

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 March 2024

I thank Judy Chevalier, Fiona Scott Morton, and Nancy Rose for thoughtful discussions of this paper and closely related work, Jimmy Roberts for suggesting the setting, Paul Eliason for helping navigate the data, John Asker, Dave Balan, Dennis Carlton, Thom Covert, Austan Goolsbee, John Kwoka, Neale Mahoney, Ariel Pakes, Chad Syverson, Haris Tabakovic, and Mike Whinston, and seminar participants at Harvard University, Yale University, the Federal Trade Commission (FTC), 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 helpful comments. Credit for the title belongs with Mike Sinkinson. (Please note that the title was adopted before the paper measured 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, nor do they 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 March 2024
JEL No. D4,D43,I11,K21,L0,L1,L11,L13,L4,L40

ABSTRACT

US antitrust authorities are only notified of large mergers, so most transactions could escape antitrust scrutiny. I study premerger notification exemptions in the dialysis industry. I find that, in sharp contrast to reportable mergers, exempt ones almost completely avoid enforcement. As a result, exempt mergers increase concentration and reduce healthcare quality, as measured by hospitalization and mortality rates. I then estimate a structural model to simulate the equilibrium response of demand, quality, and enforcement to the elimination of exemptions. I find that the benefits of eliminating exemptions in the dialysis industry far exceed the costs.

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 intend to merge 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 notification on the basis of their size, making them potentially hard to police. In the US, these exemptions are embedded within statutes from the 1970s and reflect an assumption that only large mergers merit prospective reviews [Wollmann, 2019].¹ Despite extraordinary growth in mergers, which transfer trillions of dollars in ownership rights annually each year, and despite a potentially flawed premise, which may run counter to decades of work in industrial organization economics, this assumption has never been empirically evaluated.

The potential problem is that in segmented industries, where products are differentiated to serve heterogeneous consumers, even minor deals can produce major changes in market structure and firm behavior. To illustrate, consider an industry where exactly two independently run establishments compete with one another in every city. Further, suppose that these establishments are each valued at around \$75 million. Finally, suppose that only mergers valued above \$100 million must be reported to the government. In this illustration, if premerger notifications are critical to enforcement, then a significant portion of an economically important industry could merge to monopoly without providing antitrust authorities any notice—an outcome I call "stealth consolidation."

In this paper, I study the costs and benefits of premerger notifications. If notifications facilitate enforcement, agencies negotiate effective remedies, and firms attempt to acquire competitors, then eliminating exemptions could significantly improve consumer welfare. Whether each condition holds is an empirical question. Moreover, benefits that accrue to consumers must be weighed against costs borne by the government, which needs to evaluate more notifications and investigate more mergers. Thus, I require not only primitives governing consumer demand and firm costs but also estimates of how notifications affect enforcement, whether enforcement produces lasting changes in market structure, and how agencies use resources. As far as I know, this is the first formal cost-benefit analysis of an antitrust law.

To gain traction on the problem, I narrow the scope of my analysis to a single industry and ask, "What is the effect of eliminating premerger notification exemptions for US dialysis mergers?" The answer is central to several recent policy proposals. Citing the present paper, the Senate Subcommittee on Antitrust, Competition Policy, and Consumer Rights discussed legislation that would limit exemptions,² the FTC proposed new rules that would expand premerger disclosure³ and ordered large technology platforms to divulge previously unreported acquisitions,⁴ and 23 State Attorneys General called for procedural changes aimed at strengthening enforcement.⁵ The answer to this question is also central to an ongoing debate about whether merged-induced consolidation is an important contributor to secular trends in market structure and, more importantly, income, markups, profits, and investment (see the discussion towards the end of this section).

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).

²See remarks by Senator Richard Blumenthal and Professor John Kwoka starting at 1:36 in the Senate Subcommittee hearing titled "Does America Have a Monopoly Problem?"

³See the Notice of Proposed Rulemaking at 42203 (<https://shorturl.at/IPU57>).

⁴A concise summary by Commissioners Wilson and Chopra is available at <https://shorturl.at/orTY2>.

⁵See their Public Comments at 67 (<https://shorturl.at/aLQ09>). Other changes include the restoration of prior approval.

Dialysis provides an almost ideal laboratory to examine stealth consolidation. First, it allows me to link ownership changes, premerger notifications, enforcement actions, and welfare-relevant outcomes at a facility-year (i.e., establishment-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 very close to their homes, so each location is relatively small. Thus, while the industry accounts for tens of billions of dollars in annual healthcare expenditures, individual facilities are valued at only about \$4 million each. By contrast, transaction value threshold set by Congress stands at \$81 million by the end of the sample, meaning one can acquire a substantial chain without reporting it to the agency. Fourth, enforcement is commonplace. The FTC, which has de facto responsibility for policing dialysis mergers, has prevented hundreds of ownership changes by negotiating divestitures. Fifth, entry into established dialysis markets is slow and expensive, so profits derived from combining competitors are not easily dissipated. Sixth, the cost to the government of eliminating premerger notification exemptions in the dialysis industry can be credibly estimated from variation provided by a 2001 law change.

I begin the paper by describing the setting and developing a structural model that captures its salient features. To summarize them, a federal agency enforces antitrust laws by negotiating remedies to anticompetitive transactions. The rate of enforcement depends on premerger notification and market structure. The agency chooses labor and capital to minimize its cost, subject to the constraint that it evaluates all premerger notifications it receives and investigates all mergers it finds potentially anticompetitive. In each dialysis market, patients choose a facility (or an outside option, home hemodialysis) to maximize utility, which depends on factors such as travel distance and quality of care. Throughout the model, quality is measured using risk-adjusted survival rates. Since Medicare covers most treatments and sets reimbursement rates, I assume that price does not enter utility and falls outside the provider's control. Instead, firms choose quality to maximize profits. The model highlights how acquisitions of competing facilities cause providers to internalize business-stealing externalities, thereby reducing equilibrium quality.

To describe its analysis, definitions are necessary. First, I define a "merger" as the simultaneous sale of one or more facilities from a target to an acquirer. Second, since competition and enforcement are market-level phenomena, I divide each merger up into market-specific "facilities acquisitions." Third, if a merger's transaction value meets a threshold set forth in the Hart-Scott-Rodino Antitrust Improvements (HSR) Act, then premerger notification is required. I call such mergers (and the facilities acquisitions associated with them) "reportable," and I call all others "exempt." Fourth, since agency guidelines summarize market structure changes using concentration increases that are predicted by historical figures, I compute an "HHI change" for each facilities acquisition using shares measured just prior to the transaction.⁶ Fifth, if the FTC expects a facilities acquisition to reduce competition, then it negotiates a remedy, which almost always involves the target selling off its facilities in the offending market. Thus, I use the terms "divestiture" and "enforcement" interchangeably in this setting.

To characterize divestitures, demand, and supply, I construct a dataset comprising the near-universe of US dialysis facilities that treat patients between 1996 and 2017. In this sample, I observe over 4,000 facility-level ownership changes. Only about half are reportable. In other words, more than 2,000 are never reported to US antitrust authorities. Nonetheless, the FTC negotiates the divestiture of nearly 300 facilities

⁶A previous draft of the paper used the term "simulated HHI change." This one opts for "HHI change" to ease exposition and stay consistent with guidelines issued by the agencies [US Department of Justice and Federal Trade Commission, 2010].

during this period.

Novel patterns arise in the data. *In sharp contrast to reportable facilities acquisitions, exempt ones almost completely escape antitrust scrutiny*, consistent with claims that premerger notifications are critical to effective enforcement. For reportable facilities acquisitions, divestiture rates rise alongside HHI changes from around zero to 90%. However, for exempt facilities acquisitions, divestiture rates are very close to zero regardless of HHI change. Many of these involve HHI changes exceeding 2,000 points, which is an order of magnitude higher than what agency guidelines "presume[] to be likely to enhance market power" [US Department of Justice and Federal Trade Commission, 2010]. Further analysis indicates that premerger notifications prevent consolidation. In the short run, this result is guaranteed, but in the longer run, it implies that buyers of divested assets successfully utilize them, countering recent criticism of the FTC. Finally, the data suggest that premerger notifications prevent ownership changes that would otherwise reduce quality, thereby reducing risk-adjusted hospitalization and mortality rates.

To characterize the agency's production function, I construct a dataset measuring merger-related inputs and outputs of the FTC each year from 1996 to 2017. To police mergers, the agency spends \$41 million on labor and \$16 million on capital. It also evaluates nearly 2,000 premerger notifications and investigates about 170 mergers. The data reveals that an amendment to the HSR Act, which abruptly reduced the number of mergers reported to the FTC, was followed by an equally sharp drop in the number of FTC staff tasked evaluating premerger notifications and investigating potentially anticompetitive mergers. The relative magnitudes translate into sensible estimates of productivity and underscore that the agency's expenses depend on its caseloads.

These findings beg the question of how the benefits of premerger notification compares to its costs. However, they cannot credibly answer it, as concentration indices are endogenous, one-dimensional approximations of complicated substitution and ownership patterns, to which quality responds nonlinearly. Thus, I estimate the structural model to predict market outcomes under the alternative policy in equilibrium. First, I recover parameters that determine divestiture rates. Next, I estimate demand on a panel of patient-year level facility choices. Then, I estimate marginal cost by matching moments implied by the firms' first order conditions and profit margins. Last, I estimate the productivity of agency inputs using data that measures FTC labor, capital, premerger notifications, and merger investigations. The aforementioned amendment to the HSR Act provides quasi-experimental variation in the agency's workloads over time. Once estimation is complete, I compare the costs and benefits of eliminating premerger notification exemptions in the dialysis industry. This exercise involves retroactively applying enforcement rates historically experienced by reportable mergers to exempt ones, which saves lives but also increases the costs of the FTC. I find that expected benefits of \$1.1-1.4 billion far exceed expected costs of \$10 million.

Nevertheless, my findings do not imply that all exemptions based on transaction value should be immediately eliminated. In other industries, such a policy change could have consequences that I do not consider here. As just one example, I abstract away from the possibility that firms will restructure transactions to satisfy the letter but not spirit of the HSR Act, likely by acquiring multiple minority stakes in noncorporate entities. This restriction is reasonable for dialysis mergers but may be too strong for others.⁷ Separately, my cost calculations raise but do not address nuanced political economy considerations.

⁷I thank Fiona Scott Morton for bringing this issue to my attention. I also thank two FTC attorneys for describing the problem from the agency's viewpoint. Asil et al. [2023] examines it in detail. Briefly, when interests in noncorporate entities are spread across multiple investment vehicles, the way the HSR Act and Rules define control may be misaligned with economic reality. Under certain circumstances, strategic noncompliance could occur. Typically the problem is confined acquirers backed by private equity, whose exit

Namely, due to the high cost of complying with an investigation (e.g., the time, effort, and expense of satisfying a "Second Request"), merely receiving notice of one will cause merging parties to abandon almost any modestly sized merger. On the one hand, this reduces enforcement costs, but on the other, it permits unelected officials to determine the fate of a transaction by depriving the merging parties an economical means of redress. Section 6 enumerates other considerations.

This paper contributes to several areas of research. Alongside recent work, it introduces the notion of stealth consolidation. Wollmann [2019] uses economy-wide data to study an abrupt increase in HSR thresholds in 2001, which led to a sharp drop in enforcement and a rise in mergers among rivals. However, the paper's broad scope precludes realistic measures of competition. For example, it assumes that firms compete if they occupy the same 4-digit SIC industry group, which is a wider definition than courts use in practice. It also precludes measurement of merger effects and does not attempt to measure agency costs. The present work differs in each regard. 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, Karabarounis and Neiman, 2013], various papers have related rising concentration to that decline 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.⁸

This paper also relates closely to prior work on the industrial organization of the healthcare sector and the 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.) 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.⁹ 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 threshold.

The paper proceeds out as follows. Section 2 provides institutional details and describes the data. Section 3 presents a structural model of competition and enforcement. Section 4 reports descriptive evidence. Section 5 presents estimates from the model. Section 6 compares the costs and benefits of eliminating premerger notification exemptions. Section 7 concludes.

strategies do not include public offerings or sales to public companies, which limits their use in the dialysis industry. Consistent with this claim, all but one of the mergers that I predict are reportable based on their transaction size were actually reported. Appendix D.3 provides more detail.

⁸One 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 is Peltzman [2014], who proposes the shift towards more permissive merger policy following the release of the 1982 *Merger Guidelines*.

⁹Physicians 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 antitrust law. Most of all, Section 7 of the 1914 Clayton Act prohibits transactions where the effect "may be substantially to lessen competition, or to tend to create a monopoly."¹⁰ Though it is regarded today as the cornerstone of merger-related enforcement, the Clayton Act proved flawed as originally written for reasons that are central to this paper. Even after the Celler-Kefauver amendments closed most substantive loopholes in 1950, 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, unwinding would equate to "unscrambling eggs" [Baer, 2014]. Thus, while Section 7 provided a legal basis for arresting anticompetitive deals in their incipiency, it did not provide any practical means to discover them in that state.

This practice abated in the late 1970s when Congress provided bipartisan support for the Hart-Scott-Rodino Antitrust Improvements (HSR) Act, thereby establishing the US premerger notification program. It requires merging parties that are not explicitly exempt to notify the FTC and DOJ of their proposed transaction, describe their current operations, and wait up to thirty days before completing the transaction. During this time, the merger receives what is commonly called an "ex ante" or "prospective" review: the Premerger Notification Office (PNO) provides an initial evaluation, and if the office finds the merger potentially anticompetitive, then one (but not both) of the agencies may conduct a post-PNO investigation.¹¹ Further, the act not only bars firms from structuring transactions to intentionally avoid reporting requirements but also stipulates significant penalties for failing to file, regardless of intent. Appendix D.3 provides more detail.

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 revenue, in the event the target is engaged in manufacturing).¹² Under the act as 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, the HSR threshold has nominally increased with gross national product since 2004. As a result, at present, the act exempts most transactions valued below \$111 million. For many segmented industries, these thresholds absolve all but the largest transactions from reporting responsibilities.

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

¹⁰ Other laws provide standing to sue but are utilized much less frequently. For instance, Section 5 of the FTC Act bars unfair methods of competition, while Section 2 of the Sherman Act bars monopolization, completed or attempted. In some cases, other federal agencies may have jurisdiction. Examples include mergers involving banks, airlines, motor carriers, or entities for whom a change in control might jeopardize national security. State statutes, such as California's Unfair Competition Law, may also apply.

¹¹ Since only one agency can investigate, the FTC must be "cleared" by the DOJ to investigate or vice versa. As a result, "clearance" and "investigation" are synonymous.

¹² This pared-down description is far from comprehensive but nonetheless covers most transfers of control. Supporting this claim are tight links between mergers, premerger notifications, and divestitures, which are reported in this paper.

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.¹³

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.

2.2 Kidney disease and treatment providers

The dialysis industry treats people suffering from a loss of kidney function. Kidneys are organs in the human abdomen that filter out toxins, remove excess fluids, and synthesize certain hormones, but they gradually deteriorate with prolonged exposure to conditions such as diabetes and high blood pressure. 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.¹⁴ The other option is dialysis, which involves diffusing a solution in the blood, passing the mixture across a semipermeable membrane, and removing unwanted substances.

Hemodialysis accomplishes this process outside the human body, with a machine serving as an external kidney.¹⁵ Patients typically receive 3-4-hour treatments each week, during which time they may engage in sedentary activities. Hemodialysis is by far the most common form of dialysis 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, adjusted for age, sex, ethnicity, and other factors, declined 30% by 2012, with little consensus around the cause(s), but have fallen at a slower rate since then—only about 3.5% over the following six years. Rates are still very high. Persons on

¹³For 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.

¹⁴This 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.

¹⁵Hemodialysis is ten to twenty times more prevalent than peritoneal dialysis over most of the panel, so I ignore the latter for the remainder of the paper. The reasons for its unpopularity are debated but beyond the scope of my analysis.

dialysis face 160 deaths per 1,000 person patient years [USRDS, 2017].

Since patients require frequent treatments, they prefer locations that are close to their homes. As a result, facilities are spread across the country, and patients rarely travel more than 25 miles for care.¹⁶ 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.) 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.

Geographic coverage has expanded apace with demographics, although high barriers to entry typically exist in markets with incumbent providers. Most significant is that each entrant must engage an established nephrologist to serve as its facility's medical director. According to the FTC, "locating and contracting with a nephrologist to serve as medical director is difficult because clinics typically enter into exclusive contractual arrangements with a nephrologist who is paid a medical director fee."¹⁷ Regulation raises other site-specific sunk costs associated with entry as well.¹⁸

2.3 Antitrust enforcement amid dialysis consolidation

The dialysis industry witnessed thousands of facilities acquisitions over the past two decades, and preliminary investigations into many of these raised serious competitive concerns (as shown later in this section). 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 acceptable terms under which the deal can proceed.

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 studying a single industry. For example, merger enforcement in dialysis is highly uniform. First, all cases are brought by the federal government and, in particular, the FTC (with a sole exception of a very small transaction unwound by Minnesota's 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. By extension, vertical considerations are limited.¹⁹ Third, as a nontradable service industry, imports and exports play little role in determining market structure and behavior. Fourth, economic analysis occurs at the market level, which is evident from FTC complaints, which precisely define geographic areas in which competition will be reduced. Relatedly, enforcement actions involve market-specific divestitures of target facilities, which are sold to unaffiliated third parties and are approved by the agency.

These features are shared by many US mergers. For example, the federal government brings nearly all merger challenges [Gotts, 2001]. States Attorneys General may join suits but often rely on national resources. Moreover, the FTC and DOJ police mergers very similarly. Some procedural differences arise

¹⁶Note that a previous version of this paper stated that 20, rather than 25, miles is a bound on most trips. The previous version calculated "great circle" distances, whereas this version uses road distances. Appendix E.5 provides a brief discussion of travel distances and their empirical distribution.

¹⁷See, e.g., *In the Matter of DaVita Inc. and Total Renal Care, Inc.*, FTC Matter/File 2110013 (October 25, 2021).

¹⁸All facilities must be licensed by state-level health boards and establishment a reimbursement relationship with CMS. In some areas, providers also need to prove that the surrounding community "needs" the additional capacity provided by the facility, which often involves lobbying and litigation. See, e.g., *Liberty Dialysis-Haw., LLC v. Rainbow Dialysis LLC*, SCAP-12-0000018 (June 27, 2013).

¹⁹Some providers are partially vertically integrated (e.g., Fresenius manufactures machines), but there is no evidence that this affects their operations beyond the brand of drug they administer.

for investigations arising in one or the other, but all are second order considerations. Likewise, almost all merger cases are brought under Section 7 of the Clayton Act. Similarly, most enforcement actions amount to market level remedies that require divestitures. For instance, a retrospective study by the FTC reports between 2006 and 2012, 85% of remedies involved only divestitures, with an additional 6% had a divestiture-related component [FTC, 2015].

2.4 Data sources

The Centers for Medicare & Medicaid Services (CMS) collects data from their regional offices to create Provider of Service (POS) files. Records 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 CMS-approved facility. Medicare has covered most ESRD patients since 1973, so coverage is substantially complete. For each facility, I extract the unique CMS provider number, name, address, opening and closing dates, parent company and recorded dates of any ownership changes. The sample comprises facility-year observations spanning 1996 to 2017.

CMS also collects data through its Healthcare Cost Reporting Information System (HCRIS), which is based on annual surveys that are mandatory for all certified institutions. The source consists of facility-year level records, which include CMS provider numbers. It serves to supplement the previous data sources by providing timely information about ownership changes in years when facilities did not complete form CMS-1539 and hence did not update chain affiliation details.²⁰

The United States Renal Data System (USRDS) collects detailed operational and clinical data related to kidney disease and provides access to researchers. For each facility, I observe basic attributes, including CMS provider number, and for each ESRD patient, I observe demographic and baseline health measures (e.g., age, race, gender, body mass index (BMI), and renal function at the onset of treatments, date of diagnosis, date of first dialysis session, and date of Medicare coverage), where, when, and how each was treated, and dates corresponding to hospitalizations and death [USRDS, 2017]. To construct the facility year panel, I follow prior work (e.g., Eliason [2018]) by risk-adjusting hospitalization and survival rates²¹ and then assigning patients to their modal facility in each calendar year.

The FTC publishes the complaint and consent order associated with each enforcement action. The former mostly contributes background information—it describes the operations of the target and acquirer, enumerates the relevant geographic markets, and discusses the entry conditions. The latter lists the names and addresses of each facility that the parties and the agencies agree must be divested for the deal to close.

The FTC also publishes an annual Congressional Budget Justification (CBJ) that summarizes its operations. The document reports the number of agency-wide, PNO, and post-PNO (i.e., "Merger and Joint Venture Enforcement and Compliance") employees as well as agency-wide costs by category in dollars. I equate "labor" with the number of employees, "annual compensation" with the sum of per capita salaries and benefits, and "labor costs" with the product of the two. Since the difference between total costs and labor costs comprises travel, office equipment, and payments to the GSA for the use of federal buildings, I

²⁰In about 10% of observations, cost reports suffer from aggregation across co-located providers. The problem arises most frequently when facilities are adjacent to acute care hospitals. HCRIS supplements other sources, so I ignore reports associated with them.

²¹Control variables used to risk-adjust rates are race, age, sex, BMI, indicators for diabetes, hypertension, cancer, alcoholism, drug use, and smoking, and baseline kidney function, the latter of which is measured by glomerular filtration rate (GFR). The two most widely-used equations for estimating GFR are the CKD-EPI and MDRD Study equations, so I include both estimated values.

subtract one from the other to obtain "capital rental." Following De Loecker and Eeckhout [2017], I set the "rental rate" (in real terms) equal to a depreciation rate of 12% plus the federal funds rate, and I set "capital" equal to capital rental divided by the rental rate.²² Note that since the size of the mergers that are "newly reportable" under the alternative policy I study are relatively small, litigating them is never economical (see Section 6), so I remove expert/advisory fees from costs to make the counterfactual forecasts as accurate as possible.

Finally, the FTC and DOJ jointly publish an HSR Annual Report. It reports the number and value of reportable mergers evaluated by each agency. It also reports the number of mergers each agency investigated. Note that nominal values from this source and the previous ones are converted to constant 2017 dollars using the CPI.

2.5 Data summary

Table I summarizes the data at the facility-year level. In the average year, about 4,500 facilities are active.²³ Patients face a 64% chance of hospitalization and an 86% chance of survival each year, highlighting the at-risk nature of these individuals. Chains own around three-quarters of all facilities.

Table I: *Summary of facilities*

	N	Mean	Std. Dev.	Minimum	Maximum
Patients	99781	41.8	28.6	1	290
Hospitalization rate	89239	0.64	0.19	0.097	1.33
Survival rate	82602	0.86	0.090	0.53	1.04
Machines	99129	17.8	7.93	0 [†]	174
Chain owner	99781	0.74	0.44	0	1

Each observation is a facility-year. "Patients" counts the number of individuals treated by the facility, while "Machines" counts the number of dialysis stations operated by the facility. Hospitalization and survival rates are adjusted for risk. For the purposes of this statistical summary, each facility is assigned the HHI of the market in which it is located. The sample covers 22 years. [†]denotes a mismeasured quantity (it is infeasible to operate a facility with zero machines, but only 0.0001 of facility-year observations report this number).

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 easily to infer in dialysis. Reporting requirements depend on transaction values, which one can accurately predict using information on targets. In particular, real transaction values are very closely tied to the number of facilities operated by the targets. Critically, one can then verify the accuracy of this approach using "Early Termination" notices, which are public filings that confirm a premerger notification was filed,

²²I obtain the average salary of an FTC attorney/economist and paralegal/clerk from the Office of Personnel Management. I then compute the salary of a representative FTC employee by averaging these figures and weighting by the relative number of attorneys, economists, paralegals, and clerks employed by the agency. I compute benefits by multiplying the result by the ratio of agency-wide salaries to benefits. Finally, I verify that multiplying the sum of salaries and benefits by labor is very close to the aggregate labor cost figure reported by the agency.

²³Firm-year level facility counts match those appearing in FTC complaints and orders. See Appendix Tables A1 and A2 for details.

and SEC filings.²⁴ From the combined sources, I confirm a premerger notification was filed for all but one merger that I predict is reportable, with the one discrepancy occurring in the earliest part of the sample, when Early Termination notices are not easily searchable. Similarly, I locate only one premerger notification among the mergers I predict are exempt. Thus, I conclude that exemption status is measured without any meaningful error. Appendix D describes the process of inferring and verifying reportability.

Table II summarizes ownership changes. Panel A provides a brief summary of facility level ownership changes over the sample. Over 4,300 are completed, and only around half are reportable. Panel B provides a more detailed summary at the facilities acquisition level—the relevant level to analyze competitive effects and divestitures. More than 3,200 facilities acquisitions are completed, meaning that most involve just one target facility. The average transaction date falls around the midpoint of the sample. More generally, transactions are spread evenly over the sample period (not shown). Substantial heterogeneity in predicted consolidation exists. Across all deals, mean values are between 400 and 500, but many facilities acquisitions involve no target-acquirer overlap, while others create monopolies and duopolies. (See Appendix Figure A2 for the full distribution.)

Table II: Summary of ownership changes

Panel A. Facility level					
	N	Mean	Std. Dev.	Minimum	Maximum
Reportable	4333	0.53	0.50	0	1
Divestiture	4333	0.062	0.24	0	1
Panel B. Facilities acquisition level					
	N	Mean	Std. Dev.	Minimum	Maximum
Reportable	3266	0.53	0.50	0	1
Divestiture	3266	0.066	0.25	0	1
Acquisition year	3266	2005.6	5.63	1997	2017
Δ HHI	3266	490.1	1076.1	0	5000

Table III summarizes FTC operations. To police mergers, the FTC employs 32 full-time workers in the PNO and allocates another 248 to post-PNO activities in the average year. Total annual compensation is around \$156,300. Post-PNO activities require almost an order of magnitude more resources than the PNO. This is expected. Whereas the former encompasses detailed industry analyses, negotiations with merging parties, and even litigation, the latter has a much more narrow focus. Agency-wide, the FTC employs 1,100 workers. Capital rental per employee averages \$60,000—much smaller than salaries and benefits, consistent with the idea that the FTC most intensely uses highly skilled labor. Each year, the PNO evaluates an average of 2,047 premerger notifications whose total value typically exceeds \$1 trillion. About 176 investigations are conducted annually.

²⁴When parties request expedited review, they forfeit confidentiality and agree to have a basic description of the transaction published in the *Federal Register*. Most make this request.

Table III: Summary of enforcement agency inputs and outputs

	N	Mean	Std. Dev.	Minimum	Maximum
PNO labor	21	31.9	6.53	23.5	45.2
Post-PNO labor	21	250.8	12.3	232.2	276.7
Annual compensation	21	156.3	12.5	129.1	166.8
PNO labor costs	21	4.93	0.65	3.84	6.03
Post-PNO labor costs	21	39.1	2.16	34.2	42.1
Agency-wide labor	21	1100.4	81.5	928	1211
Agency-wide capital	21	549.9	144.6	348.6	844.6
Premerger notifications	21	2047.3	1188.3	684	4749
Value of reportable mergers	21	1246.0	632.8	406.8	3000
Investigations	21	176.1	43.2	98	278

"Labor" means the number of full-time equivalent staff. "Annual compensation" means per capita salary and benefits in thousands of 2017 dollars. "Labor costs" and "capital" are measured in millions of 2017 dollars. "Premerger notifications" means the number of reportable mergers and represents the caseload of the PNO. "Value of reportable mergers" measures their value in millions of 2017 dollars. "Investigations" represents the post-PNO caseload.

3 Model

In this section, I present a model that captures the salient features of dialysis industry antitrust enforcement, patient demand, facility competition, and agency costs.

3.1 Enforcement

Each facility j owned by firm f in geographic market g and year t provides dialysis treatments. By definition, each merger m transfers ownership of one or more j from f to f' at t . Likewise, the subset of these j in a particular g constitute a facilities acquisition a . Since competition is a market-level phenomena, the FTC assesses competitive effects on a market-by-market basis, so enforcement decisions made at the facilities acquisition level. An agency enforcement function, denoted by $D(\cdot)$, maps attributes of facilities acquisitions to the probability of divestiture in the dialysis industry.

To characterize $D(\cdot)$, I first turn to the 2010 *Horizontal Merger Guidelines*, which are expressly written for this purpose.²⁵ The *Guidelines* state that "the unifying theme of [enforcement] is that mergers should not be permitted to create, enhance, or entrench market power" and "mergers that cause a significant increase in concentration and result in highly concentrated markets are presumed to be likely to enhance market power." They also state that since enforcement actions typically address incipient mergers, "market share data are normally based on historical evidence." Hence, $D(\cdot)$ may depend on expected concentration levels and changes, which are evaluated using premerger market structure. Second, I incorporate the fact that the FTC must be aware of a transaction to investigate it. Staff can learn about transactions many ways,

²⁵The DOJ and FTC state that the goal of the *Guidelines* is to "assist the business community and antitrust practitioners by increasing the transparency of the analytical process underlying the Agencies' enforcement decisions."

including trade press, statements by brokers, and SEC filings, but most mergers are never publicized, especially prior to closing [Barrios and Wollmann, 2023]. Thus, $D(\cdot)$ also depends on whether the facilities acquisition is reportable.

Formally, I assume the probability of divestiture can be written as $D(s_{g,t-1}, \Lambda_{g,t-1}, \tilde{\Lambda}_a, HSR_a)$. $s_{g,t-1}$ denotes the vector of market shares in g just prior to the merger, $\Lambda_{g,t-1}$ maps facilities to owners last period, and $\tilde{\Lambda}_a$ maps facilities to owners last period under the assumption that the target's facilities are instead owned by the acquirer. The final argument, HSR_a , takes a value of one if a is reportable and zero otherwise.

So long as $D(\cdot, HSR_a = 0) \neq D(\cdot, HSR_a = 1)$, premerger notification policies affect market structure. Such policies, of course, also affect the agency's costs. I return to this issue later in the section.

3.2 Demand

Each patient i chooses a facility or the outside option, home dialysis, to maximize total utility.²⁶ If i chooses j , then he or she receives

$$U_{ijt} = \beta_0 + \beta_x x_{jt} + \beta_n n_{jt} + \phi_t + \chi_{jt}^c + \xi_{jt} + \alpha_1 dist_{ijt} + \alpha_2 dist_{ijt}^2 + \alpha_3 dense_{g(j)t} dist_{ijt} + \alpha_4 dense_{g(j)t} dist_{ijt}^2 + \epsilon_{ijt}, \quad (1)$$

while if i chooses the outside option, then he or she receives U_{i0t} , whose mean is normalized to zero. Patients derive (dis)utility from travel, which is a function of the distance in road-miles between i 's home and j , denoted $dist_{ijt}$, and the population density of the surrounding area, denoted $dense_{g(j)t}$. Utility also depends on quality, denoted x_{jt} , and the number of machines, denoted n_{jt} . ϕ_t and χ_{jt}^c denote year and chain fixed effects, respectively.²⁷ Throughout the model, I measure quality using risk-adjusted survival rates. They are unambiguously welfare-relevant, easily interpretable by researchers and policymakers alike, and accurately measured in the data. Pragmatically, a unidimensional measure also significantly reduces the computational burdens encountered below. (For an especially transparent description of how operational decisions like under-staffing, which leaves employees inadequate time to sterilize machines, can affect patients' health, see Grieco and McDevitt [2016].)

ξ_{jt} and ϵ_{ijt} 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 that ϵ_{ijt} has an extreme value distribution. The assumption delivers closed-form "logit" individual choice probabilities, denoted s_{ijt} for each j and s_{i0t} for the outside option. Integrating the individual choice probabilities out over i yields market shares, s_{jt} , and the portion of patients that choose the outside option, s_{0gt} .

3.3 Quality-setting

Each firm f chooses quality to maximize profit, which equals

$$\pi_{fgt} = \sum_{j \in \mathcal{J}_{fgt}} (P_{g(j)t} - mc_{jt}) s_{jt} M_{g(j)t}. \quad (2)$$

²⁶Demand and marginal cost specifications borrow generously from Eliason [2018]. Since we study very different policies (e.g., I am interested in merger-related enforcement whereas he is not), my specifications depart from his in several ways. See also Footnote 28.

²⁷I omit congestion from equation 1. In its presence, increasing x_{jt} makes j more desirable to patients other than i , resulting in crowding. Thus, increasing x_{jt} increases u_{ijt} directly but decreases u_{ijt} indirectly [Richards-Shubik et al., 2021]. In this case, my estimate of β_x reflects the net effect, averaged over patients and conditional on the controls, which understates the true taste for quality. However, since the market structure changes that identify β_x are the same type of variation induced by the alternative policy I study in Section 6, this omission is unlikely to have a meaningful affect on the simulation of counterfactual outcomes.

\mathcal{J}_{fgt} denotes the set of facilities owned by f in g at t . $P_{g(j)t}$ denotes price, and $M_{g(j)t}$ denotes market size. mc_{jt} denotes marginal cost, which I parameterize as an additively separable function of quality, quality squared, the number of machines, and an unobserved shock, ω_{jt} . ω_{jt} is i.i.d. over facilities and time. γ denotes the vector of coefficients on the observable components. Prices are assumed fixed, since CMS sets reimbursement rates for a large majority of patients.²⁸

Profit maximization implies that for every facility k in \mathcal{J}_{fgt} ,

$$\sum_{j \in \mathcal{J}_{fgt}} (P_{gt} - mc_{jt}) \frac{\partial s_{jt}}{\partial x_{kt}} - \frac{\partial mc_{kt}}{\partial x_{kt}} s_{kt} = 0, \quad (3)$$

where $mc_{jt} \equiv \gamma_0 + \gamma_1 x_{jt} + \gamma_2 x_{jt}^2 + \gamma_3 n_{jt} + \omega_{jt}$. Equation 3 bears a close resemblance to the conditions that determine Nash prices under exogenously determined product characteristics. This is not coincidental: firms that set x subject to fixed P face analogous tradeoffs to firms that set P subject to fixed x .

Equation 3 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 t subscripts are omitted to ease exposition). At the margin, increasing quality 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 term captures cannibalization, i.e., that increasing x_j diverts patients away from j' , which f now owns. Thus, holding all other facility decisions fixed, x_j must fall to balance the left- and right-hand sides of the equation, meaning the incentive to provide quality at j falls following the acquisition of j' . More concisely, "when price is [regulated] above marginal cost, competition leads to more quality" [Gaynor, 2006].

3.4 Agency costs

Evaluating more premerger notifications and investigating more potentially anticompetitive mergers requires hiring labor, such as attorneys and paralegals, and adding capital, such as office space and computing/communication equipment. To forecast how agency costs are affected by policy changes, I make assumptions about the agency's objective and production functions. Since each agency's "premerger review . . . workload is not discretionary" [Department of Justice, 2023], I assume the FTC minimizes cost subject to the constraint that all premerger notifications are evaluated and all potentially anticompetitive mergers are investigated. Given the fixed proportions of labor and capital employed in antitrust enforcement (e.g., each employee requires an office, a digital work space, etc.) and their limited substitutability (e.g., faster computers do not speed up preliminary prospective merger reviews, at least under existing processes and with current technology), I assume that production is Leontief.

Rather than own the capital it employs, the FTC rents nearly all of it, typically on a short-term basis. Its largest capital item, commercial real estate, is leased from the General Services Administration under an agreement that may be amended with one month's notice. Likewise, the agency contracts for most

²⁸To be transparent, this restricts price to evolve independently of consolidation. Since some treatments are privately insured, modest price effects are possible; however, the bargaining process might also be determined at a much higher level of geographic aggregation. In any case, I lack data on insurer-provider contracts, so price effects are beyond this paper's reach. Given that the number of machine is very stable within each facility over time, I also assume that n is fixed. Note, though, that n varies in the cross section, so when Eliason [2018] studies steady-state entry decisions, he correctly assumes firms endogenously choose it.

equipment, outsources parking, and arranges commercial transportation for staff. Thus, I assume capital adjustment costs are negligible.

Since the alternative policy may affect PNO and post-PNO operations differently, I estimate separate parameters for each. Starting with the PNO, let L_t^{PNO} denote labor, K_t^{PNO} denote capital, and C_t^{PNO} denote the caseload, which are premerger notifications that must be evaluated. In each period, the FTC solves

$$\begin{aligned} & \min_{L_t^{\text{PNO}}, K_t^{\text{PNO}}} w_t L_t^{\text{PNO}} + r_t K_t^{\text{PNO}} \\ \text{s.t.} \quad & \min \left\{ \Phi_{0,L}^{\text{PNO}} + \Phi_{1,L}^{\text{PNO}} L_t^{\text{PNO}}, \Psi_{0,K}^{\text{PNO}} + \Psi_{1,K}^{\text{PNO}} K_t^{\text{PNO}} \right\} \geq C_t^{\text{PNO}}, \end{aligned} \quad (4)$$

where w_t denotes the wage, r_t denotes the rental price of capital, $\Phi_{1,L}^{\text{PNO}}$ denotes the productivity of labor, and $\Psi_{1,K}^{\text{PNO}}$ denotes the productivity of capital. Cost minimization implies that $\Phi_0^{\text{PNO}} + \Phi_1^{\text{PNO}} L_t^{\text{PNO}}$ and $\Psi_0^{\text{PNO}} + \Psi_1^{\text{PNO}} K_t^{\text{PNO}}$ exactly equal C_t^{PNO} .

The FTC faces an analogous problem when it chooses post-PNO labor and capital. Let L_t^{PP} and K_t^{PP} denote their choices, respectively, let $\Phi_{1,L}^{\text{PP}}$ and $\Psi_{1,K}^{\text{PP}}$ denote the productivity of labor and capital, respectively, and let C_t^{PP} denote the caseload, which are potentially anticompetitive mergers that must be investigated. The agency sets $\Phi_0^{\text{PP}} + \Phi_1^{\text{PP}} L_t^{\text{PP}}$ and $\Psi_0^{\text{PP}} + \Psi_1^{\text{PP}} K_t^{\text{PP}}$ exactly equal to C_t^{PP} .

4 Descriptive analysis

In this section, I describe key patterns that emerge from my data.

4.1 Divestitures

I begin with a fundamental question. How does antitrust risk vary with premerger notification and predicted consolidation? For this exercise, the appropriate unit of observation is a facilities acquisition, and the relevant outcome is a divestiture.

The analysis relies on two sources of variation. To vary predicted consolidation, I use differences in premerger market structure. Since agency guidelines emphasize concentration increases, my initial specifications are based on HHI changes. To vary premerger notification, I use differences in merger size. If a merger's transaction value is greater than or equal to the HSR threshold, then all of the facilities acquisitions associated with that merger are reportable. Otherwise, none are reportable. To fix ideas, suppose that firm A buys firm B, the HSR threshold stands at \$75 million in the year of the merger, and facilities sell for \$4 million each (in real terms, this is approximately right). For simplicity, also suppose B has at most one facility in any market. If B operates 20 facilities prior to the merger, then the FTC is alerted to 20 facilities acquisitions, but if B operates 15, then the FTC is alert to none.²⁹

I define D_a equal to one if the facilities acquisition involves a divestiture and zero otherwise.³⁰ I denote each HHI change by ΔHHI_a and set it equal to two times the product of the target's and acquirer's market

²⁹For various reasons, including fines for intentional noncompliance that may exceed \$50,120 per day, I assume that transaction values are not manipulated. Comparisons between transaction values and early termination notices strongly support this assumption. Appendix D.3 discusses this issue in detail.

³⁰This captures a very small number of cases where facilities previously owned by the acquirer rather than target were divested. Alternative approaches that avoid undercounting (e.g., proportionally allocating divestitures of acquirer facilities to target facilities in the same market) produce very similar results.

shares observed just prior to the merger. As in the previous section, I define HSR_a equal to one if the facilities acquisition is reportable and zero otherwise. I then plot D_a against ΔHHI_a separately for $HSR_a = 0$ and $HSR_a = 1$.

Figure I displays the result. The dashed line corresponds to reportable facilities acquisitions. Here, divestiture rates are closely related to HHI changes. To see this most clearly, consider the extreme points along either axis. For ΔHHI near zero, divestiture rates are also near zero, but as ΔHHI approaches 5,000, divestiture rates reach almost 90%. *The result is completely congruent with the legal standard set forth by the Clayton Act.* Enforcement is absent from ownership changes that do not "substantially lessen competition" but frequent among ones that "tend to create a monopoly."




Figure I: Exempt mergers almost completely escape antitrust scrutiny.

This figure plots an indicator for divestiture on the y-axis against HHI changes on the x-axis. The unit of observation is a facilities acquisition. To improve legibility, I bin the data into equally sized cells according to x-axis values and plot the means within each bin. The dashed line reflects reportable facilities acquisitions, and the solid line reflects exempt ones. To judge the robustness of this result, see Appendix Table A3. For confidence intervals around each plotted point, see Appendix Figure A3.

In sharp contrast, the solid line corresponds to exempt facilities acquisitions. *Here, the divestiture rate is zero regardless of the HHI change.* In other words, exempt mergers almost completely escape antitrust enforcement. In many cases, ΔHHI exceeds 2,000. For some, ΔHHI nears 5,000. Even facilities acquisitions that create monopolies and duopolies are not arrested. The lack of enforcement is even more surprising when juxtaposed with the 2010 *Horizontal Merger Guidelines*. They state, "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 of these facilities acquisitions represent concentration increases that are an order of

magnitude higher than what the agencies presume to be anticompetitive.

To be clear, these findings in no way cast blame on the FTC. If anything, they suggest the opposite: in the subset of cases for which I am certain that it is adequately informed, its staff enforce the law as Congress wrote it. Instead, the disparate treatment of exempt mergers probably reflects how difficult it is to discover private contracts shortly after they are written, echoing the even more widespread problems that preceded the HSR Act. As one of the bill's sponsors, Senator Philip Hart, stated, "The only method the [DOJ] 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 [DOJ]."

Alternative specifications discourage other interpretations. The leading candidate is prosecutorial preference over the target's size. If big transactions attract media attention and FTC staff are drawn to "headline" deals, then large mergers will face different divestiture rates irrespective of reportability. To assess this possibility, I limit the sample to reportable facilities acquisitions, which ensures that informational differences do not affect enforcement. I then plot divestiture against the size of the merger with which each facilities acquisition is associated, controlling for differences in premerger market structure. Panel A of Appendix Figure A4 reports no relationship between merger size and divestiture rate. Likewise, I test for prosecutorial preference over the size of the acquirer. Panel B of the same figure reports no relationship between acquirer size and divestiture rate.

I also investigate whether the results are sensitive to the way in which markets are defined. To do so, I redefine their boundaries entirely by drawing a uniformly sized circle around each facility. Appendix Table A3 reports that the results are unchanged. The table addresses other confounds (e.g., coincidental secular trends) as well. Finally, it reports that, conditional on HHI changes, HHI levels do not determine divestiture rates.

4.2 Consolidation

These results beg an equally fundamental question. Do the divestitures associated with premerger notifications produce lasting changes in market structure? Notice that persistence is not guaranteed. If acquirers of divested facilities struggle to operate those locations, then the facilities will shrink or even close, rendering the remedies ineffective.

To answer this question, I split the sample into reportable and exempt facilities acquisitions and plot predicted concentration changes (i.e., ΔHHI) against realized ones. Exempt facilities acquisitions are expected to lie near the 45 degree line, as they escape enforcement. However, where reportable facilities acquisitions will lie is theoretically ambiguous. If divestitures are unsuccessful, then they will also lie close to the 45 degree line, but if they prevent consolidation, then they will lie far below it.

Figure II reports the result. As expected, exempt facilities acquisitions lie along along the diagonal. In contrast, realized HHI changes exhibit almost no relationship with predicted ones. In Panel A, which measures postmerger concentration the year after the transaction, the result is almost mechanical, as divested facilities are unlikely to immediately falter. In Panel C, however, which measures postmerger concentration five years later, the result implies that premerger notifications have persistent effects on market structure.

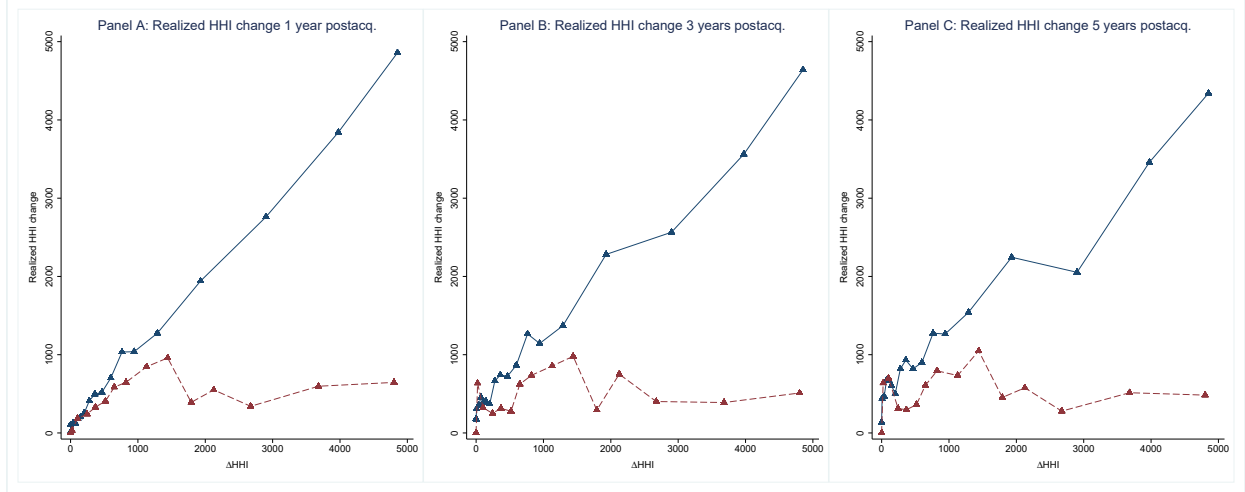


Figure II: Premerger notifications have lasting effects on market structure.

This figure plots realized HHI changes on the y-axes against predicted ones (i.e., ΔHHI) on the x-axes. The unit of observation is a facilities acquisition. To improve legibility, I bin the data into equally sized cells according to x-axis values and plot the means within each bin. Solid lines reflect exempt facilities acquisitions, whereas dashed lines reflect reportable ones. "Realized HHI change" means observed HHI after the facilities acquisition less observed HHI just prior to it. In Panel A, I measure HHI one year after the facilities acquisition. In Panels B and C, I measure it three and five years after.

4.3 Quality

The data indicate that premerger notifications prevent consolidation, and the model implies that consolidation reduces equilibrium quality. Together they predict that notifications improve quality. Are the data consistent with this prediction?

To answer this question, I specify a staggered event study. I map each facilities acquisition to the target facilities with which it is associated and the year in which it occurred. Each facility-level ownership change constitutes an event. I then compare changes in quality over time around the events. For this exercise, the unit of observation is a facility-year, and the main outcome is a risk-adjusted survival rate. For a more complete sense of clinical consequences, I also consider risk-adjusted hospitalization rates.

The comparisons are conceptually straightforward. First, recall from Section 3.3 that, all else equal, large realized concentration changes reduce quality more than small ones do. Second, recall from Sections 4.2-4.3 that the relationship between predicted and realized concentration changes is nearly one-to-one for exempt facilities acquisitions but is severely attenuated for reportable facilities acquisitions. If these conditions hold, then the extent to which large ΔHHI reduce quality more than small ΔHHI will be greater when $HSR = 0$ than $HSR = 1$.

To ease exposition, I denote "large" HHI changes by $\widetilde{\Delta HHI}$ and refer to the others as "small." Specifically, I define $\widetilde{\Delta HHI}$ equal to one if ΔHHI exceeds its mean value and zero otherwise. For the same reason, I define $EX = 1 - HSR$. The estimating equation is then given by

$$y_{jt} = \lambda_t + \lambda_j + \lambda_{jt}^c + \sum_{\tau=-7}^7 \psi_\tau d_{jt}^\tau + \sum_{\tau=-7}^7 \kappa_\tau d_{jt}^\tau EX_a + \sum_{\tau=-7}^7 v_\tau d_{jt}^\tau \widetilde{\Delta HHI}_a + \sum_{\tau=-7}^7 \mu_\tau d_{jt}^\tau EX_a \widetilde{\Delta HHI}_a + e_{jt}. \quad (5)$$

y_{jt} denotes quality. λ_t , λ_j , and λ_{jt}^c denote year, facility, and chain fixed effects, respectively. d_{jt}^τ represent

dummies for values of τ , which denotes event time.³¹ Given the sample period, I restrict attention to $\tau \in [-7, 7]$. To expedite comparisons in event time, I normalize τ -subscripted coefficients to zero just prior to each facilities acquisitions. Coefficients on the final set of interaction terms are the main interest.

Figure III reports estimates obtained from equation 5. Panel A equates quality with survival. Two features are especially striking. First, $\mu_\tau < 0$ for $\tau > 0$, consistent with the prediction that premerger notifications improve quality. In other words, the estimates suggest the HSR Act works as intended—it facilitates enforcement, which prevents consolidation that would otherwise harm consumers. Second, pre-event trends are absent, indicating that the first feature of the graph is not an artifact of persistent deterioration of certain facilities relative to others. In Panel B, risk-adjusted hospitalization rates replace survival. The fact that an empirically distinct but clinically similar measure of quality yields a similar conclusion is reassuring.




Figure III: Estimates from staggered event study

This figure plots estimates of μ_τ obtained from 5 against event time. μ_τ is normalized to zero in $\tau = -1$, which is marked with a vertical red line. Quality is measured using risk-adjusted hospitalization rates in Panel A and risk-adjusted survival rates in Panel B.

Appendix Table A4 considers alternative specifications. As in the earlier section, I redefine markets entirely by drawing a uniformly sized circle around each facility. Separately, I recompute ΔHHI using market shares that are based on the number of stations at the facility (rather than the number of patients treated by it). Finally, I re-estimate equation 5 on a restricted set of markets. Regardless of whether I measure quality using hospitalization or survival rates, estimates from the baseline specification and the alternatives are close to one another.

These findings can reconcile recent work on dialysis. Eliason [2018] relies on very different variation than what I employ here but similarly concludes that, all else equal, consolidation reduces quality. However, Cutler et al. [2016] study mergers between national chains, such as DaVita’s \$3 billion purchase of Gambro

³¹If j is involved in more than one exempt facilities acquisition, I restrict attention to the first of these ownership changes; however, this occurs infrequently, and restricting attention to the last of them does not change the results.

in 2005, and find no relationship between consolidation and quality. At first glance, their results appear at odds with one another, but upon closer inspection, they are precisely what one expects in light of the findings I reported in Section 4.1-4.2. By definition, mergers between national chains are large, so by law, they are reportable. Thus, among these ownership changes, the ones most likely to reduce quality are also most likely to require divestiture (i.e., never be completed). In other words, Cutler et al. [2016] study precisely the types of deals that face antitrust scrutiny, thereby introducing a sort of empirical puzzle that this paper solves.

These findings complement other recent work. 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 rather than concentration of ownership). 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. As alluded to earlier, my findings are consistent with their main result—acquisitions by chains reduce quality independent of competitive effects.³² Moreover, I reach the same broad conclusion about imperfect competition.³³

4.4 Agency costs

Finally, I examine the empirical relationship between agency caseloads and costs. For this exercise, I rely on variation provided by an amendment to the HSR Act. The law change, effective near the start of 2001, abruptly raised the transaction value threshold from \$15MM to \$50MM. To focus on this variation, I restrict attention to a four year window on each side of the amendment. I then plot premerger notifications and PNO employment over time.

Figure IV reports the result. PNO inputs and outputs are comparatively stable in the four years leading up to the amendment but decline sharply at that point. Premerger notifications fall around 75%, while PNO labor falls around 33%. The patterns have a natural explanation: as PNO caseloads declined, the FTC reduced the office's staff. Investigations and post-PNO labor also fell around this time (see Section 5.4 below).

5 Estimation

Concentration indices are one dimensional approximations of complicated substitution and ownership patterns, to which quality responds non-linearly in equilibrium, so they cannot reliably and accurately be used to forecast changes resulting from an alternative premerger notification policy. For this, one needs to explicitly model the response of market participants.

³²Code that produces Appendix Table A4 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. This is not a surprise—we rely on very similar data sources/samples.

³³Further along in the analysis, one of their specifications suggests consolidation does not affect quality, but it defines competition as any target-acquirer overlap within very broadly defined geographic areas. Estimates from their specifications that use more narrowly defined markets are consistent with mine. Note that the relationship between competition and quality is not central to their paper, and choosing one specification over another does not alter their main result. Appendix E.4 provides more detail.

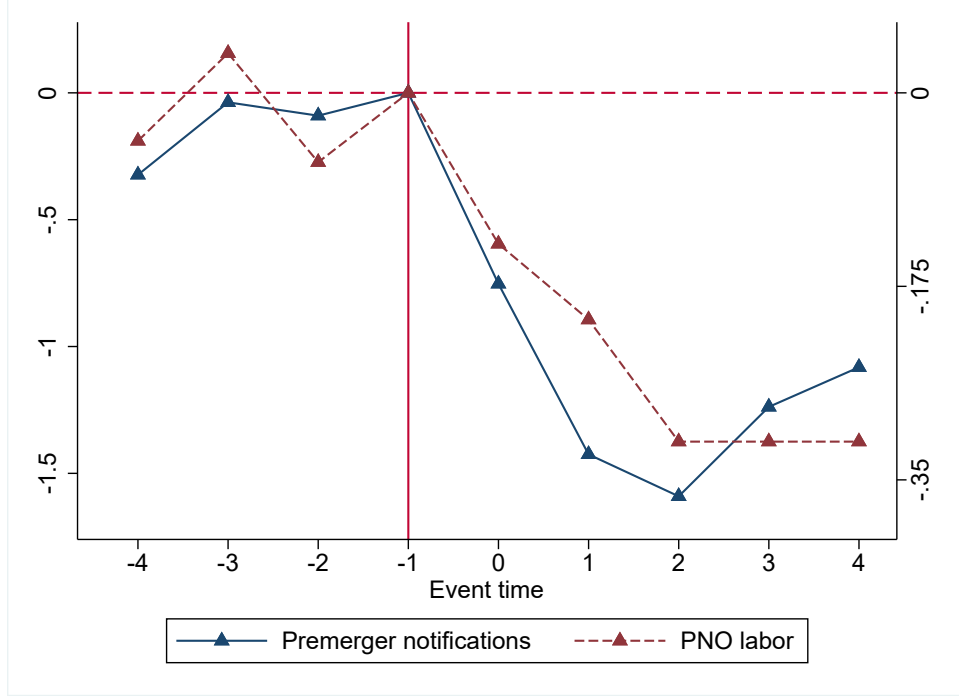


Figure IV: Premerger notifications and PNO employees decline following the HSR amendment

This figure plots the log number of premerger notifications on the primary y-axis, the log number of PNO employees on the secondary y-axis, and event time (in fiscal years relative to 2001) on the x-axis. The vertical red line corresponds to 2000, the year just prior to the one in which the HSR threshold was raised. Y-axis values are normalized to zero in that year. Discrepancies between the time path of premerger notifications reported here and in Figure I of Wollmann [2019] reflect calendar/fiscal year differences.

5.1 Enforcement

Based on the results reported in Section 4.1, I assume that the agency enforcement function can be approximated with a low-order polynomial:

$$D_a = \theta_0 + \theta_1 \text{HSR}_a + \theta_2 \Delta \text{HHI}_a + \theta_3 \text{HSR}_a \Delta \text{HHI}_a + \theta_4 \text{HSR}_a \Delta \text{HHI}_a^2 + \eta_a. \quad (6)$$

Equation 6 accounts for the concavity of divestiture rates in HHI changes and omits factors, such as HHI levels, with which divestiture rates do not vary.

Expectedly, I estimate that θ_0 , θ_1 , and θ_2 are very close to zero. Each coefficient and its corresponding standard error lies between 0 and 0.005. By contrast, I estimate that θ_3 and θ_4 equal 0.49 and -0.066, respectively. Both estimates are significant at the 1% level. In the interests of space (and in light of the fact that the graphical representation of these relationships are displayed in the previous section), I report these estimates in the Appendix (see Table A3).

5.2 Demand

Methods. To estimate demand, I begin by substituting δ_{jt} for the patient-invariant components of utility and rewriting equation 1 as

$$u_{ijt} = \delta_{ijt} + \alpha_1 \text{dist}_{ijt} + \alpha_2 \text{dist}_{ijt}^2 + \alpha_3 \text{dense}_{g(j)t} \text{dist}_{ijt} + \alpha_4 \text{dense}_{g(j)t} \text{dist}_{ijt}^2 + \epsilon_{ijt}. \quad (7)$$

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 obtain $\delta(\tilde{\alpha})$, the vector of δ_{jt} that exactly fit market shares under $\tilde{\alpha}$ [Berry, 1994], compute the individual choice probabilities at $\tilde{\alpha}$ and $\delta(\tilde{\alpha})$, and compute the log likelihood of observing the choices patients make in the data.³⁴ 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, x may depend on ζ , so regressing $\hat{\delta}$ on it (and the other observables) produces biased estimates. The solution lies in the evidence reported in Section 4: exempt facilities acquisitions involving large HHI changes reduce healthcare quality but are plausibly uncorrelated with demand shocks. To operationalize it, I define

$$z_{jgt} = \sum_a \text{EXEMPT}_a \widetilde{\Delta \text{HHI}}_a \{j \in \mathcal{J}_a\} \left(\{t \in [t_a, t_a + \rho]\} - \{t \in [t_a - \rho, t_a]\} \right). \quad (8)$$

\mathcal{J}_a denotes the set of facilities associated with a for which there are at least ρ pre- and postacquisition observations. t_a denotes the year it was completed. $\{\cdot\}$ denotes an indicator that takes a value of one if the condition within the braces is satisfied and zero otherwise. The first two terms select exempt facilities acquisitions involving large predicted concentration changes (with the precise construction of $\widetilde{\Delta \text{HHI}}_a$ is described below). Following the discussion in Section 4.2 and the fact that the number of stations are stable over time, I assume that $\mathbb{E}[z\zeta | t, c, n] = 0$.

I construct a matrix Z such that its columns include a constant vector, n , z , and vectors of year and chain dummies. I assume $\mathbb{E}[\zeta | Z] = 0$ and minimize a GMM objection function whose expected value is given by $\mathbb{E}[\zeta' Z Z' \zeta]$, which is based on moments implied by this restriction [Berry et al., 1995]. $[\hat{\beta}_0 \hat{\beta}_x \hat{\beta}_n \hat{\phi} \hat{\chi}]$ denotes the vector of parameters that minimize the empirical analog of the function.

Results. 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. Estimates related to distance (in road miles) traveled are especially precise, due at least in part to the large number of patient-year observations in the panel.

The magnitudes of the coefficients imply sensible tradeoffs. On average, patients are indifferent between

³⁴In 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.

Table IV: Demand and marginal cost estimates

Panel A. Demand estimates		
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. Marginal cost estimates		
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 road miles traveled (differing from a previous version of the paper, which used "great circle" distances). Population density is measured in 10,000s of people per square mile. Marginal costs are reported on a per-treatment basis. Robust standard errors are in parenthesis.

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 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,³⁵ the additional travel equates to \$1,248 per year. Notably, if I multiply this figure by 100, thereby translating a percentage point increase in risk-adjusted survival rates into the statistical value of a life-year, then I arrive at \$124,800, which is close to figures commonly arrived at in the literature studying "health capital." For comparison, Cutler and Richardson [1998] and Lee et al. [2009] arrive \$150,040 and \$147,800, respectively, when adjusted for inflation.

5.3 Marginal cost

Methods. To recover marginal costs, I construct two sets of moments. To obtain the first set, I take the first order condition of equation 2 with respect to quality. Substituting mc_{jt} in equation 3 with its parameterization yields that for all k in \mathcal{J}_{fgt} ,

$$\sum_{j \in \mathcal{J}_{fgt}} \left((P_{gt} - \gamma_0 - \gamma_1 x_{jt} - \gamma_2 x_{jt}^2 - \gamma_3 n_{jt} - \omega_{jt}) \frac{\partial s_{jt}}{\partial x_{kt}} \right) - (\gamma_1 + 2\gamma_2 x_{kt}) s_{kt} = 0. \quad (9)$$

To arrive at a system of equations that are additively separately in the marginal cost disturbances, I construct a matrix A_{gt} such that each (j, k) element equals $\frac{\partial s_{jt}}{\partial x_{kt}}$ if j and k are jointly owned and zero otherwise. I then multiply the linear algebra representation of equation 9 through by the inverse of A_{gt} . By doing so, I arrive at

$$P_{gt} - \gamma_0 - x_{gt} \gamma_1 - x_{gt}^2 \gamma_2 - n_{gt} \gamma_3 - \omega_{gt} - A_{gt}^{-1} s_{gt} \gamma_1 - 2A_{gt}^{-1} x_{gt} s_{gt} \gamma_2 = 0. \quad (10)$$

P_{gt} , x_{gt} , n_{gt} , ω_{gt} , and s_{gt} denote vectors of prices, quality, machines, marginal cost disturbances, and market shares, respectively, corresponding to (g, t) . I construct matrix W , whose columns consist of the following: a constant, n , $\{z = 1\}$, $\{z = -1\}$, and P . I assume $\mathbb{E}[\omega|W] = 0$ and form the first set of moments by interacting W with ω .

To obtain an additional moment, I exploit information on profit margins.³⁶ Throughout the panel, the largest US dialysis provider has published firm-wide annual average per-treatment revenues and costs of their domestic facilities.³⁷ 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 profit margins. The moment is given by

$$\sum_t \sum_g \sum_{j \in \mathcal{J}_{fgt}} P_{gt} (1 - \text{MARGIN}_{ft}) - \gamma_0 - x_{jt} \gamma_1 - x_{jt}^2 \gamma_2 - n_{jt} \gamma_3 - \omega_{jt} = 0, \quad (11)$$

³⁵According 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.

³⁶Dollar-denominated survey responses in HCRIS reports are unreliable, evidenced by implausibly low or high values. Otherwise, they could inform facility-specific cost estimates.

³⁷See, e.g., DaVita's 2017 Annual Report at <https://investors.davita.com/static-files/ca97d703-5935-4708-9fcb-4d527b258651>.

where $MARGIN_{ft}$ represents the (average) profit margin reported by DaVita in t . I construct a GMM objective function based on the full set of moments. $\hat{\gamma} \equiv [\hat{\gamma}_0 \hat{\gamma}_1 \hat{\gamma}_2 \hat{\gamma}_3]$ denotes the vector of parameters that minimize the function's empirical analog.

Results. Panel B of Table IV reports marginal cost estimates. Coefficients on x and x^2 are -2,428.8 and 1,661.9, respectively, meaning marginal cost is increasing at an increasing rate for $x > 73\%$, which applies to effectively every facility-year observation. At the median survival rate, a one percentage point increase in quality increases adds \$4.3 to the cost-per-treatment. Costs are also 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.

The average profit margin is 31.3%, and the average price-cost markup is 50.8%. These measures rise around stealth consolidation, which results in lower quality. For instance, in the three year period immediately preceding an exempt facilities acquisition involving an above-median HHI change, profit margins average 31.2%, but in the three years following one of these events, profit margins average 34.3%—an increase of just over 10%.

5.4 Agency costs

Methods. Recall from Section 4.4 that an amendment to the HSR Act, effective near the start of 2001, abruptly raised the transaction value threshold below which mergers are exempt from premerger notification. To estimate the productivity of labor, I exploit this same temporal variation in caseloads introduced by the law change. The patterns observed in Figure IV map to Φ_1^{PNO} in a transparent way: given the reduction cases, a small reduction in headcount implies highly productive labor, while a large reduction implies the opposite.

For the PNO, the first stage estimating equation is given by

$$L_t^{PNO} = \zeta_0 + \zeta_1 AMEND_t + e_t, \quad (12)$$

and the second stage estimating equation is given by

$$C_t^{PNO} = \Phi_0^{PNO} + \Phi_1^{PNO} \hat{L}_t^{PNO} + u_t^{PNO}. \quad (13)$$

$AMEND_t$ takes a value of one when $t > 2001$ and zero otherwise. \hat{L}_t^{PNO} is the predicted number of PNO employees from equation 12. u_t^{PNO} represents measurement error in C_t^{PNO} . To focus on variation derived from the amendment, I restrict attention to a four year window around the amendment (i.e., ones for which $|t - 2001| \in \{1, 2, 3, 4\}$). For post-PNO activities, analogous estimating equations yield an estimate of Φ_1^{PP} . Notice the $AMEND_t$ represents almost ideal variation for my counterfactual analysis, as the amendment and alternative policy affect similarly sized mergers.³⁸ Consequently, estimates obtained from equation 13 should produce realistic forecasts of agency labor costs will respond.

³⁸To be precise, as it pertained to dialysis mergers, the transaction value threshold was about \$15MM prior to February 2001, \$50MM between February 2001 and February 2005, and a value determined by gross national product growth thereafter. Weighted by mergers per year, the threshold averaged about \$58 million during the sample period. Thus, to a first approximation, the amendment raises the threshold from about \$15 million to \$50 million, while the alternative policy cuts it from \$58MM to \$0.

Not only does the FTC prevent anticompetitive mergers but also harmful non-merger restraints, deceptive advertisements, fraudulent sales schemes, and identity theft. Across these tasks, the number of employees required to review, investigate, and/or prosecute a case may differ, but each employee’s capital requirements are mostly the same—office space, furniture, and fittings near the National Mall as well as a modern workstation to analyze information and communicate ideas. I exploit this fact to estimate capital productivity. Specifically, I assume that the capital and labor required for each case type are proportional to one another. Under this assumption, one can regress agency-wide labor on capital and combine the result with estimates of Φ_1^{PNO} and Φ_1^{PP} to obtain estimates of Ψ_1^{PNO} and Ψ_1^{PP} , respectively. (Appendix B describes the process in detail).

Results. Table V reports the agency production function estimates. Column 1 corresponds to PNO labor productivity (Ψ_1^{PNO}). Conditional on the FTC employing a sufficient amount of capital, employees evaluate an average of 299 premerger notifications each year. Assuming 225 productive working days each year (i.e., 250 business days less three weeks paid time off and ten administrative, training, and development days), the estimate equates to 0.8 evaluations per day. To assess the reasonableness of this figure, recall that it applies to small mergers. Typically, these transactions involve targets occupying a small number of markets in a single industry, and they very rarely raise competitive concerns due to the deterrent effect of enforcement.³⁹ According to a former PNO employee with whom I spoke, it is reasonable to expect about one straightforward evaluation is completed each day.

Table V: Agency productivity estimates

VARIABLES	(1) Notifications	(2) Investigations	(3) Agency-wide capital
PNO labor	299*** (38.0)		
Post-PNO labor		2.89*** (1.00)	
Agency-wide labor			0.67*** (0.12)
Observations	8	8	22
First stage F-stat.	57.41	55.03	–
Days per notif. or investigation	.753	77.732	–
Mean capital rental per employee	–	–	79919
Productivity of PNO capital	–	–	198.92
Productivity of post-PNO capital	–	–	1.93

, **, and * denote significance at the 10%, 5%, and 1% levels, respectively. The unit of observation is a year. Capital is denominated in millions of 2017 dollars. The bottom two rows of the table report estimates of Ψ_1^{PNO} and Ψ_1^{PP} , respectively. Robust standard errors are in parentheses.*

Column 2 corresponds to post-PNO labor productivity (Ψ_1^{PP}). Employees investigate an average of 2.9 mergers each year. In other words, each merger consumes around 78 working days. As a former merger enforcement employee explained to me, one should not interpret this figure to mean one employee works

³⁹In equilibrium, small mergers are disproportionately deterred by the prospect of enforcement. The reasons are clearly described in the following section and a detailed note Appendix A.

this long on a single small merger. Rather, the typical investigation into a merger of this size last around one month but involves part-time contributors from several staff.

Despite the small number of observations, the first stage F-statistics are 57 and 55, respectively. Similarly, both second stage coefficients are significant at the 1% level. These statistics reflect sharp declines in caseloads and employment following the amendment.

Column 3 reports coefficients obtained from regressing agency-wide capital on labor. The estimate implies that the marginal FTC employee requires \$79,919 of rental capital annually, which is sensible in comparison to the sum of individual cost components. The largest of these is rent, as the agency is located in the Federal Triangle section of the nation's capital. The *Global Prime Office Occupancy Cost Report*, published by CBRE, the world's largest commercial real estate firm, states that total occupancy costs in the Washington, D.C. market average \$90 per square foot at the end of the sample period. At 395 total square feet per federal government employee [Miller, 2012], rent alone costs \$35,500. Assuming a 15% depreciation rate and using figures reported in CBRE's *Global Office Fit-Out Cost* guide, all other workplace improvements—an ergonomic workstation, a secure computing and telecommunications environment, other furnishings and equipment, and lessor markups—cost an additional \$9,000 per year. This leaves \$35,000 to cover all other expenses recruiting, supporting, and retaining a highly skilled DC Metro area office worker. These include travel, specialized software subscriptions for economic analysis, outreach, training, on-boarding, payroll services, and amenities. The last of these comprise an on-site agency-sponsored daycare center, exercise facility, and health unit as well as subsidies to park or commute using public transit.

6 Counterfactual policy analysis

In this section, I study an alternative policy that eliminates premerger notification exemptions in the dialysis industry. I retroactively apply this policy over the full sample and simulate counterfactual equilibrium outcomes at the estimated parameter values. The goal is to forecast costs and benefits.

6.1 Merger activity and market structure

The first step is to predict how market structure would evolve under the alternative policy. To simulate its effect, I apply the divestiture rates historically experienced by reportable mergers to exempt ones. In other words, under the alternative policy, an exempt facilities acquisition that would have been blocked with probability $D(\cdot, HSR = 0)$, which is nearly zero, is instead blocked with probability $D(\cdot, HSR = 1)$, which is often much greater than zero.

Since divestitures are probabilistic, market structure evolves stochastically. To see this clearly, suppose there exists a geographic market g where exactly one facility j changes ownership during the sample period. To fix ideas, suppose acquirer f buys j from target f' in year t . Further, suppose this represents an exempt facilities acquisition involving a 2,000 point HHI change. Finally, recall from Figure I that a reportable facilities acquisition involving a 2,000 point HHI change faces a 70% chance of divestiture. Thus, starting in t , there are two possible states of the world: one in which j is owned by f , which occurs with 30% probability, and one in which it does not, which occurs with 70% probability.

To accurately characterize market structure under the alternative policy, I allow for deterrence [Clougherty

and Seldeslachts, 2013, Wollmann, 2019].⁴⁰ 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 acquirer facing a high risk that a relatively small merger will be prevented will not 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 is made aware, the alternative policy deters many anticompetitive facilities acquisitions, which means the agency never actually needs to challenge—let alone review—these transactions. To account for deterrence and significantly reduce the computational burden, I assume that no firm proposes a merger where the average rate of divestiture across markets exceeds 50%. (The actual cutoff is probably much lower. See Appendix A for detailed calculations.)

The data, parameter estimates, and deterrence condition fully determine the possible states reach in each year and the probability at which each is reached. They also fully determine the expected number of mergers and investigations that occur each year. The process used to enumerate the possible states and compute the related objects is conceptually uninteresting, so I relegate it to Appendix C.

6.2 Benefits of the alternative policy

The second step is to compute benefits, which depend on the number of lives saved. Let r denote a state of the world, \bar{r} denote the state I observe in the data, and $\phi(r)$ denote the probability that r occurs. Also, let \mathcal{R}_t denote the set of possible states reached by t , and let R_t denote the cardinality of that set. Additionally, let x_{jt}^r denote the equilibrium quality offered by j at t in state r , and define s_{jt}^r analogously. Finally, let \mathcal{J}_t denote the number of facilities treating patients in t . The number of lives saved by the alternative policy each year is given by

$$\mathcal{L}_t = \sum_{j \in \mathcal{J}_t} M_{g(j)t} \left[\frac{1}{R_t} \sum_{r \in \mathcal{R}_t} \phi(r) x_{jt}^r s_{jt}^r - x_{jt}^{\bar{r}} s_{jt}^{\bar{r}} \right]. \quad (14)$$

Based on Section 5.2, I assume the statistical value of a life year equals \$125,000 in constant 2017 dollars. Using figures from Cutler and Richardson [1998] or Lee et al. [2009] would result in larger benefits. For example, Cutler and Richardson [1998] uses \$100,000 in 1998, which equals \$156,000 when adjusted for inflation.

Note that when computing equilibrium quality, one must decide whether to update the values of the chain fixed effects. Eliason et al. [2019] demonstrate that acquisitions by large chains harm quality independent of market power, and the alternative policy deters many of these transactions, updating the chain fixed effects will yield even greater benefits of the alternative policy. However, my analysis focuses on competitive effects, so I opt not to include them.

Also, note that benefits of the alternative policy do not end in 2017 simply because the 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 total value of lives saved, I let the differences in quality induced by the alternative policy persist for 3, 5, or 7 years after the sample ends. The benefits of eliminating premerger notification exemptions in the dialysis industry between 1997 and 2017 are then

⁴⁰Nocke and Whinston [2010] consider dynamic merger strategies. Their approach, which permits firms to adjust their acquisition plans in response to agency policy changes, is simply too computationally demanding in my setting.

given by

$$\$156,000 \times \left[\sum_{t=1997}^{2017} \mathcal{L}_t + \mathcal{L}_{2017} \times \Gamma \right], \quad (15)$$

where $\Gamma \in \{3, 5, 7\}$. The value of expression 15 ranges from \$1.1 billion to \$1.4 billion. The figure comprises an average per-year benefit of \$39.5 million and a terminal value of between \$235.8 and \$547.8 million.

6.3 Costs of the alternative policy

The third step is to compute costs. Under the alternative policy, the FTC must evaluate more premerger notifications and investigate more mergers, which requires more staff and capital. Let ΔC_t^{PNO} denote the difference in t between the expected number of premerger notifications evaluated under the alternative policy and the number observed in the data, and define ΔC_t^{PP} analogously. Also, let $\hat{\Phi}_1^{\text{PNO}}$ denote my estimate of Φ_1^{PNO} , and define $\hat{\Psi}_1^{\text{PNO}}$, $\hat{\Phi}_1^{\text{PP}}$, and $\hat{\Psi}_1^{\text{PP}}$ analogously.

The expected costs of eliminating premerger notification exemptions in the dialysis industry between 1997 and 2017 is then given by

$$\sum_{t=1997}^{2017} \left[\Delta C_t^{\text{PNO}} \left(\frac{w_t}{\hat{\Phi}_1^{\text{PNO}}} + \frac{r_t}{\hat{\Psi}_1^{\text{PNO}}} \right) + \Delta C_t^{\text{PP}} \left(\frac{w_t}{\hat{\Phi}_1^{\text{PP}}} + \frac{r_t}{\hat{\Psi}_1^{\text{PP}}} \right) \right]. \quad (16)$$

The value of expression 16 is \$9.9 million. PNO costs sum to \$1.2 million, while post-PNO activities account for the remainder.

To put these figures in context, I turn to the 2017 FTC Congressional Budget Justification and HSR Annual Report. In that year, the agency budgeted \$4.075 million for PNO activities evaluated 1,519 premerger notifications. Multiplying the average cost per PNO case implied by these figures by 1,511, the number of additional premerger notifications evaluated between 1997 and 2017, yields \$7.25 million. Repeating this exercise for post-PNO activities yields \$39 million. Given that the median merger currently reported to the FTC is 30-50 times larger than the ones that would be newly reportable as a result of alternative policy, expected costs of \$1.2 million and \$8.7 million, respectively, seem sensible.

6.4 Discussion

The benefits of the alternative policy are two orders of magnitude larger than the costs. Why are the latter so comparatively 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 only a few million dollars. 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 total transaction values is striking. Partly they reflect deterrence. Many of the mergers that are newly reportable under the alternative policy and likely to be investigated are never attempted in the first place.

Even so, these findings do not suggest the immediate, economy-wide elimination of exemptions, especially because net benefits may vary from one industry to another based on various environmental factors. The most pivotal ones are described below.

One factor is the burden faced by merging parties. Even absent fees, premerger notifications are costly to complete. The target and acquirer must enumerate subsidiaries, disclose revenues by NAICS industry

code (and NAPCS product code, if either engages in manufacturing), and transmit translation-related "studies, surveys, analyses, and reports" prepared by or for any officer, board member, investment banker, consultant, or advisor. These expenses may discourage mergers that generate efficiencies [De Loecker et al., 2008], which are passed through to consumers. The prevalence of welfare-improving mergers, let alone the likelihood that they are deterred by compliance costs, is not known in most industries. In dialysis, however, the matter is mostly settled. Holding fixed market structure, acquisitions by large chains are associated with lower—not higher—quality [Eliason et al., 2019].

Moreover, clever implementation of the expansion could drastically reduce the burden on merging parties. The leading example is provided by Scott Morton, who argues that a simple, interactive system for basic disclosures of small mergers would produce sufficient information. For instance, if both parties solely provide in-center dialysis treatments, then it may suffice for the system to just query the ZIP codes in which each operates. Further, such a system could lower the government's costs by automatically evaluating and screening transactions (or at least providing very concise reports for PNO staff).⁴¹

Another factor relates to the intensity of and dimensions on which firms compete. Entry barriers are an archetypal example. In markets where mergers among rivals induce rapid, robust entry that drives quality up and/or prices down, the effect of eliminating exemptions will be muted. The same goes for consumers in markets where existing firms can quickly and cheaply reposition themselves. How the speed and cost of firm- and product-level entry and exit vary across industries is a mostly open empirical question.

A third factor relates to how agencies detect mergers that are currently exempt. Customer complaints exemplify this consideration. In markets where firms closely monitor their suppliers and diligently report adverse consolidation to the agencies, the effect of eliminating exemptions will be muted. Judging by the frequency at which attorneys warn their clients not to provoke antitrust complaints from customers in the immediate wake of an exempt merger (see, e.g., Petkoski et al. [2019]), substitution almost certainly occurs in some industries. However, this mechanism fails in the dialysis industry, which may be the rule rather than exception.

Further, even when other means of detection exist, they might be highly inefficient when compared to premerger notification. In the case of customer complaints, most give rise to investigations into already completed deals—ones for which the "eggs" may already be "scrambled." Consistent with this concern, former FTC Commissioner J. Thomas Rosch estimates that consummated mergers "take on average twice as long to complete as investigations of unconsummated mergers" [Rosch, 2012].

A fourth factor relates to noncompliance. HSR Rule 801.90 states that the substance of a transaction determines its reportability—structures whose purposes are to avoid filing obligations are ignored. Historically, though, the rule has been narrowly interpreted, making noncompliance a complicated issue. Problems are unlikely in industries such as dialysis, where mostly public companies employ boilerplate purchase agreements, but the full scope of the problem is unknown. (See Footnote 7 and the detailed discussion in Appendix D.3.)

A fifth factor relates to additional exemptions that may be unaligned with economic reality. In these industries, eliminating exemptions may produce much larger net benefits than in dialysis. High-technology markets are perhaps the best candidates. Since intellectual property is incompletely capitalized under US Generally Accepted Accounting Principles until its ownership changes, extremely valuable targets may

⁴¹See Scott Morton's Testimony before the House Judiciary Committee titled "Diagnosing the Problem: Exploring the Effects of Consolidation and Anticompetitive Conduct in Healthcare Markets" and in particular her discussion of lower thresholds.

nonetheless report almost no assets. Likewise, these firms may have little to no revenue as they wait for licenses to be granted or distribution to be setup. Without assets (or sales, in the case of manufacturing), the merger fails the "size-of-persons" test, so a much higher "size-of-transaction" threshold applies. At the time of writing, if the target has less than \$22.3 million in assets (or sales, if it engages in manufacturing), then the transaction value must be at least \$445.5 million, not \$111.4 million, to require premerger notification.

7 Conclusions

In this paper, I found that premerger notifications are essential to antitrust enforcement. In their absence, mergers among close competitors that would otherwise be prevented are, in effect, completed—an eventuality I call "stealth consolidation." In the dialysis industry, this equates to facilities 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 agency guidelines.

Since premerger notifications facilitate enforcement, which prevents consolidation and preserves incentives to compete, they impact consumer welfare. In the present setting, where prices are set by the federal government, internalizing business-stealing incentives means setting quality below perfectly competitive levels. Consistent with that model of behavior, I found that exempt facility acquisitions involving large HHI increases reduce quality, as measured by risk-adjusted hospitalization and survival rates. Consistent with effective antitrust enforcement, I found that reportable facility acquisitions do not affect quality because the FTC prevents ownership changes that would reduce incentives to compete.

These facts raise the question as to whether society would benefit from eliminating premerger notification exemptions in the dialysis industry. To answer it, I estimate a structural model of antitrust enforcement, demand, and quality-setting, which allows me to predict equilibrium outcomes under any hypothetical market structure and merger activity. I then retroactively apply the divestiture rate historically experienced by reportable mergers to exempt ones. The alternative policy I propose results in less concentrated markets that offer higher quality. Over the full sample, it saves life-years worth around one billion dollars. Comparatively, the costs of the alternative policy are *much* lower. 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 prevent them. This also reflects the efficiency of the FTC, which has effectively policed billions of dollars in reportable dialysis mergers using very comparatively few resources.

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.
- Aslihan Asil, John M. Barrios, and Thomas G. Wollmann. Misaligned measures of control: Misaligned measures of control: Private equity's antitrust loophole. *Virginia Law & Business Review*, 18(1), 2023.
- 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. 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.

John M. Barrios and Thomas G. Wollmann. A new era of midnight mergers: Antitrust risk and investor disclosures. 2023.

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.

Francis P Boscoe, Kevin A Henry, and Michael S Zdeb. A nationwide comparison of driving distance versus straight-line distance to hospitals. The Professional Geographer, 64(2):188–196, 2012.

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.

Department of Justice. Fy 2023 performance budget. Technical report, Washington, DC, 2023.

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. What do we know about competition and quality in health care markets?, 2006.

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.

Paul L. E. Grieco and Ryan C. McDevitt. Productivity and quality in health care: Evidence from the dialysis industry. The Review of Economic Studies, 84(3):1071–1105, July 2016.

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.

Norm Miller. Estimating office space per worker, 2012.

Volcker Nocke and Michael D. Whinston. Dynamic merger review. Journal of Political Economy, 118(6): 1200–1251, 2010.

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.

Djordje Petkoski, William Haun, and Caroline Preel. Shearman perspectives, 2019. URL <https://www.lexology.com/library/detail.aspx?g=2639db63-3bd7-4082-8cec-152b52ad6f17>.

Jo Reynaerts, Ravi Varadhan, and John C. Nash. Enhancing the Convergence Properties of the BLP (1995) Contraction Mapping. Technical report, 2012.

Seth Richards-Shubik, Mark S Roberts, and Julie M Donohue. Measuring quality effects in equilibrium. Technical report, National Bureau of Economic Research, 2021.

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.

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.

Online Appendix

A Understanding deterrence

Firms will not pay large upfront managerial, legal, and accounting costs required to negotiate, organize, and structure their transaction unless they are reasonable sure that most of their operations will be merged. Since the alternative policy I study subjects certain ownership changes to much higher divestiture rates than they would otherwise have faced, it may deter many mergers. Deterrence can have important consequences, such as reducing the FTC's caseload, which lowers the policy's social cost. To account for its impact, I assume that mergers with an average divestiture rate across markets of 50% or more are deterred. Simply put, no one goes into a merger expecting to sell off half or more of the target's assets.

The motivation for using this cutoff rule is straightforward. Start from the premise that firms only propose a merger when the expected value it creates for them exceeds the upfront cost. Slightly more formally, merger proposal requires $\sum_a (1 - D_a) \text{VALUE}_a > \sum_a \text{UPFRONT COST}_a$.¹ For simplicity, assume that the expected value and upfront cost are constant across markets. (Doing so is unlikely to bias the conclusions of this exercise, since violations of this assumption are about as likely to affect the dialysis industry as any other.) Let A_m denote the number of markets in which the target operates, and rewrite the preceding condition as $A_m^{-1} \sum_a D_a < 1 - \text{UPFRONT COST} / \text{VALUE}$. Since $A_m^{-1} \sum_a D_a$ represents the average divestiture rate across markets, this inequality is isomorphic to the deterrence condition applied in the counterfactual simulations.

The justification for using this particular cutoff value is also straightforward. Start with the probability of divestiture in any market affected by a particular merger, which is given by $1 - \prod_a (1 - D_a)$. The least favorable case for deterrence occurs when this expression is as large as possible, so maximize its value by setting the number of markets in which the target operates as low as possible.² In other words, set $A_m = 1$, which yields $1 - \prod_a (1 - D_a) = A_m^{-1} \sum_a D_a = D_a$. In this case, firms propose mergers up to a cutoff divestiture rate, denoted by D^* , which satisfies $D^* = 1 - \text{UPFRONT COST} / \text{VALUE}$. Finally, to be even less favorable to deterrence, assume that the unconditional distribution of divestiture rates is skewed such that higher rates occur less frequently. Specifically, assume that $f(D_a) = D_a$. Under this distributional assumption, the proportion of proposed mergers with any divestiture is at least $2D^*/3$. To obtain a reasonable estimate of this figure, I turn to historical data. Over the sample period, HSR Annual Reports indicate that a total of 45,858 reportable mergers were proposed, while the DOJ's *Workload Statistics* and FTC's *Competition Enforcement Database* state that 886 reportable mergers were either restructured, involved divestiture, or were abandoned.³ These figures imply that the cutoff value, D^* , is 2.9% or less—far more conservative than the 50% figure assumed in the text.

¹I implicitly assume that firms know divestiture rates but do not know whether divestitures will actually occur. If they could perfectly predict divestiture, then one would never observe merger abandonment or restructuring resulting from antitrust enforcement. Yet, this commonly occurs—such outcomes historically account for one-quarter of all FTC remedies, according to the agency's *Competition Enforcement Database*. Hence, this is a comparatively weak assumption.

²To see this, suppose the target operates in many markets (e.g., 1000). Except when the average divestiture rate across markets is very small, $1 - \prod_a (1 - D_a) \approx 1$. In this case, unless one sees the agencies negotiate a remedy in almost every proposed merger, one concludes that the cutoff value is very low.

³The DOJ's *Workload Statistics* are available at <https://www.justice.gov/atr/division-operations>, and the FTC's *Competition Enforcement Database* is available at <https://www.ftc.gov/competition-enforcement-database>.

B Agency cost estimation details

Let Ξ index types of cases tasked by FTC staff. According to the agency, such tasks include the protection of consumers against privacy loss, identity theft, deceptive financial, marketing, and advertising practices as well as the promotion of competition through premerger notification review, merger enforcement, and nonmerger enforcement (see, e.g., the 2017 FTC Congressional Budget Justification at 41).⁴ Also, let L_t^Ξ and K_t^Ξ denote the labor and capital employed in t for Ξ . Define $L_t^{\text{FTC}} = \sum_{\Xi} L_t^\Xi$ and $K_t^{\text{FTC}} = \sum_{\Xi} K_t^\Xi$. These quantities represent the agency-wide labor and capital employed in t , respectively. Finally, let $\widehat{\Phi}_1^{\text{PNO}}$ and $\widehat{\Phi}_1^{\text{PP}}$ denote estimates of Φ_1^{PNO} and Φ_1^{PP} , respectively. (Methods used to estimate Φ_1^{PNO} and Φ_1^{PP} are described in Section 5.4 in the main text.)

At cost-minimizing input choices, $C_t^\Xi = \Phi_0^\Xi + \Phi_1^\Xi L_t^\Xi$ and $C_t^\Xi = \Psi_0^\Xi + \Psi_1^\Xi K_t^\Xi$, which implies

$$\Phi_0^\Xi + \Phi_1^\Xi L_t^\Xi = \Psi_0^\Xi + \Psi_1^\Xi K_t^\Xi.$$

Dividing each equality through by Ψ_1^Ξ yields

$$\Phi_0^\Xi/\Psi_1^\Xi + \Phi_1^\Xi L_t^\Xi/\Psi_1^\Xi = \Psi_0^\Xi/\Psi_1^\Xi + K_t^\Xi.$$

Taking the sum of the left- and right-hand sides of each task-specific equality yields

$$\sum_{\Xi} \left[\Phi_0^\Xi/\Psi_1^\Xi + \Phi_1^\Xi L_t^\Xi/\Psi_1^\Xi \right] = \sum_{\Xi} \left[\Psi_0^\Xi/\Psi_1^\Xi + K_t^\Xi \right].$$

By assumption, $\Phi_1^\Xi/\Phi_1^{\Xi'} = \Psi_1^\Xi/\Psi_1^{\Xi'}$, which implies $\Psi_1^\Xi/\Phi_1^\Xi = \Psi_1^{\Xi'}/\Phi_1^{\Xi'}$. To ease exposition, let Ψ_1^*/Φ_1^* denote the ratio of Ψ_1^Ξ to Φ_1^Ξ for any Ξ . Rearranging terms and substituting in for $\sum_{\Xi} L_t^\Xi$ and $\sum_{\Xi} K_t^\Xi$ yields

$$K_t^{\text{FTC}} = \left(\left[\sum_{\Xi} \Psi_0^\Xi/\Psi_1^\Xi \right] - \left[\sum_{\Xi} \Phi_0^\Xi/\Psi_1^\Xi \right] \right) + \frac{\Phi_1^*}{\Psi_1^*} L_t^{\text{FTC}}.$$

Regressing K_t^{FTC} on a constant and L_t^{FTC} yields $\widehat{\Phi_1^*/\Psi_1^*}$, an estimate of Φ_1^*/Ψ_1^* . Dividing $\widehat{\Phi}_1^{\text{PNO}}$ by $\widehat{\Phi_1^*/\Psi_1^*}$ yields an estimate of $\widehat{\Psi}_1^{\text{PNO}}$. Likewise, dividing $\widehat{\Phi}_1^{\text{PP}}$ by $\widehat{\Phi_1^*/\Psi_1^*}$ yields an estimate of $\widehat{\Psi}_1^{\text{PP}}$.

C Counterfactual simulation details

Let $\hat{\theta}_0, \hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4$ denote estimates of $\theta_0, \theta_1, \theta_2, \theta_3, \theta_4$. To reduce the computational burden, set $\hat{\theta}_1, \hat{\theta}_2$, and $\hat{\theta}_3$ equal to exactly zero. (Each are *very* close to zero, so this simplification will not affect the cost-benefit analysis.) To ease exposition, let \mathcal{A}_m denote the set of a in m , \mathcal{M}_t denote the set of exempt m observed in the data at t , and \mathcal{J}_m denote the set of j owned by the target in m just prior to when it occurs. Also, let A_m, M_t , and J_m denote the cardinality of these sets, respectively. Finally, let \hat{D}_a denote the expected divestiture rate for a evaluated at $\hat{\theta}_0, \hat{\theta}_1, \hat{\theta}_2, \hat{\theta}_3, \hat{\theta}_4$. Note that the process described below is tedious but straightforward.

To enumerate each possible state in each year, start at $t = 1997$. For each exempt facilities acquisition,

⁴The document may be accessed at <https://www.ftc.gov/system/files/documents/reports/fy-2017-congressional-budget-justification/2017-cbj.pdf>.

compute \hat{D}_a assuming $HSR_a = 1$. For all $m \in \mathcal{M}_{1997}$, compute the average divestiture rate weighted by the number facilities involved in it, which is given by

$$\hat{D}_m = \frac{1}{J_m} \sum_{j \in \mathcal{J}_m} \hat{D}_{a(j)}. \quad (17)$$

By assumption, m is deterred if $\hat{D}_m > 0.05$. Compute the expected PNO caseload for 1997, which is given by

$$\Delta C_{1997}^{\text{PNO}} = \left(\frac{1}{M_{1997}} \sum_{m \in \mathcal{M}_{1997}} \{\hat{D}_m < 0.5\} \right) - C_{1997}^{\text{PNO}}, \quad (18)$$

where $\{\cdot\}$ denotes an indicator that takes a value of one if the conditions inside the braces are satisfied and zero otherwise. Also, compute the expected post-PNO caseload for 1997, which is given by

$$\Delta C_{1997}^{\text{PP}} = \left[\sum_{m \in \mathcal{M}_{1997}} \{\hat{D}_m < 0.5\} \left(1 - \prod_{a \in \mathcal{A}_m} (1 - \hat{D}_a) \right) \right] - C_{1997}^{\text{PP}}. \quad (19)$$

List all a in \mathcal{A}_m such that $m \in \mathcal{M}_{1997}$, $\hat{D}_m < 0.5$, and $\Delta \text{HHI}_a > 0$. Each will lead to multiple states in the next period. Enumerate all possible states resulting from them. Denote this set by \mathcal{R}_{1998} . Compute the probability, $\phi(r)$, that each $r \in \mathcal{R}_{1998}$ occurs.

Continue on to $t = 1998$. Repeat the steps described above. To be specific, for each $r \in \mathcal{R}_{1998}$, complete the following.

1. Compute \hat{D}_m for each $m \in \mathcal{M}_{1998}$.
2. Compute equation 18 and call the result $\Delta C_{1998,r}^{\text{PNO}}$.
3. Compute equation 19 and call the result $\Delta C_{1998,r}^{\text{PP}}$.
4. List all a in \mathcal{A}_m such that $m \in \mathcal{M}_{1998}$, $\hat{D}_m < 0.5$, and $\Delta \text{HHI}_a > 0$.
5. Denote the set of states resulting from them by \mathcal{R}_{1999}^r .
6. Compute $\phi_r(r')$ for each $r' \in \mathcal{R}_{1999}^r$.
7. Compute the number of premerger notifications that must be evaluated:

$$\Delta C_{1998}^{\text{PNO}} = \sum_{r \in \mathcal{R}_{1998}} \phi(r) \Delta C_{1998,r}^{\text{PNO}}.$$

8. Compute the number of potentially anticompetitive mergers that must be investigated:

$$\Delta C_{1998}^{\text{PP}} = \sum_{r \in \mathcal{R}_{1998}} \phi(r) \Delta C_{1998,r}^{\text{PP}}.$$

Next, define \mathcal{R}_{1999} equal to the union of \mathcal{R}_{1999}^r over $r \in \mathcal{R}_{1998}$. Finally, for each $r \in \mathcal{R}_{1999}$, compute

$$\phi(r) = \sum_{r' \in \mathcal{R}_{1998}} \phi(r') \phi_{r'}(r).$$

Now, $t = 1998$ is complete. Iterate forwards until $t = 2017$ and stop once $\Delta C_{2017}^{\text{PP}}$ has been computed.

D Reportability: determination, verification, and potential avoidance

D.1 Determining reportability

I determine whether a merger is reportable or exempt in six steps. The first three steps are as follows. First, I record transaction values for 35 publicly disclosed mergers. Second, I translate nominal values into real ones. Third, I plot real transaction values against the merger size, i.e., the number of facilities acquired.

Figure A1 reports the result of the third step. 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 facilities acquisitions.

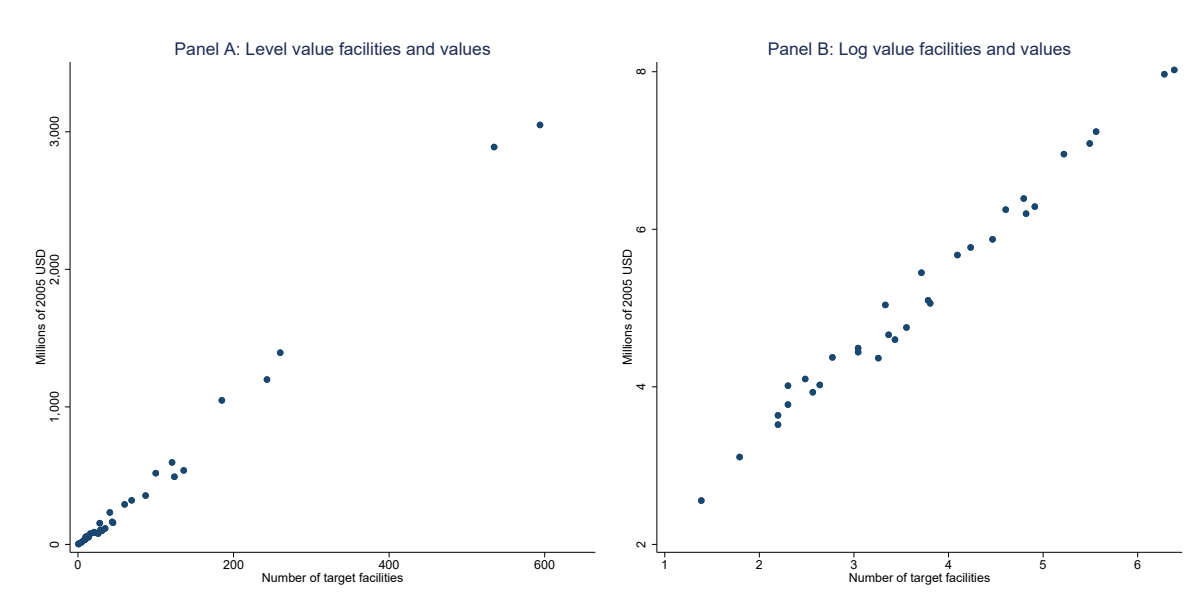


Figure A1: Real transaction values are tightly linked to the number of acquired facilities.

The unit of observation is a merger. The x-axis measures the number of facilities the target operated just prior to the merger, while the y-axis measures the transaction value in constant 2005 dollars. This currency measure facilitates comparisons to the HSR transaction value threshold, which stood at exactly \$50 million in that year, but differs from my other tables and figures, which use constant 2017 dollars. Panel A's axes measure level values, while Panel B's measure log values.

The final three steps are as follows. Fourth, I impute remaining transaction values using an average price-per-facility. Fifth, I translate nominal thresholds into real ones. Sixth, I call the merger "reportable" if the real transaction value meets or exceeds the real-valued HSR threshold, and I call the remainder "exempt."

D.2 Verifying the approach

To verify the accuracy of this approach, I rely on Early Termination Notices, which are published by the PNO. For background, when the merging parties submit a premerger notification, they can check a box that

requests "early termination" of the waiting period.⁵ If the box is checked and the prospective merger review is complete before the statutory waiting period ends, then the agencies will (at their discretion) terminate the waiting period early. However, if early termination occurs, then the PNO releases this information publicly. In particular, they report the buyer, the seller, the entity whose ownership is changing hands, and the date of the early termination. Thus, the existence of an Early Termination Notice implies a premerger notification, i.e., merger is reportable. Further, with the exception of firms wishing to keep their transactions secret for as long as possible (e.g., acquirers purchasing publicly traded stock and wishing to avoid "run ups" in the share prices), the parties request early termination. Indeed, from 1996 to 2017, 87% of merging parties checked the box.

The FTC website contains all notices published since 1999, while the text of the *Federal Register* lists earlier notices. Using the method described above, I predict 36 mergers in the sample are reportable. Of these, I confirm that at least 35 were reportable. I say "at least" because my failure to confirm reportability does not rule it out, especially since the single discrepancy occurred at the beginning of the panel, so it is not included in the FTC website's searchable database. Combined with the fact that firms' legal names frequently differ from the names under which firms do business (e.g., "FMC Corporation" versus "Fresenius Medical Care"), there is a high probability the transaction was reportable. Separately, of the large number of mergers in the sample that I predict are exempt, I find no evidence that any of them were reportable.

D.3 Understanding potential avoidance

Recognizing that merging parties may have strong incentives to escape the premerger notification program, Congress empowers the agencies to impose steep fines on firms that fail to file *for any reason*. At the time of writing, the penalty is \$50,120 per transaction per day.⁶ Equally important, the PNO aggressively pursues failure-to-file cases. For instance, it found Berkshire Hathaway liable for \$896,000 in HSR fines, despite mitigating circumstances, which include the failure being accidental, the stake equating to only 9% of the outstanding shares, and the transaction not posing any competitive concerns.⁷ As a result, the principal concern is avoidance, not outright evasion. To mitigate this concern, Congress includes a clause in the act that stipulates the filing decision must reflect the "substance" of the deal, not exemptions that arise from idiosyncratic features of how it is structured. Specifically, the act prohibits transactions or devices that are "entered into or employed for the purpose of avoiding" a filing.⁸ In the body of the main text, I claim this concern is unlikely to affect my results (even though it may limit their external validity). Support for the claim comes in two forms: technical reasons why avoidance in the dialysis industry is impractical, and patterns in the data that confirm dialysis providers are complying with the spirit as well as letter of the act. Each is described below.

The first technique that the merging parties could, at least in theory, employ to avoid filing involves manipulating transaction values downwards by way of contingent payments. If part of the consideration paid to the seller involves uncertain payments such as earn-outs based on operational achievements, then

⁵Early termination requests have not been granted since around February 2021, when the agency suspended the option, ostensibly due to the surge in pandemic-era filings. The suspension occurs years after my sample ends, so it has no effect on my analysis whatsoever.

⁶See 15 USC 18a(g)(1).

⁷Berkshire Hathaway already owned 19% of USG Corp., but it held notes, which converted to equity, increasing its share of the voting securities. It failed to file a notification related to the increase.

⁸See 16 CFR 801.90.

a portion of the acquisition price will be based on fair market value. When nascent targets offering novel products are acquired, sellers face an adverse selection problem and uncertainty exists about future outcomes, so contingent payments are large, and the merging parties have leeway in assigning probabilities to the milestones being met. Dialysis mergers do not involve meaningful contingent payments, which reflects stable operations and symmetric information. As a result, structuring a dialysis merger to include large contingent payments, marking down the probabilities they are realized, and eventually making all the payments would almost surely constitute an avoidance device.

The second technique involves manipulating asset values to avoid filing, which is even tougher. The transaction value threshold is exactly five times larger than the asset value threshold. In most industries including dialysis, the ratio of market value to assets is much less than five, so passing the size-of-transaction test guarantees passing the size-of-persons test (i.e., the asset test).⁹ Thus, to earn an asset-based exemption, firms would need to significantly undercount assets. This is made even more challenging by the fact that the act stipulates that the figure used to determine reportability should be read off the most recently generated balance sheet.

The third and perhaps best-known technique that merging parties could employ involves acquisitions of noncorporate interests (NCIs). Historically, acquisitions of NCIs such as limited partnerships were reportable only if 100% of the interests were transferred, creating a so-called "partnership loophole." If the seller retained, say, 1% interest in the acquired entity, then the transaction was exempt. The loophole grew alongside extraordinary growth in the use of limited liability companies, which were treated as NCIs. The PNO amended the HSR Rules in April 2005 such that control was determined by majority ownership. However, nearly as convoluted exemptions persist so that many NCI acquisitions are exempt or provide an incomplete picture of the transaction to the agencies. Exemptions related to NIC acquisitions are limited to private equity funds whose investors require nontraditional investment structures (e.g., ones that employ co-investment vehicles) and whose exit strategies do not include public offerings or sales to public acquirers. Thus, they are unimportant to understanding dialysis mergers. (See Asil et al. [2023] for a detailed analysis of why acquisitions that would otherwise be reportable may be exempt when sponsored by private equity.)

The data are consistent with these claims. To see this, suppose the opposite is true—the dialysis providers successfully employ these techniques. The most immediate consequence will be mergers that I determine are reportable (based on their size) but for which no premerger notification was filed. However, as stated in Appendix D.1, every merger I determine is reportable was indeed reported under the act. An even more obvious consequence is that acquirers will escape costly enforcement actions. Yet, Figure I reflects nearly 300 facility divestitures, which not only involved millions of dollars in legal fees, managerial attention, and delays but also, according to the structural estimates, curtailed profits.

The final concern is that firms avoid filings using partial acquisitions. Ostensibly, one could break apart large mergers into small ones, all of which fall below the applicable threshold. This approach is typically infeasible. To ensure that artificially divided targets receive equal treatment, the act requires merging parties to aggregate asset acquisitions that occur within 180 days of one another.¹⁰ However, the

⁹Industries such as software and biotechnology are the exceptions. The reason is that under US accounting standards, costs incurred to generate intangible assets are usually not capitalized, so intellectual property resulting from in-house R&D efforts rarely creates an asset. Hence, firms that derive most of their value from patents, trade secrets, trademarks, and/or copyrights typically report low asset values on their balance sheets. As a result, they might be able to dispose of marginal physical assets and avoid filing.

¹⁰See 16 CF 801.13(b). The parties could agree to sequential acquisitions spaced out over more than one year, but this arrangement would almost surely constitute a prohibited device.

aggregation requirement does not stop the target from carving out a part of its operations and selling them to the acquirer. Depending on the specifics of the deals (e.g., the proximity of the transaction values to the thresholds), these arrangements could easily be construed as a prohibited devices, although they should not be immediately dismissed. Instead, I point out that very few facilities acquisitions—less than 10% of nonreportable mergers involving more than one facilities acquisitions—involve carved out operations, and the majority of the carve outs do not involve any preacquisition competition with the acquirer. Consistent with the idea that these instances do not drive the results, I re-run the main analysis and confirm the main results do not change when eliminating carved out acquisitions from the sample.

E Supplementary tables and figures

E.1 Comparison of FTC and CMS data

To confirm that the Medicare and FTC data are congruent with one another, I compare the number of target and acquirer facilities reported in FTC complaints. I then compare the number of divested facilities reported in FTC orders with tabulations from Medicare. Tables A1 and A2, respectively, reports the relevant firm-by-merger comparisons.

Table A1: Facility counts extracted from complaints closely match CMS.

Acquirer 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

The data sources closely align with one another. For example, when Fresenius acquired Renal Care Group in 2005, the FTC reported that the acquirer and target owned 1,155 and 450 facilities, respectively, and that 103 facilities would be divested to National Renal Institute. CMS data indicates that in 2005 the acquirer 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.

E.2 Distribution of observed HHI changes

Figure A2 plots the density of HHI changes by reportability. Panel A reflects reportable facility acquisitions, while Panel B reflects exempt ones. Both reflect the fact that the majority of facilities acquisitions involve

Table A2: Facility counts extracted from divestiture orders closely match CMS.

Acquirer	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

no target-acquirer overlap, as evidenced by large point masses at $\Delta\text{HHI} = 0$. Since these features of the graphs make it difficult to judge how potentially anticompetitive facilities acquisitions are distributed, I replot the densities excluding HHI changes less than 500 points. Again, the graphs look similar.

E.3 Divestiture-related

E.3.1 Standard errors for Figure I

Figure I plots divestiture rates against ΔHHI separately for reportable and exempt facilities acquisitions. For legibility, I bin the data according to x-axis values and plot the mean values within each bin, but I do not provide a sense of dispersion within the bins. Figure A3 supplements the result by plotting 95% confidence intervals around each mean value.

E.3.2 Divestiture rates vs. target and acquirer size

Panel A of Figure A4 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 D_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 D_a on ΔHHI_a , and plot the residual against the number of facilities the acquirer owns just prior to the merger. Panel B of Figure A4 reports the result. Again, there's no meaningful relationship between enforcement and target size.

E.3.3 Alternative specifications

Table A3 assesses the sensitivity of equation 6 to adding controls or changing specifications. In column 1, I regress D_a on a constant, HSR_a , ΔHHI_a , $\text{HSR}_a \times \Delta\text{HHI}_a$. In column 2, I add HHI_a and $\text{HSR}_a \times \text{HHI}_a$

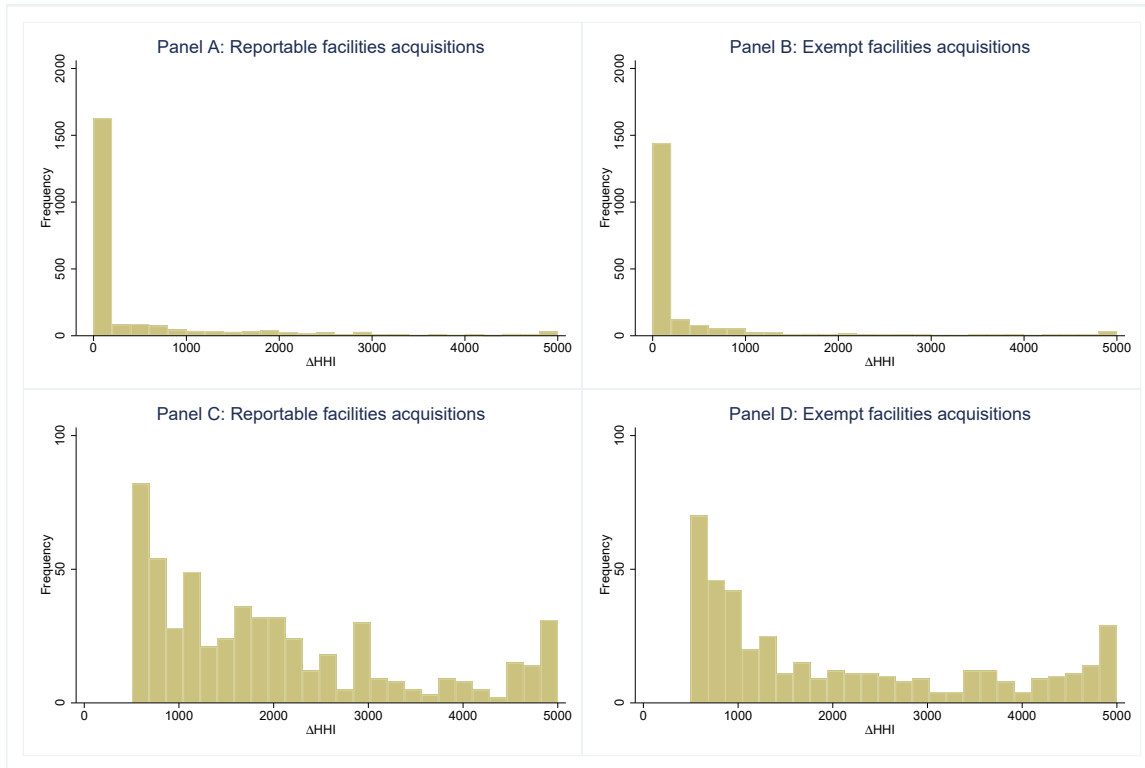


Figure A2: HHI changes associated with reportable and exempt facilities acquisitions are similarly distributed.

This figure plots the density of facilities acquisitions on the y-axis against HHI changes on the x-axis. Left-hand side panels reflect reportable facilities acquisitions, while right-hand side panels reflect exempt ones. In the bottom panels, HHI changes less than 500 points are omitted, which improves legibility.

to the right-hand side of the estimating equation. Coefficients obtained from this specification show that concentration levels do not explain divestitures conditional on HHI changes. In columns 3 and 4, I control for the size of the target (i.e., size of the merger) and size of the acquirer, respectively, and in columns 5 and 6, I add fixed effects for the identity of the acquirer and the acquisition year, respectively. These specifications produce very similar coefficients to the ones reported in column 1. Column 7 shows that the coefficients do not change in a meaningful way when I recompute HHI using machine counts rather than patient counts. Likewise, column 8 shows that the coefficients do not change when I redefine markets entirely, using a 20 mile radii around each facility to draw market boundaries. Column 9 shows that when I exclude Los Angeles (CA) Cook (IL), and San Diego (CA) counties, the results also do not change (see Section F below). Despite these additions and changes, the main coefficient (i.e., the coefficient on the interaction between ΔHHI_a and HSR_a) is stable across specifications. Column 10 reports estimates from estimating equation 6 verbatim.

E.3.4 Quadratic functional form fit

Figure I indicates that divestiture rates are concave in ΔHHI (for reportable facilities acquisitions). Likewise, the final column of Table A3 reports that R^2 rises when I include a term that interacts HSR_a with ΔHHI_a^2 . For a clearer sense of how equation 6 fits the data, I plot predicted values of the left-hand side against

Table A3: Robustness of divestiture-related estimates

VARIABLES	(1) Divest.	(2) Divest.	(3) Divest.	(4) Divest.	(5) Divest.	(6) Divest.	(7) Divest.	(8) Divest.	(9) Divest.	(10) Divest.
REPORTABLE \times Δ HHI	0.22*** (0.011)	0.23*** (0.011)	0.22*** (0.012)	0.22*** (0.0100)	0.22*** (0.011)	0.22*** (0.011)	0.21*** (0.011)	0.23*** (0.010)	0.22*** (0.011)	0.47*** (0.051)
REPORTABLE	0.025** (0.0100)	0.098*** (0.031)	0.013 (0.023)	0.047** (0.023)	0.031*** (0.010)	0.019** (0.0081)	0.033*** (0.011)	0.028*** (0.0100)	0.024** (0.010)	0.0053 (0.0048)
Δ HHI	0.0013 (0.0012)	0.0017 (0.0016)	0.0013 (0.0012)	0.0013 (0.0012)	0.000096 (0.00099)	0.00089 (0.0013)	0.00026 (0.00024)	0.0031 (0.0028)	0.0013 (0.0012)	0.0013 (0.0012)
REPORTABLE \times HHI		-0.012*** (0.0034)								
HHI		-0.00023 (0.00023)								
REPORTABLE \times Δ HHI ²										-0.060*** (0.011)
Constant	0.00044 (0.00045)	0.0017 (0.0017)	0.00031 (0.00047)	0.0078* (0.0045)	-0.0010 (0.0010)	0.0042* (0.0022)	0.00095 (0.00095)	-0.00015 (0.00011)	0.00047 (0.00048)	0.00044 (0.00045)
Observations	3,266	3,266	3,266	3,266	3,265	3,266	3,266	3,266	3,148	3,266
R-squared	0.549	0.560	0.550	0.554	0.553	0.560	0.446	0.535	0.555	0.604
Number of target fac.	No	No	Yes	No	No	No	No	No	No	No
Number of acquiror fac.	No	No	No	Yes	No	No	No	No	No	No
Chain FE	No	No	No	No	Yes	No	No	No	No	No
Acquisition year FE	No	No	No	No	No	Yes	No	No	No	No

*, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Column 2 incorporates predicted concentration levels as well as changes. Columns 3 and 4 include controls for the number of target and acquirer facilities, respectively. Columns 5 and 6 include chain and acquisition year fixed effects, respectively. In column 7, I redefine markets entirely by drawing a 20 mile radius around each facility. In column 8, I recompute HHI changes using the number of machines operated by each facility (rather than the number of patients). In column 9, I exclude Los Angeles (CA) Cook (IL), and San Diego (CA) counties (see Section F below). In column 10, I include a term that captures concavity of divestiture rates in HHI changes. Note that "Number of target facilities" is exactly equal to the variable measured by the x-axis in Panel A of A4, while "Number of acquirer facilities" is exactly equal to the variable measured by the x-axis in Panel B of A4. Standard errors are clustered at the merger level.

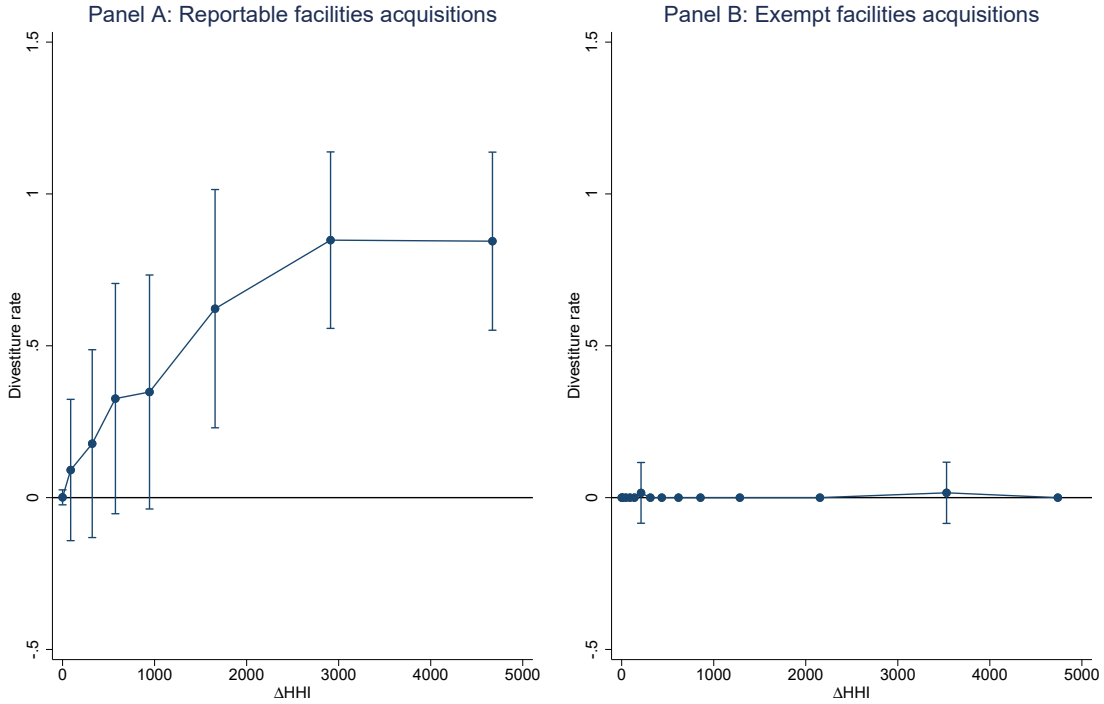


Figure A3: Relationship between divestiture rates and HHI changes with confidence intervals

This figure replicates Figure I except that error bars representing 95% confidence intervals around each plotted point are added to summarize dispersion within each bin.

actual values. Figure A5 reports the result. The fit is excellent.

E.4 Quality-related

To assess the robustness of the conclusions reached in Section 4.2, I estimate alternative specifications. To simplify the output, I collapse event time into pre- and post-acquisition periods. Recall that in the body of the main text d_{jt}^τ terms are dummies for particular values of τ , which denotes event time. Here, I set $d_{jt}^\tau = 0$ for $\tau < 0$ and $d_{jt}^\tau = d_{jt}^{\text{Post}}$ if $\tau \geq 0$, and I collapse the coefficients accordingly. In other words, I define d_{jt}^{Post} equal to one if $\tau \geq 0$ and zero otherwise. The estimating equation is given by

$$y_{jt} = \lambda_t + \lambda_j + \lambda_{jt}^c + \psi d_{jt}^{\text{Post}} + \kappa d_{jt}^{\text{Post}} EX_a + \nu d_{jt}^{\text{Post}} \widetilde{\Delta\text{HHI}}_a + \mu d_{jt}^{\text{Post}} EX_a \widetilde{\Delta\text{HHI}}_a + e_{jt}. \quad (20)$$

As in the main text, y_{jt} denotes quality. λ_t , λ_j , and λ_{jt}^c denote year, facility, and chain fixed effects, respectively. The sample period is restricted to a seven year window on each side of each event.

Table A4 reports the result of estimating equation 20. Columns 1-4 and 5-8 measure quality using risk-adjusted hospitalization and survival rates, respectively. Columns 1 and 5 correspond to the baseline specification—the "pre/post" version of the estimating equation that appears in the main text. For columns 2 and 6, I recompute HHI changes using the number of machines operated (rather than patients treated) by each facility. For columns 3 and 7, I redefine markets, using a 20 mile radii around each facility to draw

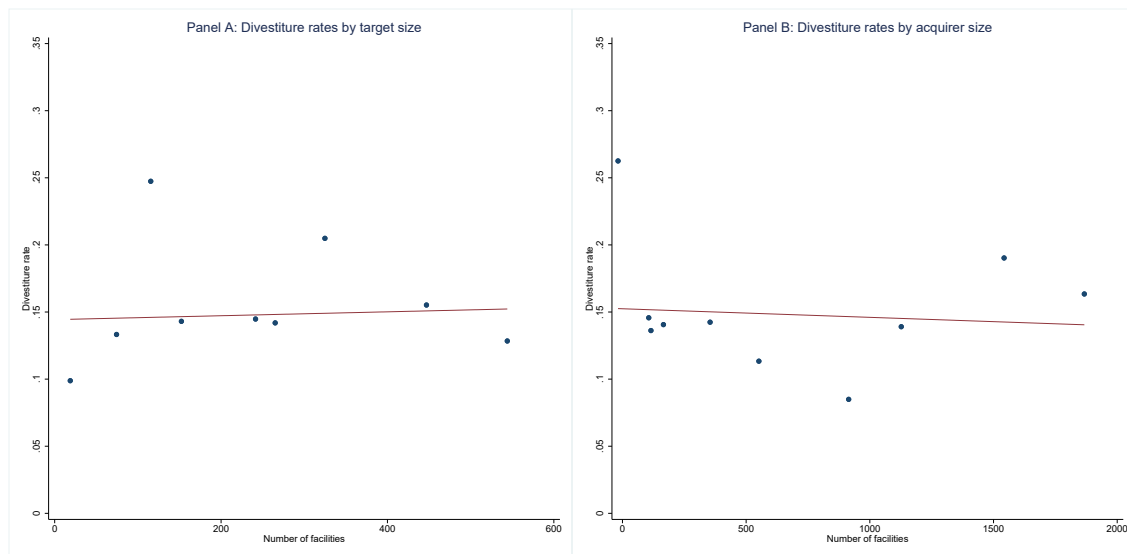


Figure A4: Target or acquirer size is unrelated with divestiture rates.

This figure plots divestiture rates, controlling for ΔHHI and ΔHHI^2 , on the y-axis against merger size on the x-axis. The sample consists of reportable facilities acquisitions. The data are sorted into equal-sized bins according to x-axis values, and averages within the bins are reported along the axes. In Panel A, merger size is measured by the number of acquired facilities, while in Panel B, it is measured by the number of facilities owned by the acquirer in the year before the merger.

boundaries. For columns 4 and 8, I exclude Los Angeles (CA) Cook (IL), and San Diego (CA) counties (see Section F). Estimates from the baseline and alternative specifications are close to one another, regardless of how quality is measured.

In Section 4.3, I state that one of the specifications in Eliason et al. [2019] suggests the relationship between consolidation and quality is limited. This point is not central to their paper but may nonetheless be useful to look at in more detail, especially since it is easy to reconcile with my results.¹¹ There are two main differences between that specification and the ones I consider. First, whereas I define markets using counties (or, in alternative specifications, by drawing circles around each facility), the specification in question bases market definitions on health service areas. The latter are large large catchments that typically extend far beyond the points from which patients travel. Second, whereas I measure market structure changes using an indicator for greater-than-average HHI changes, the specification in question uses an indicator for positive HHI changes. Due to right skew in the distribution of ΔHHI , the latter measure captures a large number of insignificant HHI changes.

To assess the impact of the first modeling choice, I turn to Online Appendix Table I5 of Eliason et al. [2019], which reports the result of redefining markets by drawing a circle with radius ten miles around each facility. (Ten miles is the largest radii considered by the paper.) Dividing columns 9's row 3 by row 5

¹¹Following Eliason [2018], I measure quality using risk-adjusted hospitalization and survival rates. To study the transfer of management practices, Eliason et al. [2019] study other measures as well. Some 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, the authors 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 the authors 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.

Table A4: Robustness of quality-related estimates

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Risk adj. survival rate	Risk adj. survival rate	Risk adj. survival rate	Risk adj. survival rate	Risk adj. hospital'n rate	Risk adj. hospital'n rate	Risk adj. hospital'n rate	Risk adj. hospital'n rate
$\text{POST} \times \text{EXEMPT} \times \widehat{\Delta\text{HHI}}$	-0.022*** (0.0063)	-0.023*** (0.0068)	-0.019*** (0.0065)	-0.019*** (0.0065)	0.035*** (0.0081)	0.022*** (0.0083)	0.025*** (0.0087)	0.034*** (0.0083)
$\text{POST} \times \widehat{\Delta\text{HHI}}$	0.0075* (0.0039)	0.0067* (0.0040)	0.0063 (0.0040)	0.0058 (0.0041)	-0.0045 (0.0052)	-0.00049 (0.0052)	-0.0052 (0.0052)	-0.0043 (0.0055)
$\text{POST} \times \text{EXEMPT}$	0.0075* (0.0043)	0.0069 (0.0042)	0.0062 (0.0043)	0.0081* (0.0045)	-0.012** (0.0055)	-0.0081 (0.0055)	-0.0088 (0.0054)	-0.013** (0.0057)
POST	-0.0050 (0.0038)	-0.0049 (0.0038)	-0.0045 (0.0037)	-0.0047 (0.0040)	0.0014 (0.0044)	-0.000081 (0.0044)	0.0014 (0.0044)	0.0014 (0.0046)
Observations	37,917	37,917	37,917	35,910	41,798	41,798	41,798	39,694
R-squared	0.489	0.490	0.489	0.487	0.510	0.510	0.510	0.506
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

* **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. Columns 1 and 5 are baseline specifications. For columns 2 and 6, target and acquirer shares are computed using the number of machines operated (rather than patients treated) by each facility. In columns 3 and 7, markets are defined by circumscribing a circle with a 20 mile radii around each acquired facility. In columns 4 and 8, I exclude all Los Angeles (CA), Cook (IL), and San Diego (CA) county markets. Standard errors are clustered at the facilities acquisition level.




Figure A5: Predicted and empirical enforcement rates

The x-axis measures 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.

implies that, in the language and notation of my paper, exempt facilities acquisitions involving $\Delta\text{HHI} > 0$ increase risk-adjusted hospitalization rates by 2.2 percentage points. To assess the impact of the second modeling choice, I re-estimate 20 but replace the indicator $\Delta\text{HHI} > \overline{\Delta\text{HHI}}$ with an indicator for $\Delta\text{HHI} > 0$, which reduces my estimate of μ by 34% to 2.3 percentage points—close to the 2.2 percentage point figure arrived at above. Minor differences in estimation account for the small, remaining gap.

E.5 Distances traveled by patients

An earlier version of the paper measured travel using great circle ("as the crow flies") distances, which are shorter than the distances patients actually travel along the road. The present version uses an approximation to road mileage proposed by Boscoe et al. [2012], who employ a sample of 66,011 trip distances between patients' homes and their nearest community hospitals.¹² The authors find that road mileage and great circle distance are highly correlated (e.g., $R^2 = 0.94$) and that the former averages 41.7% longer than the latter. Thus, I calculate great circle distances from the latitude and longitude corresponding to each ZIP5 centroid and scale up by 1.417.

¹²Until about 2015, drive times were easy to obtain from various commercial websites, including Google, making repeated calls to their map APIs, which is how Boscoe et al. [2012] constructed their sample. Since that time, all sites have eliminated this service. Currently, users must now obtain a site key and pay for each query over 250 per day. Given the R^2 value that results from regressing road mileage on great circle distances is near one, this is unnecessary—scaling great circle distances should suffice to provide accurate estimates.

Figure A6 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 ZIP5 centroids.) It shows that almost all patients travel less than 20 miles to their provider.

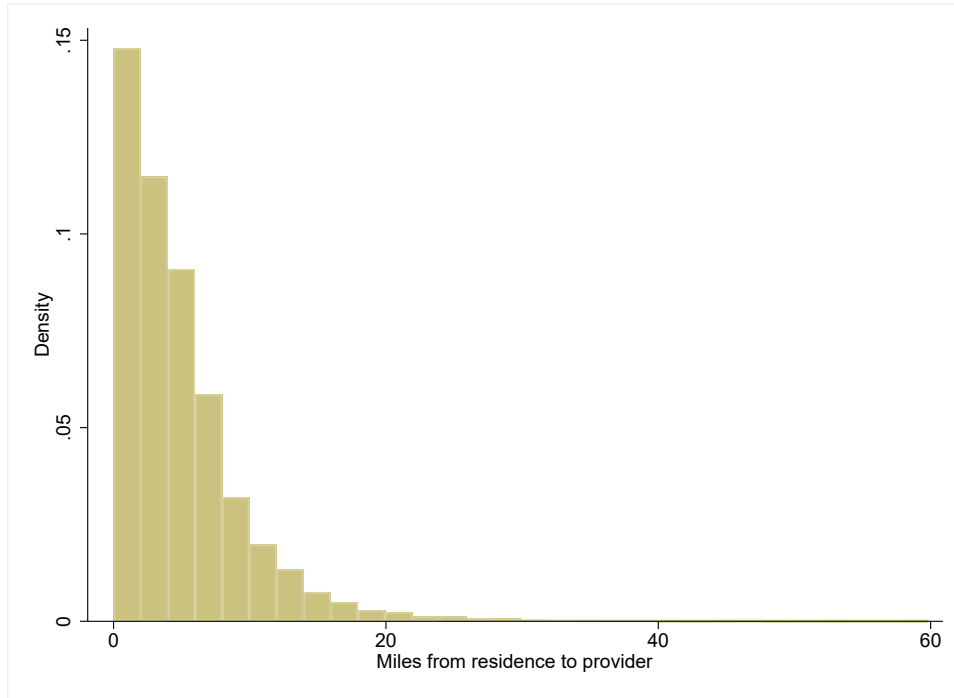


Figure A6: *Patients visit facilities close to home.*

This graph is a histogram of the distances from the centroids of the ZIP5s of patients' residences to the centroids of the ZIP5s of the facilities they visit. The unit of observation is the patient-year.

F Market definitions

To map facilities to geographic markets, I start with location information from the POS files, which report ZIP5, 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 ZIP5s and city names. Comparing the CMS information with provider websites indicates ZIP5s are more accurate, so I replace city names with the modal city name in the ZIP5. 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.¹³ Typographical errors, zip-city/city-county

¹³To be specific, I calculate the distance between the centroid of the ZIP5s 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.

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.¹⁴ 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.

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 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 Tables A3 A4, 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.

¹⁴I 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.