the higher transaction value threshold is irrelevant, since targets that are valued at \$92 million or more are practically guaranteed to have more than \$18.4 million in assets (or sales). This is the case in dialysis. However, in other industries the higher threshold may apply often. Most commonly this occurs when the target mainly consists of intellectual property (IP). One reason is that IP is not capitalized under US accounting standards and therefore not an "asset" for the purposes of the US premerger notification program.⁶⁶ Another reason is that these targets are often

⁶⁴Nearly every major law firm publishes guidelines around exempt mergers. (See, e.g., Shearman and Sterling's August 2019 memorandum titled "Risks for Consummated Deals Even Where No Notification Requirements.") Nine of the first 10 I located warned that customer complaints could alert the agencies to mergers that "raise significant antitrust concerns."

⁶⁵See Table 2 of Wollmann [2019], which tallies the number of horizontal exempt mergers by industry in the years following the HSR Amendment, which raised thresholds. For an even cleaner measure of what industries were affected, one can tally the change in mergers per year before and after the threshold was revised. (This information is not reported in the paper.) On this basis, dental laboratories, ambulatory health care services, home health care services, medical devices, hospitals including surgical centers, and pharmaceutical all rank higher than dialysis centers.

⁶⁶To be precise, under Generally Accepted Accounting Principles, research and development is recorded as an expense, not a capital expenditure. Thus, intellectual property is not recorded as an asset on the balance sheet until it is acquired by another firm, at which point it becomes an "intangible."

"pre-revenue"—they have innovative technology that is not yet in the market, often because of delays in regulatory approval or long development cycles. In the healthcare sector, biotechnology companies [Cunningham et al., 2017] and medical device startups are probably most affected .

8 Conclusions

This paper found premerger notifications are essential to antitrust enforcement. In their absence, mergers among close competitors that would otherwise be blocked are, in effect, allowed to proceed—an eventuality I refer to as "stealth consolidation." In the dialysis industry, this equates to facility acquisitions that create local duopolies and even monopolies. Several deals increase HHI 2,000 points or more—an order of magnitude higher than what is "presumed to be anticompetitive" in the Guidelines published by the US antitrust enforcement agencies, the FTC and DOJ. This paper also found that the transactions are, in fact, anticompetitive. In the dialysis industry, prices are fixed, so firms compete on quality. When facility acquisitions consolidate ownership, providers internalize business-stealing effects, and they lower quality as a result. As evidence, hospitalization rates rise while survival rates fall.

These facts raised the question as to whether society would benefit from eliminating premerger notification exemptions in this context. To answer it, I estimated a structural model, which allowed me to predict equilibrium quality under any hypothetical market structure. Next, I estimated enforcement rates among reportable mergers and retroactively applied them to exempt mergers. I found that the counterfactual policy produced less concentrated markets that offered higher quality healthcare. Over the panel, it saved thousands of life-years worth billions of dollars. Finally, I gathered historical data on FTC operations to help forecast the added burden imposed on the agency by the counterfactual policy. I found the added costs are very small in comparison to the expected benefits. Partly this reflects the deterrence effect of law enforcement—under the alternative policy, many harmful mergers are not attempted in the first place, so the agency never expends resources to block them. Partly this reflects the relative efficiency of the agency, at least as viewed through the lens of my analysis.

References

Nikhil Agarwal, Itai Ashlagi, Michael Rees, Paulo Somaini, and Daniel Waldinger. An empirical framework for sequential assignment: The allocation of deceased donor kidneys. 2019.

Antitrust Modernization Commission. Report and Recommendations to Congress and the President (Chapter II: Enforcement Institutions and Processes). Technical report, Antitrust Modernization Commission, Washington, DC, 2007.

David Autor, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. The fall of the labor share and the rise of superstar firms. NBER Working Paper No. 23396, 2017.

William J Baer. Reflections on Twenty Years of Merger Enforcement under the Hart-Scott-Rondino Act. Antitrust Law Journal, 65:825, 1996.

William J Baer. The American Bar Association Clayton Act 100th Anniversary Symposium, 12 2014. Remarks by Assistant Attorney General Bill Baer.

Simcha Barkai. Declining labor and capital shares. Working paper, 2016.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. Econometrica: Journal of the Econometric Society, pages 841–890, 1995.

Steven T. Berry. Estimating discrete-choice models of product differentiation. RAND Journal of Economics, 2(25):242–262, 1994.

Cory Capps, David Dranove, and Christopher Ody. Physician practice consolidation driven by small acquisitions, so antitrust agencies have few tools to intervene. Health Affairs, 36(9):1556–1563, 2017.

Carlos Cardelino. Daily variability of motor vehicle emissions derived from traffic counter data. Journal of the Air Waste Management Association, 48(7):637–645, 2012.

Joseph A Clougherty and Jo Seldeslachts. The deterrence effects of us merger policy instruments. The Journal of Law, Economics, & Organization, 29(5):1114–1144, 2013.

Colleen Cunningham, Florian Ederer, and Song Ma. Killer acquisitions. Working paper, 2017.

David Cutler, Leemore S. Dafny, and Christopher Ody. How does competition impact quality of care? a case study of the u.s. dialysis industry. 2016.

David M Cutler and Elizabeth Richardson. The value of health: 1970-1990. American Economic Review, 88(2):97–100, 1998.

Jan De Loecker and Jan Eeckhout. The rise of market power and the macroeconomic implications. NBER Working Paper No. 23687, 2017.

- Jan De Loecker, Jozef Konings, and Patrick Van Cayseele. Merger review: How much of industry is affected in an international perspective? Journal of Industry, Competition and Trade, 8(1):1–19, 2008.
- Paul Eliason. Market Power in Health Care: Congestion and Spatial Competition in the Dialysis Industry. 2018.
- Paul J Eliason, Benjamin Heebsh, Ryan C McDevitt, and James W Roberts. How Acquisitions Affect Firm Behavior and Performance: Evidence from the Dialysis Industry. The Quarterly Journal of Economics, 135(1):221–267, 11 2019. ISSN 0033-5533.
- Michael WL Elsby, Bart Hobijn, and Ayşegül Şahin. The decline of the US labor share. Brookings Papers on Economic Activity, 2013(2):1–63, 2013.
- Federal Trade Commission. The FTC’s Merger Remedies 2006-2012. Technical report, January 2015.
- Jason Furman and Peter Orszag. A firm-level perspective on the role of rents in the rise in inequality. 2015.
- Martin Gaynor, Kate Ho, and Robert J Town. The industrial organization of health-care markets. Journal of Economic Literature, 53(2):235–284, 2015.
- Ilene K. Gotts. The Merger Review Process. American Bar Association., 2001.
- Germán Gutiérrez and Thomas Philippon. Investment-less growth: An empirical investigation. NBER Working Paper No. 22897, 2016.
- Carla Anderson Hills. Antitrust Adviser 3rd Edition. 1985.
- Loukas Karabarbounis and Brent Neiman. The global decline of the labor share. The Quarterly Journal of Economics, 129(1):61–103, 2013.
- Thomas G. Koch, Brett W. Wendling, and Nathan E. Wilson. The effect of physician group mergers on the health outcomes of medicare beneficiaries. 2020.
- Chris P. Lee, Glenn M. Chertow, and Stefanos A. Zenios. An empiric estimate of the value of life: Updating the renal dialysis cost-effectiveness standard. Journal of the International Society for Pharmacoeconomics and Outcomes Research, 12(1):80–87, 2009.
- Volcker Nocke and Michael D. Whinston. Dynamic merger review. Journal of Political Economy, 118 (6):1200–1251, 2010.

- Volcker Nocke and Michael D. Whinston. Merger policy with merger choice. Journal of the International Society for Pharmacoeconomics and Outcomes Research, 103(2):1006–1033, 2013.
- Maureen Ohlhausen. Interview with ftc commissioner maureen ohlhausen. By Joe Tipograph. The Capitol Forum. Conducted on September 30, 2014., 2014.
- Leslie Overton. Non-reportable transactions and antitrust enforcement. Transcript of remarks at 14th Annual Loyola Antitrust Colloquium, Institute for Consumer Antitrust Studies, Chicago, April 2014, 2014.
- Sam Peltzman. Industrial Concentration under the Rule of Reason. The Journal of Law and Economics, 57(S3):S101–S120, 2014.
- PriceWaterhouseCoopers. A tax on mergers? Surveying the time and costs to business of multi-jurisdictional merger reviews (A study commissioned by the International Bar Association and the American Bar Association). 2003.
- Jo Reynaerts, Ravi Varadhan, and John C. Nash. Enhancing the convergence properties of the blp (1995) contraction mapping. Technical report, 2012.
- J. Thomas Rosch. Consummated merger challenges – the past is never dead. ABA Section of Antitrust Law Spring Meeting, 2012.
- Irving Scher and Scott Martin. Antitrust Adviser 5th Edition. Thomson Reuters, 2015.
- Fiona Scott Morton. Diagnosing the problem: Exploring the effects of consolidation and anticompetitive conduct in health care markets. Testimony before the House Judiciary Committee, 116th United States Congress, 2019.
- US Department of Justice and Federal Trade Commission. Horizontal merger guidelines. Technical report, Washington, DC, 2010.
- United States Renal Data System (USRDS). 2017 USRDS annual data report: Epidemiology of kidney disease in the United States. Technical report, National Institutes of Health, National Institute of Diabetes and Digestive and Kidney Diseases, Bethesda, MD, 2017.
- Thomas G. Wollmann. Stealth consolidation: Evidence from an amendment to the Hart-Scott-Rodino Act. The American Economic Review: Insights, 1(1):pp. 77–94, 2019.

Figures

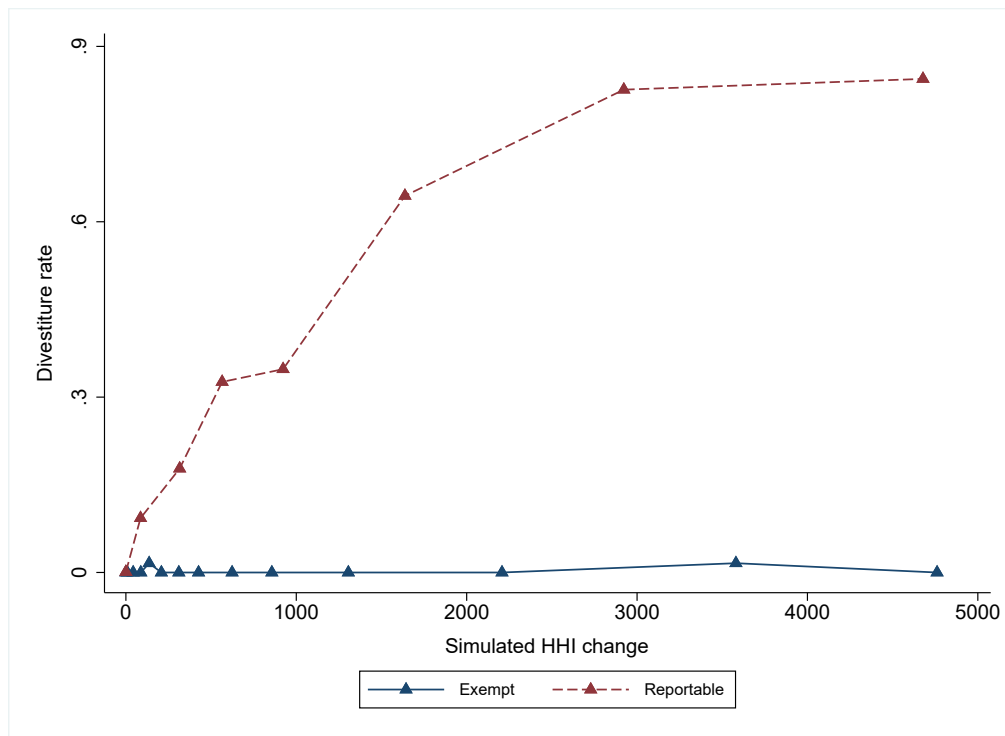


Figure I: The FTC blocks acquisitions of close competitors if—but only if—it is notified in advance.

The x-axis measures simulated HHI changes (ΔHHI), while the y-axis measures divestiture rates ($DIVEST$). Data are sorted into equal-sized bins according to x-axis values, and averages within the bins are reported along the axes. The dashed line reflects reportable facility acquisitions, and the solid line reflects exempt ones.

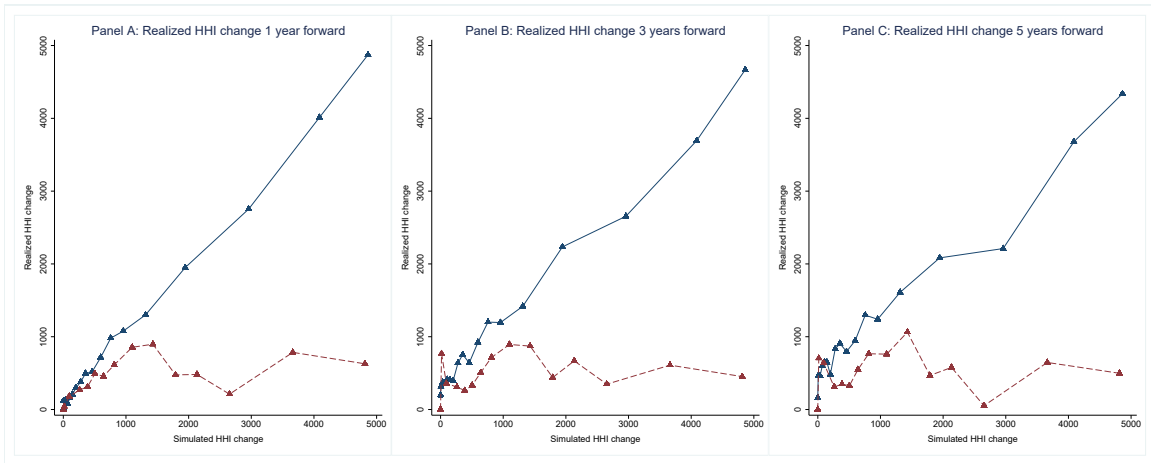


Figure II: Notifications have long-lasting effects on market structure.

The x-axes plot simulated HHI changes, while the y-axes plot realized HHI changes. In each panel, the base year for calculating realized HHI changes is the year prior to the acquisition. In Panel A, realized HHI changes are calculated between the base year and one year later. In Panels B and C, they are calculated between the base year and 3 and 5 years later, respectively.

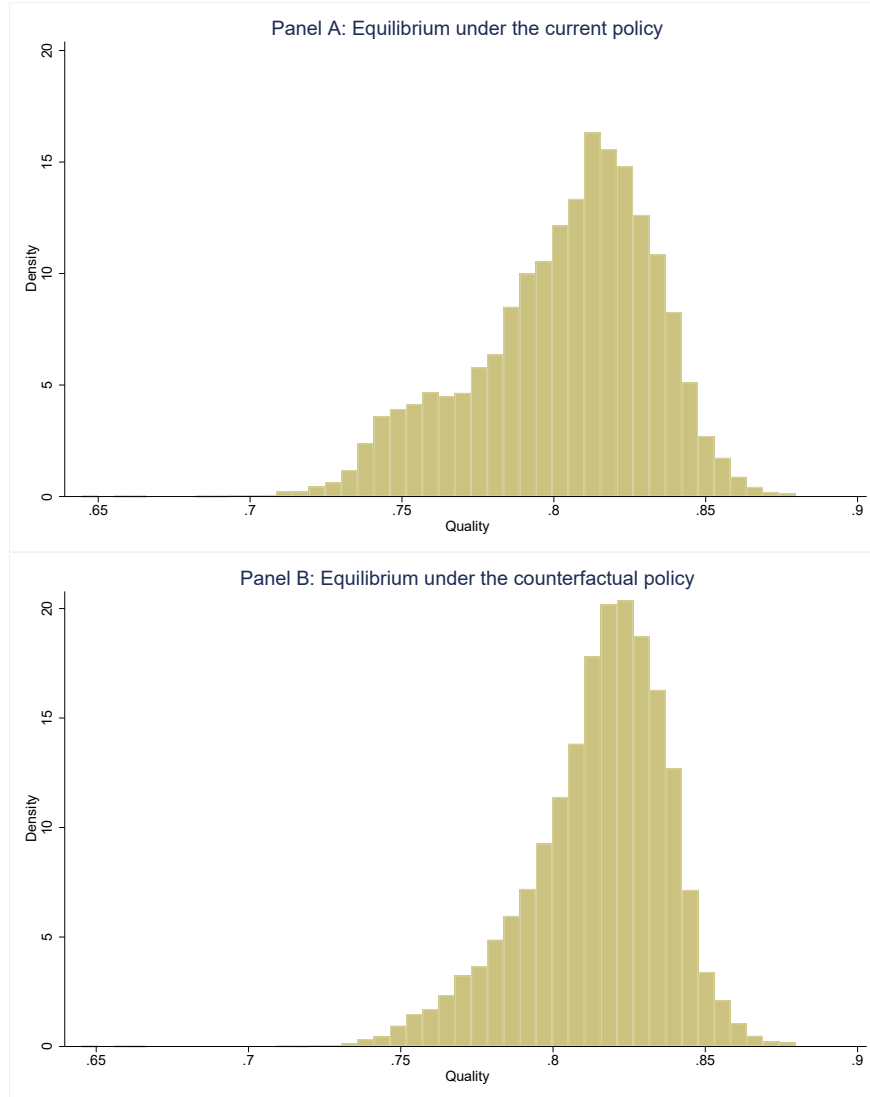


Figure III: Equilibrium quality is higher under the counterfactual policy.

The unit of observation is a facility-year. The sample comprises all facilities in all geographic areas that have witnessed at least one exempt facility acquisition involving a large simulated HHI change. Panel A reports equilibrium quality under the current policy. Panel B reports expected equilibrium quality under the counterfactual policy, where the expectation is taken over possible market structures, r . (Formally, Panel A plots the distribution of x_{jgt}^r , while Panel B plots the distribution of $\sum_{r \in \mathcal{R}_{gt}} \psi(r) x_{jgt}^r$.)




Figure IV: Merger enforcement costs and merger activity over time.

Panels A and B report PNO and Post-PNO merger costs, respectively. Panel C and D report the total value of reportable mergers and the total value of investigated mergers, respectively. All figures are in constant 2017 dollars.

Tables

Table I: *Summary of facilities.*

	N	Mean	Std. Dev.	Minimum	Maximum
Patients	99635	39.3	27.9	1	274
Hospitalization rate	88823	0.69	0.18	0.24	1.40
Survival rate	82472	0.85	0.093	0.50	1.05
Machines	98985	17.8	7.93	0	174
HHI	99635	5608.3	2907.1	554.5	10000
Chain ownership	99635	0.74	0.44	0	1

Table II: *Summary of facility acquisitions.*

	N	Mean	Std. Dev.	Minimum	Maximum
Reportable	4326	0.53	0.50	0	1
Exempt	4326	0.47	0.50	0	1
Divested	4326	0.062	0.24	0	1
Acquisition year	4326	2005.6	5.63	1997	2017
Acquisition year (Reportable only)	2310	2005.8	5.27	1997	2017
Acquisition year (Exempt only)	2016	2005.5	6.02	1997	2017
Change in simulated HHI	4326	490.7	1081.2	0	4999.9
Change in simulated HHI (Reportable only)	2310	503.1	1070.4	0	4999.2
Change in simulated HHI (Exempt only)	2016	476.6	1093.5	0	4999.9
Change in HHI realized	4326	334.7	1114.1	-5714.3	6347.8
Change in HHI realized (Reportable only)	2310	149.6	805.8	-4898.0	4977.8
Change in HHI realized (Exempt only)	2016	546.9	1354.8	-5714.3	6347.8

Table III: *Effects on hospitalization and survival rates.*

VARIABLES	(1) Hospital'n rate	(2) Survival rate	(3) Hospital'n rate	(4) Survival rate
POST × EXEMPT × ΔHHI*	0.036*** (0.0088)	-0.018*** (0.0065)		
POST × EXEMPT × ΔHHI			0.0081** (0.0031)	-0.0047** (0.0023)
POST × ΔHHI*	-0.0059 (0.0055)	0.0046 (0.0040)		
POST × ΔHHI			-0.0020 (0.0022)	-0.00046 (0.0016)
POST × EXEMPT	-0.0075 (0.0058)	0.0019 (0.0044)	-0.0037 (0.0056)	0.000089 (0.0043)
POST	0.0020 (0.0040)	0.00066 (0.0034)	0.0016 (0.0039)	0.0023 (0.0033)
Observations	52,762	47,990	52,762	47,990
R-squared	0.392	0.460	0.392	0.460
Facility acquisition FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Chain FE	Yes	Yes	Yes	Yes

, **, and * denote significance at the 10%, 5%, and 1% levels, respectively. In odd numbered columns, quality is measured by hospitalization rates; in even numbered ones, it's measured by survival rates. All specifications include facility acquisition, year, and chain fixed effects. Standard errors are clustered at the facility acquisition level.*

Table IV: Demand and marginal cost estimates.

Panel A. Demand-side parameters		
Variables	Parameters	Estimates
Distance	α_1	-0.1520 (0.0001923)
Distance ²	α_2	-0.00077589 (0.00009217)
Distance \times Population density	α_3	-0.4364 (0.00052885)
Distance ² \times Population density	α_4	0.01005 (0.0002985)
Quality	β_x	6.66899 (3.33062)
Machines	β_n	0.01708 (0.00311)
Constant	β_0	-3.0537 (2.5573)
Year fixed effects	ϕ_t	Yes
Chain fixed effects	χ^c	Yes
Number of patient-year observations		2,885,565
Number of facility-year observations		61,576
Panel B. Cost-side parameters		
Variables	Parameters	Estimates
Constant	γ_0	990.4099 (392.5425)
Machines	γ_1	1.07727 (0.48819)
Quality	γ_2	-2428.8486 (1088.1939)
Quality ²	γ_3	1661.8535 (719.0894)
Number of facility-year observations		61,576

Distance is measured in miles. Population density is measured in 10,000s of people per square mile. Robust standard errors are reported in parenthesis directly below coefficients. Marginal costs are reported on a per-treatment basis.

Appendix

A Comparing FTC and CMS facility tabulations

To verify that the Medicare and FTC data are congruent with one another, I compare the (a) number of target and acquiror facilities reported in FTC complaints and (b) the number of divested facilities reported in FTC orders with tabulations from Medicare. Tables H.1 and H.2 report firm-by-merger comparisons.

[Table V about here.]

[Table VI about here.]

The data sources closely align with one another. For example, when Fresenius acquired Renal Care Group in 2005, the FTC reported that the acquiror and target owned 1,155 and 450 facilities, respectively, and that 103 facilities would be divested to National Renal Institute. Medicare data indicates that in 2005 the acquiror and target owned 1,105 and 449 facilities, respectively, and that 103 would be divested. Discrepancies between the sources reflect differences in the timing of the reports. FTC ownership snapshots are taken at the time of the merger, while CMS tabulations are taken at year's end.

B Determining HSR exemptions

Inferring exemptions from facility counts

I take the following 6 steps to determine the reportability of mergers.

1. Gather data from 35 mergers for which the transaction value is publicly disclosed.

(These mergers involve 2,500 facility acquisitions.)

2. Translate nominal transaction values into real ones.

3. Plot real transaction values against the number of facilities involved in the merger.

(The motivation for this approach is briefly described in the main text. The result is reported below.)

4. Impute remaining transaction values using an average price-per-facility.

5. Translate real transaction values back into real ones.

6. Compare real transaction values to thresholds:

mergers are reportable if and only if the former exceed the latter.

Figure H.1 reports the result of step 3. Consistent with trade press that indicates dialysis providers are valued on a per-facility basis, the relationship is highly linear. Notably, the distribution of per-facility values are so tight that even if I construct two sets of predict transaction values using the upper and lower bounds of the interquartile range, I predict the same exemption status for over 99% of facility acquisitions.

[Figure V about here.]

C Distribution of simulated HHI changes

Panels A and B of Figure H.2 plots the frequency distribution of ΔHHI for reportable and exempt facility acquisitions, respectively. The two distributions are fairly similar. Both report a large point mass at $\Delta\text{HHI} = 0$, reflecting the fact that facility acquisitions involve no target-acquiror overlap prior to the merger.

[Figure VI about here.]

The mass around zero makes it hard to read off the frequency of facility acquisitions with large simulated HHI changes, so Figure H.3 replicates the prior one but omits facility acquisitions for which $\Delta\text{HHI} \leq 0$.

[Figure VII about here.]

D Observed distances traveled by patients

Figure H.4 reports the distribution of the distances traveled by patients to facilities. (I assume that trips begin at home and that distances between locations are approximated by distances between zip code centroids.) It shows that almost all patients travel less than 20 miles to their provider.

[Figure VIII about here.]

E Market definitions

To map facilities to geographic markets, I start with location information from the Provider of Service files, which report 5-digit zip code, city, county, and state. Then, I correct typographical errors in city names. For example, "Co Bluffs" means "Counsel Bluffs," and "Atl" means "Atlanta." Next, I address discrepancies between zip codes and city names. Comparing the CMS information with provider websites indicates zip codes are more accurate, so I replace city names with the modal city name in the zip code. I also address discrepancies between county codes and city names. For example, Hollywood and North Hollywood are not and have never been in Fresno County, CA. Here, comparing the CMS information with provider websites indicates city names are more accurate, so I replace county codes with the modal county code in the city-state pair. Note that some towns/cities do not correspond to any county, mainly due to idiosyncrasies in state-specific systems of local government. For instance, all Virginia cities are formally distinct from counties since the Civil War. When an unassigned town or city is entirely surrounded by a county, I assign it to the surrounding county, and when it is not surrounded, I assign it to the closest county.⁶⁷ Typographical errors, zip-city/city-county discrepancies, and unassigned towns/cities are few, so these are not meaningful restrictions.

Market definitions depend on county borders. However, three counties are far too large to be treated as single markets: Los Angeles (CA) Cook (IL), and San Diego (CA). I asked my faculty assistant, James Kiselik, to divide Los Angeles County and San Diego County, and he enlisted the help of a friend who grew up in Southern California.⁶⁸ I divide Cook County simply by placing Evanston, Orland Park, and Elk Grove Village in their own markets, since they are far from downtown. Note that these divisions do not in any way drive the results. For example, dropping these markets entirely has an almost imperceptible effect on the descriptive evidence reported in Section 4. For instance, it increases the estimate of θ_3 , the main coefficient of interest in equation 1, from 0.250 to 0.252, and it decreases the estimate of λ_4 , the main coefficient of interest in equation 2, from 0.357 to 0.350 (where y is measured by hospitalization rates). Also, note that while I did not test the effect of dropping these markets from the structural estimation, see the immediately below for a discussion as to why structural estimates are even less likely to be effected.

To be clear, the decision to equate counties to markets in the manner described was, at least initially, haphazard. However, subsequent analysis strongly supports the choice. First, the FTC frequently

⁶⁷To be specific, I calculate the distance between the centroid of the zip codes in the unassigned town/city and the surrounding counties. I then calculate the mean distance by county. The "closest" county has the minimum mean distance.

⁶⁸I handed this off partly because I am unfamiliar with the area and partly because I wanted to stay out of the process as much as possible. I was sufficiently familiar with Cook County that I decided to simply split this myself.

defines geographic markets using county borders in their divestiture orders. These documents are drafted by the agency, but they typically reflect negotiation with the merging parties, which means county-based market definitions are likely to reflect some consensus between enforcement officials and dialysis providers about how "locally" facilities compete. Second, I assess the robustness of Section 4 estimates by redrawing markets entirely, circumscribing a circle around each acquired facility. I report the results in Panel A of Figure H.5 and Columns 1-2 of Table H.4, which are similar to the results reported in the body of the main text. Hence, the descriptive evidence derived from my market definitions resembles descriptive evidence derived from sensibly constructed distance-based market definitions. Third and perhaps most importantly, I let the distance from patient residences to facility locations directly enter the utility function, so the model endogenously determines how intensely the facilities compete with one another for each patient. Thus, accurate structural estimates and, by extension, counterfactual policy predictions do not require precise market definitions. They only require markets to be drawn large enough so as to not separate patients and the facilities they might visit, conditional on distance. This condition is satisfied in the data. USRDS data reveals that the vast majority—over 90%—of patients visit facilities in the market in which they reside. Moreover, many (or even most) exceptions are clear errors due to outdated residence information. An example would be a patient who, according to the data, (a) originally resides in Pennsylvania and visit dialysis facilities nearby, (b) starts visiting dialysis facilities in Indiana, (c) relocates to Indiana a year or two later.

F Supplementary analysis of effects on enforcement

I consider various alternatives to the specification that derives Figure I. First, I redefine markets using a 20 mile radius around each facility. Second, I recalculate HHI using machine rather than patient counts. Later in the body of the main text (i.e., subsection 4.3), my results implicitly reflect a sample of the data for which I observe at least one premerger and postmerger measure of hospitalization and survival rates (separately). This excludes a relatively small number of acquisitions—mostly ones that occur in the first and last years of the data. Nonetheless, I explore robustness to these restrictions as well. I narrow the sample to facility acquisitions for which I observe at least one premerger and a postmerger measure of hospitalization rates and survival rates separately. Figure H.5 report the results, which mirror those of Figure I.

[Figure IX about here.]

Panel A of Figure H.6 addresses the question of whether the FTC directs enforcement actions towards large mergers. Here, size is measured by the number of facilities operated by the target just prior to the merger, which very closely relates to the transaction value of the merger. I limit the sample to reportable mergers, regress $DIVEST_a$ on ΔHHI_a , and plot the residual against the number of facilities the target owns just prior to the merger. There is no meaningful relationship between enforcement and target size. Panel B addresses the related question of whether the FTC directs enforcement actions towards "big business" (i.e., major chains). To answer this, I limit the sample to reportable mergers, regress $DIVEST_a$ on ΔHHI_a , and plot the residual against the number of facilities the acquiror owns just prior to the merger. Panel B of Figure H.6 reports the result. Again, there's no meaningful relationship between enforcement and target size.

[Figure X about here.]

Column 1 of Table H.3 reports the result of estimating equation 1. For column 2, I additionally include simulated HHI levels, denoted HHI^* , and their interaction with `REPORTABLE`. Neither one predicts divestitures, and their inclusion does not meaningfully affect the main coefficients. For columns 3-7, respectively, I include the number of target and acquiror facilities, chain fixed effects, acquisition year fixed effects, and all four controls together. The coefficients are stable across the specifications. For example, estimates of θ_3 —the coefficient on `REPORTABLE` \times ΔHHI^* —moves 20% or less in either direction across the various specifications.

[Table VII about here.]

G Supplementary analysis of effects on quality

To assess the robustness of the main quality-related result, I consider various alternatives to the specification that derives Figure III. First, I redefine markets using a 20 mile radii around each facility. Second, I calculate HHI using machine counts rather than patient counts. Third, I narrow the observation window around each acquisition from 10 years to 5 and 7 years (separately). Table H.4 reports the results, which are similar in size and significance to those reported in Table III.

[Table VIII about here.]

Acquisitions of rival facilities affect the incentives to provide quality at facilities owned by acquirors as well. Thus, to further evaluate the results, I examine the evolution of quality at acquiror facilities

located in the same markets as target facilities. The goal is to recover a parameter analogous to those in equation 2 and compare it to preceding estimates. To do so, I first map facility acquisitions, a , into acquiror-market level observations, (f, g, t) .⁶⁹ Next, I map these acquiror-market level observations back onto facilities owned by the acquiror in the year of the acquisition, each of which are indexed by a' . Finally, I construct a panel using an up to 10 year observation window around each a' . To illustrate, consider a market (g, t) where firm f_1 owns facilities h and i , where firm f_2 owns facilities j and k , and f_1 acquires f_2 . In this instance, a corresponds to (j, g, t) and (k, g, t) and a' corresponds to (h, g, t) and (i, g, t) . I map j and k into (f_1, g, t) , map (f_1, g, t) onto h and i , and record the quality offered by h and i on $[t - 10, t - 1]$ and $[t + 1, t + 10]$ (where available).⁷⁰

Note that there are no acquiror-related observations involving $\Delta\text{HHI}_a = 0$. Moreover, most facility acquisitions for which $\Delta\text{HHI}_a^* = 0$ involve no overlap at all. Thus, to study acquiror facilities, I narrow the sample to $\Delta\text{HHI}_a^* = 1$. The estimating equation is given by

$$y_{a't} = \tilde{\lambda}_1\text{POST}_t + \tilde{\lambda}_2\text{POST}_t\text{EXEMPT}_{a'} + \tilde{\mu}_{a'} + \tilde{\kappa}_t + \tilde{v}_{a't} \quad (14)$$

for a' such that $\Delta\text{HHI}_{a'}^* = 1$. $\tilde{\mu}_{a'}$ and $\tilde{\kappa}_t$ represent fixed effects specific to a' and t , respectively. $\tilde{\lambda}_1$ is analogous to $\lambda_1 + \lambda_3$ in equation 2, while $\tilde{\lambda}_2$ is analogous to $\lambda_2 + \lambda_4$.

Table H.5 reports the results of estimating equation 14, which shows that exempt facility acquisitions involving large simulated HHI changes reduce quality at acquiring facilities. In fact, the estimates are surprisingly close to previously reported coefficients. When health outcomes are measured with hospitalization rates, Table H.4 reported $\lambda_1 + \lambda_3$ is near zero while $\lambda_1 + \lambda_3$ equals 2.2 percentage points. In this exercise, I find the analogous estimates are approximately zero and 2.0 percentage points, respectively. When health outcomes are measured with survival rates, Table H.4 reported $\lambda_1 + \lambda_3$ is, again, near zero while $\lambda_1 + \lambda_3$ equals 1.66 percentage points. In this exercise, I find the analogous estimates are approximately zero and 1.30 percentage points, respectively.

[Table IX about here.]

Long run trends in hospitalization and survival rates may confound the main result. To explore this

⁶⁹ g denotes the geographic area and t denotes the year, as in Sections 5-7.

⁷⁰In rare cases, acquiror-market level observations reflect two mergers. For example, in year t a firm might acquire facility j through one merger and facility k through another merger, where j and k are located in the same geographic area g . This does not affect panel construction, but it does create multiple values of EXEMPT_a and ΔHHI_a^* at the acquiror-market level. In these cases, I use the maximum values of the indicator. Incidentally, using the minimum produces similar results.

possibility, I plot the main coefficient of interest in event time. The estimating equation is given by

$$y_{at} = \sum_{\tau=-10}^{10} \lambda_{1,\tau} \text{POST}_t \{ \tau \neq -1 \} + \sum_{\tau=-10}^{10} \lambda_{2,\tau} \text{POST}_t \text{EXEMPT}_a \{ \tau \neq -1 \} + \sum_{\tau=-10}^{10} \lambda_{3,\tau} \text{POST}_t \Delta \text{HHI}_a^* \{ \tau \neq -1 \} + \sum_{\tau=-10}^{10} \lambda_{4,\tau} \text{POST}_a \text{EXEMPT}_a \Delta \text{HHI}_a^* \{ \tau \neq -1 \} + \mu_a + \kappa_t + \zeta_{at}^c + \nu_{at}, \quad (15)$$

where λ parameters are exactly zero at $\tau = -1$, the year in event time that directly precedes the facility acquisition.

Figure H.7 reports the result of estimating equation 15. Left-hand side panels measure outcomes using hospitalization rates, while right-hand side panels use survival rates. Extreme values of τ are based on relatively fewer observations, resulting in much larger standard errors. Thus, to better assess trends in the parameter of interest, I re-estimate equation 15 over $[-7, 7]$ rather than over $[10, 10]$. There is no obvious trend in $\hat{\lambda}_4$ leading up to the facility acquisitions, indicating that long run trends are not driving these results.

[Figure XI about here.]

In Subsection 4.3, I highlight that one of the specifications in Eliason et al. [2019] (hereafter EHMR) indicates little if any relationship between competition and quality. Although it is not central to their paper, the result conflicts with estimates I report in Table III. There are two main reasons for this discrepancy, which is easy to reconcile.⁷¹ First, whereas I typically base market definitions on county borders, the EHMR specification in question bases market definitions on health service areas, which are very large and usually extend far beyond the distance any patient travels for care. Consistent with this fact, other specifications in EHMR that rely on market definitions produce estimates much closer to those reported in Table III. For example, when EHMR redraw markets using 10 mile radii around acquired facilities, they find that acquisitions increase hospitalization rates by 1.53 percentage points. (See their Online Appendix I5.)

Second, whereas I measure preacquisition competition using simulated HHI changes, the EHMR specification in question measures competition as any overlap at all between the target and acquiror.

⁷¹I follow Eliason [2018] in measuring quality using hospitalization and survival rates. To study the transfer of management practices, Eliason et al. [2019] study other measures as well, some of which impact provider profitability but do not affect competition for patients. For example, the authors show how chain ownership abruptly changes the administration of anemia medications. As a clear illustration, EHMR point out that the two leading drugs, Ferrlecit and Venofer, are "essentially substitutable" but for idiosyncrasies in how they are reimbursed, which makes the latter much more profitable to administer. Dispensing these drugs to ESRD patients consistently ranks among Medicare's largest prescription drug expenditures, so their relevance cannot be understated. Since the drugs are substitutable, switching between Ferrlecit and Venofer only affects reimbursements, leaving quality unaffected. Note, however, that EHMR also provide evidence that large chains over-prescribe these drugs, which can adversely affect the heart. If this is salient to patients, it would certainly affect quality, broadly defined.

Right skew in the distribution of concentration changes means that their definition includes many deals that amount to trivial HHI increases, especially under geographically broad market definitions. To see the effect of this change, I estimate the effect of exempt facility acquisitions using their competition measure (i.e., that simulated HHI changes are non-zero) and compare it to estimates that obtain under my competition measure. The latter are nearly 50% higher.

H Counterfactual policy

Predicting the evolution of market structure

To determine the evolution of market structure under the counterfactual policy, I need to estimate the relationship between simulated HHI changes and divestiture rates among reportable facility acquisitions, which I then retroactively apply to exempt facility acquisitions. Figure H.8 shows that a second order polynomial provides a good approximation to the shape of the relationship between simulated HHI changes and divestiture rates.

[Figure XII about here.]

The estimating equation associated with this functional form is given by

$$\text{DIVEST}_a = \theta_0 + \theta_1 + \Delta\text{HHI}_a + \theta_2\Delta\text{HHI}_a^2 + \eta_a. \quad (16)$$

This yields

$$\widehat{\text{DIVEST}}_a = 0.00508 (0.0491) + 0.480 (0.0520) \Delta\text{HHI}_a - 0.0635 (0.0115) \Delta\text{HHI}_a^2, \quad (17)$$

where standard errors are clustered at the merger-market level and reported in parentheses. Note that in practice I set $\theta_0 \equiv 0$ when calculating enforcement rates. This has no effect on the counterfactual exercise except to significantly reduce the computational burden.

Deterrence

Deterrence is an importance part of antitrust enforcement. Firms will not suffer large up-front managerial, legal, and accounting costs unless they are reasonably certain the transaction will close. In the counterfactual exercise, I assume that no firm proposes a merger where the average probability of

divestiture across target facilities exceeds 50%. In the associated footnote, I then claim that the current cutoff is arbitrarily conservative, as the true cutoff is probably below 10% (i.e., the probably the merger "succeeds" is greater than 90%).

To arrive at this conclusion, I begin with the fact that a firm proposes a merger if and only if $Expected\ benefits \times Pr.(Success) \geq Up\text{-}front\ costs$. The least favorable case for deterrence is where this condition holds exactly (i.e., $Pr.(Success)$ is as low as possible). To obtain a reasonable estimate of break-even $Pr.(Success)$, I turn to historical enforcement data. Over the sample period, about 1% of reportable mergers were abandoned or faced any enforcement action. To be as unfavorable to deterrence as possible, I assume that only horizontal mergers face enforcement/abandonment. Over the sample period, about one-third of reportable mergers were horizontal (i.e., the target and acquiror operated in the same SIC2 or NAICS3 industry). Thus, $(33\%-1\%)/33\% = 97\%$ of horizontal reportable mergers were successful. If mergers are uniformly distributed with respect to their probability of success, then the cutoff value of $Pr.(Success)=94\%$.⁷²

Calculating policy benefits

Benefits of the policy are given by

$$\sum_{t=1997}^{2017} \Lambda_t \times \$100,000 \times CPI_{2017,t} + \Lambda_{2017} \times \$100,000 \times \frac{1}{\Theta - \Gamma} \left(1 - \left(\frac{1+\Gamma}{1+\Theta} \right)^{YEARS} \right). \quad (18)$$

The first term represents real benefits over the panel, while the second term represents benefits that accrue after the panel ends. Λ_t denotes the expected number of lives saved in year t , which equal $\sum_g \sum_j \sum_{r \in \mathcal{R}_{gt}} M_{gt} x_{jgt}^r s_{jgt}^r \psi(r) / R_{gt} - \sum_g \sum_j M_{gt} x_{jgt}^{\bar{r}} s_{jgt}^{\bar{r}}$, where $x_{jgt}^{\bar{r}}$ denotes the predicted equilibrium quality chosen by j in g at t under market structure \bar{r} , x_{jgt}^r denotes an analogous object under market structure r , \mathcal{R}_{gt} denotes the set of r at (g, t) , and R_{gt} counts the number of elements in \mathcal{R}_{gt} .

$CPI_{2017,t}$ denotes the CPI adjustment factor that translates nominal currency into constant 2017 dollars. $\frac{1}{\Theta - \Gamma} \left(1 - \left(\frac{1+\Gamma}{1+\Theta} \right)^{YEARS} \right)$ denotes the annuity factor, where Γ denotes the growth rate (1%), Θ denotes the real discount rate (4%), and YEARS denotes the number of years that benefits continue past the end of the panel.

⁷²Note that this assumes the cost-benefit-enforcement relationship is similar for historical reportable mergers and newly reportable mergers. The latter are smaller by definition. Thus, if there are fixed costs at the merger level, then making this comparison is even more unfavorable to deterrence in my context.

Calculating policy costs

Bounds on PNO costs are given by

$$\left[\sum_{t=1997}^{2017} \Phi^{PNO} \Psi_t^{Dialysis, All}, \sum_{t=1997}^{2017} \Phi^{PNO} \Psi_t^{Dialysis, All} / 2 + \sum_{t=1997}^{2017} \Xi^{PNO} \Omega_t^{Dialysis, All} / 2 \right]. \quad (19)$$

Φ^{PNO} denotes the real cost to the FTC of the PNO reviewing \$1 of real transaction value. The numerator of Φ^{PNO} equals the sum over the years from 1997 to 2017 of the product of (a) the average salary of a full-time FTC staff attorney (code GS905 in the US Office of Personnel Management), (b) 1.462, which inflates base salary to total compensation (from the US Bureau of Labor Statistics), (c) a CPI adjustment factor that translates nominal currency into constant 2017 dollars, and (d) the number of full-time employees in the PNO. The denominator of Φ^{PNO} equals the sum over the years from 1997 to 2017 of the value of all reportable US mergers. $\Psi_t^{Dialysis, All}$ denotes the real transaction value of dialysis mergers at t that are reported only due to the counterfactual policy. The real transaction value of a dialysis merger equals the product of the number of facilities it involves and the real transaction value per facility, which is discussed in Subsection B. Ξ^{PNO} denotes the real cost of the PNO reviewing one additional merger. The numerator of Ξ^{PNO} equals the numerator of Φ^{PNO} . The denominator of Ξ^{PNO} equals the number of reportable US mergers from 1997 to 2017. $\Omega_t^{Dialysis, All}$ denotes the number of additional dialysis at t that are reported only due to the counterfactual policy.

Bounds on Post-PNO costs are constructed in exactly the same way, except for the following differences. First, in the numerator of Φ , I replace the number of full-time employees in the PNO with the number of full-time employees assigned to Post-PNO merger-related activities. Second, in the denominator of Φ , I replace the value of all reportable US mergers with the value of all investigated US mergers. This amount is not directly reported, so I assume that it equals the product of (a) the total value of reportable US mergers and (b) the ratio of FTC investigations ("clearances") to all reportable US mergers over the panel. HSR Annual Reports indicate that larger mergers face higher enforcement rates (unconditionally), so if this assumption is wrong, it means I overstate the costs of evaluating additional mergers. Third, in the denominator of Ξ , I replace the number of reportable US mergers with the number of investigated US mergers. Fourth, for Ψ , I replace the real transaction value of dialysis mergers that are reported only due to the counterfactual policy with the expected real transaction value of dialysis mergers that are investigated only due to the counterfactual policy. Fifth, in Ω , I replace the number of dialysis mergers that are reported only due to the counterfactual policy with the expected

number of dialysis mergers that are investigated only due to the counterfactual policy.

In the body of the main text, I assert that by assuming half of agency costs are proportional to the number of deals, with the remainder proportional to the value of those deals, I produce an upper bound for the burden imposed by the counterfactual policy. The principal reason relates to the filing of the Request for Additional Information and Documentary Material (a.k.a. "Second Request"). In the merger review process, if the initial antitrust assessment by the PNO indicates competitive concerns, the agency—more specifically, the litigation division or a regional office—will conduct initial investigation, and if concerns persist, the agency will issue a Second Request. All of the truly "expensive" work one associates with enforcement occurs after this point—the receipt and review of thousands if not millions of documents, the drafting of the formal complaint, the negotiation of remedies and/or litigation Gotts [2001]. (One FTC staff member stated to me that approximately two-thirds of all post-PNO expenses are incurred *after* the Second Request is issued.) Importantly, it is safe to assume that no merger that is reportable only under the counterfactual policy would get to that point, since the mere issuance of a Second Request would cause the parties to abandon the transaction. The reason is that the cost of complying with a Second Request would far exceed the benefits from any small merger, especially were it subject to antitrust scrutiny. Reports indicate that responding typically costs \$5 million and take six months, with some deals costing \$20 million and taking 18 months [PriceWaterhouseCoopers, 2003, Antitrust Modernization Commission, 2007]. There is even evidence to directly support this claim. Prior to the abrupt increase in HSR thresholds in 2001, mergers were frequently abandoned upon receipt of a Second Request. In 1999, for example, 5 of 45 transactions receiving Second Requests were immediately terminated. (See the 1999 FTC *Congressional Budget Justification*.)

Figures

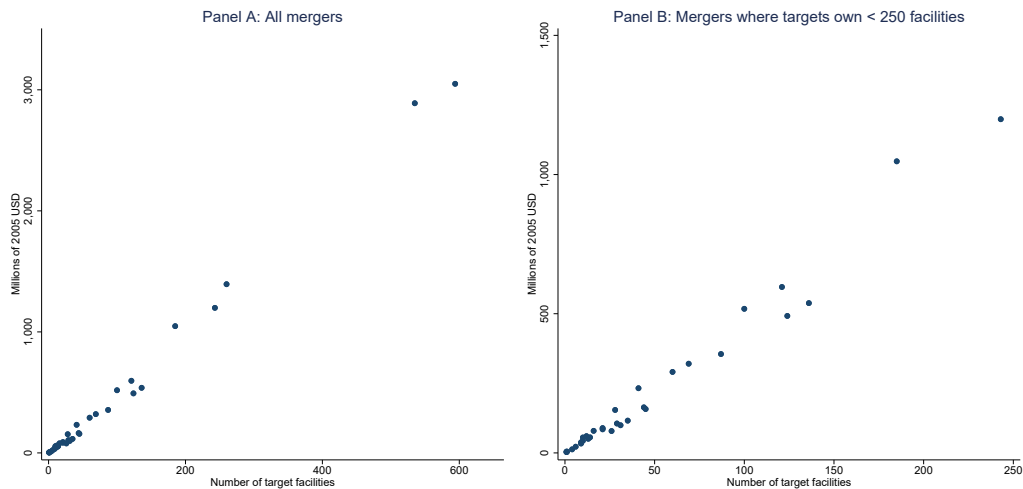


Figure H.1: Real transaction values are tightly linked to the number of target facilities.

The unit of observation is a merger. The x-axis measures the number of facilities the target operated just prior to the merger. The y-axis measures the transaction value in constant 2005 dollars. (Other figures in the paper are presented in constant 2017 dollars; I use constant 2005 dollars here to facilitate comparisons to the exemption threshold, which stood at exactly \$50 million that year.)

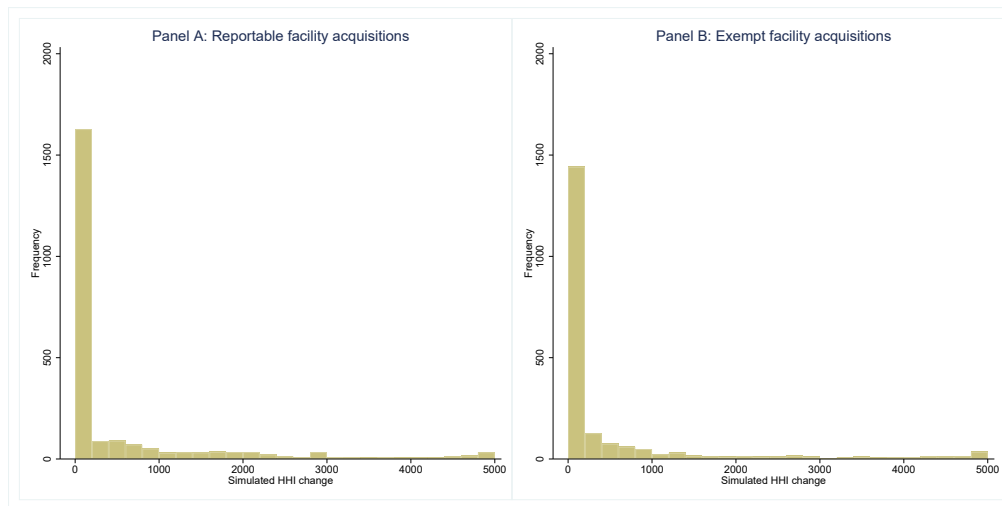


Figure H.2: Simulated HHI changes of reportable and exempt facility acquisitions.

This figure plots the frequency of facility acquisitions by simulated HHI change, which is measured on the x-axis. Panel A reflects reportability facility acquisitions, and Panel B reflects exempt ones.

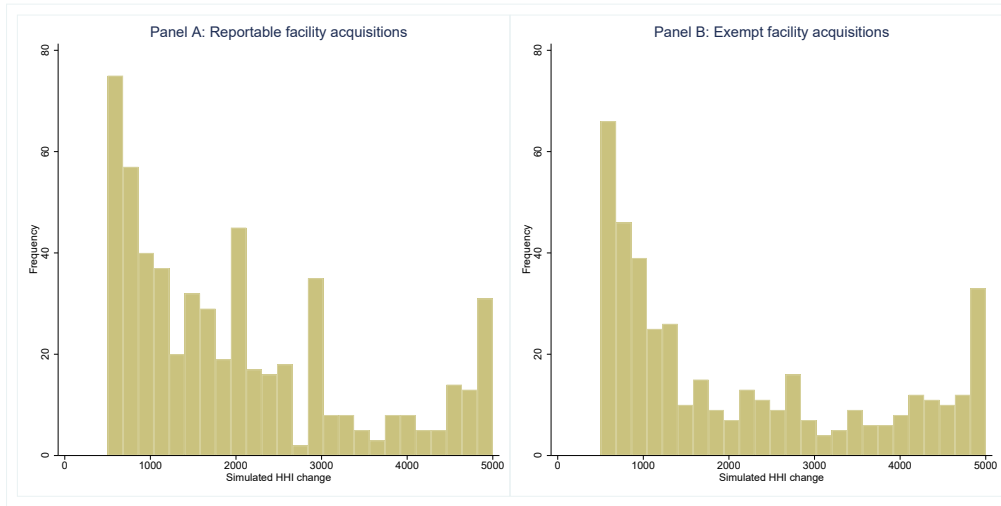


Figure H.3: Simulated HHI changes above 500.

For the sake of readability, this figure replicates Figure H.2 but omits facility acquisitions involving simulated HHI changes less than or equal to 500 points.




Figure H.4: Patients visit facilities close to home.

This graph is a histogram of the distances from the centroids of the zip codes of patients' residences to the centroids of the zip codes of the facilities they visit. The unit of observation is the patient-year.

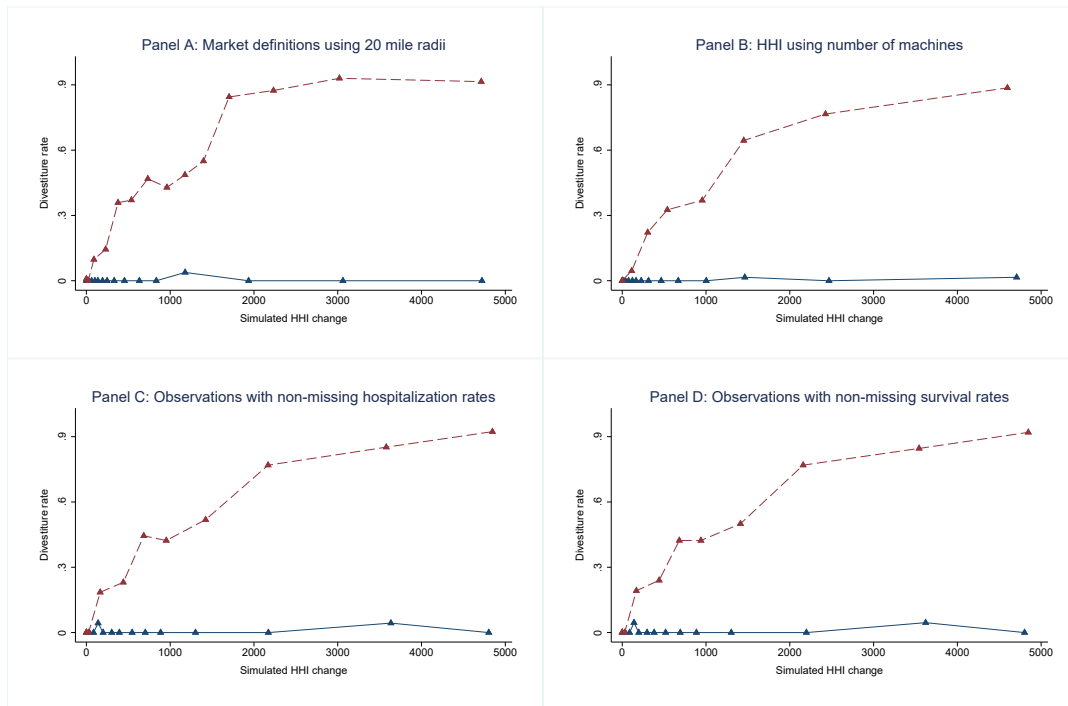


Figure H.5: Enforcement-related results are robust to alternate specifications.

Panel A redefines markets using a 20-mile radius around each target facility. Panel B uses machines rather than patients to compute market shares. Panel C restricts the sample to acquisitions for which facilities have non-missing hospitalization rates in the three years before and after the acquisitions. Panel D is analogous to Panel E except that it restricts the sample to non-missing survival rates rather than hospitalization rates.

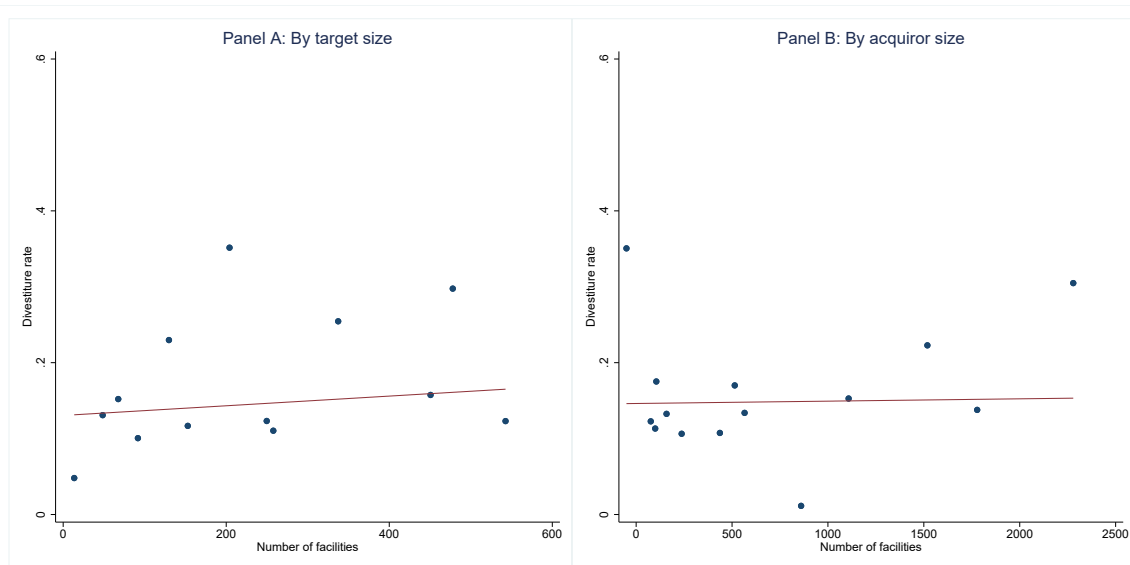


Figure H.6: Conditional on premerger market structure, the FTC does not target large mergers or acquirors.

In Panel A, the x-axis measures the number of facilities owned by the target just prior to the merger, while in Panel B, it measures the number of facilities owned by the acquiror just prior to the merger. The y-axis measures enforcement rates after controlling for premerger market structure (i.e., I regress $DIVEST_a$ on ΔHHI and plot the residual along this dimension). The data are sorted into equal-sized bins according to x-axis values, and averages within the bins are reported along the axes.

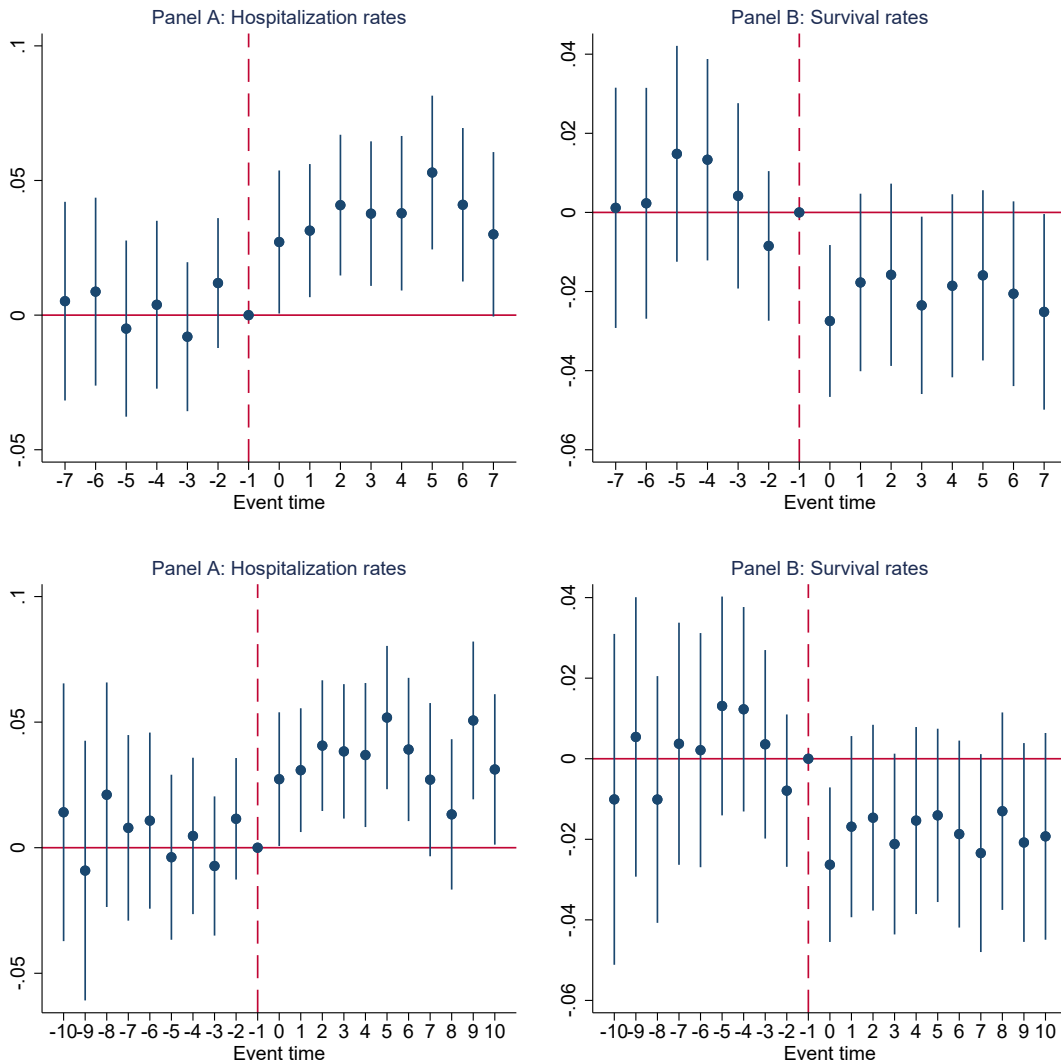


Figure H.7: Evolution of estimates of the main parameter of interest in event time.

This graph plots estimates of the main parameter of interest, λ_A , in event time from ten years before to ten years after the facility acquisitions. Vertical bars represent 95% confidence intervals around each point estimate. Standard errors are clustered at the facility acquisition level. A dotted vertical line marks the year prior to each facility acquisition.




Figure H.8: Predicted and empirical enforcement rates.

The x-axis measures simulated HHI changes, while the y-axis measures divestiture rates. The solid line corresponds to enforcement observed in the data, while the dashed one corresponds to predicted enforcement. The data are sorted into equal-sized bins according to x-axis values, and averages within the bins are reported along the axes.

Tables

Table H.1: Facility counts extracted from FTC complaints closely match CMS.

Acquiror or target	Year	FTC	CMS	CMS/FTC
DaVita	2004	665	548	82
DaVita	2010	1612	1511	94
DaVita	2016	2251	2254	100
Diversified Specialty	2010	106	141	133
Diversified Specialty	2015	100	81	81
Fresenius Medical Care North America	2005	1155	1105	96
Fresenius Medical Care North America	2011	1800	1783	99
Gambro Healthcare	2004	565	553	98
Liberty Dialysis	2011	260	250	96
Renal Care Group	2005	450	449	100
Renal Ventures	2016	36	35	97
US Renal Care	2015	200	181	91

Table H.2: Facility counts extracted from FTC divestiture orders closely match CMS.

Acquiror	Year	FTC	CMS	CMS/FTC
Alaska Investment Partners from Liberty merger	2012	1	1	100
Dallas Renal Group from Liberty Dialysis merger	2012	5	5	100
Dialysis Newco from Diversified Specialty merger	2011	29	24	83
Diversified Specialty from Liberty Dialysis merger	2012	51	49	96
Fresenius from Rhode Island/Fall River merger	2006	2	2	100
NRI from Renal Care Group merger	2006	103	103	100
PDI from Gambro Western Michigan mergers	2002	3	3	100
Physicians Dialysis from Renal Ventures merger	2017	7	7	100
Renal Advantage from Gambro Healthcare merger	2005	70	69	99
Satellite from Diversified Specialty merger	2016	3	3	100
Satellite from Gambro Healthcare merger	2005	2	2	100
Satellite from Liberty Dialysis merger	2012	2	2	100

Table H.3: Enforcement-related results are robust to various controls.

VARIABLES	(1) DIVEST	(2) DIVEST	(3) DIVEST	(4) DIVEST	(5) DIVEST	(6) DIVEST	(7) DIVEST
REPORTABLE \times Δ HHI*	0.25*** (0.0088)	0.25*** (0.0088)	0.25*** (0.011)	0.25*** (0.0086)	0.21*** (0.016)	0.24*** (0.011)	0.20*** (0.015)
REPORTABLE	0.036** (0.016)	0.074** (0.037)	0.014 (0.032)	0.055* (0.030)	0.041*** (0.014)	0.029** (0.012)	0.0085 (0.021)
Δ HHI*	0.0011 (0.00099)	0.0011 (0.0010)	0.0011 (0.00099)	0.0011 (0.0010)	0.00061 (0.0013)	-7.0e-06 (0.0015)	0.0010 (0.0018)
REPORTABLE \times HHI*		-0.0068 (0.0041)					
HHI*		-0.000074 (0.000079)					
Constant	0.00098 (0.00098)	0.0014 (0.0014)	0.00080 (0.00099)	0.0061 (0.0040)	0.0091** (0.0035)	0.0075** (0.0037)	0.026*** (0.0060)
Observations	4,331	4,331	4,331	4,331	4,330	4,331	4,330
R-squared	0.565	0.567	0.567	0.567	0.626	0.581	0.640
Number of target fac.	No	No	Yes	No	No	No	Yes
Number of acquiror fac.	No	No	No	Yes	No	No	Yes
Chain FE	No	No	No	No	Yes	No	Yes
Acquisition year FE	No	No	No	No	No	Yes	Yes

, **, and * denote significance at the 10%, 5%, and 1% levels, respectively. HHI* denotes realized HHI in the year prior to the merger. Note that "Number of target facilities" is exactly equal to the variable measured by the x-axis in Panel A of H.6, while "Number of acquiror facilities" is exactly equal to the variable measured by the x-axis in Panel B of H.6. Standard errors are clustered at the merger-market level.*

Table H.4: *Quality-related results are robust to alternate specifications.*

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Hosp. rate	Survival rate	Hosp. rate	Survival rate	Hosp. rate	Survival rate	Hosp. rate	Survival rate
POST \times EXEMPT \times Δ HHI*	0.022** (0.0092)	-0.018*** (0.0067)	0.023** (0.0090)	-0.018** (0.0071)	0.039*** (0.0088)	-0.020*** (0.0070)	0.040*** (0.0092)	-0.019** (0.0076)
POST \times Δ HHI*	-0.0037 (0.0054)	0.0066* (0.0039)	-0.0019 (0.0055)	0.0035 (0.0040)	-0.0048 (0.0055)	0.0063 (0.0043)	-0.0058 (0.0058)	0.0042 (0.0048)
POST \times EXEMPT	-0.0041 (0.0057)	0.0021 (0.0044)	-0.0039 (0.0057)	0.0015 (0.0044)	-0.016*** (0.0058)	0.0071 (0.0048)	-0.017*** (0.0063)	0.0065 (0.0054)
POST	0.0014 (0.0040)	0.000018 (0.0034)	0.00076 (0.0040)	0.0010 (0.0034)	0.0035 (0.0046)	-0.0043 (0.0040)	0.0051 (0.0053)	-0.0039 (0.0048)
Observations	52,762	47,990	52,762	47,990	41,677	37,968	32,222	29,370
R-squared	0.392	0.460	0.392	0.460	0.423	0.479	0.457	0.497
Facility acquisition FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Chain FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Specification	Radii	Radii	Stations	Stations	Window-7	Window-7	Window-5	Window-5

*, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. In columns 1-2, markets are defined circumscribing a circle with a 20 mile radii around each acquired facility. In columns 3-4, HHI is calculated machine counts rather than patient counts. In columns 5-6 and 7-8, the observation window is narrowed to 7 and 5 years, respectively. Standard errors are clustered at the facility acquisition level.

Table H.5: *Effects on acquiror facilities are congruent with those on target facilities.*

VARIABLES	(1) Hospitalization rate [Sample is Δ HHI* only]	(2) Survival rate [Sample is Δ HHI* only]
POST \times EXEMPT	0.020*** (0.0070)	-0.013** (0.0063)
POST	-0.0077 (0.0063)	0.0031 (0.0056)
Observations	8,709	8,709
R-squared	0.443	0.341
Facility acquisition FE	Yes	Yes
Year FE	Yes	Yes

, **, and * denote significance at the 10%, 5%, and 1% levels, respectively. The sample comprises acquiror facilities that are located in markets with facility acquisitions involving large simulated HHI changes. Standard errors are clustered at the facility acquisition level.*