Consolidation Under Regulated Prices: Capacity, Quality, and Patient Health in the Dialysis Industry[§]

Anwita Mahajan[*]
Job Market Paper

Adrian Rubli[†] **Francisco Garrido**[‡]

October 2026
The latest version of the paper is available here.

Abstract

How do mergers affect consumers when prices are regulated and firms compete instead on quality and capacity? We study consolidation in the U.S. outpatient dialysis industry, where Medicare sets prices, competition is local, and treatment is life-sustaining. Using 30 years of data covering more than 850 mergers and a stacked event-study design, we document contrasting short- and long-run effects. Mergers induce facility closures that, despite offsetting entry, reduce market capacity. Closures displace patients, disrupt treatment, and increase mortality by roughly 650 deaths per 100,000 patients in the merger year. Over the subsequent five years, however, health improves: hospitalizations, ICU days, and blood transfusions decline. These effects are unrelated to changes in market concentration and do not differ across Medicare payment regimes. We develop a model that identifies a novel mechanism rationalizing these patterns. Before a merger, switching costs leave patients captive, allowing low-quality facilities to survive. After a merger, the acquirer closes the target, eliminating its fixed costs and realizing scale economies in variable costs from serving more patients. Closure imposes switching costs on displaced patients, generating the short-run mortality spike, but also makes these patients contestable, strengthening incentives for surviving facilities and entrants to compete for them by raising quality. A cost-benefit analysis suggests that, under the status quo, short-run mortality costs moderately exceed long-run benefits from reduced hospitalizations, but policies that mitigate treatment disruption could reduce these costs while preserving the benefits.

[*]University of California, San Diego (anmahajan@ucsd.edu).
[†]Instituto Tecnológico Autónomo de México (adrian.rubli@itam.mx).
[‡]Charles River Associates (fagarrido@gmail.com).

[§]We thank Jeffrey Clemens, David Cutler, Paul Eliason, John Rust, Krista Ruffini, Douglas Staiger, and Thomas Wollmann for their mentorship and advice. We are grateful to Charles Ginsberg and Eugene Lin for guidance on nephrology and the dialysis industry. We thank Nikhil Agarwal, David Arnold, Jack Berger, Eric Budish, Julie Cullen, Itzik Fadlon, Craig Garthwaite, Alex Gelber, Jonathan Gruber, Sarah Kotb, Katherine Meckel, Nathan Miller, Pauline Mourot, Kate Pennington, Nancy Rose, Michael Ricks, Jonathan Roth, and Christopher Snyder for helpful comments, as well as participants at seminars and conferences where this work was presented. Work on this paper was not supported by external grants or funding sources. 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 authors and in no way should be seen as an official policy or interpretation of the U.S. government.

1 Introduction

How do mergers affect consumer welfare when prices are regulated? When firms have price-setting power, higher prices are the conventional channel through which mergers reduce consumer welfare.[1] Under price regulation, firms can instead increase profitability by cutting costs. If these cuts reduce capacity or quality, consumer welfare falls despite unchanged prices; if mergers generate scale economies or strengthen incentives to improve quality, welfare may rise. The effect of mergers on consumer welfare under price regulation is therefore an open empirical question. We study this question in the kidney dialysis industry, focusing on a novel mechanism through which mergers affect welfare—facility closures impose transitional costs, but improve quality through competition for newly contestable patients.

The U.S. outpatient kidney dialysis industry provides an ideal laboratory for several reasons. First, dialysis provides life-sustaining treatment to approximately 500,000 individuals annually in a highly concentrated market; just two firms provide roughly 70% of treatments (Rosner et al. 2023). Second, the high frequency of treatment creates strong geographic constraints, making patients particularly sensitive to local changes in facility access and quality.[2] Third, Medicare heavily regulates dialysis prices through administratively determined reimbursement rates. Together, these features let us isolate quality and capacity effects in a setting where they have immediate and substantial welfare consequences.

Our analysis draws on three decades (1991–2021) of administrative panel data from the United States Renal Data System (USRDS) and the Centers for Medicare and Medicaid Services (CMS). The USRDS provides facility-level utilization and staffing data, patient-level demographics, treatment modality, and mortality, linked Medicare claims on hospitalizations and dialysis sessions, and quality-reporting data on clinical quality and blood biomarkers. CMS Provider of Services (POS) files identify facility ownership changes and locations. Together, these data allow us to identify more than 850 mergers.

We use a stacked event-study design at the market (county) level and find contrasting short- and long-run health effects. In the short run, mergers impose substantial welfare costs. Acquired facilities close and, despite offsetting entry by non-merging firms, total market capacity declines. This capacity churn displaces patients from closing targets, disrupting

[1] This prediction follows directly from standard models of market power and is well supported empirically. Meta-analyses by Kwoka (2014) and Olsen, Orchinik, and Remer (2026) find that mergers increase prices by approximately 7% on average. By contrast, evidence on mergers' overall welfare effects is more limited, largely because evidence on offsetting effects through non-price dimensions is mixed (Kwoka and Kilpatrick 2018), as is evidence on merger-induced efficiencies (Ashenfelter, Hosken, and Weinberg 2015; Demirer and Karaduman 2024; Kwoka and Pollitt 2010).

[2] Wollmann (2020) shows that 90% of dialysis patients seek treatment within their county.

care and reducing dialysis sessions across the market. The disruption increases mortality by 3.8% in the merger year, or roughly 650 deaths per 100,000 patients.

In the long run, mergers generate meaningful welfare gains. Over the five years following a merger, hospitalizations, ICU days, and inpatient blood transfusions decline. These improvements are not explained by changes in patient risk composition over time, such as healthier patients surviving beyond the merger year, or by differences in patient risk across markets or facilities. Although mergers increase market concentration, the magnitude of this increase is not meaningfully related to either the short- or long-run effects.

We develop a model to characterize how mergers under price regulation can generate contrasting short- and long-run welfare effects. In the model, low-quality facilities survive in the pre-merger equilibrium by serving a captive patient population facing high switching costs. Following a merger, the acquirer pursues profitability by eliminating duplicate fixed costs through target closures and realizing scale economies in variable costs.[3] Closure imposes substantial switching costs on affected patients, reflected in our reduced-form mortality estimates. At the same time, it releases these patients into competition between the acquirer and potential entrants, strengthening incentives to provide quality. The same closures thus generate short-run welfare costs and long-run gains. An empirical decomposition provides evidence for three channels of long-run quality gains—the exit of low-quality facilities, the entry of higher-quality facilities, and quality improvements by acquirers.

In a cost-benefit analysis, we estimate $848 million in costs from short-run mortality and $688 million in benefits from reductions in hospitalization over the merger year and the five subsequent years. While costs therefore outweigh benefits over this period, our model has a straightforward policy implication—antitrust regulation that minimizes post-merger treatment disruptions and facilitates smooth patient transitions has the potential to reduce costs without sacrificing benefits.[4]

Finally, we provide evidence specific to consolidation in the dialysis industry, beyond price-regulated industries more broadly. First, comparing merger effects before and after two simultaneous Medicare reforms—one shifting traditional fee-for-service reimbursement to a per-treatment bundle, and another tying reimbursement to quality targets—shows that these reforms neither mitigated adverse health effects nor amplified beneficial ones. Second, merger effects on biomarkers collected contemporaneously with treatment are mixed; measures directly incorporated into quality incentives improve, consistent with targeted compliance, while other measures show no response or deteriorate. Overall, the variation in

[3]Scale economies in variable costs could arise, for example, through volume discounts or greater bargaining power in input markets (Fee and Thomas 2004).

[4]We discuss potential policy responses to reduce treatment disruptions in Section 8.

biomarker responses, despite robust improvements in inpatient outcomes, limits the extent to which these measures can be interpreted as reliable predictors of long-run health.

This paper contributes to four literatures spanning industrial organization, health economics, and public economics. First, we contribute to the industrial organization literature on the non-price effects of mergers, particularly in markets with regulated prices. While a large literature documents the price effects of consolidation (Kwoka 2014; Olsen, Orchinik, and Remer 2026), evidence on how mergers affect quality and capacity remains limited and theoretically ambiguous, owing both to the difficulty of isolating non-price mechanisms (Kwoka and Kilpatrick 2018) and to the challenge of measuring non-price welfare. We study a setting in which prices are largely fixed, allowing us to identify non-price responses to consolidation and measure its welfare consequences directly through patient health outcomes.

Our central contribution is to identify a mechanism through which mergers can reduce capacity while raising long-run quality, reconciling our evidence of contrasting short- and long-run effects on consumer welfare.[5] Before a merger, switching costs leave patients effectively captive, and low-quality facilities survive.[6] Merger-induced closures reduce capacity but also break this captivity, raising quality as surviving facilities and entrants compete for newly contestable patients. This mechanism can be interpreted as market-level "creative destruction," in which long-run quality gains come at the cost of short-run disruption.[7]

Second, we contribute to the health economics literature on consolidation and patient outcomes. Existing work focuses primarily on hospital mergers, generally finding that consolidation raises negotiated prices but has ambiguous and often modest effects on quality (Ho and Hamilton 2000; Mariani et al. 2022; Beaulieu et al. 2020; Ferrier and Valdmanis 2004). Hospitals, however, may be less informative for studying the quality and capacity effects of mergers—they retain substantial pricing power through negotiations with private insurers, compete over broader geographic markets, and treat highly heterogeneous condi-

[5] Grieco and McDevitt (2017) document a quantity-quality tradeoff in dialysis provision. How can this be reconciled with our finding that acquirers serve more patients after a merger while providing higher quality? Their framework holds capital and labor fixed. Our short-run results, in which capacity-constrained surviving facilities experience operational strain and mortality increases, are directly consistent with their findings. In the longer run, however, both acquirers and the market as a whole adjust capital and capacity, generating efficiencies that can improve quality despite higher patient volume. We describe how scale can generate quality gains in Section 6.

[6] A growing literature shows that consumer inertia can dampen competition and reduce consumer welfare, consistent with our description of the pre-merger equilibrium. In health insurance, switching costs have been shown to sustain higher prices and exacerbate adverse selection (Ericson 2014; Ho, Hogan, and Scott Morton 2017; Polyakova 2016; Handel 2013). Consumer inertia has also been shown to reduce consumer welfare in other price-regulated industries, including telecommunications (Viard 2007) and residential electricity (Hortaçsu, Madanizadeh, and Puller 2017).

[7] The term "creative destruction" is used primarily to explain how growth is driven by innovations that render older technologies obsolete (Aghion and Howitt 1992). A more general interpretation applies to our findings, where long-run quality gains are achieved at the cost of one-time reductions in quantity, measured by the number of facilities.

tions, making health outcomes a noisier measure of quality.[8,9] Dialysis, by contrast, is a standardized, life-sustaining treatment with administratively set prices, delivered to a relatively homogeneous population, providing a setting in which quality and capacity channels can be effectively isolated.

Within dialysis, prior work documents adverse effects of acquisitions among patients who remain at acquired facilities (Eliason et al. 2020) and within balanced panels of facilities observed before and after mergers (Wollmann 2020). Our key departure is to allow for facility entry and exit and to estimate effects at the market level (Dafny, Duggan, and Ramanarayanan 2012; Erickson et al. 2018, 2019). This margin accounts for both the short-run harms of target closures and the long-run efficiency and quality gains we document, which have rarely been identified in the merger literature (Dafny and Lee 2015). In other words, incorporating facility turnover fundamentally changes the welfare implications of consolidation.[10]

The market-level approach also has a methodological implication—non-acquired facilities in affected markets are not valid counterfactuals for acquired facilities, so studies relying on within-market controls may misstate the health effects of consolidation because of violations of the Stable Unit Treatment Value Assumption (SUTVA). Additionally, our findings suggest one possible explanation for the coexistence of secular improvements in dialysis outcomes and substantial industry consolidation over the past decades—consolidation may induce endogenous quality improvements.[11]

Third, we contribute to the literature on pay-for-performance and the measurement of health care quality (Jha et al. 2012; Petersen et al. 2006; Bertuzzi et al. 2023). Using a three-decade panel, we are the first to examine how the effects of dialysis mergers differ across major payment regimes, including the 2011 introduction of the bundled Prospective

[8] Hospital and tertiary care markets are typically defined using Hospital Referral Regions (HRRs), which are substantially larger than counties and are based on patient referral patterns. There are approximately 300 HRRs in the United States, compared with more than 3,000 counties (Dartmouth Atlas of Health Care 2024).

[9] We do not take a position on whether the limited quality effects documented in the hospital literature reflect genuinely small underlying effects or challenges in identifying such effects empirically.

[10] We do not view our long-run health gains as incompatible with the long-run declines documented by Eliason et al. (2020), whose results apply to patients at acquired facilities that survive. Consistent with their findings, our model and results indicate that quality remains low when targets survive, as discussed in Section 6. Our results instead speak to patients in the exposed market at large, where closures and post-merger entry raise average quality. Our sample also includes a broader set of mergers.

[11] Figure 5.1 in the United States Renal Data System 2024 Annual Data Report (NIH 2024) shows that adjusted hospitalization events per hemodialysis patient per year declined from 1.71 in 2012 to 1.48 in 2022.

Payment System (PPS) and the 2012 Quality Incentive Program (QIP).[12,13] We show that these reforms neither mitigated the adverse short-run effects of mergers nor amplified their benefits. Further, although providers respond to quality incentives by improving specific biomarkers to which penalties are tied, these intermediate measures remain poor proxies, when considered jointly, for the substantial long-run health improvements we document.

Fourth, we contribute to the literature on merger efficiencies and antitrust policy (Ashenfelter, Hosken, and Weinberg 2015; Demirer and Karaduman 2024; Kwoka and Pollitt 2010). Evaluating the efficiency defense in merger review requires weighing often-speculative cost synergies against immediate consumer harms. We provide a monetized cost-benefit analysis of mergers' non-price effects, weighing short-run mortality costs against long-run savings from reduced hospitalization. Because short-run harms arise from the transitional frictions of a process that improves quality in the long run, rather than from increased concentration, our findings point toward a different regulatory emphasis. In markets with regulated prices, consumer welfare may be better protected by conduct remedies that ease patient transitions than by blocking mergers that generate long-run efficiencies.

The rest of the paper proceeds as follows. Section 2 describes the institutional context, covering dialysis provision, provider consolidation and its regulation, and payment and price-setting. Section 3 describes the data, outcomes, study population, and descriptive statistics. Section 4 presents the empirical methods, and Section 5 presents the results. Section 6 develops a model of consolidation under price regulation to interpret the empirical findings. Section 7 presents a cost-benefit analysis comparing the monetary value of short-run health costs and long-run health benefits. Section 8 discusses the implications for antitrust policy, and Section 9 concludes.

2 Background

2.1 What is Dialysis?

End-Stage Renal Disease (ESRD) is the final stage of Chronic Kidney Disease (CKD), occurring when the kidneys' ability to filter waste and excess fluids has been severely diminished. As of 2022, 815,896 individuals were living with ESRD (NIH 2024). Treatment options for ESRD include kidney transplants and dialysis. Although transplants can restore kidney

[12] Sloan et al. (2021) examine post-reform mergers and find that smaller chains and independent facilities were more likely to be acquired, but do not study the health effects of these mergers.

[13] Wollmann (2020) analyzes health outcomes during 2012–2017, after these reforms were implemented, but does not examine whether merger effects differ from the pre-reform period.

function, they are feasible for only a limited subset of patients and frequently involve prolonged waiting periods. For the majority of ESRD patients, dialysis is the only long-term management option.

Hemodialysis is the most common modality of dialysis, and is a medical procedure that involves taking the patient's blood out of their body, passing it through the dialysis machine (dialyzer) which cleans the blood, and returning the clean blood to the patient's body.[14] Hemodialysis is predominantly performed in outpatient facilities or centers, but some patients may opt for home hemodialysis after a sustained period of training.[15] Patients receive hemodialysis three to four times per week, with each session lasting between three and five hours. Because of this high treatment frequency, patients are sensitive to facility location.[16] One implication of the localized nature of hemodialysis care is that patients have limited ability to mitigate the negative effects of a dialysis merger by seeking treatment in another market. A second implication is that, when patients are able to switch facilities, awareness of merger effects across all facilities in the market becomes relevant for both patient choice and health outcomes.

Another form of dialysis is peritoneal dialysis which uses the patient's own abdominal lining (peritoneum) for blood filtration, and is generally performed at home. Together, home hemodialysis and peritoneal dialysis patients account for a small share of the ESRD population with in-center hemodialysis patients making up 58% of all ESRD patients and around 85% of those receiving dialysis (NIH 2024). Accordingly, we focus on in-center hemodialysis—which we henceforth refer to generically as dialysis—patients and assess how mergers influence the quality of care they receive.

2.2 Dialysis Providers and Pre-Merger Review

Dialysis care is primarily delivered by for-profit entities (Medicare Payment Advisory Commission 2021). Dialysis firms have undergone rapid consolidation over the last 30 years. In 2023, two large for-profit chains DaVita Inc. and Fresenius Medical Care provided 70% of dialysis treatments in the U.S. (Rosner et al. 2023).

[14] The dialysis machine contains a filter filled with a solution called dialysate, containing water and chemicals at levels tailored to a patient's needs. The patient's blood flows through the machine, separated from the dialysate by a semi-permeable membrane. Through diffusion, water and waste products like urea and creatinine, and minerals like calcium and potassium are exchanged across the membrane until optimal levels are achieved. The cleansed blood is then returned to the patient's body.

[15] Inpatient hemodialysis is typically reserved for patients undergoing other inpatient treatments.

[16] The median travel distance to a facility for patients receiving in-center hemodialysis is 5.7 miles (Prakash et al. 2014), with patients in rural counties traveling farther than those in metropolitan areas (Velázquez et al. 2022).

Like all mergers, dialysis firm mergers fall under the purview of the Hart-Scott-Rodino (HSR) Act of 1976, or the pre-merger notification program, under which parties to proposed mergers exceeding a specified value threshold must file the deal with the Federal Trade Commission (FTC) and the Department of Justice (DOJ), after which one agency is assigned to review the case.[17] Once the merger is filed, the parties are subject to a 30-day waiting period during which the agency reviews the proposal to determine, based on data and documents they can solicit from the parties, whether the merger has potential for competitive harm.

Antitrust agencies determine whether a merger poses competitive harm through a structured process that involves defining the relevant product and geographic market in which the merging firms compete, and then measuring market concentration before and after the merger using the Herfindahl-Hirschman Index (HHI). If concentration crosses certain thresholds, the merger is presumed anticompetitive, though this presumption can be rebutted. Agencies also evaluate specific theories of harm—whether the merger would allow the combined firm to unilaterally raise prices or degrade quality, whether it would facilitate coordination among remaining competitors, and whether new entry or merger-specific efficiencies could offset any resulting harm (Federal Trade Commission 2023). Historically, this framework centered on the consumer welfare standard, under which a merger's legality turned primarily on its likely effect on price and output.

The 2023 Merger Guidelines depart from this prior approach in several respects. First, they expand the analysis beyond price effects to give greater weight to non-price dimensions of competition, such as quality, choice, and innovation (Athey, Susan and Nevo, Aviv 2026). Second, they place greater emphasis on market structure itself—rather than the consumer welfare standard's focus on likely price and output effects—as an independent basis for antitrust concern, reflected in materially lower HHI thresholds (down from 2,500 under the 2010 Guidelines to 1,800) for presuming illegality (Bolema, Ted 2026). Third, the 2023 Guidelines de-emphasize efficiencies and pro-competitive justifications as a defense, making it harder for merging parties to rebut a structural presumption of harm on those grounds (Athey, Susan and Nevo, Aviv 2026). We return to the 2023 Merger Guidelines in Section 8 to discuss their application and implications in light of our empirical findings.

Upon completing pre-merger review, if the antitrust agency determines that the risk of competitive harm is low, it either terminates the waiting period early or allows it to expire, after which the parties are free to transact the deal. Alternatively, the agency can ask for more information, imposing another waiting period, and begin an additional round of review. At the end of this round, the agency clears the deal, or clears it with caveats—primarily

[17] Prior work by Wollmann (2020) shows that mergers that escape reporting requirements are also harmful and that blocking these mergers can save lives.

ordering the acquirer to divest facilities in some markets—that aim to mitigate competitive harm, or seeks to stop the entire transaction (Federal Trade Commission 2023). In this paper, divestitures are recorded as ownership changes, with divested facilities treated as *acquired*, since mergers are identified solely through ownership changes.[18]

2.3 Payment for Dialysis and the Scope for Price-Setting

To evaluate competitive dynamics in the outpatient dialysis market, it is necessary to examine how treatments are financed and how reimbursement rules constrain provider behavior. The dialysis market operates under a bifurcated payment structure—a dominant, highly regulated public payer regime and a small, lucrative commercial segment.

Prior to 1972, Medicare coverage was restricted to individuals aged 65 and older or those with qualifying disabilities. The establishment of the End-Stage Renal Disease (ESRD) Program in 1972 extended Medicare eligibility to all Americans with ESRD who meet work-history requirements, regardless of age, with Medicaid covering individuals without work-based entitlement.

For patients under age 65, coverage follows a statutory timeline. After a three-month waiting period, Medicare acts as a secondary payer for 30 months while the patient's commercial group health insurance serves as the primary payer. After this 30-month coordination period, Medicare becomes the primary payer, and the private plan shifts to secondary status (CMS 2023).[19] Medicare Part A covers inpatient care and transplant costs without a premium. Outpatient dialysis and physician services fall under Medicare Part B, which requires a monthly premium (approximately $175 in 2024), an annual deductible ($240), and a 20% coinsurance for the Medicare-approved amount (CMS 2023). Importantly, ESRD-based coverage encompasses all Medicare-eligible services for beneficiaries with ESRD, beyond those specific to dialysis.

Prior to 2011, Medicare reimbursed dialysis facilities under a fee-for-service (FFS) system that paid a base composite rate for each treatment (i.e., time on the dialysis chair) while separately reimbursing dialysis-related drugs and laboratory tests. Beginning in 2011, Medicare replaced this system with a case-mix-adjusted prospective payment system (PPS) that bundled the dialysis treatment, routine laboratory tests, and dialysis-related injectable

[18] Our data cannot distinguish facilities acquired in a merger from those divested, as we define mergers as any ownership change. External divestiture data are available for only a small subset of mergers. That chain-to-chain mergers occurred following antitrust enforcement and involved divestitures is supported by FTC rulings. See, for example, the FTC's order in Fresenius' 2012 acquisition of Liberty: `https://tinyurl.com/yyz4j4a3`

[19] ESRD-based Medicare coverage terminates 12 months after a patient ceases dialysis or 36 months post-transplant (CMS 2023).

and oral drugs into a single payment. Under the bundled PPS, outpatient facilities receive a fixed per-treatment base rate—set at $273.82 in 2025 (CMS 2024)—adjusted for patient case-mix, facility volume, nurse training, wage indices, and outlier costs. Concurrently, CMS established the Quality Incentive Program (QIP), which penalizes facilities up to 2% of reimbursement if they fail to meet clinical and reporting benchmarks.

Because Medicare FFS sets the per-treatment rate administratively for the majority of patients—between 51% and 70.6% of patients in the 2013-2021 period—price competition within the traditional Medicare segment is thus completely suppressed by policy design.[20] Among the remaining patients, around three-fifths are covered by Medicare Advantage plans and two-fifths by commercial insurers, for whom prices are theoretically negotiable.[21] Do the payment dynamics in this non-Medicare segment contribute to providers' price-setting power?

Although commercial payers account for a relatively small share of dialysis volume, the magnitude of commercial markups is large. Childers et al. (2019) note that in 2017, DaVita received $1,041 per treatment for commercially insured patients, who comprised around 10.5% of its patients, compared with $248 per treatment from government payers. Mechanically, this implies an average reimbursement of approximately $331 per treatment, or about 34 percent above the level that would prevail if all treatments were reimbursed at the government rate. Thus, even a modest privately insured volume can substantially raise average reimbursement when commercial prices very high.

Similarly, a 2021 Medicare Payment Advisory Commission (MedPAC) report using 2018 Medicare Advantage (MA) encounter data found that, after accounting for age and wage-index differences, MA plans paid about 14 percent more for dialysis services than FFS Medicare rates, though contract-level prices varied widely.[22] Lin et al. (2022) provide more detailed evidence using 2016–17 claims from three large MA insurers, estimating counterfactual government-rate payments and finding that MA plans paid 27 percent more than FFS Medicare for the median outpatient dialysis treatment. At the same time, out-of-network dialysis prices were close to FFS rates, indicating that the departure from FFS pricing appears to arise primarily through negotiated in-network contracts, especially with large dialysis organizations.

[20] The statistic accounts for coverage under Medicare FFS and dual Medicare-Medicaid. See Figure 9.1b in NIH (2024).

[21] In 2021, 28.4% of dialysis patients were covered Medicare Advantage plans and 17.5% by commercial plans. See Figure 9.1b in NIH (2024).

[22] Available at `https://www.medpac.gov/document/chapter-12-the-medicare-advantage-program-status-report-march-2021-report/`, last accessed May 18, 2026.

We analyze 23 years of annual Form 10-K filings from DaVita, a leading U.S. dialysis provider, submitted to the U.S. Securities and Exchange Commission to trace the evolution of the firm's revenues and costs over this period. We report the findings in Appendix Figure A.1. Between 2002 and 2024, DaVita's nominal revenue per treatment increased from $291 to $391. In real terms, however, normalizing to 2021 dollars using the Personal Consumption Expenditures healthcare index, revenue per treatment declined from $451 to $364. Thus, despite the high markups paid by commercial and Medicare Advantage plans, DaVita did not obtain a higher real price per treatment during a period of increasing consolidation.[23] This pattern is also consistent with the declining share of commercially insured patients. At the same time, the figure shows that DaVita's real patient-care costs per treatment declined from $316 in 2002 to $240 in 2024. Real costs per treatment therefore fell by 24%, compared with a 19% decline in real revenue per treatment. Taken together, these trends are consistent with our claim that price regulation in dialysis constrains firms' ability to raise prices following consolidation, and that cost reduction is the primary margin through which consolidation can increase profitability.

3 Data and Descriptive Statistics

3.1 Data

The primary data source is the United States Renal Data System (USRDS), which links longitudinal facility, patient, claims, treatment, payer, laboratory, and clinical data for U.S. dialysis patients.[24] We supplement the USRDS with the CMS Provider of Services (POS) files, which provide facility-level information that is unavailable or insufficiently detailed in the USRDS. In particular, POS records identify facility ownership changes, allowing us to construct merger events, while ownership-type information distinguishes acquisitions of independent facilities from transactions involving dialysis chains. Facility location information allows us to assign facilities to markets. We link the POS and USRDS using Medicare facility identifiers. We construct outcomes at the facility and market levels, assigning patients to

[23] DaVita acquired three large dialysis firms over this period. Gambro Healthcare was acquired in October 2005, DSI Renal was acquired in 2011, and HealthCare Partners was acquired in 2012.

[24] The USRDS is a registry funded jointly by the Centers for Medicare and Medicaid Services (CMS) and the National Institute of Diabetes and Digestive and Kidney Diseases (NIDDK) of the National Institutes of Health (NIH). It includes longitudinal data on all outpatient dialysis facilities in the United States, as well as all Medicare Part A and B claims beginning on the first day of dialysis for all Medicare Fee-for-Service (Original Medicare) patients, payer information, treatment history, laboratory measurements from CROWNWeb submitted in compliance with Medicare QIP requirements, and patients' information from the CMS Form 2728, which is required within 45 days of dialysis initiation. The USRDS is not a public dataset and can be accessed upon Institutional Review Board (IRB) review, proposal submission and a Data Use Agreement (DUA).

facilities based on where they receive the largest share of their annual dialysis treatments; the assignment algorithm is described in detail in Appendix J.

In all analyses, we include market-level covariates intended to capture county demographics and proxy for local dialysis demand. County median income and population by age (in 10-year bins), sex, and race/ethnicity (non-Hispanic Black, non-Hispanic White, and Hispanic) are drawn from publicly available annual U.S. Census vintages. Data on the presence of state-level certificate-of-need (CON) laws governing renal care facilities are obtained from the American Health Planning Association (2026).[25]

3.2 Outcomes and Study Population

We organize our outcomes into six categories, each capturing a distinct margin along which mergers could affect markets, facilities, and patients. Appendix C provides, for each outcome, the clinical background needed to interpret it, the rationale for its inclusion (including where it relates to CMS's Quality Incentive Program), a discussion of how it is constructed from the underlying claims and clinical data (with full variable-construction details in Appendix J), and a review of the relevant empirical and clinical literatures.

First, we measure facility entry and closure using market-level counts of total, acquired, and non-acquired facilities. We classify facilities targeted in the reference merger—those undergoing an ownership change—as *acquired*. All remaining facilities are classified as *non-acquired*, encompassing both acquirer and other facilities, as our data do not permit reliable identification of parent firm ownership. These outcomes speak directly to antitrust concerns over whether entry or expansion by incumbent facilities can offset the competitive effects of a merger. To our knowledge, ours is the first study to empirically examine how mergers affect entry and closure in the dialysis industry, or among health care providers more broadly.

Second, we examine treatment disruptions using dialysis sessions per patient. Mean travel distance and the patient-to-station ratio additionally capture the transition and capacity frictions faced by patients displaced by facility closures. To our knowledge, the empirical literature has not documented these frictions following dialysis facility closures.

[25] Under Certificate of Need (CON) laws, entities proposing the new provision of a health service in a market must prove the public need for the said service in the market and obtain a permit from the state. Every state covers different medical services under the CON umbrella and the existence of a dialysis CON in a state is relevant to this analysis since it restricts facility entry for reasons unrelated to mergers. AHPA data on the status of dialysis CON laws are available only for the years 2004, 2008, 2012, and 2016. Except for one state, no state changes its CON law status across these years. For the intervening years, the CON law status is extrapolated to cover the full sample period. If a state is not present in the AHPA data, its CON law status is assumed to be 0.

Third, we measure market concentration using the facility-level Herfindahl-Hirschman Index (HHI) within county-level markets.[26] We define markets at the county level, following Wollmann (2020), who finds that more than 90% of dialysis patients receive care within their county of residence. Following their approach, we further subdivide three large counties—Cook County, IL; Los Angeles County, CA; and San Diego County, CA—into distinct markets based on observed patient-flow patterns. We calculate facility market shares based on the number of patients receiving treatment at each facility.

Fourth, we examine three facility decisions that serve as observable, upstream indicators of care quality—staffing levels of nurses, technicians, and social workers; anemia management, specifically the share of patients prescribed an erythropoiesis-stimulating agent (ESA) and average monthly ESA dosage; and vascular access modality, specifically the share of patients using arteriovenous fistulas, the lowest-risk access type.

Fifth, we study *biomarkers*—objective physiological measures collected monthly through CROWNWeb and observable in real time by facility staff—including dialysis adequacy (Kt/V), hypercalcemia, ultrafiltration rate (UFR), hemoglobin, albumin, and iron saturation. Kt/V and hypercalcemia are QIP measures tied to reimbursement through explicit performance thresholds; UFR is a QIP reporting measure that does not affect reimbursement. Hemoglobin was a reimbursement-linked QIP measure only from 2012–2014, while albumin and iron saturation are not part of the QIP. Comparing QIP and non-QIP biomarkers allows us to distinguish mergers' effects on explicitly incentivized measures from their effects on patient care more broadly.

Sixth, we examine *clinical endpoints*—outcomes that capture how patients ultimately feel, function, and survive, following the definition of Atkinson Jr et al. (2001). We include hospitalizations, ICU days, and all-cause mortality, as well as inpatient blood transfusions, which provide a measure of anemia management and, together with ESA prescribing behavior, help interpret ESA use as a measure of quality.

The study horizon and population vary by outcome as per USRDS data availability. Appendix Table B.1 provides detail regarding the unit of analysis, data availability years, eligible merger cohorts and patient population characteristics for each outcome evaluated in the primary analyses. To summarize, most outcomes are observed for the full population of in-center hemodialysis patients, including facility capacity, travel distance, patients per station, staffing, mortality, biomarkers. Claims-based outcomes—dialysis sessions, hospitalizations, ICU days, and blood transfusions—are available for Medicare patients. CROWNWeb-based outcomes, including biomarkers, as well as ESA prescription and fistula use are observed for

[26] Because our data do not reliably identify parent firms, we construct market shares at the facility level.

patients regardless of insurance coverage, although they are constructed from the subset of patients observed by facilities on a single day each month.[27] In a set of robustness checks, we assess the sensitivity of our conclusions to varying the unit of analysis (e.g., aggregating facility-level outcomes to the market level).

3.3 Descriptive Statistics

Table 2 reports summary statistics for market-level outcomes and covariates, while Tables 3 and 4 present the corresponding statistics at the facility level.[28] Our analyses encompass 724 treated markets (counties) exposed to mergers and 1,784 control markets, with a mean of 1.79 facilities and 20.98 dialysis stations per market. These markets serve approximately 76 patients on average, of whom 66% are Medicare beneficiaries.

It follows from these aggregates that the mean ratio of patients to stations is 3.65. Patients travel approximately 15 miles from their residence to their dialysis clinic, and receive 116 dialysis sessions annually, corresponding to a schedule of approximately two to three sessions per week. A majority of patients (72%) have an ESA prescription, and 53% have a fistula implanted for vascular access. Facilities employ 11.5 staff members, including approximately five nurses, five technicians, and one social worker.

The dialysis population in our sample performs comparably to the broader U.S. population on some health measures, such as hypercalcemia, but fares substantially worse on others, including low hemoglobin and hospitalizations. Most notably, the crude all-cause death rate among dialysis patients is 0.18, or approximately 18,000 deaths per 100,000 patients, nearly 20 times the corresponding U.S. death rate.[29] Dialysis patients are also disproportionately likely to be Black and to be aged over 70.[30]

[27] Patients receiving peritoneal or home dialysis are excluded from all outcomes. Patients covered by Medicare Advantage, Medicaid, or private insurance are additionally excluded from analyses relying on Medicare claims data.

[28] Our summary statistics are computed using the stacked dataset, so the same market may appear in multiple merger cohorts as a treated or control market. We do not regard this as problematic because our clean-window restriction separates these observations by several years, over which substantial turnover in the dialysis patient population means that the later observation reflects a largely new set of patients and market conditions.

[29] The mean share of a facility's patients with hypercalcemia in the sample is 2%, compared with 1–2% in the broader U.S. population as of 2022 (Cleveland Clinic 2022). The share of patients with low hemoglobin in the sample is 19%; by comparison, the prevalence of anemia, measured by low hemoglobin, in the broader population between 2021 and 2023 was 9.3% (Centers for Disease Control 2024). In the U.S. population, average hospital admissions per person, which we calculate as total hospital admissions (American Hospital Association 2023) divided by the 2023 U.S. population, are approximately 0.1, compared with 1.74 in our sample. Similarly, the crude 2023 U.S. death rate is approximately 923 per 100,000 (Murphy et al. 2024), compared with a death rate of 18% among dialysis patients in our sample, nearly 20 times higher.

[30] In 2023, Black individuals accounted for 14.4% of the U.S. population (Pew Research Center 2025), compared with 28% in our sample. In 2023, 17.8% of the U.S. population was aged 65 and older (KFF 2025), compared with 36% of patients in our sample aged 70 and older.

4 Empirical Methods

To estimate the effects of dialysis firm mergers on capacity, quality, and health outcomes, we use a stacked event-study design implemented through a two-step estimation procedure. This approach captures the dynamic evolution of merger effects, allowing us to distinguish short- and long-run impacts and to account for market adjustments that unfold over time, including firm entry, exit, and scale efficiencies.

Our empirical strategy compares outcomes within markets (counties) that experience a merger with those in markets that do not, and have at least one dialysis facility during the same period. We condition on market-level controls that proxy for demographics and dialysis demand and, in analyses particularly susceptible to threats from patient risk selection, we additionally condition on facility-level controls. Importantly, we treat all facilities in a market experiencing a merger as treated, regardless of whether they are directly acquired. This allows us to compare the trajectories of acquired and non-acquired facilities within the same market and capture market-wide effects of the merger.

4.1 Is merger exposure endogenous?

A first-order consideration motivating our relatively "agnostic" approach to control-group selection is whether mergers are endogenous—that is, whether markets experiencing mergers are selected in ways that require a more narrowly tailored control group. To assess this possibility, we examine which pre-merger observables distinguish markets that experience a merger from those that do not. Appendix D provides a detailed description of this analysis; we summarize the results here. Figure 2 reports coefficients from linear probability models estimating the determinants of merger exposure separately by year. We highlight coefficients significant at the 10% level or higher, with color indicating the direction and shading intensity the magnitude of the coefficient.

Overall, the results reveal relatively sparse and inconsistent predictive power. Most observables do not significantly predict merger exposure in most cohorts, and no characteristic exhibits a stable relationship with merger probability from 1996–2018. The number of facilities is the clearest exception; an additional facility is associated with a 1–13 percentage-point increase in merger probability across cohorts, although even this relationship varies in magnitude and significance across years. Other market, county, facility, and patient-composition characteristics—including county population, demographics, facility age, and our measure of exogenous patient risk (which we describe in further detail in Section 4.2)—show little systematic relationship with merger exposure. Consistent with this pattern, the adjusted

R^2 ranges from 0.07 to 0.42 and clusters around 0.20, indicating that even a rich set of pre-merger observables explains only a limited share of the cross-sectional variation in merger exposure. These findings provide reassurance that treated and control markets are broadly comparable along observable dimensions, while offering little basis for constructing a more narrowly tailored control group. We therefore do not condition the control group on characteristics with weak or unstable predictive power.

Additionally, to gain insight into the factors firms themselves cite when describing merger decisions, Appendix E provides an AI-assisted qualitative analysis of publicly disclosed statements from joint press releases, SEC filings, and regulatory agencies for 10 dialysis-provider mergers. These statements identify recurring themes among disclosed merger rationales, including quality and patient care, strategic and geographic positioning, market expansion, economies of scale and cost efficiencies, and shareholder value. However, these broadly stated objectives generally do not identify the specific economic mechanisms, competitive conditions, or market characteristics that made particular transactions attractive, making them difficult to map directly onto the observable selection patterns. Taken together, these patterns provide little evidence to warrant a more restrictive empirical approach.

4.2 Stacked Event-Study Design

We identify a merger in a market when at least one facility undergoes an ownership change. A market is classified as treated in calendar year t if it experiences a merger in t, and as a control if no facility in the market changes ownership in t or during the surrounding years, as described below, while at least one facility remains in operation throughout that period.[31]

The first-step estimating equation is

$$y_{i,e,d} = \sum_{\substack{k=-\mathcal{L}' \\ k\neq -1}}^{k=\mathcal{L}''} \beta_e^{unadj} \left(\mathrm{D}_{i,d} \times \mathbb{1}\{e=k\}\right) + \mathrm{X}'_{m(i),e}\chi + \mathrm{Q}'_{f(i),e}\theta + \lambda_{i,d} + \Omega_{e,d} + \epsilon_{i,e,d}. \qquad (1)$$

To estimate equation 1, we construct a separate sub-experiment d for each merger year, restricting observations to $\mathcal{L}'$ years before and $\mathcal{L}''$ years after the merger and including markets treated in year d and markets not treated in that year. $\mathcal{L}'$ and $\mathcal{L}''$ vary by outcome, as detailed in Appendix Table B.1. We define event time as $e = t - d$ and stack these sub-experiments, aligning observations around the merger year ($e = 0$). Some markets may

[31] Mergers span three classes, determined by comparing a facility's chain-ownership status one year before and two years after the ownership change (the two-year lag allows time for data updates to be reflected): chain-to-chain (acquisitions of chains by other chains), independent-to-chain (acquisitions of independent facilities by chains), and independent-to-independent (acquisitions of independent facilities by other independent firms). The sample is 47% chain-to-chain, 20% independent-to-chain, and 28% independent-to-independent (Figure 1).

therefore appear as controls in multiple sub-experiments. We describe the stacking procedure in detail in Appendix J.

For some outcomes, the unit of analysis is the market m; for others, it is a facility f within a market, with patient outcomes aggregated to the facility level. Treatment, however, is consistently defined at the market level—all facilities in markets experiencing a merger are considered treated, regardless of whether they are directly acquired.

$y_{i,e,d}$ denotes the outcome for unit i at event time e in sub-experiment d, and $\mathrm{D}_{i,d}$ indicates treatment status. We interact treatment with event-time indicators for $k \in [-\mathcal{L}', \mathcal{L}'']$, excluding $e = -1$ as the reference period. The regression includes unit-by-sub-experiment fixed effects $\lambda_{i,d}$, which absorb time-invariant unit characteristics, and event-time-by-sub-experiment fixed effects $\Omega_{e,d}$, which absorb shocks common to all units in a sub-experiment at a given event time.

To account for time-varying differences across markets m, in all analyses, we include a vector of market-level covariates $\mathrm{X}_{m(i),e}$. These are median household income, populations in 10-year age groups starting from 0 to 9 years and up to 80 years and older, male population, population in each of the racial groups Hispanic, non-Hispanic White and non-Hispanic Black and an indicator for whether the market's state had a Certificate of Need law for dialysis facilities.[32,33]

In analyses of health outcomes, we additionally include facility-level covariates $\mathrm{Q}_{f(i),e}$ including facility age, and variables capturing patient demographics, specifically the share of patients who are female, Non-Hispanic White, Non-Hispanic Black, Hispanic, and those in 10-year age groups ranging from 10–19 up to 80-89. A key additional covariate included in $\mathrm{Q}_{f(i),e}$ is the patient case-mix. Including this covariate accounts for patient selection, so that the estimated merger effects reflect changes due to the merger itself rather than differences in patient health profiles across facilities or over time.[34]

The challenge in constructing the patient case-mix from patients' observed health outcomes is that patient health profiles are endogenous to the merger. To address this challenge we construct a plausibly exogenous metric of the case-mix in the following way. In each cohort-specific dataset, before aggregating outcomes at the facility level or stacking the

[32] Under Certificate of Need (CON) laws, entities proposing the new provision of a health service in a market must prove the public need for the said service in the market and obtain a permit from the state. Every state covers different medical services under the CON umbrella and the existence of a dialysis CON in a state is relevant to this analysis since it restricts facility entry for reasons unrelated to mergers.

[33] Market-level covariates are available at the county-level. Therefore, for almost all markets, the covariates describe the market uniquely but for the three large counties that were further split into smaller markets, covariate values are common within the county in which these markets lie.

[34] Including this covariate also aligns the measurement of hospitalization and transfusion with the Medicare QIP procedure, which standardizes these outcomes to an expected value based on the facility's case-mix as described in Appendix C.6.

data, we regress hospitalizations among Medicare patients on age, dialysis vintage (years on dialysis), sex, race, Hispanic origin, and urbanicity, using only patients in the control group and pooling across event-times.[35] We then use the estimated relationships to predict hospitalizations for all patients, including those in treated facilities, and non-Medicare patients. A patient's predicted risk is fixed based on her characteristics and "travels" with her as she moves across facilities, such that she contributes to the risk composition of each facility she attends over time.

Then, the facility-level case-mix measure is defined as predicted hospitalizations per patient, or total predicted hospitalizations divided by the total number of patients. Because the prediction is based solely on predetermined patient characteristics and estimated using the untreated group, it is not affected by merger-induced changes in hospitalization rates. It therefore captures *exogenous* differences in patient composition across facilities, rather than differences arising from the merger.

The stacked event study approach addresses an important limitation of traditional event study designs in settings with staggered treatment timing, where different units experience treatment (e.g., mergers) in different years. Standard event studies can produce biased or misleading estimates under staggered treatment timing because they impose inappropriate weights, conflate heterogeneous treatment effects, and use contaminated control groups that are either already treated or eventually receive treatment themselves. The stacked event study, which has gained prominence in the recent applied economics literature (Wing, Freedman, and Hollingsworth 2024; Cengiz et al. 2019, and others), resolves these issues by estimating treatment effects separately for each treatment cohort, aligning treatment events as if they occurred simultaneously across cohorts, and comparing treated units only to never-treated or not-yet-treated units during the same relative time window, producing credible estimates of dynamic treatment effects.

To ensure that estimated dynamic effects reflect the reference merger rather than other merger activity, we restrict our sample to markets with a sufficiently long period without merger activity preceding the reference merger and no subsequent merger during the post-merger estimation window.[36] These restrictions yield an analytic sample of 867 merger events across 724 markets from 1996 to 2018. Appendix J provides full details of the sample construction.

[35] Hospitalization data are only available for Medicare patients since they are obtained using Medicare claims.

[36] For cohorts between 2001 and 2016, we require 10 years without merger activity preceding the merger, corresponding to approximately five pre-periods plus five post-periods following any prior merger. For earlier and later cohorts, we impose the longest feasible restriction given data availability from 1991 to 2021.

A potential concern is that the clean-window restrictions exclude markets with consecutive merger activity, so our estimates may reflect merger effects among markets with moderate rather than high merger activity. To assess the importance of this selection, we compare the 3,805 merger events in the full population (Appendix Figure A.3) with those in our analytic sample (Figure 1), which contains 23% of all events. The distributions are broadly similar across samples, both over time and by merger type, with no noteworthy change in merger composition after the 2011 introduction of Medicare PPS. As a complementary and more formal check, we estimate effects on our main outcomes—capacity, dialysis sessions, mortality, and hospitalization—using the full population of merger events and find no qualitative differences from our primary estimates (Appendix Figure A.4). We likewise find no meaningful differences in mortality or hospitalization effects before versus after the Medicare payment policy changes within our primary sample (Figure 17).

4.3 Trend Adjustment

For the outcome measuring market capacity using the total number of facilities in a market, unadjusted estimates of equation 1 reveal a statistically significant negative pre-trend at $e \in [-5, -4]$ (Appendix Figure A.5), consistent with merger-exposed markets experiencing increases in facility counts prior to the merger. This pattern is consistent with the evidence in Section 4.1, which shows that the pre-merger facility count is the only variable with consistent predictive power for merger exposure. We therefore implement a second-stage adjustment, applied uniformly across outcomes, following Freyaldenhoven, Hansen, and Shapiro (2019) and Rambachan and Roth (2023), and used, for example, by Dobkin et al. (2018), Bhuller et al. (2013) and Goodman-Bacon (2018) which subtracts a linear extrapolation of the pre-merger trend from each estimated unadjusted coefficient.

Specifically, after first estimating the unadjusted coefficients $\hat{\beta}e^{unadj}$ using equation 1, we fit a linear trend to the pre-merger coefficients:

$$\hat{\beta}_e^{unadj} = c + m\, e, \qquad e \in [-\mathcal{L}', -1]. \tag{2}$$

We then extrapolate this fitted trend over the full event-time horizon, yielding the predicted value of the trend at each event time:

$$\hat{\delta}_e = \hat{c} + \hat{m}\, e, \qquad e \in [-\mathcal{L}', 5]. \tag{3}$$

We define the adjusted coefficients by subtracting the extrapolated trend from the corresponding unadjusted coefficients:

$$\hat{\beta}_e^{adj} = \hat{\beta}_e^{unadj} - \hat{\delta}_e. \tag{4}$$

Because the extrapolated trend is netted out of every coefficient, including pre-period ones, $\hat{\beta}_{-1}^{adj}$ is not mechanically zero as it would be in a standard event-study normalization. We present only adjusted estimates $\hat{\beta}_{e}^{adj}$ in the main exhibits.

The standard errors associated with the adjusted estimates are bootstrapped, and thus estimated non-parametrically.[37] The bootstrap procedure repeats the full two-step procedure multiple times, sampling unit-sub-experiment pairs from the stacked dataset with replacement, and derives standard errors and confidence intervals for each event-time from the distribution of the resulting bootstrapped adjusted coefficients.[38]

The trend-adjustment procedure rests on the identifying assumption that, absent the merger, the pre-merger linear trend would have continued unchanged. Although this assumption cannot be tested directly, we assess its plausibility by examining the linearity of the unadjusted pre-period coefficients and the magnitude and precision of the *adjusted* pre-period coefficients. For the latter, estimates close to zero across pre-merger periods indicate that the linear fit captures the pre-period trajectory well, while narrow confidence intervals indicate that the pre-merger coefficients are tightly estimated around the fitted line. Observing this pattern provides greater confidence that the linear specification adequately captures the pre-merger trajectory and, consequently, that its extrapolation into the post-merger period is credible. In Section 5, we consider additional threats to the credibility of this extrapolation, including the possibility that the estimated post-merger effects reflect mean reversion rather than the causal effect of the merger.

We note that, for most outcomes, the unadjusted pre-trend is close to and statistically indistinguishable from zero, such that adjustment is unlikely to materially affect the estimates. We nonetheless apply the two-step procedure uniformly across outcomes for consistency, accepting the additional imprecision it introduces where adjustment is unnecessary in exchange for protection against pre-trend-driven bias where it is.

Our choice to conduct pre-trend adjustment also relates to our earlier analysis of merger determinants (Section 4.1). There, we show that pre-merger observables have limited predictive power for merger incidence, making it unclear whether any set of controls could fully account for the observed pre-trends. To the extent that unobserved factors explaining merger probability contribute to these trends, our procedure provides a way to account for their implications without requiring us to observe or specify those factors directly.

[37] Equation 1 clusters standard errors at the unit-sub-experiment level. These are the standard errors associated with unadjusted coefficients that are not presented in any exhibits except Appendix Figure A.5.

[38] Specifically, we use the 2.5th and 97.5th percentiles of the bootstrapped coefficient distribution to construct the lower and upper bounds of the 95% confidence interval, respectively. Because these percentile-based confidence intervals do not impose symmetry around the point estimate, the interval need not be symmetric when the bootstrap distribution is skewed or otherwise non-normal.

4.4 Multiple Hypothesis Testing

For our three primary health outcomes of interest—sessions per patient, mortality, and hospitalizations per patient—we perform multiple hypothesis testing to address both the increased probability of a false positive when testing each outcome separately at the conventional 5% significance level, and the fact that the outcomes are estimated on overlapping samples. We perform the correction separately for the adjusted merger-year (post-0) and long-run (post-5) effects, controlling the family-wise error rate at 5% within each post-merger period across all three outcomes jointly.

We use the Holm step-down procedure, which orders p-values within each family from smallest to largest and tests them sequentially against increasingly permissive thresholds ($\alpha/3$, $\alpha/2$, and α), stopping at the first failure to reject. Unlike procedures that exploit the joint distribution of test statistics across outcomes (e.g., Romano-Wolf), Holm requires no assumptions about the dependence structure between tests, making it appropriate given that our three outcome-specific regressions are estimated on overlapping but non-identical samples. As shown in Appendix Table B.2, both the reduction in sessions ($\hat{\beta} = -4.48$, adjusted $p = 0.00024$) and the increase in the death rate ($\hat{\beta} = 0.0065$, adjusted $p = 0.0467$) remain statistically significant in the merger year, while the effect on hospitalizations remains statistically insignificant (adjusted $p = 0.4465$), as in our baseline estimates. Five years post-merger, only the decline in the hospitalization rate remains statistically significant ($\hat{\beta} = -0.076$, adjusted $p = 0.0287$); neither the sessions nor mortality effects are distinguishable from zero. Thus, accounting for multiple hypothesis testing leaves our conclusions unchanged.

5 Results

5.1 Effects on Facility Closures and Entry

Figure 3 presents the coefficients for the market-level effects on total facilities from equation 4. We find that mergers lead to an initial increase in the total number of facilities in a market, followed by a decline in subsequent years (panel A). By five years after the merger, mergers reduce net market capacity by 0.15 facilities, or by 8.7% relative to the baseline mean.

Decomposing this total effect into acquired and non-acquired facilities reveals a sharp decline in the number of acquired facilities—five years after the merger, there are 0.4 fewer acquired facilities relative to a baseline of 1.03, corresponding to roughly a 40% decline in

the probability of survival (panel B). This decline is partially offset by entry among non-acquired facilities, which increases by 0.18 facilities over the same period (panel C).[39] The relative timing of these responses is also informative—in the early post-merger years, entry by non-acquired facilities more than offsets the exit of acquired facilities, resulting in a net increase in market capacity. By five years after the merger, however, the cumulative decline in acquired facilities exceeds the compensating entry response, yielding a net reduction in market capacity.[40] We observe an analogous merger effect on dialysis stations in the market as shown in Figure 4—a net long-run reduction in station capacity, reflecting a decline in stations at acquired facilities that is partially offset by an increase at non-acquired facilities.

We next consider several additional features of the facility-count results. First, do facility closures reflect mean reversion? Markets experiencing mergers are on an upward trajectory in facility counts prior to the merger, raising the possibility that this pre-merger growth would have leveled off or reversed even in the absence of the merger. To assess this possibility, we exclude markets in the top percentile of pre-merger facility growth, measured as the change in facility count between the fifth and first pre-merger periods, and re-estimate our effects.[41] As shown in Appendix Figure A.6, the resulting estimates are remarkably similar in magnitude to our baseline estimates. This similarity suggests that our findings are not driven by markets with unusually strong pre-merger growth and, consequently, a greater propensity for mean reversion.

Second, are facility closures and entry driven by particular types of mergers or markets? Appendix Figure A.13 shows that mergers trigger both closures of acquired facilities and entry by non-acquired facilities across all merger types, indicating the generality of our findings. The relative magnitude of these responses, however, varies across merger types, particularly for entry. As a result, chain-to-chain mergers generate a net reduction in total capacity, independent-to-chain mergers produce little to no net change, and independent-to-independent mergers generate a modest net increase in total capacity. Panel A of Appendix Figure A.4 shows that mergers also reduce net market capacity when we consider the full

[39] We note again that data limitations only allow for the distinction between a acquired facility and other facilities, but preclude the assignment of granular firm identity to these facilities. Therefore it cannot be determined whether the entry of non-acquired facilities is by the acquiring firm, which could open new facilities in addition to acquiring the target in the same market, or by other non-merging firms.

[40] The post-merger coefficients from panels B and C of Figure 3 sum to values that are close to, but not exactly equal to, the post-merger coefficients in panel A. This discrepancy arises from the inclusion of market-level controls, which alter the partial correlation between each outcome and the treatment in distinct ways, as well as the additional step of adjusting coefficients relative to a fitted line, which is estimated separately for each outcome.

[41] We use the 99th percentile as the cutoff because pre-merger facility growth takes on relatively few distinct values. As a result, lower percentile thresholds often map to the same underlying growth value and do not provide a meaningfully distinct sample restriction.

population of mergers, without imposing the baseline sample restrictions on the frequency of consecutive merger activity within markets.[42]

Third, given barriers to entry, is the expansion by non-acquired firms plausible? Startup costs for new dialysis clinics can be substantial, and entrants must obtain a Certificate of Need before opening a facility. Yet, the supply of dialysis facilities has grown over time (Velázquez et al. 2022). Using anecdotal estimates of facility startup costs and projected annual profits (Arnold, Carrie 2020), our back-of-the-envelope calculation suggests a conservative break-even horizon of approximately six years.[43] This calculation does not imply that entry is infeasible, however, since potential entrants may base their decisions on expected profits net of startup and other entry costs. We discuss this mechanism more formally in Section 6. For now, we note that the entry response is driven primarily by chain operators (Appendix Figure A.7), rather than by independent firms. Chain operators may face lower effective expansion costs due to economies of scale, making entry more feasible for these firms. The observed entry patterns are therefore consistent with sufficiently attractive profit opportunities to induce new market participation, particularly among chain operators.

Fourth, why does entry by non-acquired facilities, which is observed as early as the merger year, precede the exit of acquired facilities? Several explanations are plausible. First, the pattern may reflect the annual frequency of our facility-level data. Because facility status is recorded only once per year, an acquired facility that closes earlier in the year and a non-acquired facility that opens later in the same year will both appear as active, making entry appear to precede exit even when the closure occurs first. Second, the timing may reflect firms' forward-looking responses. Non-acquired firms may anticipate subsequent closures of acquired facilities and enter preemptively to capture the resulting demand, potentially drawing on information from prior market experience. Third, the timing may reflect administrative delays—facility closures may take time to be fully implemented and recorded, such that the observed exit date lags the underlying decision to close.

[42] This broader sample comes with the caveat that event-time observations may overlap across mergers within a market; for example, period 1 following one merger may coincide with period 2 following an earlier merger. Thus, the estimates should not necessarily be interpreted as isolating the effect of a single merger in markets experiencing closely spaced merger activity. Nevertheless, the persistence of the capacity reduction in this unrestricted sample provides additional evidence that mergers reduce net market capacity.

[43] Arnold, Carrie (2020) reports startup costs of approximately $3 million for a new facility. Annual revenues are also roughly $3 million, with an estimated net profit margin of 18%, implying annual profits of slightly more than $500,000. Assuming the initial startup costs are recovered evenly over time, a facility would require approximately six years to recoup its initial investment.

5.2 Effects on Patient Volume, Travel Distance and Utilization

A natural question that arises from the analysis of facility closures is how mergers affect patient volume and displacement. Figure 5, presenting the effect of a merger on the market's patient volume, shows that the total patient volume in the market remains largely unaffected by the merger. This evidence supports the view that patients remain local. Under the assumption that mergers do not affect the underlying incidence of kidney disease, there is no net patient entry or exit from the market.[44,45] At the same time, patient volume declines at acquired facilities, offset by a corresponding increase at non-acquired facilities. The relative magnitudes of the period-by-period effects on patient volume are informative about placement frictions—gains at non-acquired facilities lag losses at acquired facilities, suggesting delays in patients' transitions to new clinics.

Patient displacement following facility closures also increases travel distances. Figure 6 uses a facility-level measure of road travel distance between patients' residence ZIP codes and their clinic ZIP codes and shows an increase of nearly 0.4 miles.[46] While this increase in travel distance may appear modest, its clinical relevance may be greater for dialysis patients, who often face substantial mobility constraints (Bogataj et al. 2023). For patients who rely on public transportation or others for transportation to treatment, even modest increases in distance can translate into substantially greater travel burdens. Figure 7 also suggests that the ratio of patients to stations increases following the merger, indicating greater utilization of facility capacity, although the estimate is somewhat imprecise.

5.3 Effects on Market Concentration

As discussed in Section C.3, we construct our measure of market concentration—the HHI—using facility-level shares of patients. Changes in market concentration therefore capture the post-merger effects of facility closures, entry by non-merging firms, and the resulting reallocation of patients across facilities within the market. Figure 8 shows that mergers consistently increase the HHI throughout the post-merger period, with an increase of approximately 6.5 percentage points five years after the merger. This magnitude would raise

[44] Minor long-run declines in the total patient count may reflect changes due to mortality, as will be discussed in Section 5.4, among prevalent patients within the market.

[45] Appendix Figure A.8 shows a modest shift in insurance composition following the merger, with fewer non-Medicare and more Medicare patients by five years post-merger.

[46] We calculate great-circle ("as the crow flies") distances of up to 100 miles between ZIP-code centroids using the ZIP Code Distances database from National Bureau of Economic Researc (2025) and scale them by 1.417 to approximate road distance, following Wollmann (2020). This factor corresponds to the average ratio of road mileage to great-circle distance estimated by Boscoe, Henry, and Zdeb (2012) using 66,011 patient-to-hospital trips, for which the two measures are highly correlated.

the HHI from 85% at baseline to approximately 91.5% five years after the merger, or, using the scale employed by the Department of Justice and Federal Trade Commission, from 8,500 to 9,150 points. This increase is large by the standards of antitrust enforcement—the 2023 Merger Guidelines characterize markets with HHI above 1,800 as highly concentrated and treat an increase of more than 100 points in such markets as sufficient to trigger a structural presumption of anticompetitive harm (U.S. Department of Justice and the Federal Trade Commission 2023). The estimated 650-point increase is therefore more than six times this threshold, highlighting the extent to which mergers further concentrate an industry that is already highly concentrated.

5.4 Short-Run Effects on Treatment and Mortality

Figure 9 presents the facility-level effects of mergers on annual dialysis sessions per patient. Mergers reduce sessions by approximately 4.5 per patient in the merger year, with no detectable effects in subsequent years. If concentrated within a short period, such as a single month, this reduction would represent a clinically meaningful 37% decline relative to the typical prescribed frequency of 12 sessions per month, potentially exposing patients to acute risks associated with missed dialysis treatments (Leggat et al. 1998; Saran et al. 2003). Given its lack of persistence, we interpret this effect as a disruption in treatment.

Disruptions in treatment are consistent with both temporary gaps in care as patients are displaced by facility closures and have yet to secure placement elsewhere, and congestion at surviving facilities as displaced patients are absorbed. Even absent eventual closure, reduced operations, renovations, or resource reallocation in anticipation of closure may temporarily constrain treatment capacity. The decline in treatment delivered at surviving facilities in the merger year may therefore reflect these transitional disruptions, with care stabilizing in subsequent years as patients are reassigned and facilities adjust their operations.

Next, we present the evidence on mergers' effects on a facility's death rate. Figure 10 shows that mergers cause a sharp increase in the mortality rate in the year of the merger by about 0.65 percentage points and 3.8% of the baseline mean of a 17% mortality rate. For ease of interpretation, this represents an increase in the death rate among the dialysis population by 650 per 100,000 patients, which is a little over two-thirds of the full-population crude U.S. death rate in 2023 (Murphy et al. 2024), around three times the crude death rate from heart disease, and around 3.5 times the crude death rate from cancer in 2023 (Ahmad, Cisewski, and Anderson 2024).[47]

[47] Crude death rates are most comparable to our measure of mortality, defined as the number of deaths divided by the relevant population. Our specifications additionally include demographic controls for the age and race composition of counties and

The strong temporal coupling between mergers' effects on sessions per patient and mortality, together with clinical evidence on the consequences of missed dialysis sessions, leads us to conclude that these outcomes are mechanistically linked. Interpreted alongside the effects on facility closures and patient displacement, the evidence points to a mechanism in which mergers induce closures, displace patients, and generate short-run frictions in placement at new clinics that disrupt treatment, thereby increasing short-run mortality.

Before turning to a fuller discussion of the results and their underlying mechanisms, we first examine the robustness of the estimated effects on sessions and mortality. First, we show that the qualitative pattern of effects is robust to aggregating outcomes at the market rather than facility level, and to estimating effects over the full population of mergers. Appendix Figure A.10 shows that, even when the unit of analysis is the market rather than the individual facility, mergers reduce sessions per patient (panel A) and increase mortality (panel B).[48] The magnitudes of both effects are somewhat smaller at the market level but remain clinically meaningful. Market-level analysis is also useful in addressing concerns about how patients are assigned to facilities during periods of substantial transition, which could otherwise introduce spurious variation into facility-level estimates. Panels B and C of Appendix Figure A.4 additionally show that mergers reduce sessions per patient in the merger year and increase mortality in the full population of mergers, suggesting that our baseline effects are not driven by particulars of the markets which become a part of our sample.[49]

Second, we address a potential concern that constructing the death rate using the full patient count in the denominator may distort the measure when deaths occur early in the calendar year. Appendix Figure A.9 presents estimates using a denominator that includes only patients who survive at least the first half of the calendar year. We again find a pronounced increase in mortality in the merger year.

---

facilities, providing an implicit adjustment for differences in underlying mortality risk across markets. Thus, the comparison is not driven simply by treated markets having older or otherwise higher-risk patient populations. We calculate the 2023 U.S. crude death rate by dividing the 3,090,964 deaths reported in Murphy et al. (2024) by the 2023 U.S. population of 334,914,895 from the Census Bureau's population estimates. We obtain the approximate numbers of cancer and heart disease deaths in 2023 from Figure 2 of Ahmad, Cisewski, and Anderson (2024) and divide each by the U.S. population to calculate the corresponding crude death rates.

[48] For market-level analyses of sessions, mortality, and hospitalizations per patient, we weight observations by the number of facilities in the market, aligning conceptually with our baseline facility-level analyses, where each facility receives equal weight.

[49] The magnitudes are somewhat smaller, however, and the mortality effect appears to persist for one to two years following the merger. This pattern is unsurprising given that, in the unrestricted sample, event-time observations from successive mergers may overlap. As a result, estimates for a given merger may be contaminated by the dynamic effects of other mergers occurring in the same market, including the longer-run stabilization in treatment intensity and mortality effects associated with earlier mergers. Such overlap would tend to attenuate the estimated merger-year effects and spread the estimated response across event time.

We next return to a discussion of the baseline results, subjecting the causal interpretation of the relationship between closures, treatment disruptions, and mortality to further scrutiny. Figure 11 presents the effects of mergers on these outcomes separately for markets with fewer than three facilities and those with three or more facilities in the year before the merger. Markets with smaller baseline capacity experience larger facility closures in both levels and percentage terms—five years after the merger, these markets have 0.42 fewer facilities relative to a baseline of 0.945, compared with 0.21 fewer facilities relative to a baseline of 1.21 in larger markets. A larger closure response in smaller markets is consistent with greater scope for acquirers to consolidate and recapture displaced demand, a closure motive we discuss in detail in Section 6.[50]

Further, patients in markets with smaller baseline capacity experience larger treatment disruptions—nearly 7 fewer sessions per patient, compared with 3.75 fewer sessions in markets with larger baseline capacity, despite similar baseline treatment volumes of 118–121 sessions per patient. Mortality also increases more in smaller markets, by 0.73 percentage points relative to a baseline of 0.180, compared with a 0.45 percentage-point increase from a baseline of 0.167 in larger markets. Thus, the same markets that experience larger closure responses also experience larger treatment disruptions and mortality increases, providing further evidence for the hypothesized causal pathway from closures to treatment disruption to mortality.

A final question concerning the interpretation of the mortality effects is whether they reflect "harvesting," or mortality displacement, whereby mergers disproportionately hasten the deaths of patients who would have died in the near term regardless. Two pieces of evidence argue against this interpretation. First, if the increase in mortality in the merger year reflected harvesting, we would expect to observe offsetting declines in mortality in subsequent years, as deaths that would otherwise have occurred later were shifted forward into the merger year. Figure 10 shows no evidence of such offsetting declines.

Second, Appendix Table B.3 examines the dynamic effects of mergers on the risk composition of patients at affected facilities. We estimate an unadjusted facility-level event study using the specification in Equation 1, without additional covariates. If the merger-year mortality effect reflected harvesting, we would expect the subsequent patient pool to become systematically less risky as patients whose deaths were accelerated by the merger would no longer be present. We find little evidence of such a compositional shift. Mergers have no effect on the gender composition (column 2) or racial composition (columns 3–5) one or two

[50] Closures are more attractive when acquirers can recapture displaced patients. In smaller markets, overlapping catchment areas make patients from a closed facility more likely to shift to the acquirer's other facilities rather than competitors, increasing the profitability of closure.

years after the merger. Moreover, rather than reducing exogenous risk, as would be expected under harvesting, mergers modestly increase predicted hospitalization rates over this period (column 1). We likewise find no significant decline in the share of older patients in the post-merger years (columns 10–13 for patients aged 50 and older), despite the fact that these patients would be among those whose deaths we would expect to be disproportionately accelerated under a harvesting interpretation. Taken together, neither the dynamic pattern of mortality nor the evolution of patient risk composition exhibits the patterns that harvesting would predict.

5.5 Effects on Facility Decisions

The next set of results describes effects of mergers on facility decisions. First, we examine whether mergers affect facility staffing, namely the sum of nurses, patient care technicians and social workers. Panel A of Figure 12 shows that mergers result in a reduction in staff in the short-term, with approximately 0.3 fewer personnel in the year of the merger. The timing of the response to staffing mirrors the timing of the decline in treatment, consistent with short-term operational disruptions. Appendix Figure A.11 shows that reductions are concentrated among nurses and technicians, with no effect on social workers and no meaningful changes in the full-time versus part-time composition. This pattern again suggests a temporary disruption to employment rather than a persistent change in workforce composition or employee quality.

Second, we consider whether mergers affect anemia drug utilization. Panel B of Figure 12 shows that mergers result in a reduction in the share of patients with an erythropoietin stimulating agent (ESA) prescription.[51] Appendix Figure A.12 reflects another dimension of merger-induced changes in ESA prescription practices, and shows a reduction in the average monthly dosage of ESAs. Our results, which show an unambiguous decrease in ESA prescriptions, contrast with the findings in Eliason et al. (2020) who find that that ESA dosage increases when independent firms are acquired by chains. However, we note that the study sample in Eliason et al. (2020) does not directly overlap with ours. Their study examines acquisitions by chains of independent firms occurring before the 2011 Medicare bundled payment reform, whereas ours focuses on the post-reform period, using ESA dosing data from 2012 onward from CROWNWeb, and encompasses a broader set of merger types. The timing of the reform is especially relevant for interpretation. Prior to 2011, dialysis facilities were reimbursed on a fee-for-service basis for ESA doses, whereas after the reform,

[51] We acknowledge that the pre-period coefficients for this outcome are not as close to zero as we would prefer. Results for dosing, however, provide reassurance regarding the interpretation of this result.

ESAs were incorporated into the larger dialysis bundle, fundamentally flipping the profit incentives for ESA prescriptions.[52] Our finding is consistent with Eliason et al. (2020) in suggesting that consolidation amplifies the prevailing financial incentive—in this case, by enhancing firms' ability to extract per-treatment profits from reducing ESA use. Effective ESA dosing is crucial for anemia management, but neither high nor low doses alone reveal whether the dosing is appropriate. To assess its impact, we examine the need for blood transfusions, which we discuss in Section 5.6.

Finally, the panel C of Figure 12 suggests that mergers do not have any meaningful effects on the share of patients utilizing fistulas for vascular access. Because higher fistula use indicates higher quality, this result suggests no change in quality along this dimension.

5.6 Long-Run Effects on Hospitalization

As discussed in Section 5.4, mergers cause a short-run disruption to treatment and a corresponding spike in mortality. In the long run, however, patient health improves. Figure 13 shows that, across facilities, hospitalizations per patients decline in the post-merger period. Overall, patients experience around 0.075 fewer hospitalizations, which, relative to the baseline of 1.75, represents a 4.3% decline. We show in Section 7 that nationally, the reductions to hospitalization translate to around 1% in Medicare inpatient savings for ESRD patients. In Panel C of Appendix Figure A.10, we show that, even after aggregating hospitalizations and patients to the market level to conduct a analogous market-level analysis, hospitalizations still appear to decline, with point estimates approaching the magnitude of our baseline estimates, although the effects are estimated with substantially greater imprecision. Panel C of Appendix Figure A.4 likewise shows that mergers reduce hospitalizations when we consider the full population of mergers, without imposing the baseline sample restrictions on the timing of successive merger events.

To assess whether mergers affect the severity of hospitalizations, we examine their impact on ICU days associated with a hospitalization. We show in Figure 14 that mergers result in reductions in ICU days per patient. They also reduce transfusions per patient as shown in Figure 15.[53] Our evidence suggests that the reductions in ESA prescriptions that we discussed in Section 5.5 do not compromise health, as they are not linked to increases in severe, transfusion-requiring anemia.

[52] See the report "How the Bundle Has Changed Clinical Care at Fresenius Kidney Care" (Fresenius Medical Care N.d.) from a large dialysis chain that provides anecdotal evidence for the relevance of this mechanism.

[53] We emphasize that our outcome is transfusion events, and not units of blood transfused.

To summarize, while the short-run effects of mergers point to health harms, patients experience unambiguous improvements in health in the long run, as the market stabilizes. A potential concern is that the long-run gains are an artifact of risk selection over time. If the merger-year mortality spike disproportionately removes high-risk patients from the sample, the remaining patient pool would be healthier on average, potentially generating the subsequent reductions in hospitalizations through compositional change rather than improvements in care. We find little evidence of such selection. Column 1 in Appendix Table B.3 reports the dynamic effects of mergers on our exogenous measure of patient risk—predicted hospitalizations per patient. We find no reduction in this measure in the post-merger period. Similarly, we find no evidence that the post-merger years are characterized by a lower share of patients in higher-risk groups—the coefficients for years 1 through 5 are not significantly negative for the share of patients aged 50 and older (columns 10–13) or for the share of Hispanic or Black patients (columns 4–5). The post-merger patient pool, therefore, does not appear to become systematically lower risk.

Moreover, our hospitalization estimates condition on a rich set of patient-level controls, including predicted hospitalization risk and other patient demographics, further reducing the scope for compositional changes to account for the estimated reductions in hospitalizations. Finally, the dynamics of the hospitalization effects are inconsistent with a one-time compositional shift. If the merger-year mortality increase removed high-risk patients from the market's patient pool, we would expect a discrete decline in hospitalizations immediately after the merger, followed by stabilization once the patient composition adjusted. Instead, hospitalizations decline progressively over the subsequent years. This pattern is more consistent with accumulating improvements in care than with a one-time change in the risk composition of the patient pool.

What, then, explains these long-run gains? We provide a detailed exposition of the mechanisms in Section 6. In brief, the long-run gains arise through two channels—within-firm improvements in quality at facilities that survive the merger, particularly those operated by the acquiring firm, and market composition effects resulting from the exit of lower-quality target facilities and the entry of higher-quality facilities. Importantly, the same mechanism that generates the short-run harm—facility closures—also plays a role in generating the long-run benefits.

5.7 Effects on Biomarkers

Figure 16 shows the effect of mergers on biomarkers. For ease of interpretation the quadrant with green shading in the panels within this figure represent a clinically beneficial outcome.

We note that the figure reflects the analysis of outcomes measured starting in 2012 when Medicare began collecting detailed biomarker data as a part of the Quality Incentive Program (QIP). Panel A shows that the share of a facility's patients with recommended dialysis adequacy ($\geq$ 1.2 mg/dL) increases after a merger, though the effect is noisily measured. Panel B shows the effect of mergers on the share of patients with hypercalcemia ($\geq$10.2 mg/dL), finding that mergers significantly reduce this share.[54]

Next we consider patients' ultrafiltration rate (UFR), where both very high and very low UFRs are considered to affect patient health negatively. Panel C in Figure 16 shows that mergers lead to a sustained increase in the UFR in years following the merger. The magnitude of this increase reaches at most 0.3, which, when evaluated in relation to the baseline mean, would result in a UFR of at most 8 mL/kg/hr. This level remains below the thresholds associated with adverse endpoints in clinical literature (Flythe, Kimmel, and Brunelli 2011; Kim et al. 2018), yet represents an adverse outcome following a merger.

Probing further, Panel D reports the effects of mergers on patients' hemoglobin levels. The effects on the share of patients with hemoglobin below 12 g/dL are mixed—the share increases in the merger year but subsequently declines. With respect to albumin, panel E shows that by three years post-merger the mean level declines enough relative to the baseline mean of 3.66 g/dL to fall just below the recommended minimum of 3.5 g/dL. For iron saturation, the results in panel F indicate a decline from the baseline mean of about 30% to about 25%. This level remains within the normal range for the average dialysis population (National Kidney Foundation 2025).[55]

Taken together, the evidence in Figure 16 is somewhat ambiguous. Three biomarkers—Kt/V (panel A), hypercalcemia (panel B), and hemoglobin (panel D)—improve by three years after the merger, while three others—UFR (panel C), albumin (panel E), and iron saturation (panel F)—deteriorate. An interesting pattern emerges; the biomarkers that improve are those that are consistently included in the set of measures to which QIP directly attaches penalties (Kt/V and hypercalcemia), or that were previously included in that set (hemoglobin). In contrast, the biomarkers that do not improve are not directly tied to penalties, even though they are required to be reported or collected. This pattern suggests that merger-induced improvements in biomarkers may reflect targeted responses to QIP incen-

[54] Appendix Figure A.2 shows that this is effect is not coupled with increases in hypocalcemia (<8.5 mg/dL), suggesting an unambiguous improvement in calcium levels among dialysis patients.

[55] A 25% iron saturation falls below recommended standards for EPO-resistant patient populations (Fishbane et al. 1996). The reductions in iron saturation shown in panel F align with the declines in facilities' ESA dosages reported in Figure 12. While facilities often offset ESA reductions with increases in intravenous iron to manage anemia, the evidence on iron saturation suggests that such compensatory adjustments are not occurring here.

tives rather than broad-based improvements in clinical management, consistent with mergers strengthening firms' ability or incentives to comply with targeted quality requirements.

At the same time, we find unambiguous improvements in downstream health outcomes, including hospitalizations, ICU days, and blood transfusions. This contrast raises two possible clinical interpretations, particularly given the still-evolving evidence on the clinical relevance of these biomarkers and the appropriate policy thresholds. One possibility is that, taken as a bundle, these biomarkers are imperfect predictors of downstream inpatient health. Another is that the QIP may be identifying clinically meaningful dimensions of care—the biomarkers prioritized by the program improve alongside downstream measures of patient health. We do not take a position on these interpretations; rather, the results suggest that the relationship between targeted quality metrics and broader health outcomes warrants further investigation.

5.8 Quality Incentives and Merger Effects

In 2011-2012, Medicare implemented two major policy changes. First, it introduced the Quality Incentive Program (QIP), which requires facilities to report biomarkers and other quality measures on a monthly basis and imposes penalties of up to 2% of reimbursement for failure to meet specified standards. Second, Medicare introduced the Prospective Payment System (PPS), which replaced the incumbent itemized payment regime with an all-inclusive bundle. An important question is whether the introduction of explicit quality incentives through QIP changes the effects of mergers on patient health.

To investigate this question, we estimate merger effects on clinical endpoints separately for the periods 1991–2011 (merger cohorts 1996–2008) and 2012–2021 (merger cohorts 2015–2018). Before discussing the results, we highlight three considerations. First, because QIP and the PPS were implemented contemporaneously, we cannot cleanly attribute any change in merger effects to QIP alone. Indeed, at least one of our outcomes—anemia drug prescribing—appears to respond to the PPS.

Second, some clinical endpoints may not be expected to respond directly to QIP. Most notably, QIP does not directly incentivize mortality, and our evidence suggests that the merger-year mortality effect reflects short-run operational disruptions rather than changes in long-run quality. Nevertheless, if QIP broadly improves clinical quality, we would expect mortality to co-move with outcomes that are directly incentivized by the program.

Third, the two policy changes could have endogenously altered the composition of mergers, for example by changing which facilities or firms choose to merge. In that case, a comparison of merger effects before and after 2012 could reflect changes in the types of

mergers being undertaken rather than changes in the effects of a given merger type. We return to this compositional concern after presenting the results.

Figure 17 reports point estimates and confidence intervals for the *difference* between the estimated merger effects in the pre- and post-QIP samples for two key clinical endpoints—mortality and hospitalizations. To account for the possibility that the same facilities appear in both samples, we implement a joint bootstrap in which the pre-QIP and post-QIP estimates, and hence their difference, are computed using the same resampled dataset in each iteration. This procedure does not require the two estimates to be independent—any covariance arising from facilities appearing in both estimation samples is preserved automatically through the joint resampling procedure.

Panel A shows no statistically significant difference in the merger effect on mortality across the two periods. The estimated merger-year mortality difference is close to zero in the pooled sample, although mergers after 2012 appear to have a somewhat larger mortality effect in the subsequent year. Panel B likewise shows no statistically significant difference in the effect on hospitalizations per patient. Thus, we find no evidence that the introduction of QIP mitigated the adverse health effects of mergers, as measured by merger-year mortality, or amplified their long-run effects on hospitalizations. If anything, the estimated differences are positive for both outcomes in some periods, implying somewhat higher mortality and hospitalization effects after 2012, although neither difference is statistically distinguishable from zero.

We next consider whether changes in the merger composition could obscure an effect of QIP. If the policy changes altered the types of mergers taking place, and merger effects differ systematically across types, then the pre- versus post-QIP comparison could confound changes in merger composition with changes in the effect of a given merger type. For example, QIP could improve health for some types of mergers while having no effect, or even adverse effects, for others. If the post-QIP period contained a larger share of mergers in the latter group, this compositional shift could attenuate the aggregate effects and make QIP appear less effective than it actually is.

Two patterns in our data make this explanation unlikely. First, Figures 1 and 2 provide little evidence of a substantial change in merger composition after 2012. The types of mergers occurring after 2012 do not appear qualitatively different from those occurring before 2012, nor do the characteristics of markets in which mergers occur appear to shift substantially. Thus, we find little evidence that the policy changes materially altered the selection of merger types.

Second, the direction of the observed pre- versus post-QIP differences is difficult to reconcile with a compositional explanation based on the heterogeneity documented in Appendix

Figure A.15. Consider, for example, a shift toward independent-to-independent mergers. These mergers are associated with modest increases in long-run hospitalizations but uniquely exhibit declines in long-run mortality. If a greater share of such mergers after 2012 were responsible for attenuating the long-run hospitalization gains associated with QIP, we would also expect the post-QIP mortality effect to become more negative. We do not observe such a pattern—the difference in mortality effects instead points, if anything, toward a more positive post-QIP effect. Conversely, suppose the post-QIP period contained a larger share of chain-to-chain mergers. These mergers exhibit little effect on long-run mortality, which could mechanically push the mortality difference upward. But chain-to-chain mergers also generate larger reductions in hospitalizations, implying that such a compositional shift should make the hospitalization difference more negative. Again, this is not what we observe.

The fact that the two outcomes move in directions that are difficult to reconcile with the merger-type heterogeneity in Appendix Figure A.15 makes it unlikely that changes in merger composition are driving the pre- versus post-QIP comparison. Taken together, the evidence therefore provides little indication that the introduction of QIP materially changed the effects of mergers on the clinical outcomes we study.

5.9 The Mechanistic Role of Market Concentration

In Section 5.3, we showed that dialysis firm mergers increase market concentration. We next present evidence suggesting that, while changes in concentration are a mechanical consequence of mergers, their role as a mediator of mergers' health effects is more limited.

To test the role of market concentration changes, we conduct our outcome analyses separately for mergers that ex post generate larger versus smaller increases in market concentration. Because market concentration is itself an endogenous outcome, we stratify markets based on pre-merger HHI rather than on the realized change in concentration. Specifically, for each market, we calculate its pre-merger HHI as the average HHI over the five pre-merger periods and classify markets as either above or below the median of this market-level measure.[56] We then conduct our two-step event-study analyses separately for the two groups and test whether the resulting coefficients differ across groups. We calculate the difference between the two outcome coefficients and divide it by the square root of the sum of their bootstrapped standard errors, following Clogg, Petkova, and Haritou (1995) and Paternoster et al. (1998). Because the two groups consist of mutually exclusive sets of markets and facilities, the estimates are based on independent samples, making this test appropriate.

[56] We perform this calculation in the raw, unstacked market-level dataset.

We present the results in Figure 18, plotting the coefficients and confidence intervals separately for each subsample and reporting both the difference between the coefficients and the corresponding t-statistic.[57] Before turning to the health effects, we first examine whether mergers generate differential changes in market concentration across markets with below- versus above-median baseline HHI. This analysis establishes the relevant treatment intensity for interpreting the subsequent heterogeneity in health outcomes. Ex ante, it is ambiguous whether mergers would generate larger or smaller HHI increases in markets with higher pre-merger concentration. Mechanically, the change in HHI from a merger is roughly proportional to the product of the merging facilities' baseline market shares. This creates two offsetting effects. In concentrated markets, if the merging facilities are themselves large, this could generate large HHI increases. But high-concentration markets are usually dominated by one or two large incumbents—as is the case in our sample—while the remaining, smaller facilities compete for the residual share, and these smaller facilities are the ones more likely to be involved in a given merger or closure. As a result, concentrated markets often have less "room" for the specific facilities involved in a merger to generate large HHI changes.[58] We find this second effect dominates—panel A of Figure 18 shows that markets with above-median baseline HHI see smaller post-merger increases in HHI than markets with below-median baseline HHI.

Panel B shows that mergers do not have significantly different effects on treatment disruption across markets with larger versus smaller concentration gains, as measured by the merger-year effect on sessions per patient.[59] Panel C likewise shows no statistically significant difference in merger-year mortality effects between markets with larger and smaller concentration gains. Finally, panel D shows that the effects on hospitalizations are also not significantly different across the two groups, though the effects are more precisely estimated among markets with higher concentration changes.

[57] In Figure 18, we report the baseline mean HHI for each group, defined as the mean HHI in period -1 among the markets assigned to that group. Thus, the baseline mean describes the average concentration of markets within each group and is distinct from the criterion used to assign markets to groups.

[58] For example, consider a below-median-HHI market with three facilities holding patient shares of 0.667, 0.167, and 0.167, yielding a baseline HHI of 0.50. If one of the smaller facilities merges with the largest, the resulting shares are 0.834 and 0.167, increasing HHI to 0.72—a 0.22 increase. Now consider an above-median-HHI market with two facilities. If their shares are 0.816 and 0.184, baseline HHI is 0.70, and a merger raises HHI to 1.00, a 0.30 increase. In this case, higher baseline concentration predicts a larger concentration gain. However, if the two facilities instead have shares of 0.947 and 0.053, baseline HHI is 0.90, and the merger again raises HHI to 1.00, but the increase is only 0.10. Thus, depending on the initial distribution of patient shares, either above-median or below-median HHI markets can experience larger merger-induced increases in concentration.

[59] We find some evidence of differential long-run effects on sessions; however, the pattern runs counter to the hypothesized mechanism, with markets experiencing smaller concentration gains exhibiting more persistent declines in treatment.

Overall, our analyses lead us to conclude that mergers' harmful effects, particularly the treatment disruptions and mortality effects, are not driven by changes in market concentration. We find some suggestive evidence that larger concentration gains are associated with larger long-run reductions in hospitalizations, consistent with a mechanism in which acquirers use the scale efficiencies associated with a larger patient base to deliver higher-quality care. We discuss this mechanism in Section 6. Our broader conclusion regarding the limited role of concentration is consistent with prior studies of the dialysis industry (Eliason et al. 2020; Cutler, Dafny, and Ody N.d.), which find no systematic relationship between market concentration and patient outcomes.

6 A Model of Consolidation, Capacity and Quality Under Price Regulation

We develop a simple model of how consolidation affects capacity and quality when prices are fixed. In markets where prices are flexible, mergers increase firms' ability to exercise market power by raising prices—the main margin for rent extraction and profit maximization. When prices are fixed, however, firms cannot increase profits by raising prices, making cost reductions through capacity and quality choices the primary margin for profit maximization.

The model rests on three key features. First, facilities receive a fixed payment per patient. Second, patients face switching costs. Third, variable costs exhibit economies of scale, so an additional patient is worth more to a larger facility.

Before the merger, switching costs make patients captive, so each facility provides its minimum quality. After the merger, the acquirer decides whether to close the target and a potential entrant decides whether to enter; active facilities then choose quality, and patients choose a facility. Closure forces the target's patients to switch, after which they bear a portion of the switching cost and become contestable. Where entry follows, the acquirer and the entrant compete for these patients on quality. The model delivers our key results: a net reduction in the number of facilities, a short-run health cost, and higher average quality after the merger. In this section, we describe the model's key intuition; all proofs are in Appendices G and H.

6.1 Setup

There are three potential facilities in a market: an acquirer A, a target T, and a potential entrant E. Before the merger, A and T operate, facility $i \in \{A, T\}$ serves $n_i^0 > 0$ patients, and the market has $N \equiv n_A^0 + n_T^0$ patients.

Costs. Facility i's total cost is $C_i(n_i, q_i) = F_i + c_i(n_i, q_i)\, n_i$, where $F_i > 0$ is a fixed cost and c_i is the variable cost per patient, with

$$\frac{\partial c_i}{\partial q_i} > 0 \quad \text{(quality is costly)}, \qquad \frac{\partial c_i}{\partial n_i} < 0 \quad \text{(scale economies)}. \tag{1}$$

The marginal cost of an additional patient, $\mathrm{MC}_i \equiv \partial C_i / \partial n_i = c_i + (\partial c_i / \partial n_i)\, n_i$, therefore lies strictly below the average variable cost c_i. We assume $p > \mathrm{MC}_i$ at all relevant volumes and qualities.

Prices and quality. Each patient yields a fixed reimbursement p. Facility i chooses quality $q_i \in [\underline{q}_i, \bar{q}]$, where the floor $\underline{q}_i$ is the lowest quality it can deliver given its equipment, staffing, and clinical protocols. Patient utility $u(q)$ is strictly increasing in quality.

Assumption 1 (Quality differentials within the merging pair)**.** *The merging pair consists of a high-quality acquirer and a low-quality target: the target's quality floor lies strictly below that of the acquirer and of a potential entrant; the acquirer is at least as cost-efficient per patient as the target at their pre-merger scales; and the acquirer could serve the whole market at no greater total cost than the target:*

$$\underline{q}_T \;<\; \underline{q}_A = \underline{q}_E \;\equiv\; \underline{q}, \qquad c_A(n_A^0, \underline{q}) \;\le\; c_T(n_T^0, \underline{q}_T), \qquad F_A + c_A(N, \underline{q})\, N \;\le\; F_T + c_T(N, \underline{q}_T)\, N. \tag{2}$$

Every market contains facilities of differing quality, whether or not it experiences a merger. Our analysis of merger determinants shows that markets with and without mergers are similar on average (Section 4.1). What Assumption 1 adds is that, within a market, a merger joins a higher-quality acquirer with a lower-quality target. The second and third conditions require that the acquirer, despite delivering higher quality, is not less efficient. Let $\Delta \equiv u(\bar{q}) - u(\underline{q}_T) > 0$ denote the largest utility gain attainable over the quality space.

Switching costs. A patient k who moves from facility i to facility j incurs

$$s_k^{ij} = b + \varepsilon_k^{ij}, \qquad \varepsilon_k^{ij} \sim F, \tag{3}$$

where $b > 0$ is a common cost of leaving an existing provider relationship—the delay in finding a new placement and the associated health risks, including missed or shortened sessions, and the costs of transferring records and learning how a new provider operates. A patient avoids b only by remaining at her current facility. The term ε_k^{ij} captures patient-facility specific fit, such as distance, scheduling, or insurance networks; it is drawn independently for each patient and destination from a continuous distribution F with support $[\underline{\varepsilon}, \bar{\varepsilon}]$. Patient k switches from an open facility i to j if and only if

$$u(q_j) - s_k^{ij} \ge u(q_i). \tag{4}$$

Assumption 2 (Treatment within the market)**.** *Leaving the market—forgoing in-center dialysis or seeking it elsewhere—yields utility V_0, with*

$$V_0 \;<\; u(\underline{q}_T) - b - \bar{\varepsilon}. \tag{5}$$

The right-hand side of (5) is the lowest utility a patient can obtain at any active facility—the lowest feasible quality, net of the switching cost and the worst possible fit. The assumption therefore requires that even a patient with the worst fit prefers some active facility to leaving the market. Fit determines *which* facility a patient chooses, not *whether* she is treated in the market. Because dialysis is life-sustaining and delivered locally, we view this as natural, and consistent with it, mergers leave total patient volume in the market largely unchanged (Figure 5).

6.2 Pre-Merger Equilibrium

Proposition 1 (Pre-merger equilibrium)**.** *If inertia dominates the largest attainable quality advantage,*

$$b + \underline{\varepsilon} \;>\; \Delta, \tag{6}$$

then no patient switches and each facility sets quality at its own floor, $q_A^0 = \underline{q}$ and $q_T^0 = \underline{q}_T$. In particular, the target provides strictly lower quality than the acquirer.

When inertia is large, no feasible quality gap can overcome the switching cost, so patients are captive. Because quality cannot attract patients but raises costs, each facility chooses the minimum quality it can deliver, and captivity sustains a low-quality target alongside a higher-quality acquirer.

6.3 The Post-Merger Game

We take the merger as given and model the post-merger market as a three-stage game.[60]

- **Stage 1 (Network).** A chooses to close T ($c = 1$) or keep it open ($c = 0$). Simultaneously, E chooses to enter ($e = 1$) or stay out ($e = 0$).
- **Stage 2 (Quality).** Observing (c, e), every active facility simultaneously chooses its quality.

[60] The model does not take a stand on where mergers occur. Our design therefore does not require random mergers, but that treated markets would have evolved like control markets absent the merger. Within-market differences in the level of quality do not threaten our identification.

- **Stage 3 (Demand).** Patients observe qualities and their idiosyncratic fit and choose a facility.

Let $\pi_A^*(c,e)$ denote the operating profit (revenue net of variable cost) of the acquirer's network—A and, if open, T—at the Stage-2 equilibrium qualities, and $\pi_E^*(c)$ the entrant's operating profit conditional on entering. The acquirer pays F_A, and also F_T if T remains open; the entrant pays $F_E > 0$ if it enters. When indifferent, the entrant enters and the acquirer keeps T open.

6.4 Equilibrium

We solve by backward induction. Four network structures $(c,e) \in \{0,1\}^2$ are possible; we begin with the focal one, closure with entry.

6.4.1 Quality Competition After Closure

Assumption 3 (Forced switching of target patients)**.** *Closing T removes the option of remaining at T. Each displaced patient k must switch, and bears the switching cost at whichever active facility m she moves to: $s_k^{Tm} = b + \varepsilon_k^{Tm}$.*

All patients face the same switching cost; what differs after closure is the alternative they compare it against. A patient at A can stay at no cost, so she moves to E only if E's quality advantage exceeds $b + \varepsilon_k^{AE}$. Since $b + \varepsilon_k^{AE} \geq b + \underline{\varepsilon} > \Delta$, no feasible quality gap is large enough, and she stays. A patient at T cannot stay. She must choose between A and E and pays b at either, so she prefers A if and only if

$$u(q_A) - b - \varepsilon_k^{TA} \geq u(q_E) - b - \varepsilon_k^{TE} \iff D_k \equiv \varepsilon_k^{TA} - \varepsilon_k^{TE} \leq u(q_A) - u(q_E). \tag{7}$$

Because b is paid at either destination, it cancels. b still affects her welfare, which is the source of the short-run mortality increase, but no longer affects her choice. Closure thus makes the target's patients contestable, while the acquirer's patients remain captive. With G and g the CDF and density of D_k, the acquirer captures a share $r(q_A,q_E) = G\big(u(q_A) - u(q_E)\big)$ of the displaced pool, and volumes are

$$n_A^C = n_A^0 + r(q_A,q_E)n_T^0, \qquad n_E^C = \big[1 - r(q_A,q_E)\big]n_T^0. \tag{8}$$

The acquirer chooses q_A to maximize $\Pi_A^C = \big[p - c_A(n_A^C, q_A)\big]\, n_A^C$. Its first-order condition is (Appendix G)

$$\underbrace{\big(p - \mathrm{MC}_A\big)\, n_T^0\, g\big(u(q_A) - u(q_E)\big)\, u'(q_A)}_{\text{margin on contested patients won}} = \underbrace{\frac{\partial c_A}{\partial q_A}\, n_A^C}_{\text{cost of quality across all patients}}. \tag{9}$$

Higher quality wins displaced patients, each worth the margin $p - \mathrm{MC}_A$, but raises the cost of treating every patient at A. Scale economies enter here: because $\mathrm{MC}_A < c_A$, each patient the acquirer absorbs is worth more than its average margin, which strengthens the incentive to compete for the contestable pool—and, at the network stage, raises the value of closing T when the acquirer expects to absorb many of its patients.

The entrant solves the analogous problem on n_E^C, with first-order condition

$$\big(p - \mathrm{MC}_E\big)\, n_T^0\, g\big(u(q_A) - u(q_E)\big)\, u'(q_E) \;=\; \frac{\partial c_E}{\partial q_E}\, n_E^C. \tag{10}$$

A Stage-2 Nash equilibrium (q_A^*, q_E^*) solves (9)–(10); we assume it exists and is unique.[61]

Lemma 1 (Strict quality upgrades)**.** *Suppose that at the floor each firm's margin on contested patients exceeds its cost of raising quality across its patient pool:*

$$\begin{aligned} \big(p - \mathrm{MC}_A\big)\, n_T^0\, g(0)\, u'(\underline{q}) \;&>\; \tfrac{\partial c_A}{\partial q_A}\big|_{\underline{q}}\, \big[n_A^0 + G(0)\, n_T^0\big], \\ \big(p - \mathrm{MC}_E\big)\, n_T^0\, g(0)\, u'(\underline{q}) \;&>\; \tfrac{\partial c_E}{\partial q_E}\big|_{\underline{q}}\, \big[1 - G(0)\big]\, n_T^0. \end{aligned} \tag{11}$$

Then both firms choose quality strictly above the floor: $q_A^*(1,1) > \underline{q}$ and $q_E^*(1,1) > \underline{q}$.

Closure creates a contestable pool, and because marginal patients are profitable, each firm pays for quality to win them. This margin is absent in the captive pre-merger market.

6.4.2 The Other Subgames

Lemma 2 (No entry without closure)**.** *If the acquirer keeps T open ($c = 0$), the entrant attracts no patients for any qualities, so $\pi_E^*(0) = 0$ and E does not enter.*

The fourth network structure, entry without closure, therefore never arises. While T is open, every patient in the market is at an open facility and can avoid b by staying. To attract any patient, the entrant would have to offer a utility gain larger than $b + \varepsilon_k \geq b + \underline{\varepsilon}$, but by condition (6) no feasible quality gap exceeds $\Delta < b + \underline{\varepsilon}$. The entrant would serve no patients and lose $F_E > 0$. Closure is the only event that creates demand an entrant can win, so by keeping T open the acquirer always deters entry.

Lemma 3 (Closure without entry)**.** *If T closes and E stays out, the acquirer serves all N patients and sets $q_A^*(1,0) = \underline{q}$.*

With no entrant, displaced patients have no alternative to A (Assumption 2), so A's volume does not respond to quality and A stays at its floor. If T remains open and E stays out, the pre-merger qualities $(\underline{q}, \underline{q}_T)$ persist.

[61] Existence follows from continuity on the compact set $[\underline{q}, \bar{q}]^2$ and quasi-concavity of each firm's objective, which we assume via standard regularity conditions on c_i, u, and G. Uniqueness ensures that $\pi_A^*(1,1)$ and $\pi_E^*(1)$ are well defined.

6.4.3 The Network Stage

Stage-1 payoffs, with operating profits evaluated at the Stage-2 equilibrium qualities, are:

Stage 1 net payoffs (Acquirer, Entrant)

		Entrant E	
		Enter ($e = 1$)	Stay out ($e = 0$)
Acquirer A	Close ($c = 1$)	$\pi_A^*(1,1) - F_A, \quad \pi_E^*(1) - F_E$	$\pi_A^*(1,0) - F_A, \quad 0$
	Keep open ($c = 0$)	$\pi_A^*(0,1) - F_A - F_T, \quad -F_E$	$\pi_A^*(0,0) - F_A - F_T, \quad 0$

Let $\overline{F}_E \equiv \pi_E^*(1) > 0$ denote the entrant's operating profit after closure, the highest entry cost it will pay.

Proposition 2 (Post-merger equilibrium)**.** *Under condition* (6)*, Assumptions 1–3, and the conditions of Lemma 1, the unique subgame-perfect equilibrium is*

$$(c^*, e^*) = \begin{cases} (1,1) & \text{if } F_E \le \overline{F}_E \text{ and } F_T > \pi_A^*(0,0) - \pi_A^*(1,1), \\ (0,0) & \text{if } F_E \le \overline{F}_E \text{ and } F_T \le \pi_A^*(0,0) - \pi_A^*(1,1), \\ (1,0) & \text{if } F_E > \overline{F}_E. \end{cases} \tag{12}$$

Quality rises strictly above the floor in $(1,1)$*; in* $(1,0)$ *and* $(0,0)$ *every active facility remains at its floor.*

Closure is deterred only by the threat of entry. When entry is blocked ($F_E > \overline{F}_E$), closing T always pays: it saves F_T, and because displaced patients have no alternative, the acquirer treats all of them at its own, lower cost without competing on quality. When entry would follow closure ($F_E \le \overline{F}_E$), keeping T open preserves its operating profit and deters entry, so the acquirer closes only if the saved fixed cost exceeds the operating profit it forgoes, $\pi_A^*(0,0) - \pi_A^*(1,1)$. This forgone profit is smaller the more displaced patients the acquirer expects to absorb and the lower the cost at which it treats them—the scale economies in its first-order condition (9). Closure with entry thus requires intermediate entry costs—low enough that an entrant would enter once the target's patients become contestable, but positive, so it would never enter against an open target—together with a target whose fixed cost is large relative to the profit forgone by closing it.

6.5 Empirical Footprints

The model rationalizes the two central patterns in our data: net capacity falls despite entry, and quality rises even though entry does not fully offset closures.

Footprint 1: Closure, entry, and quality gains. Where entry costs are intermediate, the equilibrium features closure of the target, entry by a new facility, and quality above the floor at both the acquirer and the entrant (Proposition 2 and Lemma 1). Closure makes the target's patients contestable, and because price is fixed, the acquirer and the entrant compete for them on quality. Because every displaced patient must switch and bear b, closure also imposes a transitional health cost. Our estimates match each element of this equilibrium: acquired facilities close, non-acquired facilities enter (Figure 3), mortality rises sharply in the merger year (Figure 10), and hospitalizations per patient decline in the long run (Figure 13).

Footprint 2: Lower capacity and higher quality. Entry costs differ across markets. Where they are high, entry is blocked, but the acquirer still closes the target: with no rival to compete for the target's patients, it absorbs all of them at its own, lower cost. In these markets, capacity falls by one facility. Every entry is therefore matched by a closure, but not every closure is matched by an entry, so capacity falls on average (Appendix Proposition 3). Empirically, five years after a merger there are 0.4 fewer acquired facilities and 0.18 more non-acquired facilities per market, so roughly 45% of closures are followed by entry, and total capacity falls by 0.15 facilities.

Quality nonetheless rises, even in markets where no entry follows. There, the target's patients move from a facility providing $\underline{q}_T$ to the acquirer, which provides $\underline{q} > \underline{q}_T$: the weakest facility exits and its patients receive better care. Where entry does follow, quality rises further, because the acquirer upgrades and the entrant provides higher quality than the facility it replaces. Quality therefore rises in every market in which the target closes (Appendix Corollary 1), so on average the model predicts lower capacity together with higher quality—the pattern we observe.

Table 1: Post-merger equilibria: capacity and quality

Equilibrium	**Δ Facilities**	**Quality**	**Source of quality gain**
Closure + entry $(1,1)$	0	$\uparrow$	exit of target, acquirer upgrade, higher-quality entrant
Closure, no entry $(1,0)$	-1	$\uparrow$	exit of target
Status quo $(0,0)$	0	unchanged	—

The model thus identifies two sources of quality gains: *selection*, through the exit of low-quality targets and the entry of higher-quality facilities, and *within-firm improvements* at surviving facilities. We assess both in Figure 19. Restricting the sample to facilities that

survive five years after the merger isolates within-firm improvements; adding facilities that exit should attenuate the decline in hospitalizations, and adding entrants should strengthen it. Panel A shows a decline in hospitalizations among survivors, panel B shows that adding exiting facilities attenuates it, and panel C shows that adding entrants modestly strengthens it. Quality gains therefore arise both from selection and from improvements among surviving facilities.

Footprint 3: Market sparsity. Closures should be larger in markets with fewer facilities, where the acquirer is the natural destination for the target's patients and so recaptures more of them after closure, making closure less costly (Appendix Proposition 4). Consistent with this, markets with fewer than three facilities before the merger experience larger closures in levels and in percentage terms: five years after the merger, they have 0.42 fewer facilities relative to a baseline of 0.945, compared with 0.21 fewer relative to a baseline of 1.21 in larger markets (Appendix Figure 11, panel A).

6.6 Scope of the Model

Three features of the post-merger market lie outside the model. First, the model is static: quality moves to its new level immediately upon closure, whereas hospitalizations decline progressively over five years. We interpret q^* as the long-run level toward which facilities converge as investments in staffing, protocols, and equipment accumulate. Second, the model abstracts from the transitional frictions that accompany closure—reduced sessions, staffing disruptions, congestion, and longer travel—and reallocates displaced patients instantly; the switching cost b is a reduced-form summary of these frictions. Third, the model does not consider what happens once displaced patients settle in. As they build relationships with their new facility, inertia may re-form and make them captive again, which in a dynamic setting could lead facilities to revert toward their quality floors once the contestable pool has been absorbed. We do not model this possibility; empirically, hospitalizations continue to decline through the fifth year after the merger, so any such reversion does not occur within the horizon we observe.

7 Cost-Benefit Analysis

Drawing on standard frameworks for monetizing health effects, we weigh the *health-related* benefits of mergers, reflected in long-run reductions in hospitalizations, against short-run mortality costs over the five years following the merger. We emphasize that this exercise

is intentionally limited to patient health outcomes and does not constitute a full welfare analysis. In particular, we do not directly value changes in patients' travel and time costs, treatment convenience, or other dimensions of access and quality of life. Some of these factors may affect health outcomes and are therefore indirectly reflected in our estimates, but their direct welfare consequences are not captured. The goal is to answer a narrower, evidence-based question—on balance, do mergers improve or worsen patient health over the span of five years following the merger? In doing so, we provide a framework for evaluating consumer welfare on non-price margins, for which conventional antitrust tools for assessing the effects of consolidation are less directly applicable. Table 5 presents the estimates from our analysis, and a description is provided below.

We evaluate the health-related costs and benefits of mergers over a five-year period. This horizon provides a natural endpoint for both our event-study analysis and our calculation of costs and benefits. In the context of ESRD, patient survival is often measured over five years following the onset of kidney failure or initiation of dialysis, after which survival typically declines by half (NIH 2024). Conceptually, this suggests that, for an average patient who begins dialysis shortly before a merger, the most consequential health effects are likely to occur within five years.

We begin by estimating the number of dialysis patients exposed to mergers. Our sample contains 724 markets with mergers (Table 2), representing slightly more than 23% of the 3,143 U.S. counties. Applying this share to the total population of 474,468 patients receiving in-center hemodialysis in 2021 (NIH 2024) yields an estimated 109,295 patients plausibly exposed to mergers.[62]

To estimate merger benefits through long-run reductions in hospitalizations, we utilize the cumulative of coefficients over the 5-year post-merger horizon from the panel A of Figure 13, which shows that mergers reduce hospitalizations by nearly 0.248 per patient over 5 years. To translate this effect into monetary terms, we use Medicare Part A spending data from the Kaiser Family Foundation (KFF). From KFF (Kaiser Family Foundation 2021*a*), we obtain the 2021 figures for total traditional Medicare Part A enrollees (35,337,335) and program payments per enrollee ($5,292), yielding aggregate Part A payments of $187,216 million. Because our hospitalization outcome data reflect Medicare Part A claims for traditional Medicare, these populations are directly comparable. We use 2021 data to reflect the most up-to-date value of benefits, although our regression estimates reflect effects of mergers spanning 1991–2021.

[62] We obtain a similar estimate by directly counting patients in exposed markets in our stacked sample, treating patients in different, time-separated merger cohorts as distinct observations—104,946 patients.

Program payments per enrollee do not directly capture spending per hospitalization, since each enrollee may experience multiple hospitalizations. To obtain this measure, we draw on a separate KFF report (Kaiser Family Foundation 2021*b*) that provides the total number of Part A discharges. Dividing aggregate payments by total discharges yields an estimated spending of $25,433 per hospitalization. Importantly, both the KFF and our own estimates are calculated over the full population of enrollees, including those with no hospitalizations, ensuring conceptual consistency.

Multiplying the per-patient reduction in hospitalizations (0.248) by the number of exposed patients (109,295) and the spending per hospitalization ($25,433) gives an aggregate spending reduction of $688.22 million, which we interpret as the monetary value of the health benefits of mergers. If we assume that the estimated hospitalization savings are distributed evenly across the six-year post-merger period, the resulting annual savings of roughly $115 million correspond to a little over 1% of annual Medicare inpatient spending of around $11 billion.[63,64]

Our estimates of costs comes from the short-term increase in mortality from the coefficient at event-time zero in Figure 10, which indicates a 0.65 percentage point rise in mortality following mergers. Applied to the merger-affected population, this implies roughly 710 additional deaths.

We translate merger-related deaths into years of life lost (YLL) using life expectancy for the dialysis population. The 2021 USRDS estimates (NIH 2023) report an average remaining life expectancy of 4.3 years for dialysis patients aged 65–69, the age group encompassing the average age at death in our sample (68.4 years). Applying this estimate to the 710 merger-related deaths yields 3,052 years of life lost.

We then monetize these life-years using the value of a statistical life year (VSLY). We derive a central VSLY of $278,000 by dividing the central value of a statistical life (VSL) of $11.4 million reported by U.S. Department of Health and Human Services (2021) by average U.S. life expectancy. The resulting value of the estimated 3,052 life-years lost is $848.5 million. Equivalently, this implies a cost of approximately $1.20 million per merger-related death.

Our baseline analysis therefore shows that, under the status quo, the health costs of mergers moderately outweigh their benefits, yielding an estimated negative net value (-$160.24 million). Importantly, however, this imbalance is sufficiently narrow that a modest

[63] In 2021, the annual nominal spending on inpatient care (Medicare Parts A and B) for Medicare patients was approximately $11 billion (NIH 2024). See NIH (2024) Figure 9.5.

[64] The magnitude of savings as a percentage of Medicare spending would be smaller the relevant denominator were *total*, rather than inpatient, Medicare spending on ESRD patients, which was about $33 billion in 2021 (NIH 2024). See NIH (2024) Figure 9.2.

reduction in merger-related mortality would reverse it. A policy intervention that averts at least 135 of the estimated merger-related deaths would reduce mortality costs to $687.08 million, bringing the total health costs below the estimated benefits and making the merger net-beneficial even under the status quo.[65] Such an intervention would therefore need to avert only about 19% of merger-related deaths to tip the cost-benefit balance in favor of mergers. We discuss potential policy interventions capable of achieving such a reduction in Section 8. In Panel B of Appendix I, we further show that it would take just one additional year for benefits to exceed costs under the status quo, whether we assume that the reduction in hospitalizations observed five years post-merger persists into year 6 or linearly extrapolate the estimated coefficients to year 6.

In Panel A of Appendix I, we assess the sensitivity of our cost-benefit analysis to alternative assumptions, varying one aspect of our baseline methodology at a time. First, we account for the fact that our estimates of hospitalization spending are based on Medicare fee-for-service data, whereas our exposure estimates encompass all dialysis patients. We therefore provide an upper bound on benefits by assuming that all patients are privately insured and would generate the larger hospitalization savings due to higher private reimbursement rates. Second, rather than using a central VSLY to monetize mortality costs, we consider a low VSLY of $130,000 and, separately, a high VSLY of $423,000. Third, rather than using the life expectancy of dialysis patients to calculate YLL, we use the life expectancy—16.3 years—of a typical U.S. adult at the average age of death in our sample. The VSLY and life-expectancy assumptions affect costs only. Fourth, we apply a 3% discount rate to both costs and benefits. Fifth, we use the market-level estimates from Figure A.10 to calculate merger-induced reductions in hospitalization and mortality. The latter two modifications affect both costs and benefits.

Across these six variants, three yield a positive net value of mergers, while three continue to yield a negative net value. The assumption about life expectancy is the most consequential source of variation relative to our baseline estimate—using the life expectancy of the general U.S. population substantially increases the monetized cost of mortality. By contrast, discounting, as well as switching from facility-level to market-level estimates produce small changes in net value relative to the baseline. Still, the most realistic variants preserve the relative magnitudes of benefits and costs in a range where even modest reductions in mortality from policy interventions would generate positive net value.

[65] With 134 deaths averted, mortality costs would remain at $688.28 million and therefore continue to exceed the estimated benefits.

8 Antitrust Policy

Our results yield two sets of implications for antitrust regulation. The first concerns how authorities should evaluate proposed mergers. In light of the 2023 Merger Guidelines, we highlight that mergers may generate substantial benefits and efficiencies that, if not appropriately considered, could lead to a mischaracterization of their overall effects. We also provide new insights for antitrust analysis in settings characterized by non-price competition, showing that mergers can simultaneously reduce capacity and improve quality, and we offer a simple methodology for quantifying consumer benefits and harms along non-price margins (Section 7). Finally, our results suggest de-emphasizing blunt measures such as market concentration: in our setting, changes in concentration play little role in explaining either the harms or the benefits of mergers.

The second set of implications concerns policy tools that mitigate merger-related harms when mergers proceed. Under the current policy environment, the estimated costs of mergers moderately exceed their estimated benefits (Section 7). Yet the mortality increase we document is concentrated in the merger year rather than reflecting a persistent decline in care quality. Through the lens of our model, it arises from the switching cost b that patients bear when their facility closes. This raises a natural question: can a regulator reduce this cost without preventing closure, and without sacrificing the long-run quality gains that closure and entry generate?

The model separates the two components of switching costs. The common component b sustains pre-merger captivity (condition (6)) and is the cost displaced patients bear at closure, generating the mortality spike in Figure 10. But b does not affect the quality competition that follows closure: displaced patients pay it at whichever facility they choose, so their choice between the acquirer and the entrant—and hence firms' incentives to compete on quality—is governed by quality and idiosyncratic fit ε_k^{ij} alone. Consequently, a policy that lowers b for patients whose facility closes reduces the mortality cost of closure without changing quality, closure, or entry (Appendix Proposition 5). Two conditions are essential for this neutrality. First, the policy must be local in time and across patients—for example, expedited medical-record transfer or transition case management triggered by a closure notice—so that b remains at its baseline level before the merger and for patients at facilities that stay open. Second, displaced patients must continue to be treated in the market, so that the policy changes how costly their transition is, not whether they remain in the contestable pool.

Neutrality also depends on who bears the cost of transition support. If the cost falls on the acquirer and is incurred only when it closes a facility—for example, a requirement

to fund case management for displaced patients—it acts as a tax on closure: it reduces the net savings from closing the target and can deter closure in markets where those savings are small, forgoing the associated quality gains. Neutrality is preserved when the cost is borne by a third party, such as CMS, or when the acquirer's obligation does not depend on whether it closes a facility, such as a commitment fixed at the time of merger approval. A formal characterization of how the incidence of such policies affects closure, entry, and quality is beyond the scope of this paper, and we leave it to future work.

Three other policies that target switching costs need not preserve the quality gains. First, an economy-wide reduction in b, in place before any merger, weakens pre-merger captivity: if patients become responsive to quality differentials, facilities compete on quality even without a merger, so baseline quality is higher and the within-market quality gap that closure eliminates is smaller. Such a policy lowers the mortality cost of closure but also shrinks the merger's long-run quality gains, so its net welfare effect is ambiguous. Second, a post-merger reduction in b for all patients in the market makes the acquirer's own patients responsive to the entrant's quality. This changes the acquirer's quality problem, the profit it forgoes by closing the target, and the entry threshold; it may strengthen incentives to compete on quality, but it can also shrink the set of markets in which an equilibrium with closure and quality gains arises. Third, a policy that assigns a share of displaced patients to a designated facility regardless of quality severs the link between quality and demand for those patients. As this share grows, the acquirer's incentive to upgrade disappears and entry is deterred, eliminating the mechanism that generates the long-run quality gains (Appendix Proposition 6).

The policy implication is therefore clear. Policies that reduce the common cost of switching for patients whose facility closes can mitigate the mortality cost of closure without affecting closure, entry, or equilibrium quality, provided they are narrowly targeted to the closure event. Existing antitrust and CMS tools offer several approaches consistent with this principle. Pre-merger notification could require disclosure of the capacity and location of nearby non-merging facilities, helping ensure that displaced patients retain a meaningful choice set. Merger approval could also be conditioned on enforceable commitments—such as expedited medical-record transfer, minimum advance notice of closure, or transition case-management support—that lower the common component of switching costs without steering patients toward a particular facility or altering switching frictions economy-wide. To avoid discouraging efficient closures, such commitments are best funded publicly or fixed at the time of approval rather than levied on the acquirer only when it closes a facility.

9 Conclusion

In markets where prices are heavily regulated, the effects of firm consolidation may manifest primarily through changes in quality and capacity. Our analysis of the U.S. outpatient kidney dialysis industry reveals that mergers in such environments generate a stark intertemporal trade-off. In the short run, acquiring firms close target facilities and consolidate operations, disrupting patients' treatment and generating a temporary but substantial increase in mortality. We estimate the resulting short-run health costs at $848 million. Over the subsequent five years, however, these effects reverse—hospitalizations decline, generating $688 million in long-run health benefits. We develop a model that links these two effects. Before a merger, high switching costs create a captive patient base for each facility, weakening firms' incentives to invest in quality. With prices fixed, an acquiring firm can increase profitability after a merger by closing a redundant facility and reducing fixed costs. Closure forces displaced patients to bear a large share of the switching costs, but also renders some of them contestable. This newly contestable patient pool creates opportunities for entry and strengthens firms' incentives to compete on quality, leading to quality improvements following the merger.

We find that the health effects of mergers are stable across Medicare payment regimes. We also find that mergers increase market concentration, but changes in concentration do not systematically explain their effects on facility closures or patient health.

These findings suggest new foundations for merger policy in price-regulated markets. Because the long-run quality efficiencies associated with consolidation are substantial, relying solely on structural presumptions—such as lower concentration thresholds—may inadvertently prevent mergers that generate meaningful improvements in patient welfare. Instead, welfare may be better protected through targeted conduct remedies. Our model shows that antitrust stipulations that facilitate patient transitions after closures can mitigate the short-run harms associated with displacement while preserving the long-run benefits of scale and efficiency.

References

Aghion, Philippe, and Peter Howitt. 1992. "A model of growth through creative destruction." *Econometrica* 60(2): 323–351.

Ahmad, Farida B., Jodi A. Cisewski, and Robert N. Anderson. 2024. "Mortality in the United States — Provisional Data, 2023." MMWR Morb Mortal Wkly Rep 2024;73:678–681. https://www.cdc.gov/mmwr/volumes/73/wr/pdfs/mm7331a1-H.pdf Centers for Disease Control and Prevention.

American Health Planning Association. 2026. "Archives: Certificate of Need Matrix of Service Coverage*." https://www.ahpanet.org/matrix_copn.html.

American Hospital Association. 2023. "Fast Facts on U.S. Hospitals, 2023.".

American Renal Associates Holdings, Inc. 2020. "American Renal Associates Enters Into Definitive Agreement to Be Acquired by Nautic Partners (Press Release)." Business Wire. https://www.businesswire.com/news/home/20201002005131/en/American-Renal-Associates-Enters-Into-Definitive-Agreement-to-be-Acquired-by-Nautic-Partners.

Arnold, Carrie. 2020. "Should Your Kidney Doctor Have a Financial Stake in Dialysis?" Undark.

Ashenfelter, Orley C, Daniel S Hosken, and Matthew C Weinberg. 2015. "Efficiencies brewed: pricing and consolidation in the US beer industry." *The RAND Journal of Economics* 46(2): 328–361.

Athey, Susan and Nevo, Aviv. 2026. "DOJ and FTC Chief Economists Explain the Changes to the 2023 Merger Guidelines." https://www.promarket.org/2023/12/19/doj-and-ftc-chief-economists-explain-the-changes-to-the-2023-merger-guidelines/.

Atkinson Jr, Arthur J, Wayne A Colburn, Victor G DeGruttola, David L DeMets, Gregory J Downing, Daniel F Hoth, John A Oates, Carl C Peck, Robert T Schooley et al. 2001. "Biomarkers and surrogate endpoints: preferred definitions and conceptual framework." *Clinical pharmacology & therapeutics* 69(3): 89–95.

Beaulieu, Nancy D, Leemore S Dafny, Bruce E Landon, Jesse B Dalton, Ifedayo Kuye, and J Michael McWilliams. 2020. "Changes in quality of care after hospital mergers and acquisitions." *New England Journal of Medicine* 382(1): 51–59.

Béguin, Lisa, Thierry Krummel, Nathalie Longlune, Roula Galland, Cécile Couchoud, and Thierry Hannedouche. 2021. "Dialysis dose and mortality in haemodialysis: is higher better?" *Nephrology Dialysis Transplantation* 36(12): 2300–2307.

Bertuzzi, Luca, Paul J Eliason, Benjamin Heebsh, Riley J League, Ryan C McDevitt, and James W Roberts. 2023. Gaming and effort in performance pay. Technical report National Bureau of Economic Research.

Bhattacharya, Vivek, Gastón Illanes, and David Stillerman. 2023. Merger effects and antitrust enforcement: Evidence from us consumer packaged goods. Technical report National Bureau of Economic Research.

Bhuller, Manudeep, Tarjei Havnes, Edwin Leuven, and Magne Mogstad. 2013. "Broadband internet: An information superhighway to sex crime?" *Review of Economic studies* 80(4): 1237–1266.

Bogataj, Špela, Jernej Pajek, Blaž Slonjšak, and Vanja Peršič. 2023. "Prevalence of impaired physical mobility in dialysis patients: a single-centre cross-sectional study." *Journal of Clinical Medicine* 12(20): 6634.

Bolema, Ted. 2026. "Decoding the 2023 FTC and DOJ Merger Guidelines: Insights into Shifting Antitrust Enforcement." `https://www.mercatus.org/research/policy-briefs/decoding-2023-ftc-and-doj-merger-guidelines-insights-shifting-antitrust`.

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

Bougette, Patrice, Kai Hüschelrath, and Kathrin Müller. 2014. "Do horizontal mergers induce entry? Evidence from the US airline industry." *Applied Economics Letters* 21(1): 31–34.

Brunelli, Steven M, Glenn M Chertow, Elizabeth D Ankers, Edmund G Lowrie, and Ravi Thadhani. 2010. "Shorter dialysis times are associated with higher mortality among incident hemodialysis patients." *Kidney international* 77(7): 630–636.

Calzada, Joan, Xavier Fageda, and Fernando Martínez-Santos. 2023. "Mergers and bank branches: two decades of evidence from the USA." *Empirical Economics* 64(5): 2411–2447.

Caradonna, Peter, Nathan H Miller, and Gloria Sheu. 2025. "Mergers, entry, and consumer welfare." *American Economic Journal: Microeconomics* 17(3): 103–130.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The effect of minimum wages on low-wage jobs." *The Quarterly Journal of Economics* 134(3): 1405–1454.

Centers for Disease Control. 2024. "Anemia Prevalence: United States, August 2021–August 2023.".

Childers, Christopher P, Jill Q Dworsky, Gerald Kominski, and Melinda Maggard-Gibbons. 2019. "A comparison of payments to a for-profit dialysis firm from government and commercial insurers." *JAMA internal medicine* 179(8): 1136–1138.

Cleveland Clinic. 2022. "Hypercalcemia.".

Clogg, Clifford C, Eva Petkova, and Adamantios Haritou. 1995. "Statistical methods for comparing regression coefficients between models." *American journal of sociology* 100(5): 1261–1293.

CMS. 2023. "Medicare Coverage of Kidney Dialysis & Kidney Transplant Services.".

CMS. 2024. "Calendar Year 2025 End-Stage Renal Disease (ESRD) Prospective Payment System (PPS) Final Rule (CMS-1805-F).".

Cutler, David, Leemore Dafny, and Christopher Ody. N.d. "How does competition impact the quality and price of outpatient service facilities? A case study of the U.S. dialysis industry." *Working Paper*: Forthcoming.

Dafny, Leemore, Mark Duggan, and Subramaniam Ramanarayanan. 2012. "Paying a premium on your premium? Consolidation in the US health insurance industry." *American Economic Review* 102(2): 1161–1185.

Dafny, Leemore S, and Thomas H Lee. 2015. "The good merger." *New England Journal of Medicine* 372(22): 2077–2079.

Dartmouth Atlas of Health Care. 2024. "FAQ." `https://www.dartmouthatlas.org/faq/`.

DaVita HealthCare Partners Inc. 2015. "DaVita HealthCare Partners to Acquire Renal Ventures (Press Release)." PR Newswire.`https://www.prnewswire.com/news-releases/davita-healthcare-partners-to-acquire-renal-ventures-300132059.html`.

DaVita Inc. 2005. "DaVita Closes Gambro Healthcare Acquisition (Press Release)." Filed as Exhibit 99.3 to DaVita Inc. Form 8-K, U.S. Securities and Exchange Commission. `https://www.sec.gov/Archives/edgar/data/0000927066/000119312505199639/dex993.htm`.

DaVita Inc. 2011. "DaVita to Acquire DSI Renal, Inc. (Press Release)." Filed as Exhibit 99.1 to DaVita Inc. Form 8-K, U.S. Securities and Exchange Commission. `https://www.sec.gov/Archives/edgar/data/0000927066/000119312511030748/dex991.htm`.

Demirer, Mert, and Ömer Karaduman. 2024. Do mergers and acquisitions improve efficiency? Evidence from power plants. Technical report National Bureau of Economic Research.

Dobkin, Carlos, Amy Finkelstein, Raymond Kluender, and Matthew J Notowidigdo. 2018. "The economic consequences of hospital admissions." *American Economic Review* 108(2): 308–352.

Eliason, P, Benjamin Heebsh, R League, R McDevitt, and James Roberts. 2022. "The effect of bundled payments on provider behavior and patient outcomes: evidence from the dialysis industry." *Working Paper*: Forthcoming.

Eliason, Paul J, Benjamin Heebsh, Ryan C McDevitt, and James W Roberts. 2020. "How acquisitions affect firm behavior and performance: Evidence from the dialysis industry." *The Quarterly Journal of Economics* 135(1): 221–267.

Erickson, Kevin F, Wolfgang C Winkelmayer, Vivian Ho, Jay Bhattacharya, and Glenn M Chertow. 2019. "Market consolidation and mortality in patients initiating hemodialysis." *Value in Health* 22(1): 69–76.

Erickson, Kevin F, Yuanchao Zheng, Vivian Ho, Wolfgang C Winkelmayer, Jay Bhattacharya, and Glenn M Chertow. 2018. "Market competition and health outcomes in hemodialysis." *Health services research* 53(5): 3680–3703.

Ericson, Keith M Marzilli. 2014. "Consumer inertia and firm pricing in the Medicare Part D prescription drug insurance exchange." *American Economic Journal: Economic Policy* 6(1): 38–64.

Federal Trade Commission. 2023. "Premerger Notification and the Merger Review Process.".

Fee, C Edward, and Shawn Thomas. 2004. "Sources of gains in horizontal mergers: evidence from customer, supplier, and rival firms." *Journal of financial Economics* 74(3): 423–460.

Ferrier, Gary D, and Vivian G Valdmanis. 2004. "Do mergers improve hospital productivity?" *Journal of the Operational Research Society* 55(10): 1071–1080.

Fishbane, Steven, Edward A Kowalski, Louis J Imbriano, and John K Maesaka. 1996. "The evaluation of iron status in hemodialysis patients." *Journal of the American Society of Nephrology* 7(12): 2654–2657.

Flythe, Jennifer E, Stephen E Kimmel, and Steven M Brunelli. 2011. "Rapid fluid removal during dialysis is associated with cardiovascular morbidity and mortality." *Kidney international* 79(2): 250–257.

Fresenius Medical Care. N.d. "How the "Bundle" Has Changed Clinical Care at Fresenius Kidney Care.".

Fresenius Medical Care AG. 2005. "Fresenius Medical Care to Acquire Renal Care Group, Inc. (Press Release)." `https://www.fresenius.com/node/4677`.

Fresenius Medical Care AG & Co. KGaA. 2006. "Fresenius Medical Care Announces Closing of the Renal Care Group Acquisition (Press Release)." `https://www.fresenius.com/node/4724`.

Fresenius Medical Care AG & Co. KGaA. 2022*a*. "Fresenius Medical Care Closes Three-Way Merger and Creates Premier Value-Based Kidney Care Provider in the U.S. (Press Release)." Filed as Exhibit 99.1 to Fresenius Medical Care AG & Co. KGaA Form 6-K, U.S. Securities and Exchange Commission. `https://www.sec.gov/Archives/edgar/data/1333141/000110465922093926/tm2224458d1_ex99-1.htm`.

Fresenius Medical Care AG & Co. KGaA. 2022*b*. "Fresenius Medical Care Participates in Three-Way Merger to Create Premier Value-Based Kidney Care Provider in the U.S. (Press Release)." Filed as Exhibit 99.1 to Fresenius Medical Care AG & Co. KGaA Form 6-K, U.S. Securities and Exchange Commission.`https://www.sec.gov/Archives/edgar/data/1333141/000110465922036439/tm229941d1_ex99-1.htm`.

Freyaldenhoven, Simon, Christian Hansen, and Jesse M Shapiro. 2019. "Pre-event trends in the panel event-study design." *American Economic Review* 109(9): 3307–3338.

Goodman-Bacon, Andrew. 2018. "Public insurance and mortality: evidence from Medicaid implementation." *Journal of Political Economy* 126(1): 216–262.

Grieco, Paul LE, and Ryan C McDevitt. 2017. "Productivity and quality in health care: Evidence from the dialysis industry." *The Review of Economic Studies* 84(3): 1071–1105.

Hakim, Raymond M, Julia Breyer, Nuhad Ismail, and Gerald Schulman. 1994. "Effects of dose of dialysis on morbidity and mortality." *American Journal of Kidney Diseases* 23(5): 661–669.

Handel, Benjamin R. 2013. "Adverse selection and inertia in health insurance markets: When nudging hurts." *American Economic Review* 103(7): 2643–2682.

Hanner, Daniel, Daniel Hosken, Luke M Olson, and Loren K Smith. 2015. "Dynamics in a mature industry: Entry, exit, and growth of big-box grocery retailers." *Journal of Economics & Management Strategy* 24(1): 22–46.

Held, Philip J, Nathan W Levin, Randall R Bovbjerg, Mark V Pauly, and Louis H Diamond. 1991. "Mortality and duration of hemodialysis treatment." *Jama* 265(7): 871–875.

Ho, Kate, Joseph Hogan, and Fiona Scott Morton. 2017. "The impact of consumer inattention on insurer pricing in the Medicare Part D program." *The RAND Journal of Economics* 48(4): 877–905.

Ho, Vivian, and Barton H Hamilton. 2000. "Hospital mergers and acquisitions: does market consolidation harm patients?" *Journal of Health Economics* 19(5): 767–791.

Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven L Puller. 2017. "Power to choose? An analysis of consumer inertia in the residential electricity market." *American Economic Journal: Economic Policy* 9(4): 192–226.

Incentive AB, and Gambro AB. 1997. "Incentive's Gambro Group to Acquire Vivra Incorporated (Press Release)." Filed as Exhibit 99.(a)(1) to Vivra Inc. Schedule 14D-1, U.S. Securities and Exchange Commission. `https://www.sec.gov/Archives/edgar/data/0000850882/000091205797016522/0000912057-97-016522.txt`.

Jha, Ashish K, Karen E Joynt, E John Orav, and Arnold M Epstein. 2012. "The long-term effect of premier pay for performance on patient outcomes." *New England Journal of Medicine* 366(17): 1606–1615.

Kaiser Family Foundation. 2021*a*. "Medicare Spending Per Beneficiary." `https://tinyurl.com/3n2acxnn`.

Kaiser Family Foundation. 2021*b*. "Use of Inpatient Hospital Services in Traditional Medicare." `https://tinyurl.com/5n6txu7k`.

KFF. 2025. "Population Distribution by Age.".

Kim, Tae Woo, Tae Ik Chang, Tae Hee Kim, Jason A Chou, Melissa Soohoo, Vanessa A Ravel, Csaba P Kovesdy, Kamyar Kalantar-Zadeh, and Elani Streja. 2018. "Association of ultrafiltration rate with mortality in incident hemodialysis patients." *Nephron* 139(1): 13–22.

Kwoka, John. 2014. *Mergers, merger control, and remedies: A retrospective analysis of US policy*. Mit Press.

Kwoka, John, and Michael Pollitt. 2010. "Do mergers improve efficiency? Evidence from restructuring the US electric power sector." *International Journal of Industrial Organization* 28(6): 645–656.

Kwoka, John, and Shawn Kilpatrick. 2018. "Nonprice effects of mergers: issues and evidence." *The Antitrust Bulletin* 63(2): 169–182.

Leggat, JE, Sean M Orzol, Tempie E Hulbert-Shearon, Thomas A Golper, Camille A Jones, Philip J Held, and Friedrich K Port. 1998. "Noncompliance in hemodialysis: predictors and survival analysis." *American Journal of Kidney Diseases* 32(1): 139–145.

Lin, Eugene, Bich Ly, Erin Duffy, and Erin Trish. 2022. "Medicare Advantage Plans Pay Large Markups To Consolidated Dialysis Organizations." *Health Affairs* 41(8): 1107–1116.

Loughnan, Alice, Galil Rahman Ali, and Sumith C Abeygunasekara. 2011. "Comparison of the therapeutic efficacy of epoetin beta and epoetin alfa in maintenance phase hemodialysis patients." *Renal failure* 33(3): 373–375.

Manolis, Antonis A, Theodora A Manolis, Helen Melita, Dimitri P Mikhailidis, and Antonis S Manolis. 2022. "Low serum albumin: A neglected predictor in patients with cardiovascular disease." *European Journal of Internal Medicine* 102: 24–39.

Mariani, Marco, Leuconoe Grazia Sisti, Claudia Isonne, Angelo Nardi, Rosario Mete, Walter Ricciardi, Paolo Villari, Corrado De Vito, and Gianfranco Damiani. 2022. "Impact of hospital mergers: a systematic review focusing on healthcare quality measures." *European Journal of Public Health* 32(2): 191–199.

Medicare Payment Advisory Commission. 2021. "March 2021 Report to the Congress: Medicare Payment Policy.".

Murphy, Sheryl L., Kenneth D. Kochanek, Jiaquan Q. Xu, and Elizabeth Arias. 2024. "Mortality in the United States, 2023." NCHS Data Brief, No. 521. `https://www.cdc.gov/nchs/data/databriefs/db521.pdf` National Center for Health Statistics.

National Bureau of Economic Researc. 2025. "ZIP Code Distance Database." `https://www.nber.org/research/data/zip-code-distance-database`.

National Kidney Foundation. 2025. "Anemia and Iron Needs in Dialysis.".

National Vital Statistics Reports. 2022. "United States Life Tables, 2020." https://www.cdc.gov/nchs/data/nvsr/nvsr71/nvsr71-01.pdf.

Nephrology News and Issues. 2015. "U.S. Renal Care, DSI Renal to merge." https://www.healio.com/news/nephrology/20180227/us-renal-care-dsi-renal-to-merge.

Nguyen, Hoai-Luu Q. 2019. "Are credit markets still local? Evidence from bank branch closings." *American Economic Journal: Applied Economics* 11(1): 1–32.

NIH. 2023. "USRDS Annual Data Report: Epidemiology of Kidney Disease in the United States." National Institutes of Health.

NIH. 2024. "USRDS Annual Data Report: Epidemiology of Kidney Disease in the United States." National Institutes of Health.

Olsen, Andrew, Reed Orchinik, and Marc Remer. 2026. "Price Effects in US Merger Retrospectives: A Meta-Analytic Approach." : Forthcoming.

Paternoster, Raymond, Robert Brame, Paul Mazerolle, and Alex Piquero. 1998. "Using the correct statistical test for the equality of regression coefficients." *Criminology* 36(4): 859–866.

Petersen, Laura A, LeChauncy D Woodard, Tracy Urech, Christina Daw, and Supicha Sookanan. 2006. "Does pay-for-performance improve the quality of health care?" *Annals of internal medicine* 145(4): 265–272.

Pew Research Center. 2025. "Facts About the U.S. Black Population.".

Polyakova, Maria. 2016. "Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D." *American Economic Journal: Applied Economics* 8(3): 165–195.

Prakash, Suma, Rick Coffin, Jesse Schold, Steven A Lewis, Douglas Gunzler, Susan Stark, Matthew Howard, Darlene Rodgers, Douglas Einstadter, and Ashwini R Sehgal. 2014. "Travel distance and home dialysis rates in the United States." *Peritoneal Dialysis International* 34(1): 24–3.

Rambachan, Ashesh, and Jonathan Roth. 2023. "A more credible approach to parallel trends." *Review of economic studies* 90(5): 2555–2591.

Rosner, M. H., C. R. Manley, E. V. Hickey, and J. S. Berns. 2023. "Stakeholder Theory and For-Profit Dialysis: A Call for Greater Accountability." *Clinical Journal of the American Society of Nephrology* 18(9): 1225–1227.

Saran, Rajiv, Jennifer L Bragg-Gresham, Hugh C Rayner, David A Goodkin, Marcia L Keen, Paul C Van Dijk, Kiyoshi Kurokawa, Luis Piera, Akira Saito, Shunichi Fukuhara et al. 2003. "Nonadherence in hemodialysis: associations with mortality, hospitalization, and practice patterns in the DOPPS." *Kidney international* 64(1): 254–262.

Sloan, Caroline E, Abby Hoffman, Matthew L Maciejewski, Cynthia J Coffman, Justin G Trogdon, and Virginia Wang. 2021. "Trends in dialysis industry consolidation after Medicare payment reform, 2006-2016." Paper presented at the JAMA Health Forum,.

Total Renal Care Holdings, Inc., and Renal Treatment Centers, Inc. 1997. "Total Renal Care and Renal Treatment Centers Agree to Merge (Press Release)." Filed as Exhibit 99.1 to Total Renal Care Holdings, Inc. Form 8-K, U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/0000927066/000101706297002112/0001017062-97-002112.txt.

UCSF Hospitalist Handbook. N.d. "Hypocalcemia.".

U.S. Department of Health and Human Services. 2021. "Appendix D: Updating Value per Statistical Life (VSL) Estimates for Inflation and Changes in Real Income." https://aspe.hhs.gov/sites/default/files/2021-07/hhs-guidelines-appendix-d-vsl-update.pdf.

U.S. Department of Justice and the Federal Trade Commission. 2023. "Merger Guidelines." https://www.justice.gov/d9/2023-12/2023%20Merger%20Guidelines.pdf.

U.S. Renal Care, Inc. N.d. "U.S. Renal Care and DSI Renal Announce Definitive Merger Agreement (Press Release)." https://www.usrenalcare.com/press-releases/u-s-renal-care-and-dsi-renal-announce-definitive-merger-agreement/.

U.S. Renal Care, Inc., and Ambulatory Services of America, Inc. N.d. "U.S. Renal Care and Ambulatory Services of America Announce Definitive Merger Agreement (Press Release)." https://www.usrenalcare.com/press-releases/u-s-renal-care-and-ambulatory-services-of-america-announce-definitive-merger-agreement/.

Vanholder, Raymond, Griet Glorieux, and Sunny Eloot. 2015. "Once upon a time in dialysis: the last days of Kt/V?" *Kidney international* 88(3): 460–465.

Velázquez, Alexis F, Rebecca Thorsness, Amal N Trivedi, and Kevin H Nguyen. 2022. "County-level dialysis facility supply and distance traveled to facilities among incident kidney failure patients." *Kidney360* 3(8): 1367–1373.

Viard, V Brian. 2007. "Do switching costs make markets more or less competitive? The case of 800-number portability." *The RAND Journal of Economics* 38(1): 146–163.

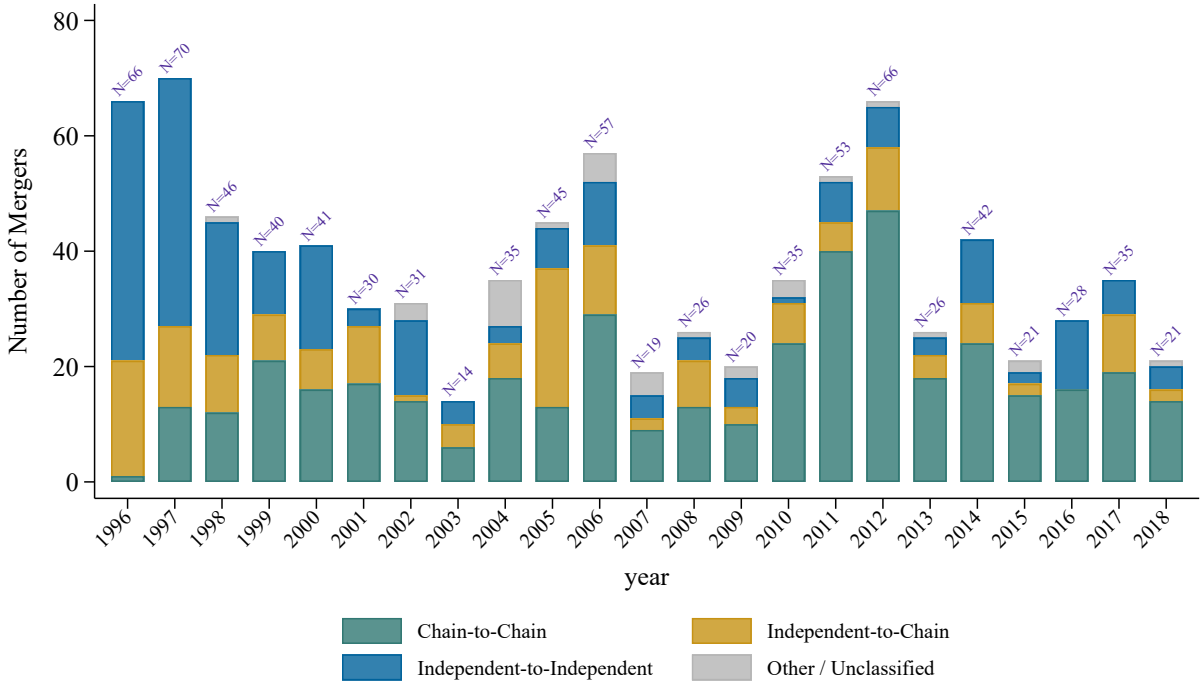
Whaley, Christopher M, Brian Briscombe, Rose Kerber, Brenna O'Neill, and Aaron Kofner. 2022. "Prices paid to hospitals by private health plans: findings from round 4 of an employer-led transparency initiative." *Rand Health Quarterly* 10(1): 5.

Wing, Coady, Seth M Freedman, and Alex Hollingsworth. 2024. Stacked difference-in-differences. Technical report National Bureau of Economic Research.

Wollmann, Thomas G. 2020. How to get away with merger: stealth consolidation and its effects on US healthcare. Technical report National Bureau of Economic Research.

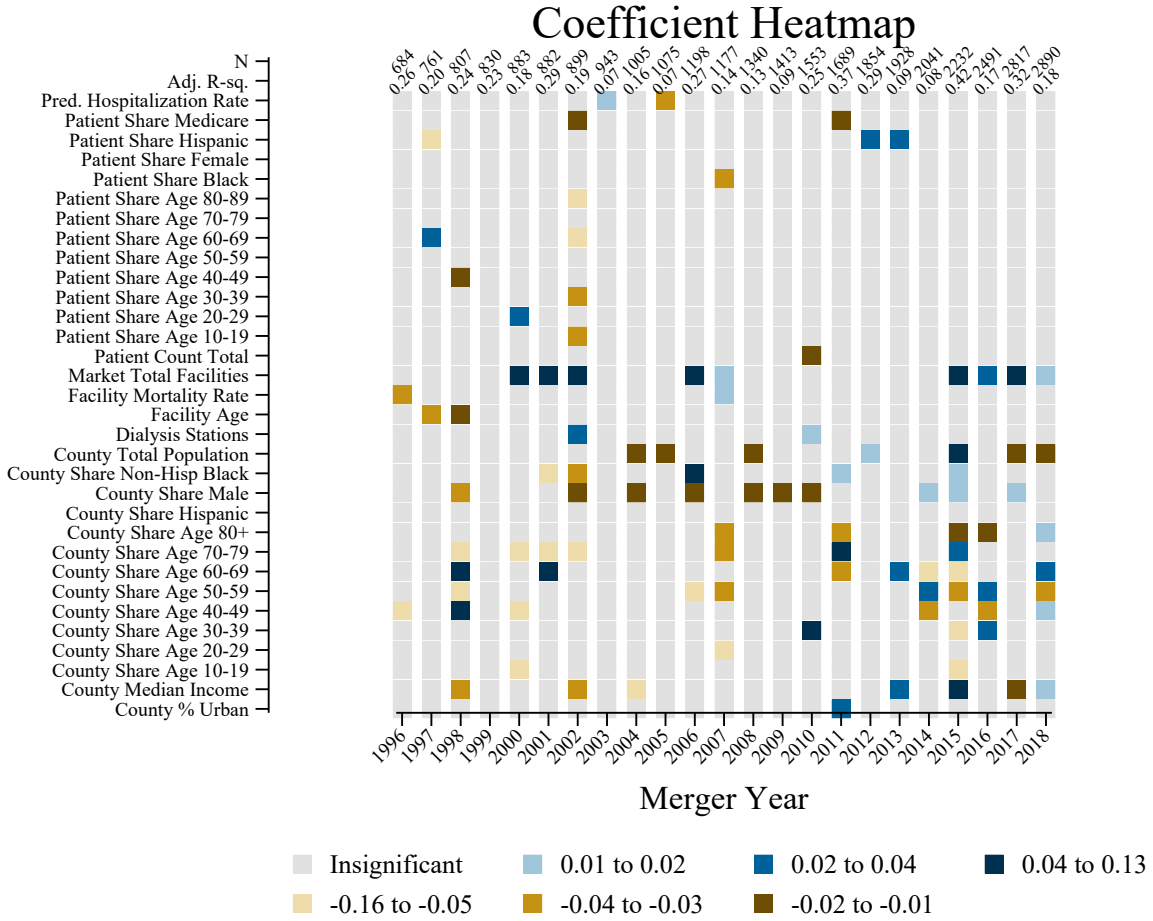
Figures

Figure 1: Merger Events in the Study Sample



Notes: This figure displays the unique merger events included in the analysis sample. The sample is constructed by applying the market-exclusion criteria described in Section 4.2 to all mergers occurring between 1996 and 2018. Bar height indicates the total number of merger events, with counts reported above each bar. Shading indicates merger subtype (target–acquirer), where green denotes chain-to-chain, yellow denoted independent-to-chain, blue denotes independent-to-independent, and gray denotes mergers that could not be categorized.

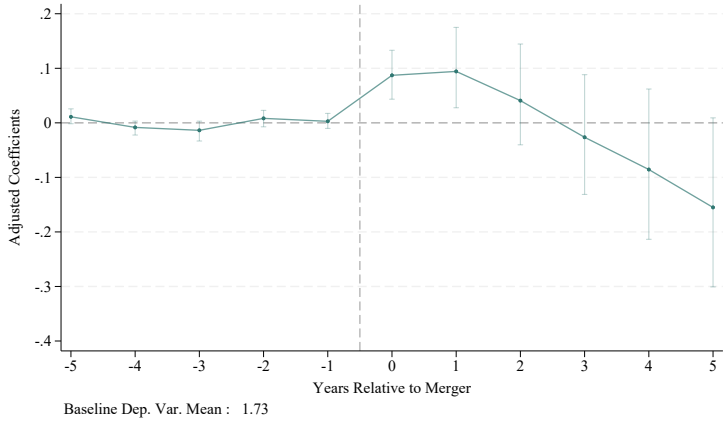
Figure 2: Predictors of Markets With Mergers



Notes: This figure displays coefficients from the linear probability regression described in equation 1. For each merger cohort, we construct a dataset containing all facilities in treated markets and facilities in control markets, defined as markets with no merger in that year. We restrict the analysis to the five years preceding the cohort-specific merger year and measure each predictor as its mean over this five-year pre-merger period for each facility. We then estimate the regression separately for each cohort year. The vector of predictors includes measures of market size, county demographics, facility characteristics, and patient composition. Cells corresponding to statistically insignificant coefficients are shaded gray. For coefficients significant at the 10% level or higher, positive estimates are shaded blue and negative estimates yellow, with shading intensity reflecting the magnitude of the coefficient.

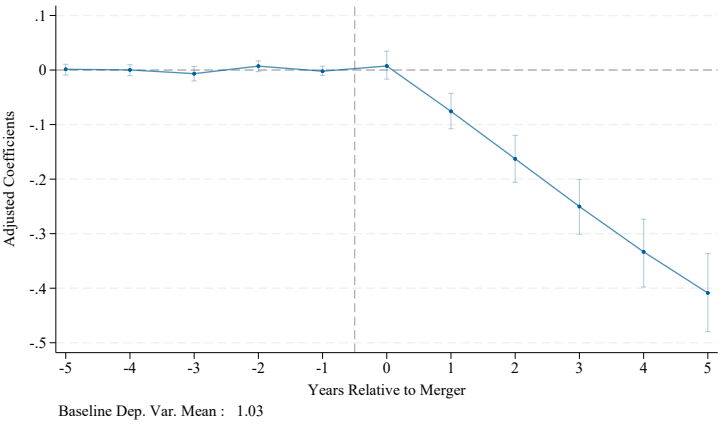
Figure 3: The Effect of Mergers on the Number of Facilities in the Market

(A) Total Facilities



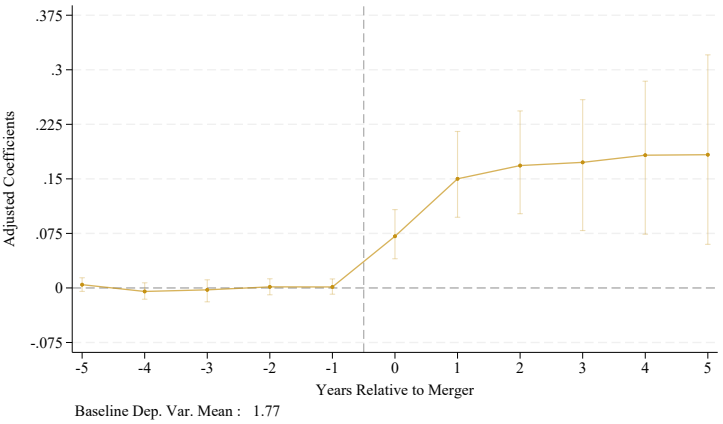
Baseline Dep. Var. Mean : 1.73

(B) Acquired Facilities



Baseline Dep. Var. Mean : 1.03

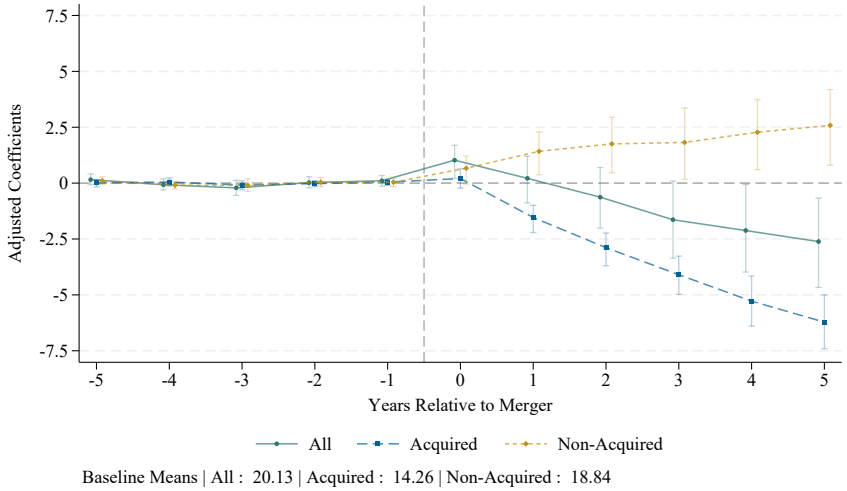
(C) Non-Acquired Facilities



Baseline Dep. Var. Mean : 1.77

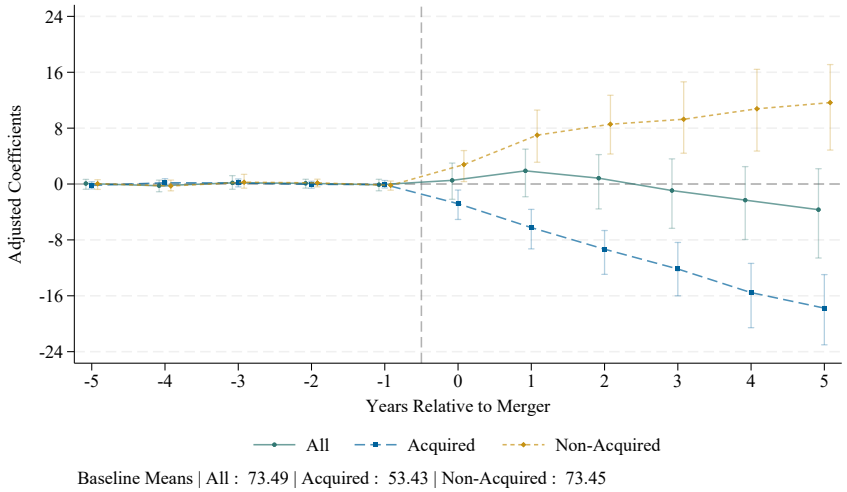
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of all facilities (panel A), acquired facilities (panel B), and non-acquired facilities (panel C). The unit of analysis is a market, with markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets in Panel A and treated markets in Panels B and C.

Figure 4: The Effect of Mergers on the Number of Dialysis Stations in the Market



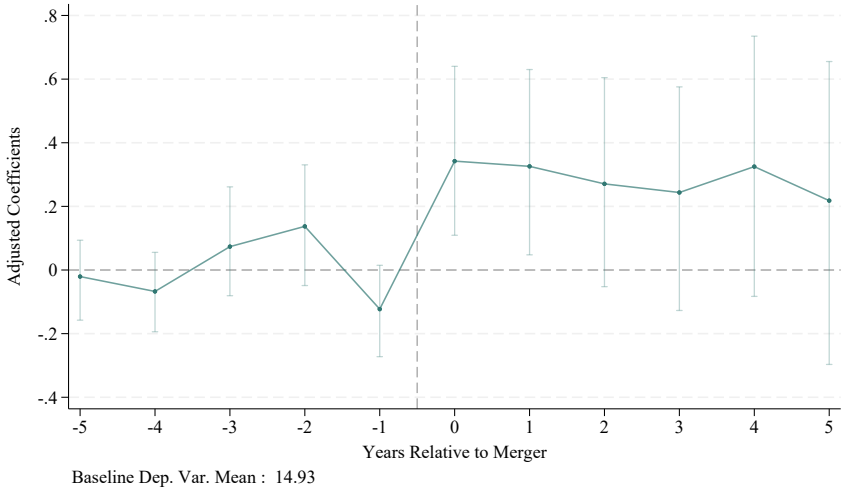
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of dialysis stations. The unit of analysis is a market, with markets experiencing a merger designated as treated. Effects are estimated separately for all facilities (green, solid line), acquired facilities (blue, dashed line), and non-acquired facilities (yellow, dotted line). The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets for all-facility stations and treated markets for acquired and non-acquired facility stations.

Figure 5: The Effect of Mergers on the Number of Patients in the Market



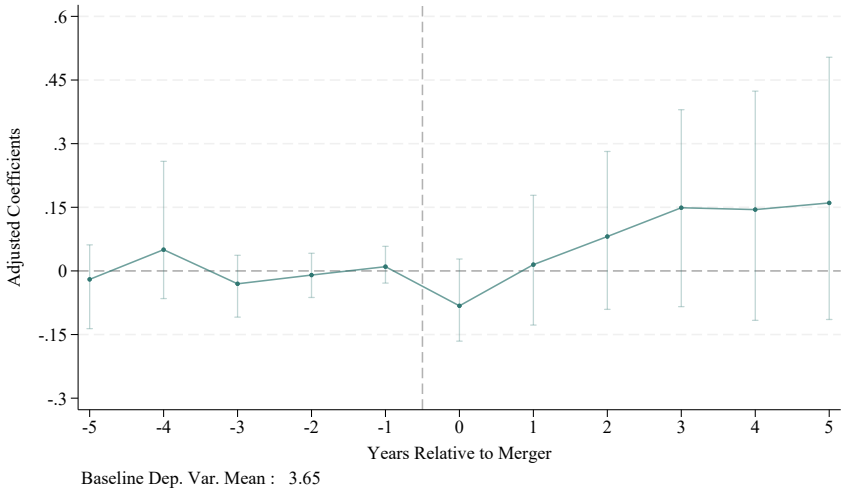
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of dialysis patients. The unit of analysis is a market, with markets experiencing a merger designated as treated. Effects are estimated separately for all facilities (green, solid line), acquired facilities (blue, dashed line), and non-acquired facilities (yellow, dotted line). The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets.

Figure 6: The Effect of Mergers on Travel Distance (Mean Miles Per Patient)



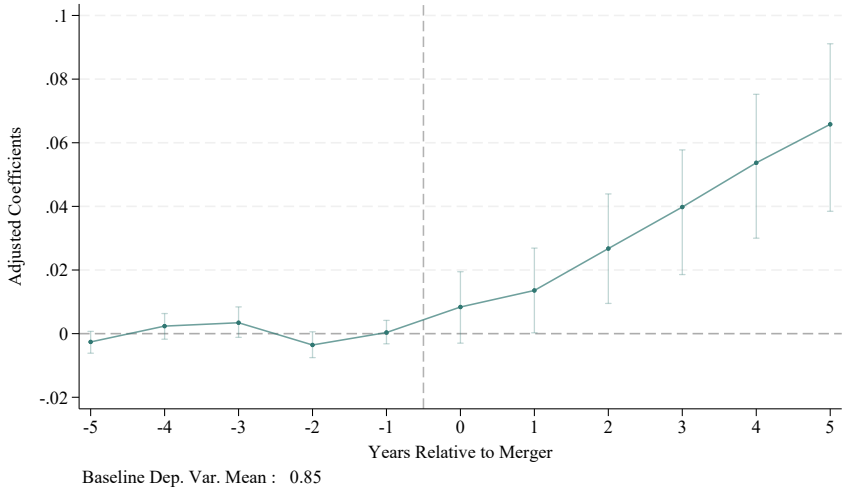
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on mean travel distance to a dialysis facility. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The outcome is the mean distance in miles between patients' residence ZIP codes and the facility's ZIP code. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time −1, calculated across all treated and control facilities.

Figure 7: The Effect of Mergers on Patients Per Dialysis Station



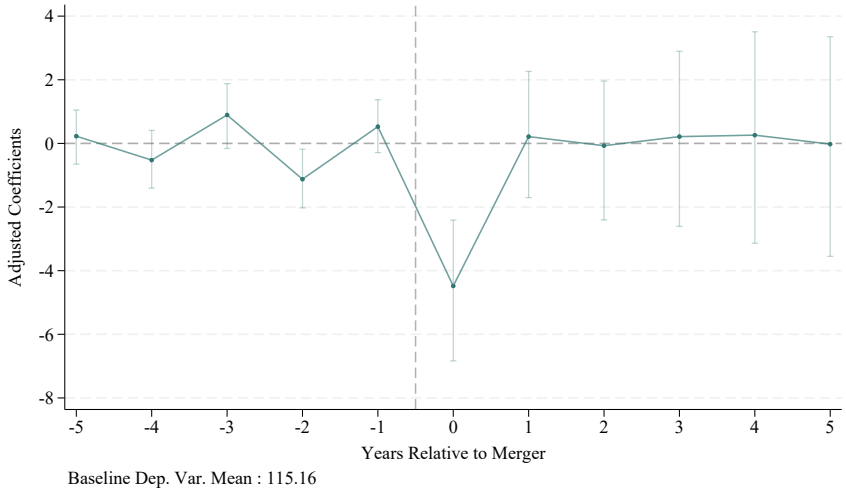
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on patients per station. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The outcome is the total number of patients divided by the total number of dialysis stations at the facility. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time −1, calculated across all treated and control facilities.

Figure 8: The Effect of Mergers on Market Concentration



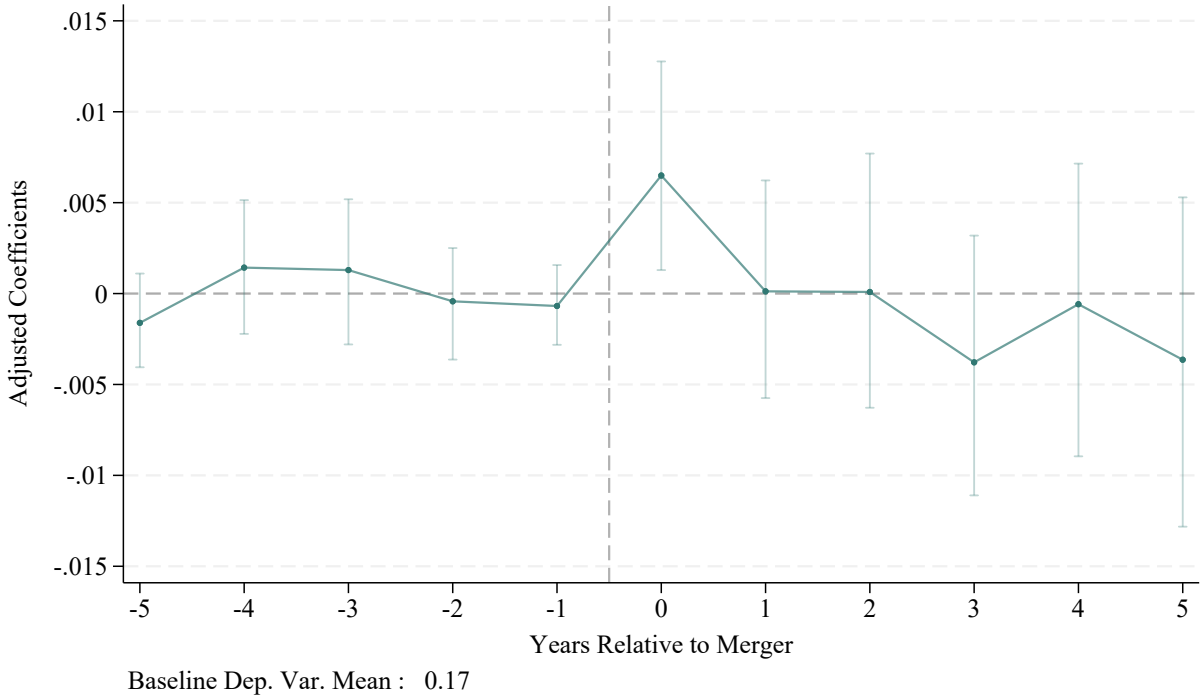
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on market concentration (HHI), measured as the sum of squared facility-level patient shares within a market. The unit of analysis is a market, with markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets.

Figure 9: The Effect of Mergers on Dialysis Sessions Per Patient



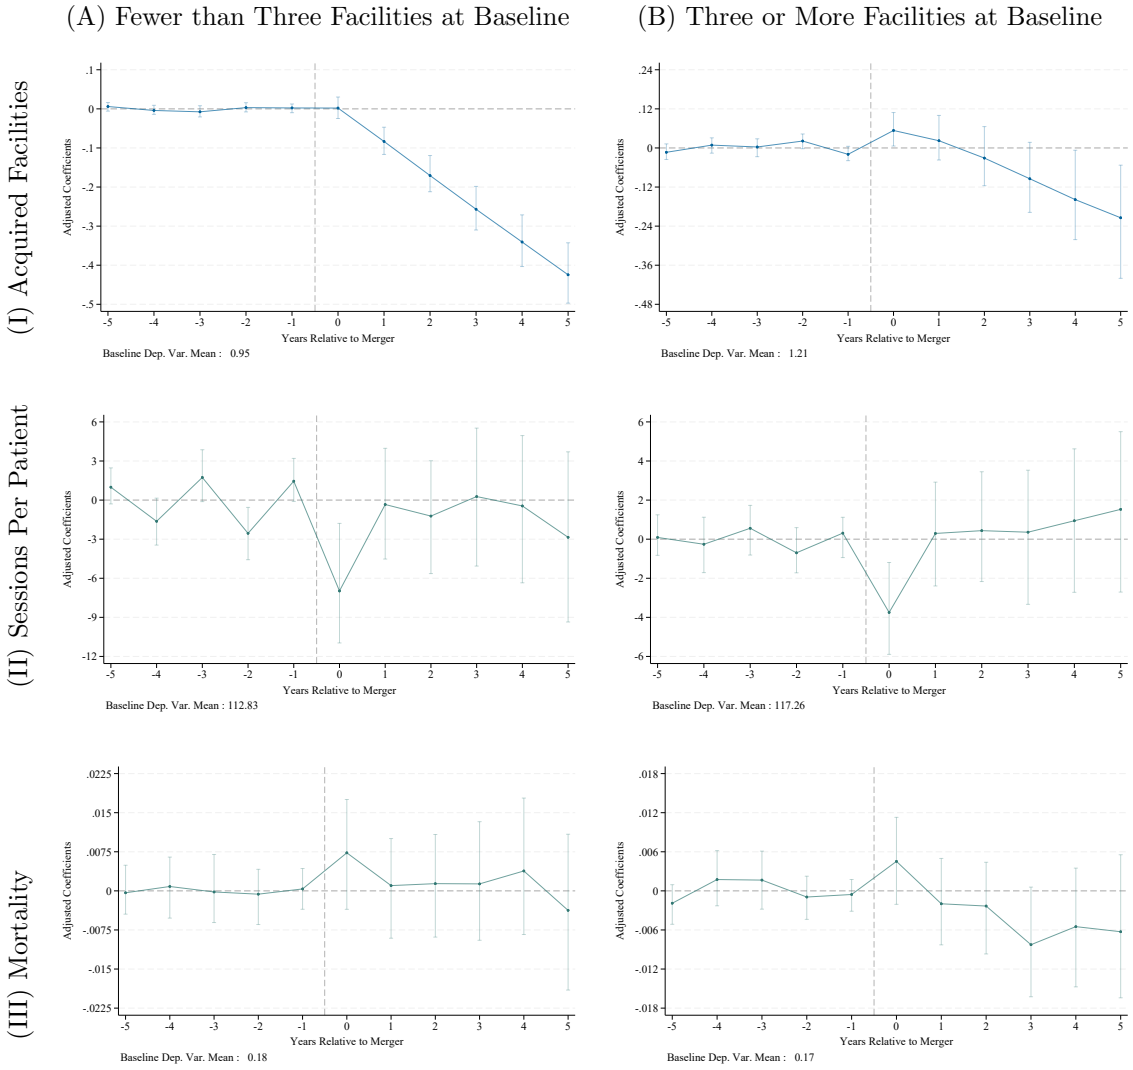
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on dialysis sessions per patient. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The outcome is the total number of sessions received annually divided by the total number of dialysis patients at the facility. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

Figure 10: The Effect of Mergers on Mortality



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on mortality. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The outcome is the total number of deaths divided by the total number of dialysis patients at the facility. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

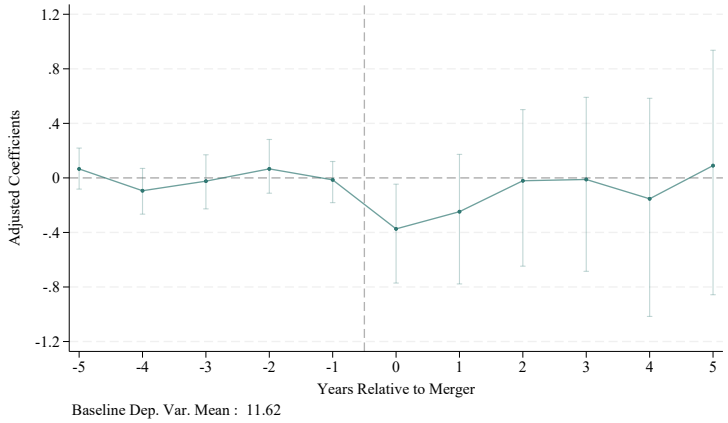
Figure 11: Heterogeneity in the Effect of Mergers on the Capacity, Sessions Per Patient and Mortality, by Baseline Capacity



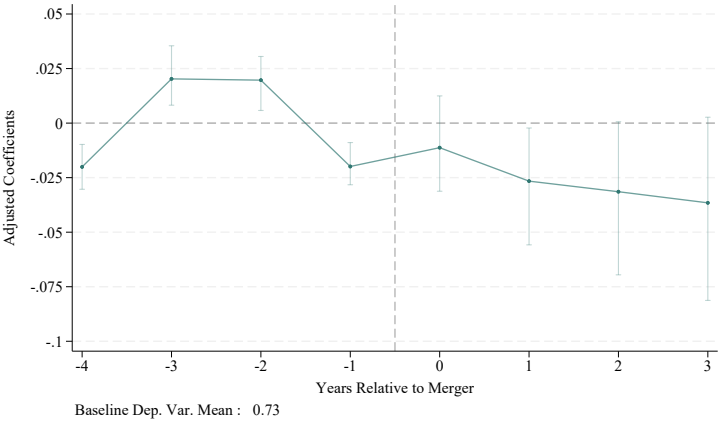
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on acquired facilities (I), sessions per patient (II), and mortality (III) for markets with two or fewer facilities (Panel A), and markets with three or more facilities (panel B) at event time −1. The unit of analysis is a market in panel A and a facility in panels B-C, with all facilities (acquired and non-acquired) in markets experiencing a merger of the reference type designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls in panel A, and market- and facility-level controls in panels B-C. The baseline dependent variable mean is the mean outcome at event time −1, calculated across all treated and control facilities.

Figure 12: The Effect of Mergers on Staffing, Anemia Drug Prescription and Fistula Usage

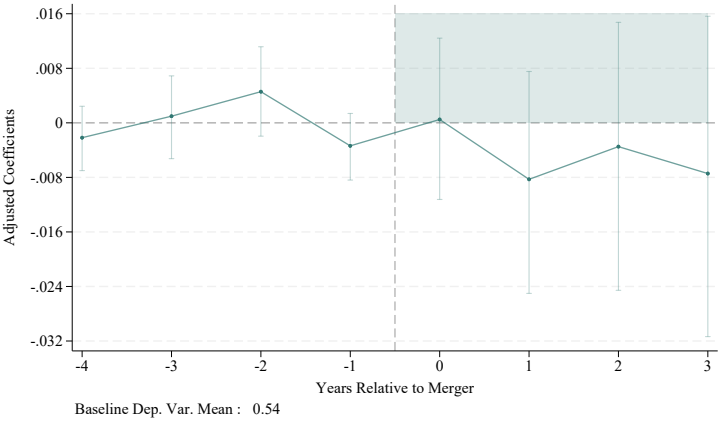
(A) Staff Volume



(B) Share of Patients With ESA Prescription



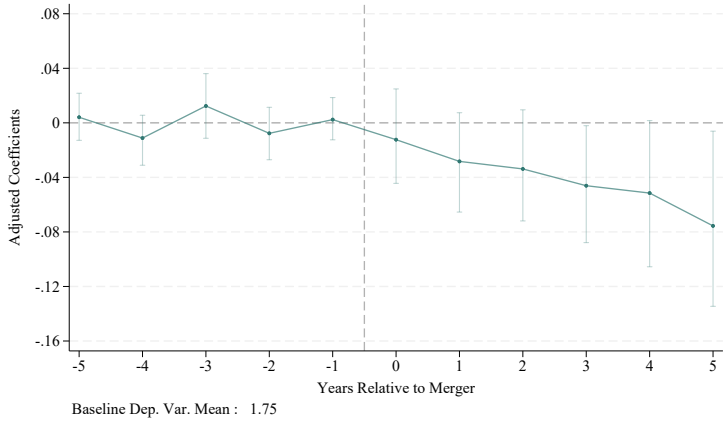
(C) Share of Patients With Fistula



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on staffing volume (panel A), the share of the facility's patients with a prescription for an erythropoiesis-stimulating agent (ESA; Panel B), and the share of the facility's patients with a fistula for vascular access (panel C). Green shading indicates the direction of clinically beneficial merger effects. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

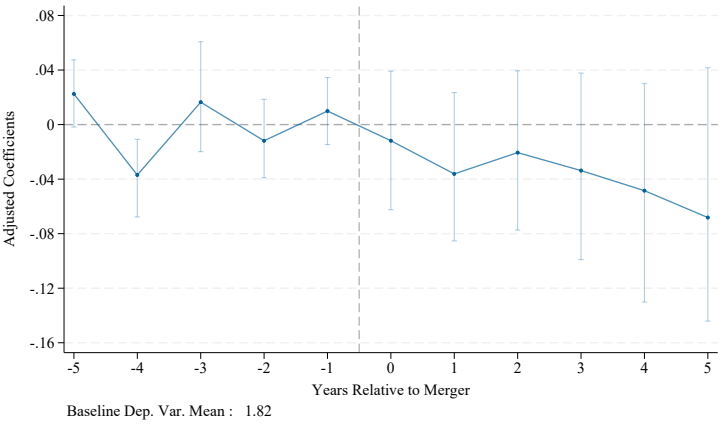
Figure 13: The Effect of Mergers on Hospitalizations Per Patient

(A) All Facilities



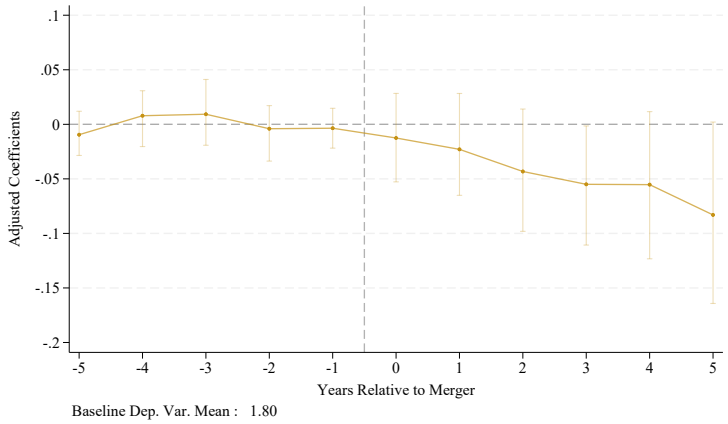
Baseline Dep. Var. Mean : 1.75

(B) Acquired Facilities



Baseline Dep. Var. Mean : 1.82

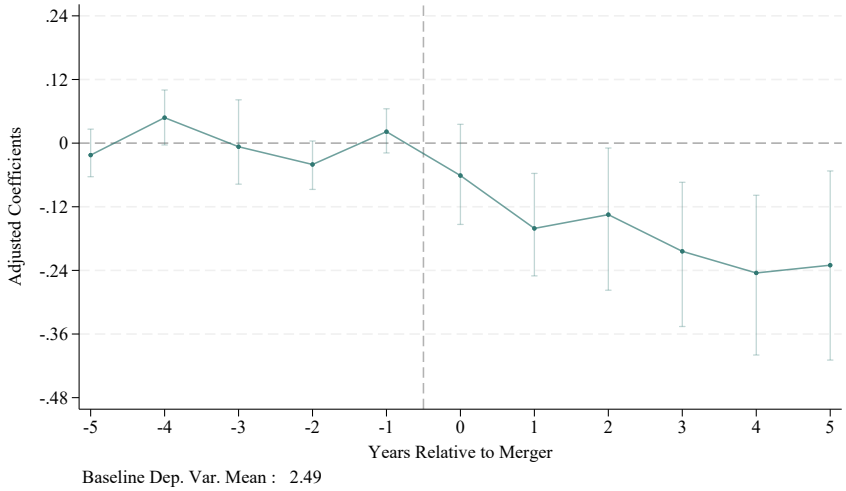
(C) Non-Acquired Facilities



Baseline Dep. Var. Mean : 1.80

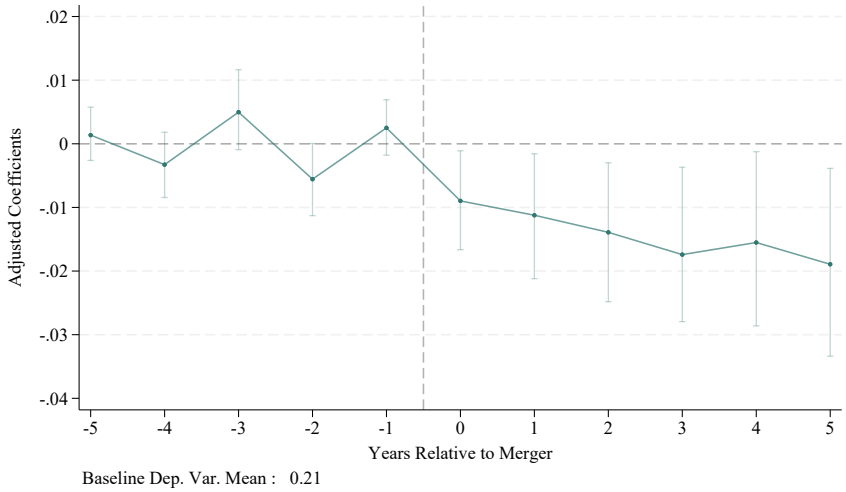
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on hospitalizations per patient at all facilities (panel A), acquired facilities (panel B), and non-acquired facilities (panel C). The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The outcome is the total number of hospitalizations (including 0) divided by the total number of dialysis patients at the facility. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets in Panel A and treated markets in Panels B and C.

Figure 14: The Effect of Mergers on ICU Days Per Patient



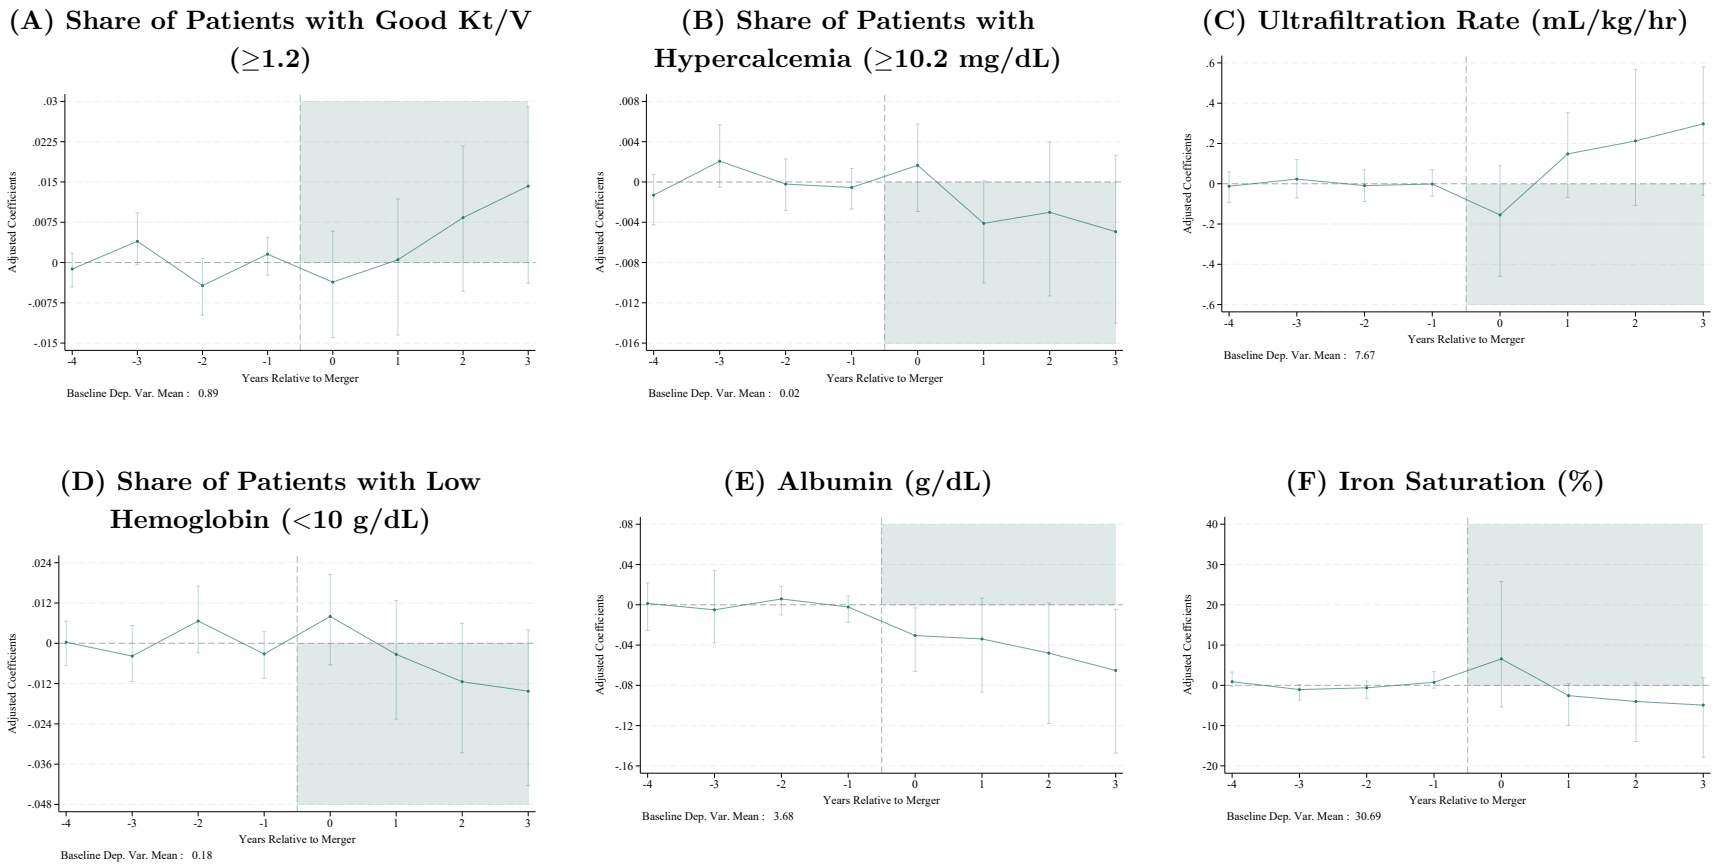
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on days in the ICU per patient. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The outcome is the total number of days in the ICU (including 0) divided by the total number of dialysis patients at the facility. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

Figure 15: The Effect of Mergers on Blood Transfusion Events Per Patient



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on inpatient blood transfusion events per patient. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The outcome is the total inpatient blood transfusion events (including 0), identified using ICD codes in inpatient claims as described in Appendix J, divided by the total number of dialysis patients at the facility. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

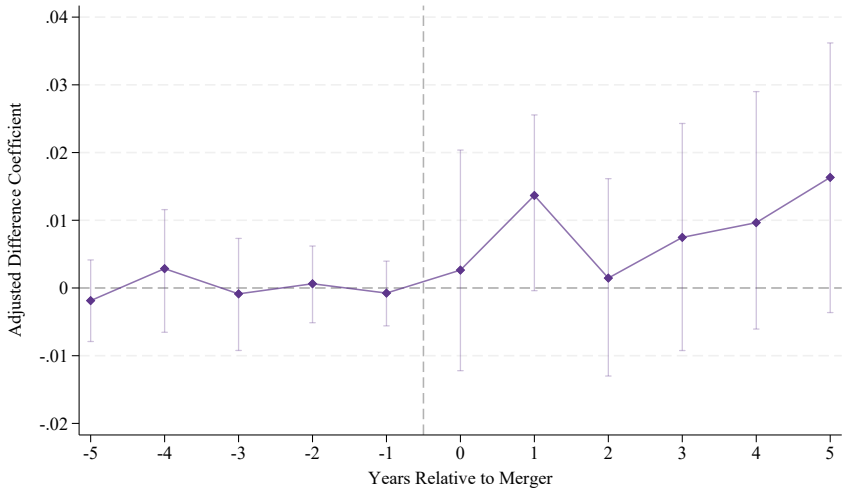
Figure 16: The Effect of Mergers on Biomarkers



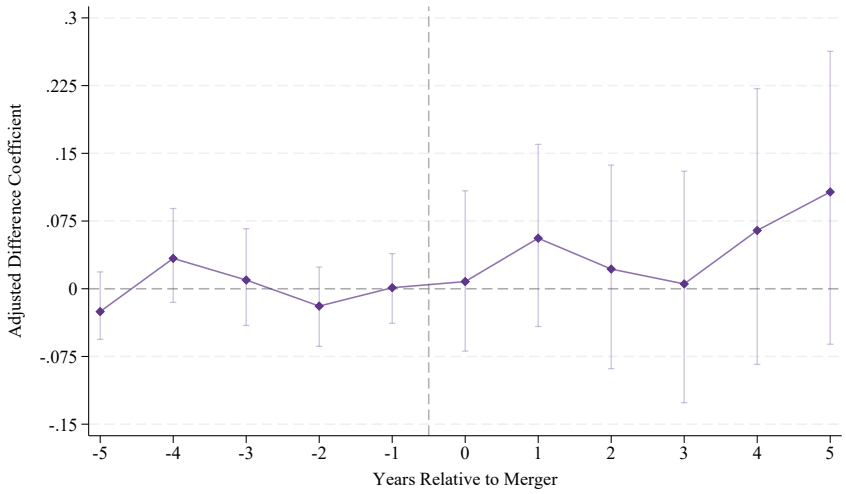
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on blood biomarkers. The outcome in panel A is the share of patients with good Kt/V ($\geq$ 1.2 mg/dL). The outcome in panel B is the share of patients with hypercalcemia defined as serum calcium levels exceeding 10.2 mg/dL. The outcome in panel C is the mean ultrafiltration rate (mL/kg/hr). The outcome in panel D is the share of patients with low hemoglobin (<10g/dL). The outcome in panel E is the mean serum albumin (g/dL). The outcome in panel F is the mean iron saturation (%). Green shading indicates the direction of clinically beneficial merger effects. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

Figure 17: Difference in the Effect of Mergers on Mortality and Hospitalizations From Pre-2011 to Post-2012

(A) Difference in the Effect on Mortality

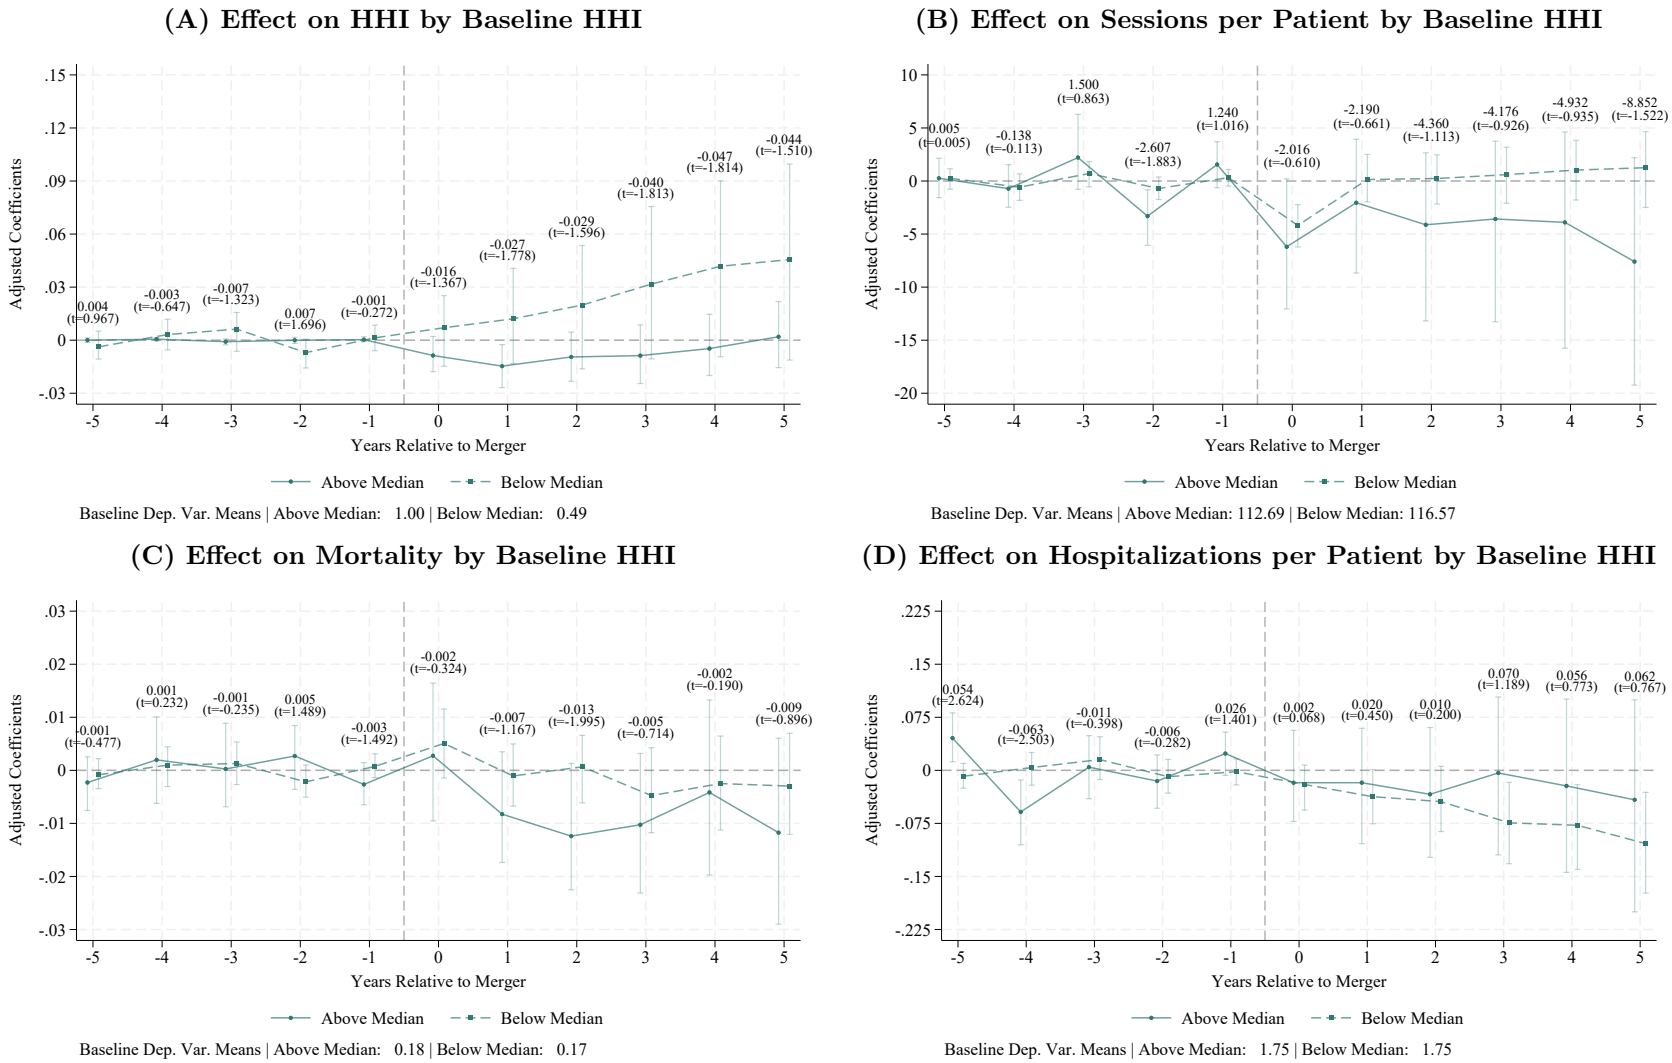


(B) Difference in the Effect on Hospitalizations per Patient



Notes: This figure presents the difference in the effect of a merger on mortality (panel A) and hospitalizations per patient (panel B) before versus after 2012, when the Medicare Quality Incentive Program and bundled Prospective Payment System were introduced. We estimate equation 4 separately for the pre-QIP sample (merger cohorts 1996–2008) and the post-QIP sample (merger cohorts 2015–2018). To account for the possibility that the same facilities appear in both samples, we implement a joint bootstrap, computing the pre-QIP and post-QIP estimates, and their difference, using the same resampled dataset in each iteration. The unit of analysis is a a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

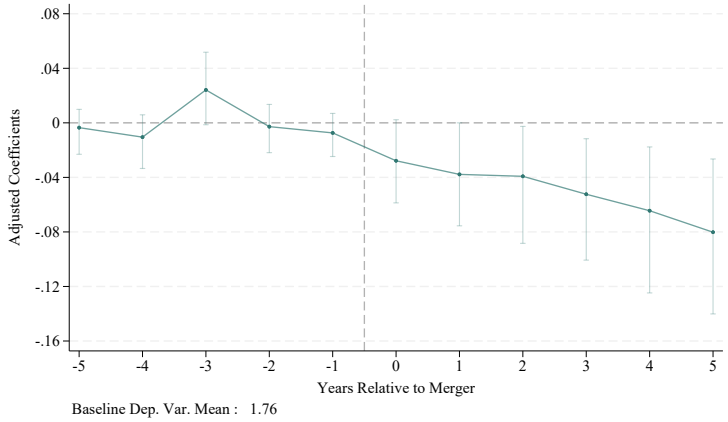
Figure 18: Heterogeneity in The Effect of Mergers by Baseline Market Concentration



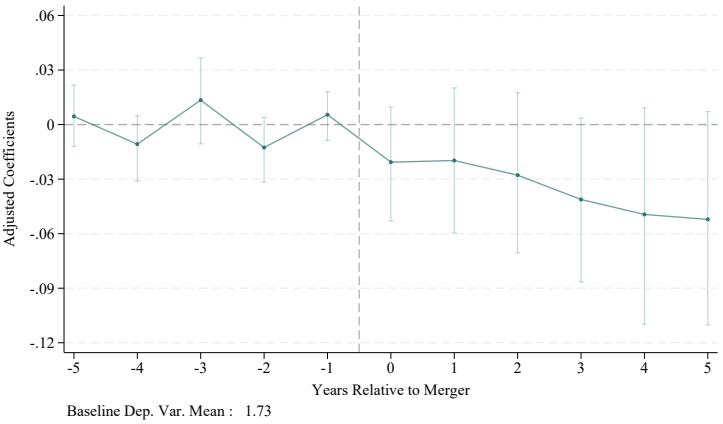
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on market concentration (HHI; Panel A), dialysis sessions per patient (Panel B), mortality (Panel C), and hospitalizations per patient (Panel D), by baseline market HHI. The coefficients denoted by the dashed line are estimated for markets with HHI below the cross-market median at event time −1, while those denoted by the solid line are estimated for markets with HHI at or above the median. The median is calculated in a stacked market-level dataset, treating duplicate instances of markets in the stack as distinct observations that plausibly correspond to unique markets sufficiently separated in time. We conduct the two-step event-study analysis separately for the two groups and test for differences in the resulting coefficients. The figure reports the difference in coefficients and the corresponding t-statistic. The unit of analysis is a market in panel A and a facility in panels B-D, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time −1, calculated across all treated and control facilities.

Figure 19: A Decomposition of the Sources of Long-Run Gains in Quality

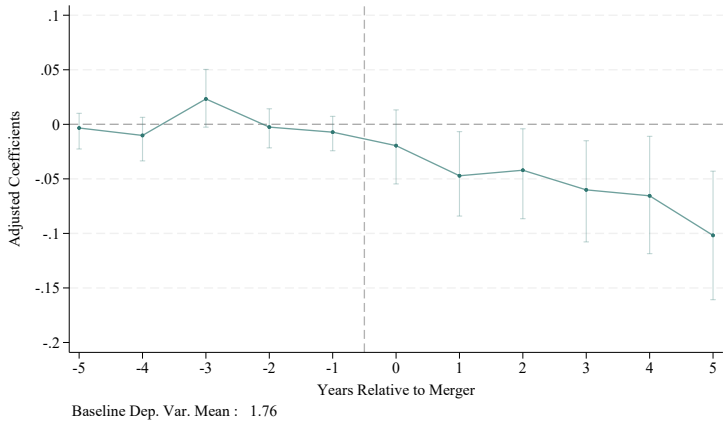
(A) Hospitalization Per Patient at Surviving Facilities



(B) Hospitalization Per Patient at Surviving + Exiting Facilities



(C) Hospitalization Per Patient at Surviving + Entering Facilities



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on hospitalizations per patient. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. In Panel A, the sample includes survivors—facilities that were in the market for at least three years prior to the merger and remained in the market throughout the entire post-merger period. In Panel B, the sample includes both survivors and exiting facilities, where exiting facilities were in the market for at least three years prior to the merger but exited at some point before the end of the full post-merger period. In Panel C, the sample includes both survivors and entrant facilities, where entrant facilities were not in the market at any point during the pre-merger period but entered at some point during the post-merger period. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

Tables

Table 2: Summary Statistics for Outcomes and Covariates in Market-Level Analyses

(A) Outcomes			
Variable	Mean	SD	Count
Markets : treated			724
Markets : control			1784
Facilities : all	1.79	2.18	
Facilities : acquired	1.74	2.05	
Facilities : non-acquired	1.76	2.16	
Stations : all	20.98	32.13	
Patients: all	76.10	132.21	
Patients: Medicare	66.27	108.77	
Patients: non-Medicare	9.83	32.27	
HHI	0.84	0.26	
(B) Covariates			
Variable	Mean	SD	
Median income	42228	13078	
Pop. ages 0-9	10675	30961	
Pop. ages 10-19	11596	31559	
Pop. ages 20-29	11336	35647	
Pop. ages 30-39	11280	35882	
Pop. ages 40-49	11660	32985	
Pop. ages 50-59	10669	28606	
Pop. ages 60-69	7963	20381	
Pop. ages 70-79	5113	13050	
Pop. ages 80+	3052	8162	
Pop. female	42416	121309	
Pop. Hispanic	9110	61064	
Pop. non-Hispanic Black	9242	53912	
Pop. non-Hispanic White	60677	120970	
Dialysis CON law	0.24	0.43	

Notes: This table presents summary statistics for the market-level covariates, and the outcomes analyzed at the market level, and reports the number of unique treated and control markets in our analytical sample. Statistics for outcomes are computed using the stacked dataset for the market-level analyses. The data sources and construction of covariates are described in Section 3.1 and Appendix J.

Table 3: Summary Statistics for Outcomes in Facility-Level Analyses

Variable	Mean	SD	Variable	Mean	SD
Patients per station	3.65	2.90	Death rate	0.18	0.09
Travel distance (miles per patient)	14.86	7.69	Hospitalizations per patient	1.74	0.61
Sessions per patient	116.01	29.79	ICU days per patient	2.53	2.13
			Transfusions per patient	0.20	0.17
Share pat. with ESA prescription	0.72	0.20			
Mean ESA dosage (micrograms)	13266.15	1269691.60	Share pat. with good Kt/V	0.89	0.15
Mean ESA dosage (units)	54641.22	655284.19	Share pat. with hypercalcemia	0.02	0.03
Share pat. with fistula	0.53	0.15	Share pat. with hypocalcemia	0.10	0.07
			Share pat. with high hemoglobin	0.08	0.06
Staff : all	11.48	7.33	Share pat. with low hemoglobin	0.19	0.11
Nurses	5.10	3.91	Mean ultrafiltration rate	7.53	4.79
Technicians	5.30	4.35	Mean iron saturation	30.40	36.40
Social workers	1.07	0.55	Mean serum albumin	3.72	0.72

Notes: This table presents summary statistics for the outcomes analyzed at the facility level. Statistics are computed using the stacked dataset used in the regression analysis of each outcome.

Table 4: Summary Statistics for Covariates in Facility-Level Analyses

Variable	Mean	SD
Facility age	11.95	9.14
Case-mix (predicted hospitalizations per patient)	1.75	0.08
Share pat. female	0.45	0.10
Share pat. non-Hispanic White	0.66	0.29
Share pat. non-Hispanic Black	0.28	0.29
Share pat. Hispanic	0.07	0.14
Share pat. ages 10-19	0.00	0.03
Share pat. ages 20-29	0.02	0.03
Share pat. ages 30-39	0.05	0.04
Share pat. ages 40-49	0.11	0.06
Share pat. ages 50-59	0.19	0.08
Share pat. ages 60-69	0.25	0.08
Share pat. ages 70-79	0.24	0.09
Share pat. ages 80-80	0.12	0.08

Notes: This table presents summary statistics for the facility-level covariates. We report here statistics computed using the stacked dataset for the mortality analysis, which is the broadest in terms of both time coverage and patient inclusion. Analogous facility-level covariates are included in the stacked datasets used for the other analyses.

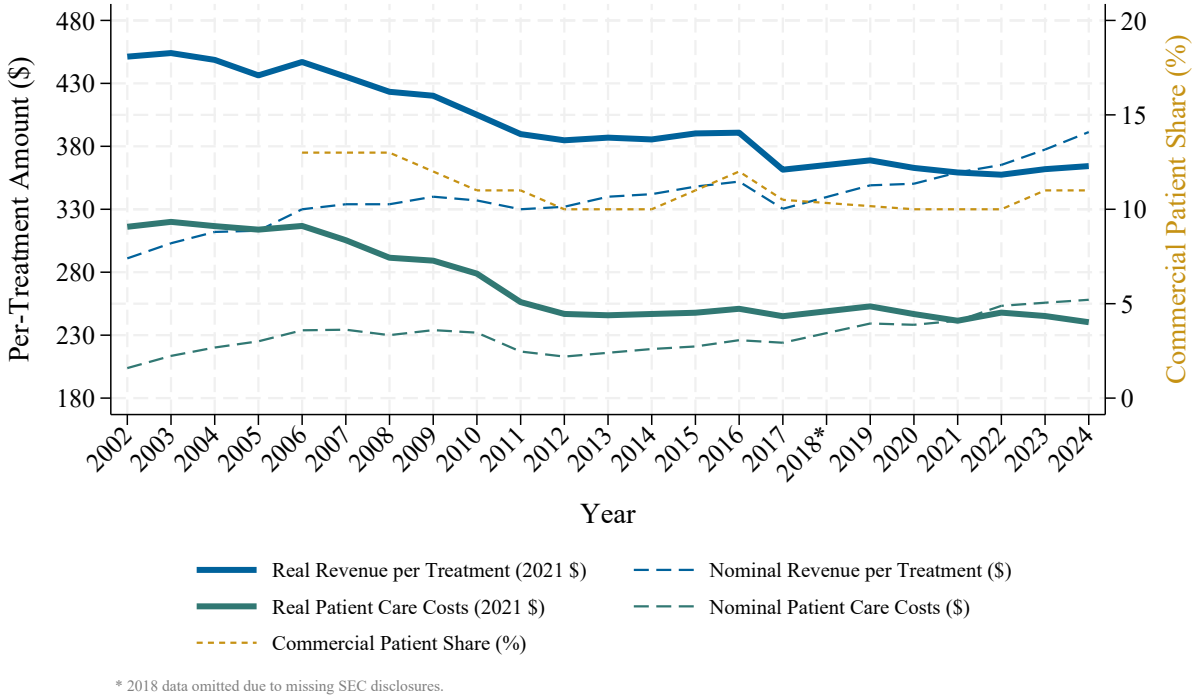
Table 5: Baseline Cost-Benefit Calculations

Component	Source	Value
1. Exposed Population & Core Parameters		
Exposed Dialysis Patients	USRDS / Market Estimate	109,295.21
Cost per Hospital Discharge	KFF / Internal Estimates	$25,433.38
Valuation of Statistical Life Year (VSLY)	Central Estimate	$278,000.00
Life Expectancy (Dialysis Population)	USRDS (Table 6.1)	4.30 Years
2. Mortality Costs (Year 0 Effect)		
Mortality Rate Coefficient ($t = 0$)	Regression Output	+0.006494
Total Excess Deaths	Exposed Patients $\times$ Coefficient	709.77
Cost per Death	LE (4.3) $\times$ VSLY ($278k)	$1,195,400.00
Total Mortality Cost	**Deaths $\times$ Cost per Death**	**$848.46 Million**
3. Hospitalization Benefits (Cumulative Years 0–5)		
Sum of Hospitalization Coefficients ($t = 0$ to 5)	Event Study Outputs	–0.247584
Total Averted Discharges	Exposed Patients $\times \sum$ Coefficients	27,055.97
Total Hospitalization Benefit	**Discharges $\times$ Cost per Discharge**	**$688.22 Million**
4. Baseline Net Balance		
Net Benefit (Benefits − Costs)		**–$160.24 Million**

Notes: This table presents the calculations underlying our cost-benefit analysis, in which we monetize the merger-year health harms from increased mortality and the cumulative long-run health gains from reduced hospitalizations. The calculation procedure and the data sources used to construct the relevant inputs are described in detail in Section 5.

A Appendix Figures

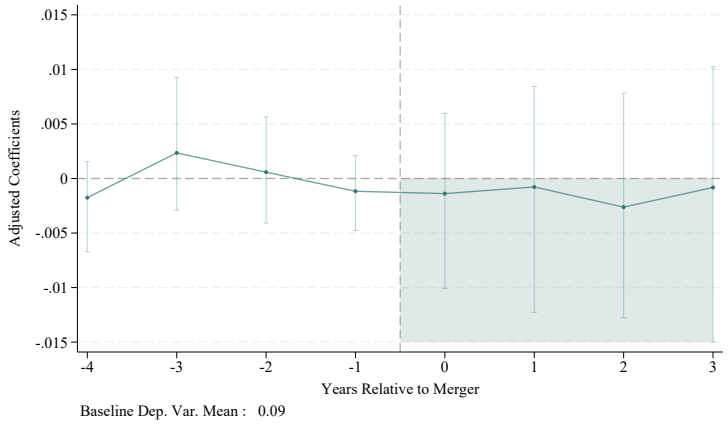
Figure A.1: DaVita Per-Treatment Revenue, Costs, and Commercial Patient Share



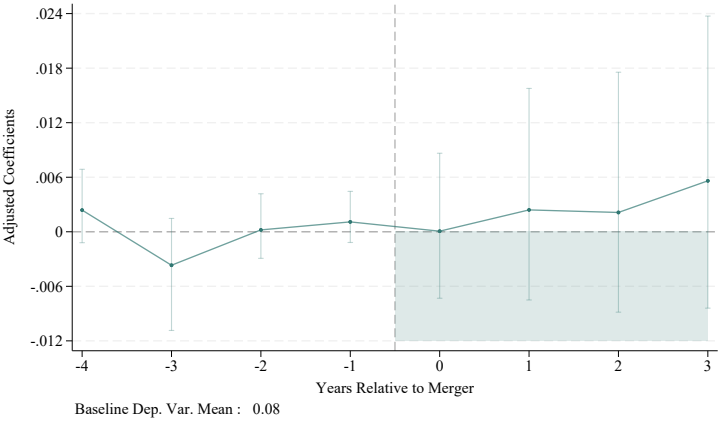
Notes: This figure presents financial information for DaVita Inc., a large dialysis provider, from 2002 to 2024, based on the firm's Form 10-K filings with the U.S. Securities and Exchange Commission (SEC). We obtain each year's filing from the SEC's public archives and use the filing for the fiscal year ending December 31. Dashed lines show nominal revenue per treatment (blue) and patient care costs per treatment (green). We use the Personal Consumption Expenditures Health Care Price Index, provided by the U.S. Bureau of Economic Analysis and obtained from the Federal Reserve Bank of St. Louis's FRED database, to convert nominal amounts to 2021 dollars, shown by the thick solid lines. The yellow dotted line represents the share of DaVita patients with commercial insurance, also obtained from the firm's filings; its values are reported on the right y-axis. Data for 2018 are unavailable.

Figure A.2: The Effect of Mergers on Hypocalcemia and the Share of Patients with High Hemoglobin

(A) Share of Patients with Hypocalcemia (< 8.5 mg/dL)

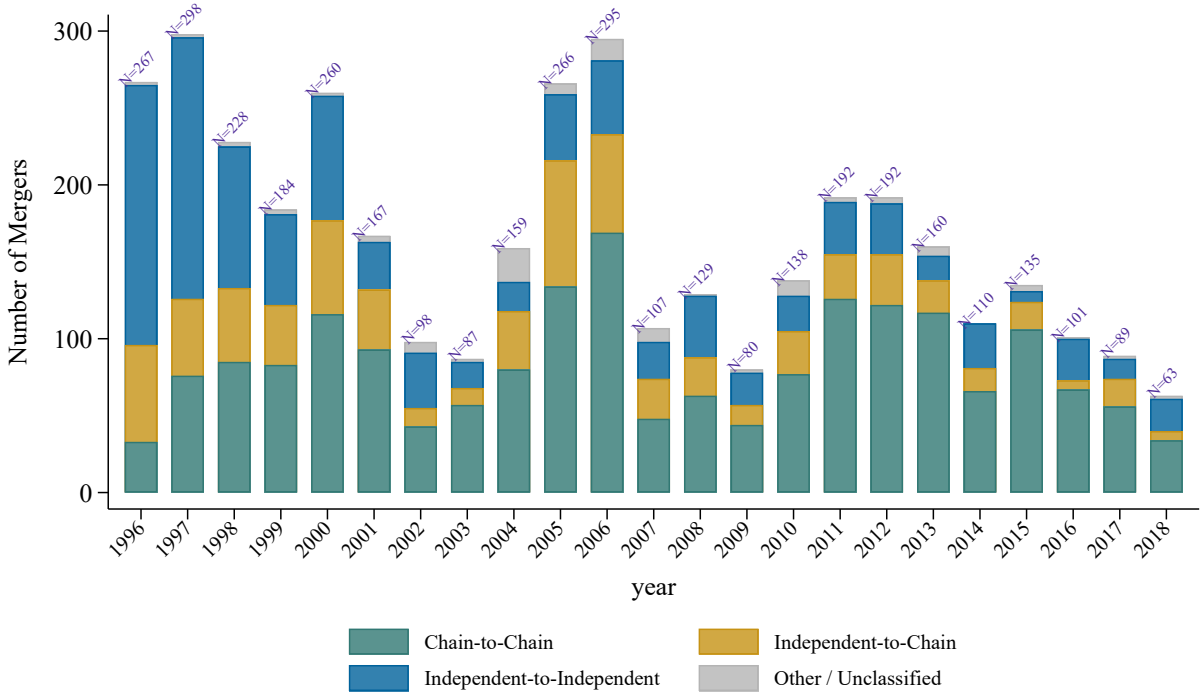


(B) Share of Patients With High Hemoglobin (> 12 g/dL)



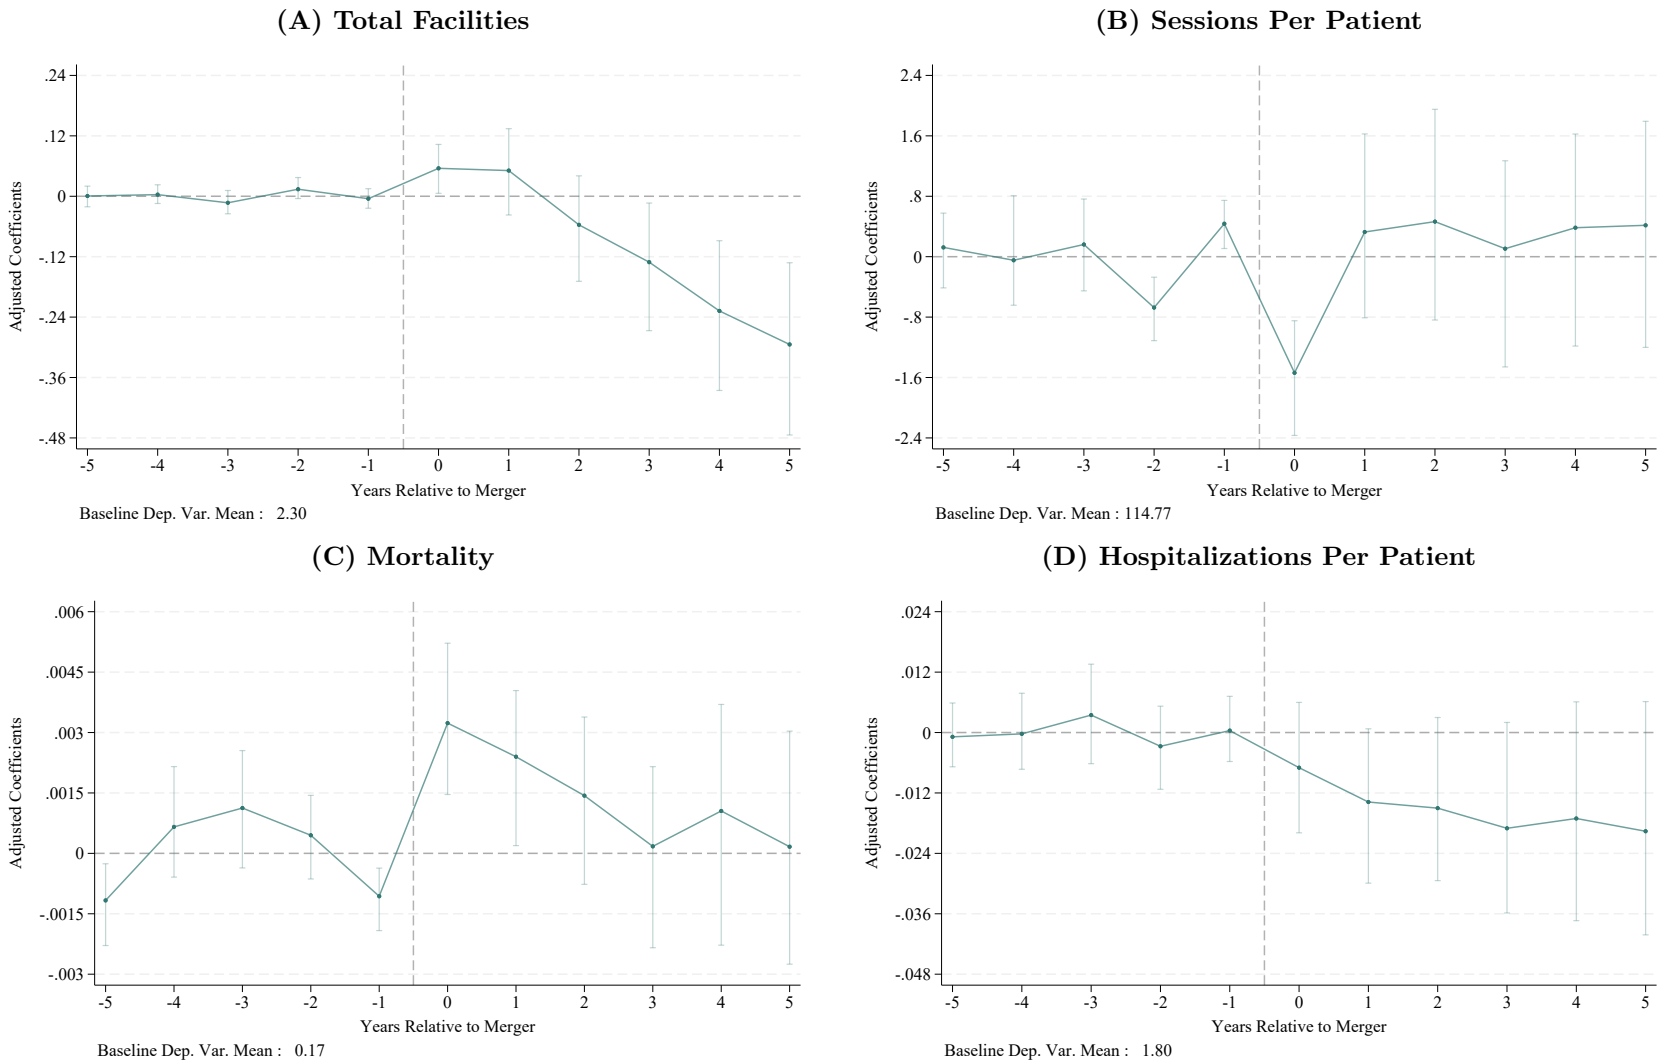
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on blood biomarkers. The outcome in panel A is the share of patients with hypocalcemia defined as serum calcium levels below 8.5 mg/dL. The outcome in panel B the share of patients with high hemoglobin (>12g/dL). Green shading indicates the direction of clinically beneficial merger effects. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

Figure A.3: Merger Events in the Population



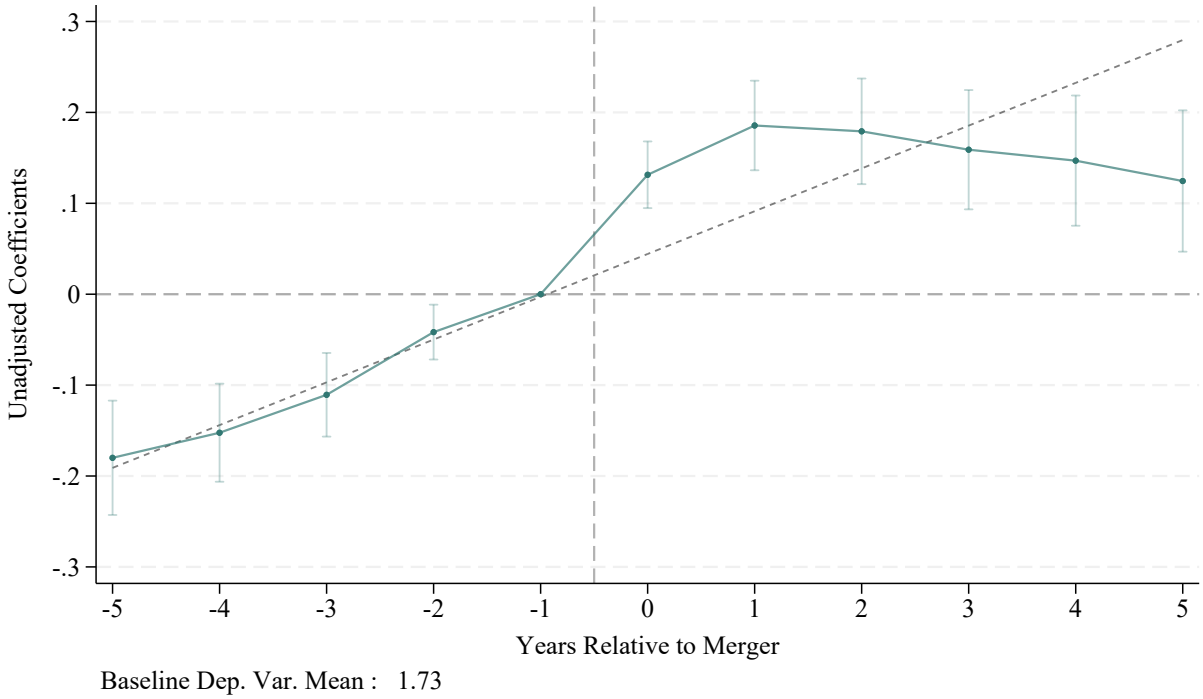
Notes: This figure displays the unique merger events occurring between 1996 and 2018 before applying the market-exclusion criteria described in Section 4.2 that determine the analysis sample. Bar height indicates the total number of merger events, with counts reported above each bar. Shading indicates merger subtype (target–acquirer), where green denotes chain-to-chain, yellow denotes independent-to-chain, blue denotes independent-to-independent, and gray denotes mergers that could not be categorized.

Figure A.4: The Effect of Mergers on the Number of Facilities, Sessions Per Patient, Mortality and Hospitalizations Per Patient in the Full Population of Mergers



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of facilities (Panel A), sessions per patient (Panel B), mortality (Panel C), and hospitalizations per patient (Panel D) in the full population of merger events, without the sample restrictions imposed to construct our analytical sample. The unit of analysis is a market in panel A, and a facility in panels B-D, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls in panel A, and market- and facility-level controls in panels B-D. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

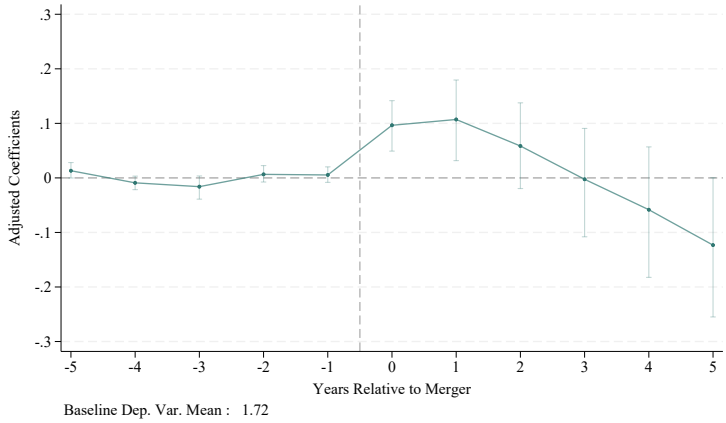
Figure A.5: The Effect of Mergers on the Number of Facilities, Without Pre-Trend Adjustment



Notes: This figure presents coefficients from equation 1, representing the effect of a merger on the number of facilities without pre-trend adjustment. The unit of analysis is a market, with markets experiencing a merger designated as treated. The slanted dotted line shows the linear trend fitted to the pre-merger coefficients and extrapolated to the post-merger event times. The trend-adjusted coefficients presented in Figure 3 are the differences between the unadjusted coefficients shown here and this fitted trend. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals with standard errors clustered at the market-cohort level. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets.

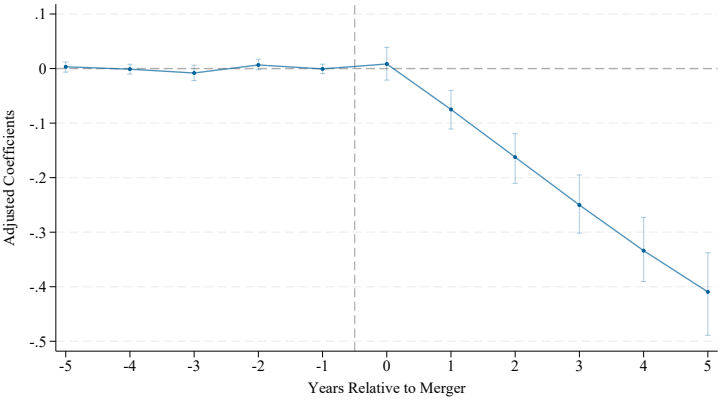
Figure A.6: The Effect of Mergers on the Number of Facilities, Excluding Markets With Very High Pre-Merger Facility Growth

(A) Total Facilities



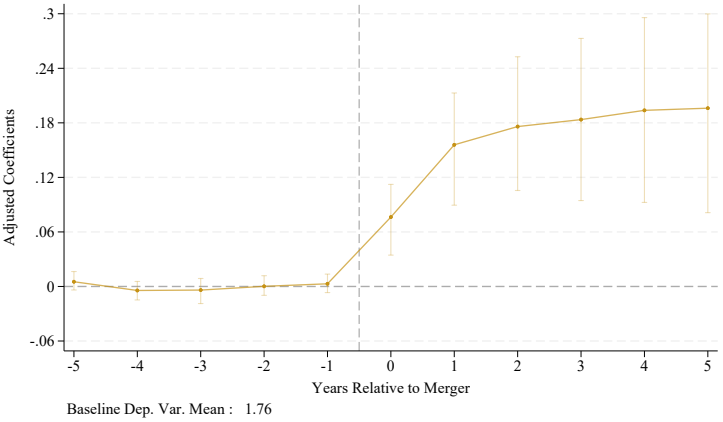
Baseline Dep. Var. Mean : 1.72

(B) Acquired Facilities



Baseline Dep. Var. Mean : 1.03

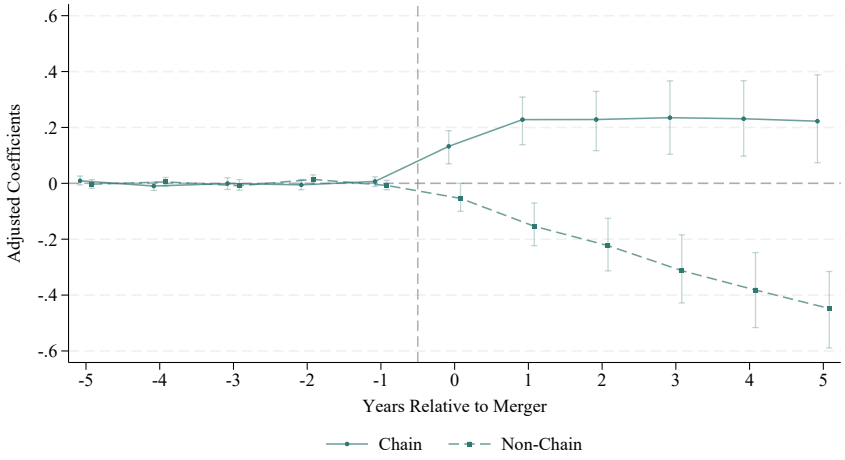
(C) Non-Acquired Facilities



Baseline Dep. Var. Mean : 1.76

Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of all facilities (panel A), acquired facilities (panel B), and non-acquired facilities (panel C), within the subsample of markets that excludes those in the top percentile of pre-merger facility growth, measured as the change in facility count between the fifth and first pre-merger periods, and re-estimate our effects. The unit of analysis is a market, with markets experiencing a merger designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets in Panel A and treated markets in Panels B and C.

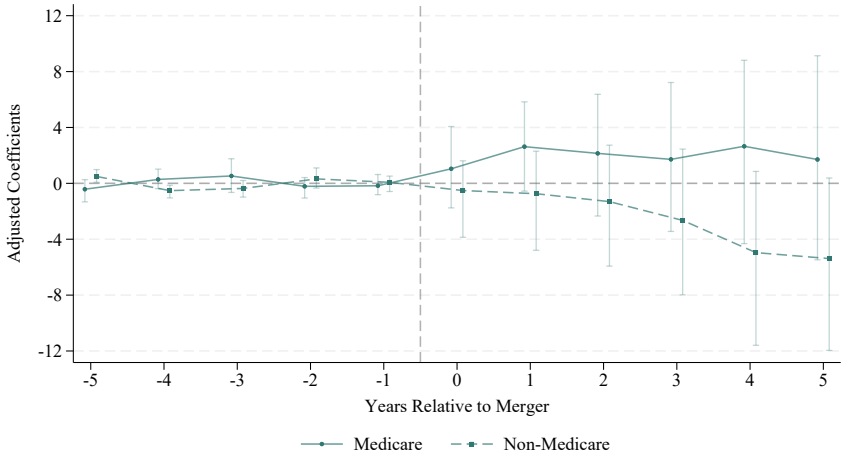
Figure A.7: The Effect of Mergers on the Number of Chain-Owned vs Non-Chain-Owned Facilities



Baseline Dep. Var. Means | Chain: 1.69 | Non-Chain: 1.69

Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of chain-owned (solid line) and non-chain-owned (dashed line) facilities. The unit of analysis is a market, with markets experiencing a merger of the reference type designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time −1, calculated across all markets.

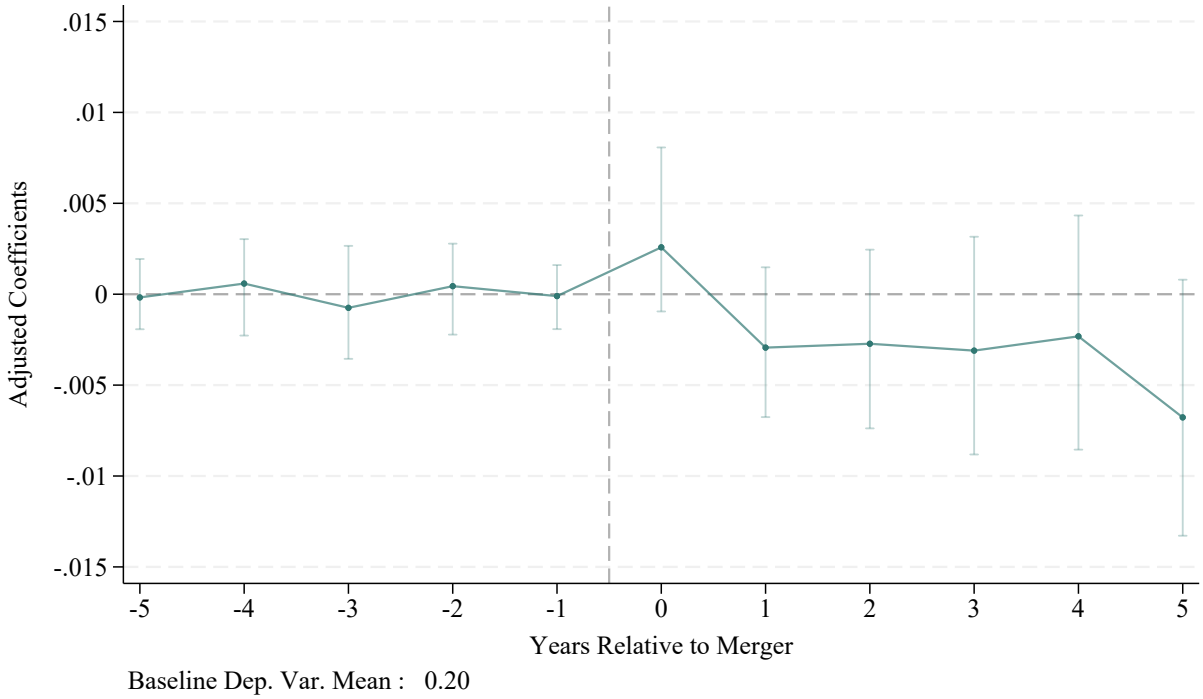
Figure A.8: The Effect of Mergers on the Number of Patients, by Insurance Coverage



Baseline Means | Medicare: 112.14 | Non-Medicare: 9.67

Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of dialysis patients with Medicare coverage (solid line) and non-Medicare coverage (dashed line). The unit of analysis is a market, with markets experiencing a merger of the reference type designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time −1, calculated across all markets.

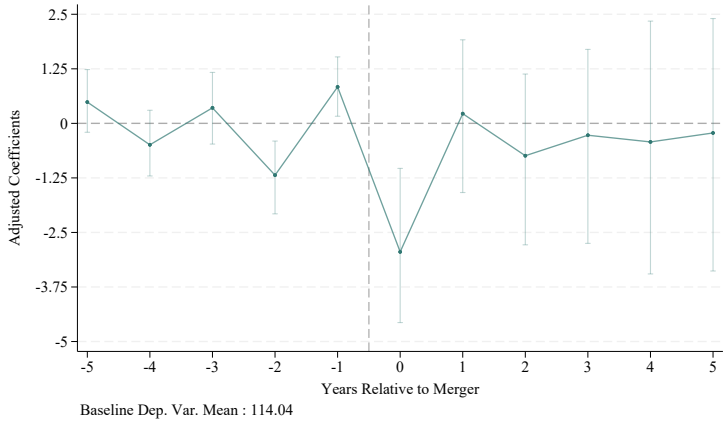
Figure A.9: The Effect of Mergers on Mortality, Excluding Patients with Early-Year Deaths from the Denominator



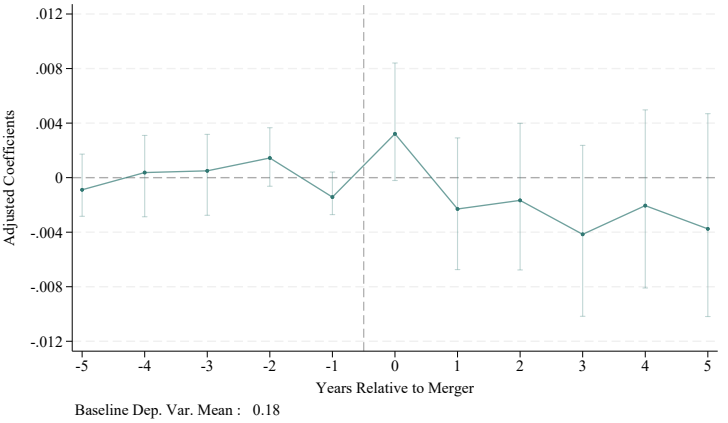
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on mortality. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. The outcome is the total number of deaths divided by the total number of dialysis patients that survived past June of the calendar year at the facility. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

Figure A.10: The Effect of Mergers on Sessions Per Patient, Mortality and Hospitalizations Per Patient, Measured at the Market Level

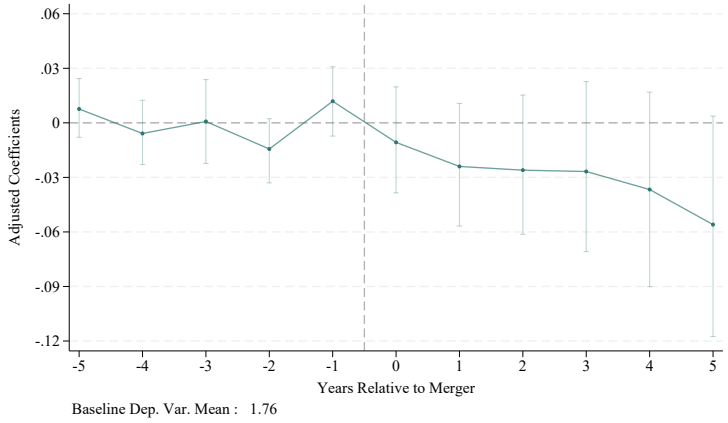
(A) Sessions Per Patient



(B) Mortality



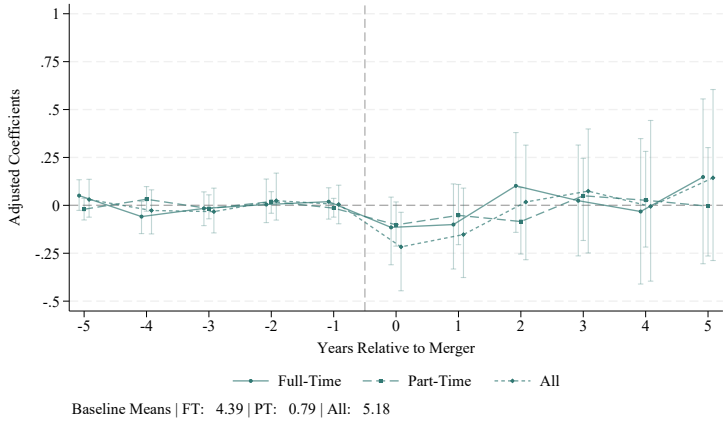
(C) Hospitalizations Per Patient



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on dialysis sessions per patient (panel A), mortality (panel B) and hospitalizations per patient (panel C). The unit of analysis is a market, with markets experiencing a merger designated as treated. The regression includes weights for the market's facility count. The outcome is the total number of sessions (panel A), deaths (panel B) or hospitalizations (panel C) in the market, divided by the total number of dialysis patients in the market. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls that capture both broader market characteristics and the characteristics of the facilities within the market. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets.

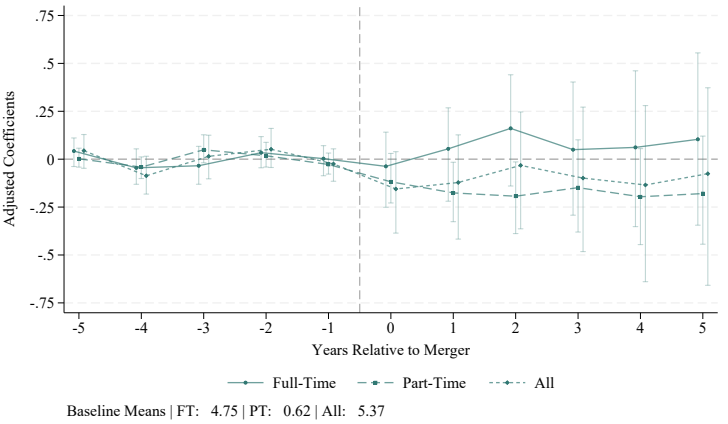
Figure A.11: The Effect of Mergers on Staffing, by Occupation and Full-Time Status

(A) Nurses



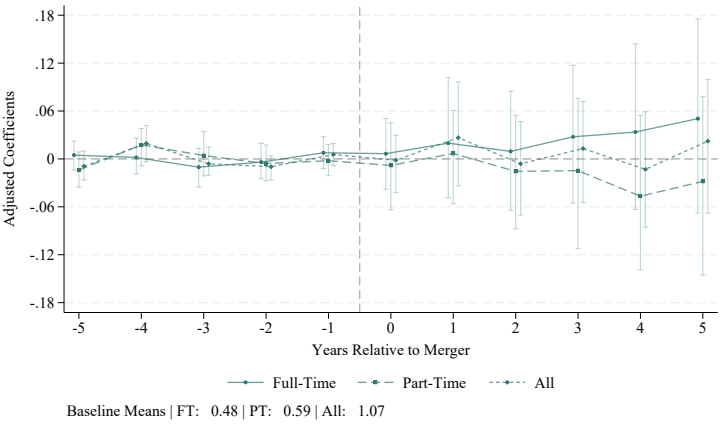
Baseline Means | FT: 4.39 | PT: 0.79 | All: 5.18

(B) Technicians



Baseline Means | FT: 4.75 | PT: 0.62 | All: 5.37

(C) Social Workers

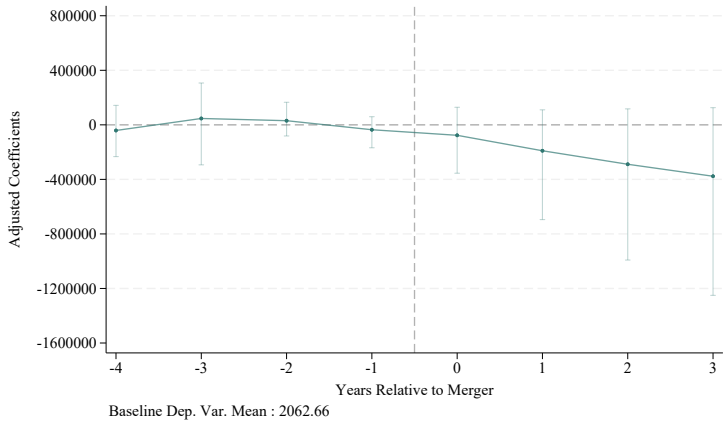


Baseline Means | FT: 0.48 | PT: 0.59 | All: 1.07

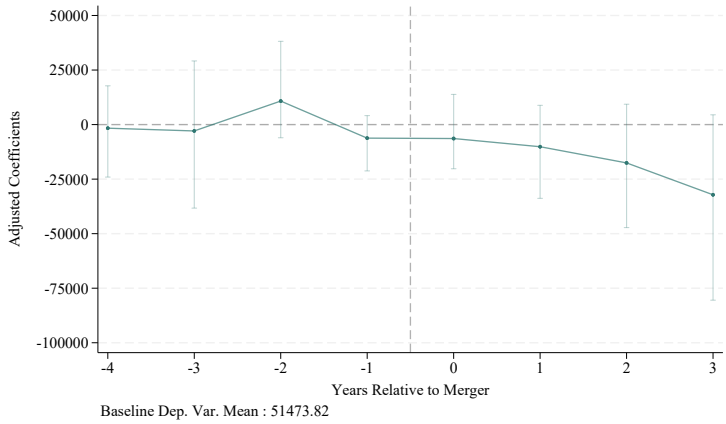
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of nurses (panel A), patient care technicians (panel B) and social workers (panel C). The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. Effects are estimated separately for full-time staff (solid line), part-time staff (dashed line), and total staff (dotted line). The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

Figure A.12: The Effect of Mergers on Anemia Drug Dosing

(A) ESA Dosage (Micrograms)



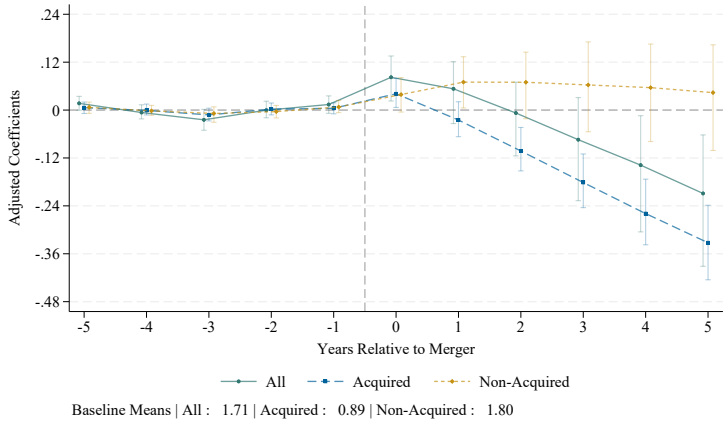
(B) ESA Dosage (Units)



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on ESA (erythropoiesis-stimulating agent) dosing. The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger designated as treated. Panel A presents effects on facility-level ESA dosage measured in micrograms, while Panel B presents effects measured in units. Units are specific to the ESA product and therefore cannot be generically converted to an equivalent microgram dose. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities.

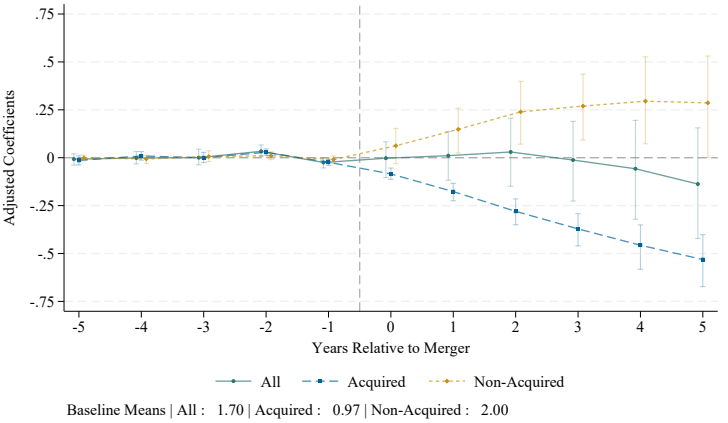
Figure A.13: Heterogeneity in the Effect of Mergers on the Number of Facilities in the Market, by Merger Type

(A) Facilities in Chain-to-Chain Mergers



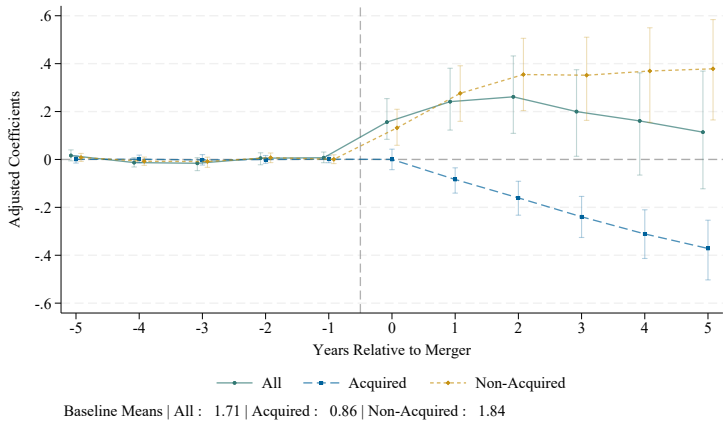
Baseline Means | All : 1.71 | Acquired : 0.89 | Non-Acquired : 1.80

(B) Facilities in Independent-to-Chain Mergers



Baseline Means | All : 1.70 | Acquired : 0.97 | Non-Acquired : 2.00

(C) Facilities in Independent-to-Independent Mergers

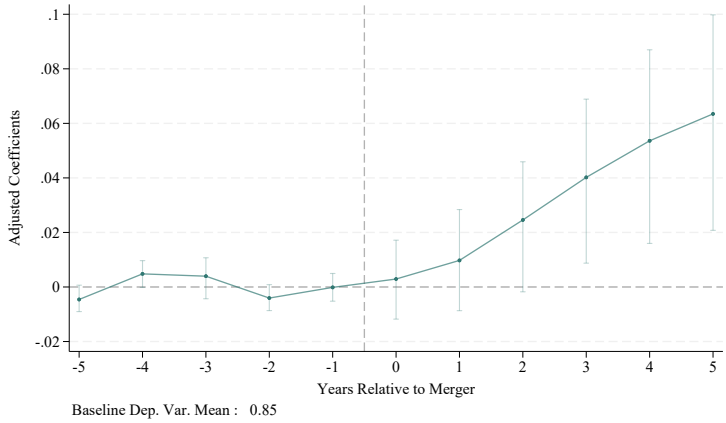


Baseline Means | All : 1.71 | Acquired : 0.86 | Non-Acquired : 1.84

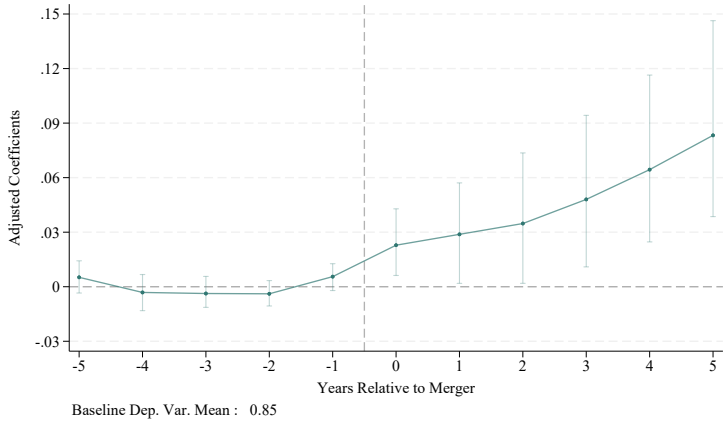
Notes: This figure presents coefficients from equation 4, representing the effect of a merger on the number of facilities separately for chain-to-chain mergers (panel A), independent-to-chain mergers (panel B) and independent-to-independent mergers (panel C). The unit of analysis is a market, with markets experiencing a merger of the reference type designated as treated. Effects are estimated separately for all facilities (green, solid line), acquired facilities (blue, dashed line), and non-acquired facilities (yellow, dotted line). The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets when the outcome is all facilities, and treated markets when the outcome is acquired or non-acquired facilities.

Figure A.14: Heterogeneity in the Effect of Mergers on Market Concentration, by Merger Type

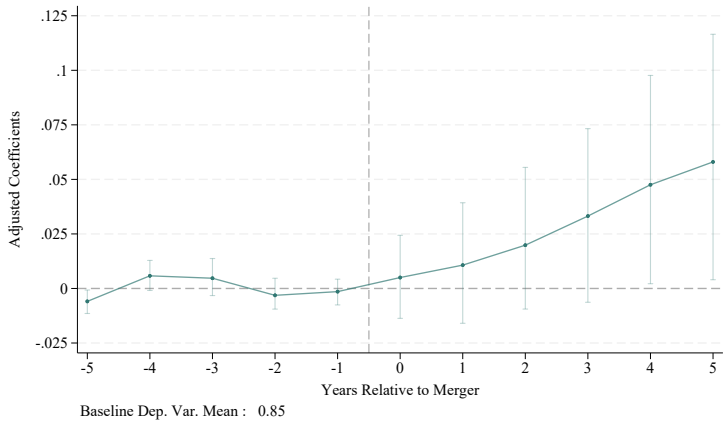
(A) HHI in Chain-to-Chain Mergers



(B) HHI in Independent-to-Chain Mergers

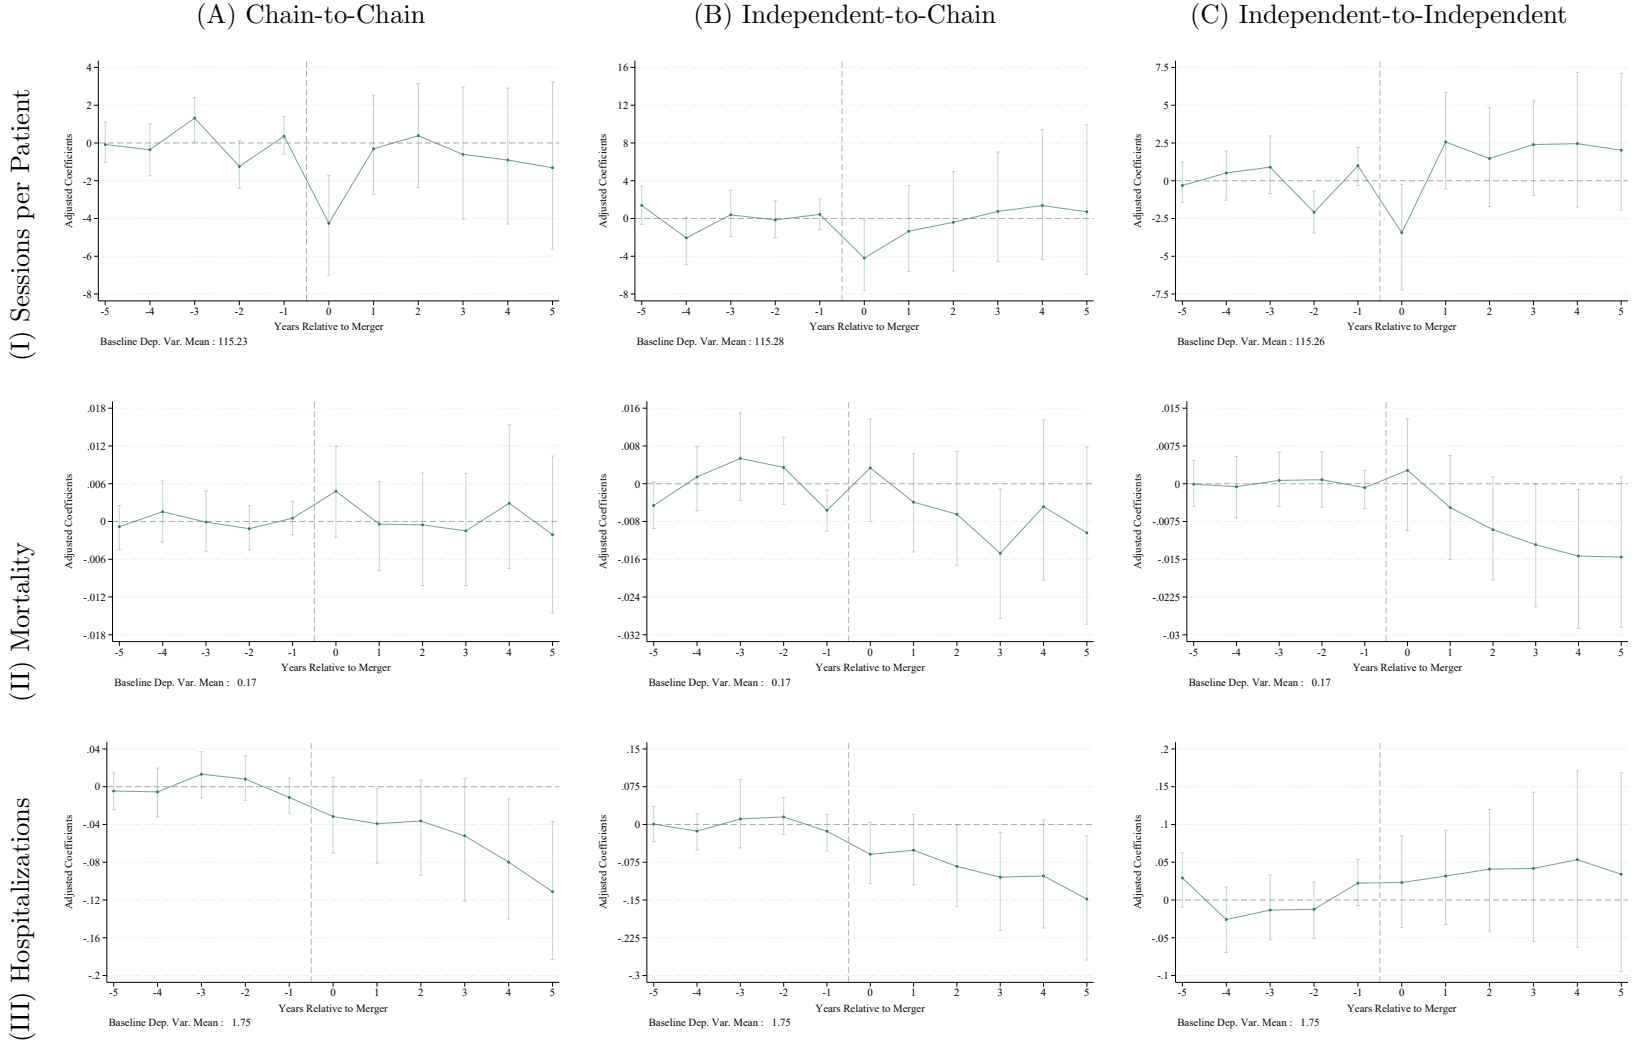


(C) HHI in Independent-to-Independent Mergers



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on market concentration (HHI) for chain-to-chain mergers (Panel A), independent-to-chain mergers (Panel B), and independent-to-independent mergers (Panel C). HHI is measured as the sum of squared facility-level patient shares within a market. The unit of analysis is a market, with markets experiencing a merger of the reference type designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market-level controls. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all markets.

Figure A.15: Heterogeneity in the Effect of Mergers on Sessions per Patient, Mortality, and Hospitalizations, by Merger Type



Notes: This figure presents coefficients from equation 4, representing the effect of a merger on sessions per patient (I), mortality (II), and hospitalizations per patient (III) for chain-to-chain mergers (Panel A), independent-to-chain mergers (Panel B), and independent-to-independent mergers (Panel C). The unit of analysis is a facility, with all facilities (acquired and non-acquired) in markets experiencing a merger of the reference type designated as treated. The vertical dashed line marks the merger event, with event time 0 corresponding to the year of the merger. Vertical lines around each coefficient represent 95% confidence intervals based on bootstrapped standard errors, estimated non-parametrically so that the intervals need not be symmetric. The regression includes market- and facility-level controls. The baseline dependent variable mean is the mean outcome at event time −1, calculated across all treated and control facilities.

B Appendix Tables

Table B.1: Summary of Outcomes, Data Availability, Merger Groups, Patient Coverage and Period Windows

Outcomes	**Data Availability (Years)**	**Merger Cohorts**	**Patient Coverage**	**Pre-Periods (incl. −1)**	**Post-Periods (incl. 0)**
MARKET-LEVEL					
Facilities	1991–2021	1996–2018		5	4 or 5
Stations	1991–2021	1996–2018		5	4 or 5
Patients	1991–2021	1996–2018		5	4 or 5
HHI	1991–2021	1996–2018		5	4 or 5
FACILITY-LEVEL					
Patients Per Station	1991–2021	1996–2018	All	5	4 or 5
Travel Distance	1991–2021	1996–2018	All	5	4 or 5
Staff	2004–2021	2008–2018		4 or 5	
Sessions[a]	2002–2021	2006–2018	Medicare	4 or 5	4 or 5
ESA Prescription[a]	2012–2021	2015–2018	All*	3 or 4	4
Fistula Use[a]	2012–2021	2015–2018	All*	3 or 4	4
Biomarkers[a]	2012–2021	2015–2018	All*	3 or 4	4
Hospitalizations	1991–2021	1996–2018	Medicare	5	4 or 5
ICU Days	1991–2021	1996–2018	Medicare	5	4 or 5
Transfusions	1991–2021	1996–2018	Medicare	5	4 or 5
Deaths	1991–2021	1996–2018	All	5	4 or 5

Notes: This table summarizes outcome variables categorized by unit of analysis (market,facility), showing outcome data availability years, merger cohorts analyzed, and period window specifications.
* Patients are included regardless of insurance coverage type, but the outcome is constructed from CROWNWeb data and represents a subset of the facility's patients treated on a single day each month.

Table B.2: Holm Step-Down Procedure Results for Post-0 and Post-5 Families

Outcome	(1) Coeff. (Adj.)	(2) Bootstrapped SE	(3) Raw p	(4) Rank	(5) Adj. Threshold	(6) Adj. p	(7) Reject?
Panel A: Effect 0 Years Post-Merger							
Sessions	−4.4809	1.1351	0.00008	1	$\alpha/3 = 0.0167$	0.00024	Yes
Death rate	0.00649	0.00286	0.02337	2	$\alpha/2 = 0.025$	0.04674	Yes
Hospitalization rate	−0.01234	0.01621	0.44648	3	$\alpha/1 = 0.05$	0.44648	No
Panel B: Effect 5 Years Post-Merger							
Hospitalization rate	−0.07557	0.02917	0.00957	1	$\alpha/3 = 0.0167$	0.02871	Yes
Death rate	−0.00364	0.00467	0.43570	2	$\alpha/2 = 0.025$	0.87140	No
Sessions	−0.02212	1.78858	0.99014	3	$\alpha/1 = 0.05$	0.99014	No

Notes: This table presents our multiple-hypothesis correction for three key outcomes—sessions per patient, mortality, and hospitalizations per patient—using the Holm step-down procedure. We perform the correction separately for the trend-adjusted merger-year coefficients from our baseline event-study analyses and for the coefficients five years post-merger, reported in Figures 9, 10, and 13. The corresponding coefficients, bootstrapped standard errors, and baseline p-values are reproduced in Columns 1-3. Within each family of hypotheses, we order the outcomes by their raw p-values from smallest to largest, producing a rank ordering denoted in column 4. The Holm-adjusted p-value is calculated as the raw p-value multiplied by $(N - \text{rank} + 1)$, where N is the number of hypotheses in the family (here, 3). We then compare the adjusted p-value with the conventional 5% (α) significance level. Equivalently, the Holm procedure compares each raw p-value with a rank-specific significance threshold of $\alpha/(N - \text{rank} + 1)$. As in our baseline analyses, sessions and mortality are statistically significant in the year of the merger but not five years post-merger, while hospitalizations are not significant in the merger year but remain statistically significant five years post-merger.

Table B.3: The Effect of Mergers on Patient Risk and Demographic Composition

	(1) Pred. Hosp. Rate	(2) Share Female	(3) Share White	(4) Share Black	(5) Share Hispanic
0 years (merger year)	0.0012	-0.0003	0.0016	0.0009	-0.0013
	(0.0014)	(0.0025)	(0.0023)	(0.0020)	(0.0014)
1 year after	0.0025*	0.0002	0.0008	0.0014	-0.0007
	(0.0013)	(0.0023)	(0.0022)	(0.0021)	(0.0014)
2 years after	0.0022*	0.0014	0.0010	0.0015	-0.0013
	(0.0013)	(0.0023)	(0.0021)	(0.0019)	(0.0013)
3 years after	0.0009	-0.0008	-0.0026	0.0044**	-0.0014
	(0.0012)	(0.0020)	(0.0020)	(0.0018)	(0.0012)
4 years after	0.0003	-0.0022	-0.0037*	0.0057***	-0.0003
	(0.0012)	(0.0021)	(0.0020)	(0.0018)	(0.0012)
5 years after	0.0008	-0.0013	-0.0017	0.0037**	-0.0006
	(0.0012)	(0.0020)	(0.0019)	(0.0017)	(0.0012)
5 years before	0.0030*	-0.0010	0.0067***	-0.0012	-0.0020
	(0.0018)	(0.0028)	(0.0026)	(0.0024)	(0.0014)
4 years before	0.0031*	0.0022	0.0058**	-0.0035	-0.0022
	(0.0018)	(0.0028)	(0.0027)	(0.0024)	(0.0015)
3 years before	0.0021	0.0015	0.0013	0.0008	-0.0023
	(0.0016)	(0.0026)	(0.0025)	(0.0023)	(0.0014)
2 years before	0.0010	-0.0011	0.0001	0.0029	-0.0016
	(0.0016)	(0.0025)	(0.0025)	(0.0023)	(0.0014)
N	1,016,777	1,016,777	1,016,777	1,016,777	1,016,777
Baseline Mean	1.7558	0.4506	0.6621	0.2844	0.0695

	(6) Share Age 10–19	(7) Share Age 20–29	(8) Share Age 30–39	(9) Share Age 40–49	(10) Share Age 50–59	(11) Share Age 60–69	(12) Share Age 70–79	(13) Share Age 80–89
0 years (merger year)	-0.0004	-0.0007	0.0008	-0.0019	-0.0005	-0.0061***	0.0054**	0.0018
	(0.0006)	(0.0006)	(0.0012)	(0.0016)	(0.0020)	(0.0020)	(0.0022)	(0.0016)
1 year after	-0.0000	0.0000	0.0010	-0.0020	0.0014	-0.0005	0.0001	-0.0013
	(0.0003)	(0.0006)	(0.0010)	(0.0015)	(0.0019)	(0.0021)	(0.0021)	(0.0016)
2 years after	-0.0003	0.0001	0.0005	-0.0027**	0.0021	-0.0013	0.0011	0.0007
	(0.0002)	(0.0006)	(0.0010)	(0.0014)	(0.0018)	(0.0020)	(0.0020)	(0.0015)
3 years after	-0.0004	-0.0016***	0.0010	-0.0015	0.0021	-0.0005	0.0000	-0.0006
	(0.0004)	(0.0006)	(0.0010)	(0.0013)	(0.0017)	(0.0018)	(0.0019)	(0.0014)
4 years after	0.0001	-0.0008	0.0019*	-0.0010	0.0004	-0.0002	-0.0010	-0.0005
	(0.0003)	(0.0008)	(0.0011)	(0.0013)	(0.0018)	(0.0017)	(0.0018)	(0.0014)
5 years after	0.0005	-0.0004	0.0026**	-0.0021*	0.0004	-0.0003	0.0007	-0.0023*
	(0.0004)	(0.0006)	(0.0011)	(0.0012)	(0.0017)	(0.0017)	(0.0017)	(0.0014)
5 years before	-0.0001	0.0002	0.0004	-0.0003	0.0001	0.0004	-0.0027	0.0021
	(0.0003)	(0.0009)	(0.0014)	(0.0017)	(0.0022)	(0.0025)	(0.0024)	(0.0019)
4 years before	-0.0003	-0.0006	-0.0002	-0.0001	-0.0008	0.0002	-0.0004	0.0023
	(0.0003)	(0.0008)	(0.0013)	(0.0019)	(0.0022)	(0.0023)	(0.0024)	(0.0018)
3 years before	-0.0002	-0.0002	0.0003	-0.0030*	0.0002	-0.0024	0.0037	0.0014
	(0.0003)	(0.0010)	(0.0013)	(0.0016)	(0.0021)	(0.0022)	(0.0025)	(0.0018)
2 years before	-0.0007**	0.0002	-0.0005	-0.0017	0.0022	-0.0013	-0.0008	0.0017
	(0.0003)	(0.0007)	(0.0012)	(0.0016)	(0.0019)	(0.0022)	(0.0022)	(0.0017)
N	1,016,777	1,016,777	1,016,777	1,016,777	1,016,777	1,016,777	1,016,777	1,016,777
Baseline Mean	0.0033	0.0211	0.0530	0.1065	0.1871	0.2541	0.2377	0.1236

Notes: This table presents estimates from equation 1 of mergers' effects on the risk composition of facilities within the market. No covariates are included. The analysis uses the stacked dataset for the mortality analysis, which is the broadest in terms of both time coverage and patient inclusion. The unit of analysis is the facility, and all facilities (acquired and non-acquired) within a market experiencing a merger are designated as treated. Standard errors are clustered at the facility-cohort level. The baseline dependent variable mean is the mean outcome at event time -1, calculated across all treated and control facilities. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Appendix: Background and Inclusion Rationale for Outcomes

In this section, we provide a clinical overview of each category of outcomes and, where required, a rationale for their inclusion in the analysis. We also briefly note how we construct each outcome, with details pertaining to variable construction provided in Appendix J.

C.1 Facility Entry and Closure

The potential for firm entry and incumbent expansion has historically been an important consideration in antitrust enforcement, particularly in assessing whether a merger's effects on market concentration are likely to be offset by competitive responses. Although the 2023 Merger Guidelines have somewhat de-emphasized pro-competitive merger effects, the potential for entry remains relevant to the FTC's assessment of whether a merger poses competitive harm. These dynamics may be especially consequential in health care, where the overall density of facilities in a market, as well as the differential responses of merging and non-merging firms, can have important implications for patient welfare given that dialysis patients seek care locally. Yet the empirical industrial organization literature on mergers' effects on entry and exit is surprisingly thin. To our knowledge, this is the first paper to empirically study the effects of mergers in the dialysis industry or, more broadly, among health care providers, although a small literature examines these dynamics in other industries.[66]

We measure the effect of mergers on market-level closures and entry using the counts of total facilities, acquired facilities, and non-acquired facilities in the market as separate outcomes. The CMS POS data identify ownership changes at the facility level—which we use to identify mergers—but do not reliably identify the parent firm that owns each facility. This distinction between facility-level ownership and firm-level identity has three implications for our empirical analysis. First, we cannot conduct analyses at the firm level or reliably trace the expansion and contraction of individual firms across markets. Second, we can distinguish facilities acquired in the merger from other facilities in the market, but we cannot determine which of the non-acquired facilities are owned by the acquiring firm and which are owned by other firms. Third, the inability to reliably identify parent firms constrains how we measure entry and exit outcomes in control markets.[67]

[66] In the retail and airline industries, mergers appear to have heterogeneous effects on entry and network expansion (Bhattacharya, Illanes, and Stillerman 2023; Hanner et al. 2015; Bougette, Hüschelrath, and Müller 2014). A structural analysis of commercial cellphone carriers finds no evidence of entry following mergers (Caradonna, Miller, and Sheu 2025). In banking, mergers appear to increase closures and reduce capacity (Calzada, Fageda, and Martínez-Santos 2023; Nguyen 2019). Although there is a substantial literature on hospital closures, we are not aware of evidence directly examining whether mergers cause such closures. Interestingly, the industries in which the effects of mergers on entry and closure appear to be nonexistent or heterogeneous are among the least price-regulated settings in this literature, with banking being the most price-regulated. This pattern is consistent with the idea that capacity responses may play a less prominent role in industries where firms have greater scope to compete through prices.

[67] First, this precludes empirical strategies that use markets in which the acquiring firm is present but does not participate in a merger as control markets. Second, the distinction between acquired and non-acquired facilities is defined only in markets experiencing a merger. In practice, therefore, our comparison is between the evolution of acquired and non-acquired facilities in merger markets and the evolution of total facilities in control markets.

C.2 Treatment

Closures require patients to transition to new dialysis facilities, a process that can involve several frictions. Patients may face information constraints or difficulty finding a facility that accepts their insurance, while available facilities may face capacity constraints, nurse shortages, or limited treatment schedules. Travel constraints can further limit patients' options. These frictions may be particularly consequential in the short run, when patients must find alternative facilities and incumbent facilities adjust to additional demand. Whether the effects of these disruptions persist in the longer run depends on the timing and likelihood of entry, the ability of new facilities to absorb residual demand, and the extent to which capacity constraints and other frictions at incumbent facilities ease over time. To our knowledge, however, there is no empirical literature documenting the extent of these transition and capacity frictions following dialysis facility closures.

These disruptions may have important consequences for patients because dialysis treatment duration is closely linked to health outcomes. A robust clinical literature finds that shorter dialysis duration is associated with higher mortality (Hakim et al. 1994; Brunelli et al. 2010; Held et al. 1991; Béguin et al. 2021). Although time spent on the dialysis machine is an important determinant of kidney health, previous studies of dialysis mergers have not, to our knowledge, examined this outcome (Eliason et al. 2020; Wollmann 2020). We therefore study how mergers affect the number of dialysis sessions per patient in every facility exposed to a merger, providing direct evidence on how the disruptions associated with facility closures translate into changes in treatment. We also examine the effects of mergers on patients' mean travel distance to their facility and on the ratio of patients to dialysis stations, which captures capacity constraints and facility utilization.

C.3 Market Concentration

We measure market concentration using the Herfindahl-Hirschman Index (HHI), defined as "the sum of the squares of the market shares" (U.S. Department of Justice and the Federal Trade Commission 2023) of firms in the relevant market.[68] We define markets at the county level, following Wollmann (2020), who finds that more than 90% of dialysis patients receive care within their county of residence. Following their approach, we further subdivide three large counties—Cook County, IL; Los Angeles County, CA; and San Diego County, CA—into distinct markets based on observed patterns of patient flows. We define market shares based on the number of patients receiving treatment at each facility.

As discussed in Section C.1, limitations in our ability to identify the ultimate parent firm preclude us from constructing firm-level market shares. We therefore construct the HHI using facility-level market shares. Facility-level concentration will generally increase when facilities exit the market, since exit removes a competitor and redistributes its patients among the remaining facilities; this mechanical effect is strongest when the exiting facility is small, whereas exit by a very large facility can instead disperse share among survivors and lower concentration. Concentration can also change independent of facility exit, however, as

[68] The delineation of the relevant market is itself a substantive antitrust judgment, for which enforcement agencies provide various guidelines in Section 4.3 of the 2023 Merger Guidelines (U.S. Department of Justice and the Federal Trade Commission 2023).

patients reallocate across incumbent facilities, thereby shifting market shares even when the set of facilities in the market is unchanged.

C.4 Facility Decisions

We consider three facility decisions that serve as intermediate, directly observable indicators of quality under facility control, before quality is ultimately reflected in patient health outcomes.

First, we consider the volume of staff employed at dialysis facilities. Mergers could affect staffing through three channels. First, facility closures may create short-run staffing shortages, mirroring the frictions that can disrupt patient treatment—staff at closing facilities may face difficulties finding employment at other facilities. Second, mergers may increase monopsony power in local labor markets, allowing firms to reduce staffing wages through greater bargaining power and thereby compressing the supply of labor in the long run. Third, even holding staffing wages fixed, mergers may lead firms to reduce staffing as a longer-run cost-cutting measure. This channel may be particularly relevant in a price-regulated industry, where firms cannot increase prices following consolidation and therefore rely on cost reductions to increase profitability.

The worker categories we include are nurses, patient care technicians, and social workers. A dialysis patient's care team is led by a nephrologist, who sets the overall care plan and prescribes medications.[69] Day-to-day treatment is managed by nurses, including registered nurses, licensed practical nurses, and advanced practice nurses, who form the primary staffing input in outpatient dialysis clinics and typically receive specialized dialysis training. Patient care technicians prepare the dialysate, set up patients, monitor vital signs, and maintain equipment. Social workers play a key role in transplant eligibility; CMS requires a psychosocial evaluation by a social worker to process transplant candidacy. In our primary analysis, we consider total staffing across all staff types and employment statuses. In supplementary analyses, we separately examine each staff type and its full- and part-time components, treating shifts from full- to part-time employment as a potential reduction in quality.

Second, we study anemia drug utilization at facilities. Anemia is highly prevalent in individuals with kidney disease, primarily due to reduced renal synthesis of erythropoietin (EPO), the hormone responsible for initiating the production of red blood cells. During dialysis, anemia is treated with the prescription of erythropoietin stimulating agents (ESAs), which are artificial replacements of EPO, tailored to patient needs and administered intravenously. When anemia management is inadequate, patients typically need red blood cell transfusions to maintain red blood cell volume. CMS is sensitive to the relevance of anemia management in measuring facility quality and monitors it using a facility's standardized transfusion ratio, which captures the number of annual transfusion events for patients at a facility relative to the expected number given the facility's case-mix. In this analysis, we study facility processes relevant to anemia management, namely the share of patients pre-

[69] The USRDS facility data do not provide information about nephrologist staffing.

scribed an ESA, and the average monthly dosage of ESAs.[70] Commonly prescribed ESAs include Epoetin Alpha, Epoetin Beta, and Darbepoetin Alpha which differ slightly in molecular structure and absorption time, but not in treatment efficacy or safety if correctly dosed (Loughnan, Ali, and Abeygunasekara 2011). So, we don't consider the choice of the ESA agent as a quality decision but pool doses of all ESAs measured in the same units in the analysis of the monthly ESA dosage.[71]

Third, we consider the utilization of arteriovenous fistulas as the modality of vascular access. To receive dialysis, blood must flow in and out of the body at high rates. Central venous catheters are flexible tubes inserted into a large vein providing this high flow rate in the neck, chest, or groin and are used mainly for short-term or emergency dialysis because they carry the highest infection and complication risks. Two other vascular methods combine the property of arteries to provide the high-pressure flow needed for dialysis, and the property of veins to be accessible and puncture easily. Of these, the first method is arteriovenous fistulas which are direct surgical connections between an artery and a vein, usually in the arm, that enlarge the vein and increase blood flow over time, making them the most durable access with the lowest infection risk. The second method is arteriovenous grafts which use a synthetic tube to connect an artery and a vein when a patient's own veins are too small or damaged for a fistula. Grafts last longer than catheters but still have higher rates of infection and clotting than fistulas. Because of their lowest infection risk, fistulas are considered the gold standard for dialysis access, but require outpatient surgery and maturation time. The Medicare QIP rewards facilities for maintaining a higher proportion of patients with fistula access. In this analysis, we examine each facility's share of patients using fistulas relative to other access types.

C.5 Biomarkers

In our analysis of health outcomes we distinguish between *biomarkers* and *clinical endpoints*. We borrow this categorization from Atkinson Jr et al. (2001) who define biomarkers as "[a] characteristic that is objectively measured and evaluated as an indicator of normal biological processes, pathogenic processes, or pharmacologic responses to a therapeutic intervention" and clinical endpoints as "[a] characteristic or variable that reflects how a patient feels, functions, or survives". In our setting, biomarkers are blood or other physiological measurements taken monthly in facilities, and clinical endpoints are hospitalization, days in the ICU, inpatient transfusions and mortality. Clinical endpoints have been the central focus of much of the literature on dialysis mergers and are important indicators of care quality on their own, but biomarkers can provide additional information on patients' clinical status and serve as more immediate targets of changes in care. Moreover, correlations between biomarkers and endpoints provide insight into the effectiveness of the Medicare QIP, which incorporates a subset of biomarkers into its quality ratings, though their predictive validity remains contested (Vanholder, Glorieux, and Eloot 2015). Biomarker data are collected by

[70] The effect of Medicare payment policy on the use of ESAs has been a focus of work by Eliason et al. (2022) and is not in itself a focus of this paper.

[71] ESAs are prescribed either in micrograms or in units. Because the orders of magnitude differ across these dosage measures, we consider them separately, pooling across ESA agents within each measure.

facilities monthly for the set of patients treated on a single day each month and submitted to CMS through an online reporting system called CROWNWeb.

The first biomarker of interest to our analysis is dialysis adequacy, an indicator of how effectively the dialysis procedure cleans the blood, measured by the single-pool Kt/V. Kt/V captures the amount of urea or other solutes cleared from the body from before to after the dialysis session.[72] QIP guidelines specify a minimum target Kt/V of 1.2 and evaluate facilities based on the proportion of reported patient-month observations that meet or exceed this threshold. Accordingly, our outcome measure is the share of a facility's patients achieving a "good Kt/V," defined as Kt/V of at least 1.2 mg/dL.

The second biomarker of interest is hypercalcemia. Impaired kidney function in patients with ESRD compromises the body's ability to maintain electrolyte balance, often leading to abnormal blood calcium levels. Hypercalcemia, or elevated calcium levels, poses a significant cardiovascular risk, as excess calcium can disrupt the electrical impulses that regulate cardiac muscle function and contribute to vascular calcification. Accordingly, QIP tracks patients with serum calcium levels exceeding 10.2 mg/dL.[73] In this paper, we study the share of a facility's patients with hypercalcemia as defined by the QIP.

The third biomarker of interest is the ultrafiltration rate (UFR), which measures the rate at which fluid is removed from a patient's blood during a dialysis session.[74] A high UFR is undesirable because rapid fluid removal can cause complications such as hypotension, cramping, and cardiac stress. Conversely, a low UFR may also be harmful, increasing cardiovascular risks. Unlike Kt/V or hypercalcemia, UFR is a reporting measure in the QIP; facilities are required only to report the components used to calculate UFR, and the program does not specify optimal thresholds. Nonetheless, clinical evidence suggests that UFR values exceeding 10 mL/kg/hr are associated with higher cardiovascular risk (Flythe, Kimmel, and Brunelli 2011; Kim et al. 2018). In this analysis, we compute each patient's UFR using CROWNWeb data elements and examine how mergers affect the facility-level annual average UFR, which captures variation in both excessively high and low fluid removal rates.

We also examine three additional biomarkers collected by CMS through CROWNWeb that are either not included or only inconsistently included in the QIP's portfolio of facility quality measures—hemoglobin, albumin, and iron saturation.[75] Specifically, we analyze the share of patients with low hemoglobin levels (<10 g/dL), the annual facility average of serum albumin (g/dL), and the annual facility average of pre-dialysis iron saturation (%). Among these, low hemoglobin and iron saturation are markers of anemia risk, while low albumin levels (<3.5 g/dL) indicate advanced kidney failure and are a known predictor of cardiovascular risk (Manolis et al. 2022). Both QIP and non-QIP biomarkers are available

[72] Specifically, Kt/V is measured as a patient's urea clearance, multiplied by the duration of the dialysis treatment in minutes, divided by the volume of total body water. Another commonly used metric of dialysis adequacy is the Urea Reduction Ratio (URR). The URR is calculated as the percentage change in the blood urea nitrogen from before to after dialysis. Kt/V and URR are mathematically related.

[73] At the same time, hypocalcemia defined as serum calcium levels below 8.5 mg/dL (UCSF Hospitalist Handbook N.d.), is also undesirable as it increases muscle dysfunction and cardiovascular risks. We report the effects of mergers on hypocalcemia in Appendix Figure A.2, finding no merger-induced changes in this outcome.

[74] UFR (mL/kg/hr) $= \frac{\text{Net fluid removal (mL)}}{\text{Dialysis time (min)} \times \text{Post-dialysis weight (kg)}} \times 60$. Net fluid removal (mL) = [Pre-dialysis weight (kg) − Post-dialysis weight (kg)] × 1000.

[75] Hemoglobin was included as a QIP measure from 2012 to 2014 but was subsequently removed.

from 2012 onward, following the rollout of CROWNWeb, but they offer complementary insights. Effects of mergers on QIP biomarkers may capture policy interactions, that is, how mergers influence facilities' performance on measures explicitly monitored by the QIP. In contrast, effects on non-QIP biomarkers more directly reflect the intrinsic impact of mergers on patient care, as these outcomes are not subject to QIP reporting or incentives.

C.6 Clinical Endpoints

As outlined in Section C.5, clinical endpoints in this analysis are hospitalization, ICU days, inpatient blood transfusions and mortality. Hospitalizations signal shifts from outpatient to inpatient care, serve as indicators of suboptimal outpatient quality, and also represent substantial ESRD-related costs that merit consideration in Medicare's assessment of the dialysis program. In the QIP, facilities are graded on the standardized hospitalization ratio, which is the ratio of total inpatient admissions for patients at a facility divided by the expected admissions given the case-mix. In this analysis, we consider a facility's hospitalizations per patient, measured from inpatient Medicare claims. Second, we study the number of days in the ICU per patient, which is informative about the severity of hospitalizations.

The third endpoint of interest is the frequency of blood transfusions, measured in the QIP by a standardized transfusion ratio and, in this study, by the number of inpatient events comprising red blood cell transfusions per patient. A higher per-patient blood transfusion rate is indicative of inadequate management of anemia. We evaluate transfusion propensity alongside ESA prescribing behavior, since changes in ESA prescribing are ambiguous in isolation but provide a clearer indication of deteriorating care when they are accompanied by increases in inpatient blood transfusions.

The final clinical endpoint of interest is all-cause mortality. We measure mortality as the facility's annual instantaneous mortality rate, which is the ratio of total patients who die in any year to the total patients treated. Both mortality and the number of ICU days are not within the QIP's portfolio of facility-quality measures.

D Appendix: Are Mergers in Dialysis Endogenous?

We examine which pre-merger observables distinguish markets that experience a merger from those that do not. We use our most comprehensive facility-level dataset, which combines CMS Provider of Services (POS) data, USRDS facility data, and outcomes observed for patients regardless of insurance status. For each merger cohort, we construct a dataset containing all facilities in treated markets—those experiencing a merger in the cohort year—and facilities in control markets, defined as markets with no merger in that year. We restrict the analysis to the five years preceding the cohort-specific merger year and measure each predictor for a facility as its mean over this pre-merger period. We then estimate, separately for each cohort year, the linear probability model

$$TreatA_{is} = \alpha + X'_{is}\beta + \gamma_s + \varepsilon_{is}, \tag{1}$$

where $TreatA_{is}$ indicates whether facility i is located in a market that experiences a merger in the cohort year, X_{is} is a vector of county- and facility-level characteristics averaged over the five pre-merger years, and γ_s denotes state fixed effects. The vector X_{is} includes measures of market size, county demographics, facility characteristics, and patient composition. Standard errors are robust to heteroskedasticity.

Figure 2 summarizes the coefficients on these characteristics across merger cohorts. We highlight coefficients that are statistically significant at the 10% level or higher, with color indicating the direction of the coefficient and shading intensity indicating its magnitude. Overall, the figure reveals relatively sparse and inconsistent predictive power. Most observables do not significantly predict merger exposure in most years, and no characteristic other than market size exhibits a stable relationship with merger probability throughout the 1996–2018 period.

The number of facilities in a market is the clearest exception. It is the most consistently predictive characteristic, with an additional facility associated with an increase in the probability of merger exposure ranging from approximately 1 to 13 percentage points across cohorts. Even for this measure, however, the magnitude and statistical significance vary across years. Other market and county characteristics exhibit substantially less consistent relationships with merger exposure. County population, for example, is associated with a lower probability of merger exposure in a number of years but has a modestly positive relationship in others. Gender composition is predictive in some cohorts, whereas age composition and median income exhibit little systematic relationship with merger exposure. Facility age is negatively predictive in several of the earlier cohorts but does not exhibit a comparable relationship in later years.

We find similarly little evidence of systematic selection on facility characteristics that capture patient composition and risk. In particular, our measure of exogenous patient risk, the predicted hospitalization rate, which we describe in more detail in Section 4, does not consistently predict merger exposure. More broadly, the limited predictive power of these characteristics suggests that, beyond market size as measured by the number of facilities, there is relatively little observable information that consistently distinguishes markets that experience a merger from those that do not.

The explanatory power of the models reinforces this interpretation. Adjusted R^2 ranges from approximately 0.07 to 0.42 across cohorts, with values clustering around 0.20. Thus,

even this rich set of pre-merger county and facility characteristics explains only a limited share of the cross-sectional variation in merger exposure. The highest adjusted R^2, approximately 42%, occurs for the 2015 cohort; in most cohorts, the models explain substantially less.

These findings have a two-sided implication for our empirical strategy. The limited systematic differences in pre-merger observables provide some reassurance that treated and control markets are broadly comparable along the dimensions we can observe. At the same time, the substantial variation in merger exposure that remains unexplained by these observables provides little basis for constructing a more narrowly tailored control group. We therefore do not condition the control group on characteristics that have weak or unstable predictive power for merger exposure, as doing so could impose restrictions that are not strongly supported by the data.

To gain further insight into the factors that firms themselves cite when describing merger decisions, Appendix E provides an AI-assisted qualitative analysis of publicly disclosed statements from joint press releases, SEC filings, and regulatory agencies for a set of 10 dialysis-provider mergers. These statements identify recurring themes, including quality and patient-care improvement, strategic and geographic positioning, market expansion, economies of scale and cost efficiencies, and shareholder value. Although they provide insight into firms' stated motivations, they generally do not identify the specific economic mechanisms or market characteristics that made particular transactions attractive.

The broad nature of these stated rationales makes it difficult to map them directly onto the observable selection patterns in Figure 2. For example, an emphasis on geographic expansion, market positioning, or scale might suggest that mergers should be more likely in larger or geographically attractive markets, while an emphasis on quality and patient care might suggest selection on observable patient or facility characteristics. Similarly, stated motivations related to cost efficiencies might imply selection into markets or facilities with greater potential for economies of scale. Yet these broad objectives need not correspond closely to the specific observable characteristics available in our data. Public statements describe firms' stated goals, but generally provide little information about the particular economic incentives, competitive conditions, or market-specific considerations that make a given transaction attractive.

E Appendix: Disclosed Rationales for Mergers Among Dialysis Providers

This appendix provides an AI-assisted qualitative analysis compiling publicly disclosed statements from joint press releases, SEC filings, and regulatory agencies that provide insight into the motivations underlying particular dialysis-provider mergers. The compilation is intended to provide an overview of the motivations underlying dialysis mergers broadly, rather than to characterize the rationale for any particular merger in our sample. Importantly, we cannot establish whether any of the mergers in our sample correspond to the transactions described in these public statements. We used an LLM, Claude Sonnet 5, to assist in identifying and describing relevant publicly available statements. Our prompt was, verbatim:

> "Quietly search the web for contemporaneous public announcements and other first-hand sources documenting why dialysis firms merged or were acquired. Focus on statements made by company executives or other company officials (e.g., press releases, investor materials, SEC filings, earnings calls, or interviews), rather than retrospective commentary or third-party accounts.
>
> For each merger, provide the names of the merging/acquiring and target firms; the merger year; whether each firm was a small independent operator or a large dialysis chain at the time of the transaction; one or two short, direct quotes from company officials that capture the stated motivation for the transaction; and the sources, with full citations in '.bib' format.
>
> Please prioritize transactions for which there is clear first-hand evidence of the parties' stated rationale. After compiling the transactions, conduct a qualitative theme analysis of the stated merger rationales. Identify the major themes that emerge from the executive quotes and other first-hand materials, and rank them from most frequent to least frequent. For each theme, briefly describe what it encompasses and indicate how many transactions exhibit the theme. Distinguish, where possible, between motivations such as geographic/patient-base expansion, economies of scale or cost efficiencies, access to capital or resources, complementary capabilities, quality/patient-care improvements, strategic positioning, and other recurring rationales. The goal is to characterize the range and relative prevalence of stated motivations for consolidation in the dialysis industry, rather than to infer firms' underlying motives."

We used an LLM to identify relevant publicly available statements in order to minimize researcher discretion in selecting which transactions and statements to highlight. In this sense, the compilation provides a more systematic and reproducible account of publicly stated merger motivations that is less conditioned on our own selection of transactions or preferred interpretations. The transaction descriptions and thematic analysis presented below are our own rephrasing and synthesis of the LLM-assisted search and coding results; they are not reproduced verbatim from the LLM's output. The direct quotations from company executives, in contrast, are reproduced from the underlying public sources and are separately cited.

Merger 1: Gambro (Incentive AB) acquired Vivra Renal Care, Inc. in 1997. Both firms were large dialysis chains at the time. Gambro was already a major global renal-care and dialysis-products company and Vivra was the second-largest U.S. dialysis provider operating around 262 centers and serving 15,800 patients.

> *"This transaction represents a major strategic move for Incentive and continues our strong thrust into the high growth health care sector... The acquisition of Vivra Renal Care is a logical next step in the development of Incentive Group and demonstrates our strong commitment to shareholder value creation."*
> *Mikael Lilius, President and CEO of Incentive AB (Incentive AB and Gambro AB 1997)*

> *"Vivra Renal Care is a well-managed company that fits well with Gambro. VRC will strengthen our position as a global leader in renal care... this transaction will generate significant operating synergies, while allowing Gambro to continue to provide a service of the highest quality."*
> *Berthold Lindqvist, President and CEO of Gambro (Incentive AB and Gambro AB 1997)*

Merger 2: Total Renal Care Holdings, Inc. (TRC) acquired Renal Treatment Centers, Inc. (RTC) in 1997. Both firms were large, publicly traded dialysis chains. TRC operated 194 U.S. centers serving approximately 14,900 patients, while RTC operated 164 centers serving approximately 11,000 patients.

> *"We will now have the critical mass necessary to ensure our position as the leading independent provider of dialysis services in the U.S. and the resources to rapidly expand into highly fragmented international markets... Renal Treatment Centers provides a strong platform for us to accelerate our aggressive, yet disciplined, growth strategy both in the U.S. and internationally."*
> *Victor M.G. Chaltiel, Chairman, President and CEO of TRC (Total Renal Care Holdings, Inc. and Renal Treatment Centers, Inc. 1997)*

> *"We are excited to be joining forces with TRC. Its quality and outcomes programs, focused managed-care strategy, extensive renal-related support programs, and exceptionally strong track record will ensure future growth."*
> *Robert L. Mayer, Jr., President and CEO of RTC (Total Renal Care Holdings, Inc. and Renal Treatment Centers, Inc. 1997)*

> *"We have identified $8.2 million and $22.8 million of annual cost savings and synergies to be achieved in 1998 and 1999, respectively, and we expect the transaction to be immediately accretive to earnings."*
> *Victor M.G. Chaltiel, Chairman, President and CEO of TRC (Total Renal Care Holdings, Inc. and Renal Treatment Centers, Inc. 1997)*

Merger 3: DaVita Inc. acquired Gambro Healthcare U.S. in 2005. Both firms were large dialysis chains. DaVita operated over 700 outpatient centers serving over 57,000, while Gambro Healthcare added roughly 520 centers and about 39,000 patients after required divestitures, in a deal valued at roughly $3.1 billion.

"The DaVita Village is excited to welcome the good people of Gambro. We continue to believe that this combination will significantly increase our ability to meet the needs of all of our constituencies over the long-term."
Kent Thiry, Chairman and CEO of DaVita (DaVita Inc. 2005)

Merger 4: Fresenius Medical Care AG acquired Renal Care Group, Inc. (RCG) in 2005 announced. Both firms were large chains. Fresenius Medical Care was already the world's largest dialysis products and services company, and Renal Care Group operated more than 425 clinics serving over 30,400 patients, with revenue of about $1.35 billion.

"This acquisition solidifies Fresenius Medical Care's position as the leader in dialysis services in the U.S. With its service network, Renal Care Group provides an excellent strategic and geographic fit to Fresenius Medical Care's operations in the world's largest dialysis market. We are convinced that this acquisition will enhance the Company's growth prospects and profitability... [and] will provide opportunities to successfully leverage the Company's cost leadership position through its combination with Renal Care Group."
Ben Lipps, CEO of Fresenius Medical Care (Fresenius Medical Care AG 2005)

"Both companies... have three key assets in common: a strong commitment to continuous quality improvement, dedicated and highly-motivated people and experienced and visionary management teams. This combination clearly gives us an opportunity to create a superior platform for innovation in the delivery of services and products and will further enhance the lives of patients entrusted to our care."
Gary Brukardt, President and CEO of Renal Care Group (Fresenius Medical Care AG 2005)

"This makes us unique and strengthens our position as the world's largest renal therapy company. As we combine the best of both companies, we do so for the benefit of all – patients, employees, physicians, customers and shareholders."
Fresenius Medical Care at acquisition close (Fresenius Medical Care AG & Co. KGaA 2006)

Merger 5: DaVita Inc. acquired DSI Renal, Inc. in 2011. DaVita was a large dialysis chain and DSI Renal was a mid-sized regional operator with 106 centers serving approximately 8,000 patients, and around 360 million in annualized revenue, acquired for approximately $690 million.

"We look forward to joining the dedicated teammates and physicians associated with DSI in continuing to provide quality care to ESRD patients. This acquisition introduces us to several new geographies and makes us a more effective competitor in selected areas. Through this acquisition we will be able to bring the broader line of DaVita chronic kidney disease services to DSI patients. These services will be beneficial for patients, physicians, payors, and taxpayers in providing more effective care and helping to reduce costs to the health care system."
Kent Thiry, Chairman and CEO of DaVita (DaVita Inc. 2011)

"DaVita has always shared DSI's commitment to improving the quality of life of our patients and to providing high quality care. We believe that combining with DaVita in today's complex healthcare environment will help us take our operations to the next level and greatly benefit our patients and staff."
Leif Murphy, President and CEO of DSI (DaVita Inc. 2011)

Merger 6: U.S. Renal Care, Inc. (USRC) acquired DSI Renal in 2015. This was the second company to operate under the DSI Renal name—DaVita had acquired the original DSI Renal in 2011 but was subsequently required by the FTC to divest dozens of clinics as a condition of the transaction. The divested clinics were reorganized as a new company, also named DSI Renal, under private equity ownership (Nephrology News and Issues 2015). U.S. Renal Care subsequently acquired the new DSI Renal. USRC was a large, growing dialysis chain and DSI Renal was a mid-sized regional operator.

"We are very excited about the opportunity to partner with DSI and its employees and physicians. DSI's commitment to providing exceptional care to patients, coupled with its focus on physician partnerships, is very consistent with USRC's approach. We look forward to working together to continue to provide high quality care to patients."
Chris Brengard, CEO of U.S. Renal Care (U.S. Renal Care, Inc. N.d.)

Merger 7: U.S. Renal Care, Inc. (USRC) acquired Ambulatory Services of America (ASA), operating as Innovative Dialysis Systems in 2013. USRC was a growing national chain, while ASA was a smaller regional independent operator.

"Innovative Dialysis has built a strong reputation for providing high quality, patient-centered care. We are pleased to partner with their physicians and employees and believe that our shared commitment to quality and our joint-venture approach will allow us to continue to build on that success."
Chris Brengard, CEO of U.S. Renal Care (U.S. Renal Care, Inc. and Ambulatory Services of America, Inc. N.d.)

"We are excited to partner with U.S. Renal Care. Their focus on physician partnerships and commitment to high-quality care make them an excellent fit for Innovative Dialysis and our physician partners."
Timothy Martin, CEO of Ambulatory Services of America (U.S. Renal Care, Inc. and Ambulatory Services of America, Inc. N.d.)

Merger 8: Innovative Renal Care (Nautic Partners) acquired American Renal Associates Holdings, Inc. (ARA) in 2021. ARA was a large, publicly traded national dialysis chain servign arounf 17,000 patients and operating a physician joint-venture model. Innovative Renal Care was a newly formed, smaller platform company backed by private equity.

"At ARA, we have created a unique platform for delivering superior quality care to patients by partnering with physicians around the country. This transaction recognizes the value of the Company and delivers a meaningful premium to shareholders."
Joe Carlucci, Chair and CEO of ARA (American Renal Associates Holdings, Inc. 2020)

"As the healthcare industry continues to evolve, recent innovation throughout kidney and renal care has allowed for dramatic and exciting improvements in care delivery and patient experience. ARA's significant infrastructure and geographic footprint provides a unique opportunity to advance the development of a comprehensive renal care model for the patients and communities we serve."
Nick Mendez, President and CEO of Innovative Renal Care (American Renal Associates Holdings, Inc. 2020)

Merger 9: Fresenius Medical Care, InterWell Health, and Cricket Health underwent a three-way merger in 2022. Fresenius was a large multinational dialysis company; InterWell Health and Cricket Health were smaller, specialized value-based-care technology and clinical-services companies, rather than traditional dialysis clinic chains.

"I am enthusiastic to see this merger being completed which will support our sustainable profitable growth. InterWell Health will be the premier value-based kidney care provider in the U.S., combining and leveraging innovative new tools, vast experience, and a deep nephrologist network. By further expanding into the strategically important chronic kidney disease stage-3-to-5-market, we will be able to support even more patients throughout the Renal Care Continuum, including earlier interventions to prevent disease prior to kidney failure. This is good for patients, the payers, and our shareholders."
*Helen Giza, Deputy CEO of Fresenius Medical Care (Fresenius Medical Care AG & Co. KGaA 2022*b,a*)*

"We are excited to see the transformative impact we will have on patient outcomes and kidney care delivery in the U.S. as we bring together three leading value-based kidney care entities... InterWell Health is well positioned to improve patient outcomes with fewer hospital admissions, slow disease progression, increase transplant referrals and rates."
*Bill Valle, CEO of Fresenius Medical Care North America (Fresenius Medical Care AG & Co. KGaA 2022*a*)*

Merger 10: DaVita HealthCare Partners Inc. acquired Renal Ventures Limited, LLC in 2015. DaVita was a large national chain, while Renal Ventures was a small regional independent operator of 36 dialysis clinics in six states, acquired for $415 million.

"We are excited to work with Renal Ventures' patients, employees and physicians... We look forward to serving Renal Ventures' patients with industry-leading outcomes and comprehensive care."
Javier Rodriguez, CEO of DaVita Kidney Care (DaVita HealthCare Partners Inc. 2015)

We then used the LLM to qualitatively code the direct executive quotations across the transactions and identify recurring themes in the stated merger rationales. The most consistently invoked rationale was *quality and patient-care improvement*. Executives on both sides of nearly every deal framed the combination as a means of improving clinical quality, patient outcomes, or continuity of care. This theme appears consistently across press releases regardless of time period, deal size, or whether the target was a large chain or a smaller operator. At the same time, many of these statements are sufficiently broad that they provide limited economic clarity about the specific mechanism through which the merger is expected

to improve quality or generate value. In many cases, the stated rationale identifies a desired outcome rather than a clearly articulated economic mechanism linking the combination of the firms to that outcome.

The second most common theme was *strategic and geographic positioning, market expansion, and leadership*. Many acquirers described the transaction as a means of extending market leadership, filling geographic gaps, or strengthening their competitive position. This theme is especially prominent in the largest chain-to-chain combinations, including Gambro/Vivra, TRC/RTC, DaVita/Gambro, and Fresenius/RCG, where executives emphasized the scale and strategic fit of the combined organizations. DaVita/DSI similarly described the acquisition as expanding its geographic reach and strengthening its competitive position. A more specific form of strategic positioning appears in the Fresenius Health Partners/InterWell Health/Cricket Health merger, where executives emphasized expanding beyond dialysis for end-stage kidney disease into earlier-stage chronic kidney disease management. This forward-looking market-expansion framing is largely unique to the more recent, value-based-care-oriented transaction in the sample.

Related to, but distinct from, market positioning, several statements frame transactions as a means of sustaining growth or creating a platform for future expansion. This framing appears in large chain-to-chain combinations such as TRC/RTC and Incentive/Gambro, as well as in the more recent ARA/IRC transaction, where the parties emphasized leveraging ARA's existing infrastructure and geographic footprint to develop a more comprehensive renal-care model. Together, these statements suggest that strategic positioning encompasses not only expansion of the existing dialysis footprint but also efforts to create platforms for broader growth within renal and kidney care.

The third most prevalent theme was *economies of scale and cost efficiencies*. Explicit references to cost savings, synergies, or scale economies appeared in several transactions, most notably TRC/RTC, Gambro/Vivra, and Fresenius/RCG. TRC/RTC provided the most concrete statement of expected cost savings, while Gambro/Vivra and Fresenius/RCG emphasized operating synergies and the ability to leverage the acquirer's existing cost advantages. DaVita/DSI likewise connected the acquisition to an objective of reducing costs to the health care system. This theme is more common in the largest chain-to-chain transactions than in smaller acquisitions of independent or regional operators, where explicit cost-synergy rationales are largely absent from the quoted material.

A smaller number of releases explicitly reference *shareholder value, earnings accretion, or premium consideration* as a rationale, rather than merely as a transaction term. Incentive/Gambro, TRC/RTC, and ARA/IRC each included statements emphasizing shareholder value, expected earnings effects, or the premium paid to shareholders. This theme is more characteristic of publicly traded companies addressing investors and shareholders directly.

Among transactions where the target was explicitly a smaller or independent operator, target executives sometimes framed joining a larger company as providing access to broader resources or a stronger organizational platform. This is particularly apparent in DSI's description of the DaVita transaction as enabling the company to take its operations to the next level. Similar language appears in other smaller-target transactions, although the quoted statements more often emphasize shared commitments to quality and patient care than specific mechanisms through which additional resources would create value.

Overall, the stated rationales emphasize improvements in patient care and clinical capabilities alongside expansion and strategic positioning, while explicit cost savings and shareholder-value considerations appear less consistently and are concentrated in particular types of transactions. More broadly, however, the executive statements often provide limited economic clarity about the underlying motivations for consolidation: they frequently describe broad objectives such as improving quality, expanding access, or strengthening the organization without specifying the economic mechanisms through which the merger is expected to achieve those objectives. Thus, the statements are informative about how firms publicly characterize their transactions, but they provide less evidence about the underlying economic incentives driving the mergers.

F Appendix: Heterogeneous Effects of Merger Types

Our primary analyses shed light on the overall effects of mergers in dialysis. Here, we undertake the analysis of key outcomes separately for three merger subtypes.[76] The first subtype is mergers in which the target and acquirer are both chains (chain-to-chain). The second subtype is mergers in which the target is independent but the acquirer is a chain (independent-to-chain).[77] The third subtype is one in which both entities are independent firms (independent-to-independent).

Figure A.13 presents heterogeneity in merger effects on facility closures and entry. Merger-induced closures are widespread, but the magnitude of closures as well as the degree of offsetting entry by non-merging facilities varies substantially by merger type. Overall, chain-to-chain mergers lead to net reductions in the number of facilities within a market, independent-to-chain mergers have no net effect, and independent-to-independent mergers yield a modest net increase that stabilizes within five years after the merger.

Figure A.14 shows that all three merger subtypes produce comparable increases in market concentration. The effects are largest for independent-to-chain mergers (panel B), which increase the HHI by just over 7.5 percentage points five years after the merger. The corresponding increases are modestly smaller for chain-to-chain mergers (panel A) and independent-to-independent mergers (panel C), at approximately 6 percentage points. [78]

Figure A.15 shows the effects of mergers across merger subtypes on our key outcomes—sessions per patient, mortality, and hospitalizations per patient. Across all three merger types, the evidence points to disruptions in treatment in the merger year followed by stabilization in the long run. A spike in merger-year mortality is also visible across merger types, although disaggregating by merger type reduces the precision of the short-run mortality estimates. Importantly, the magnitude of both the merger-year treatment disruption and the increase in mortality is largest for chain-to-chain mergers, followed by independent-to-chain mergers and then independent-to-independent mergers. This ordering closely mirrors the magnitude of the associated declines in net capacity, consistent with the relationship between capacity reductions, treatment disruptions, and mortality documented earlier.

We find that long-term gains in hospitalization appear most strongly in chain-to-chain and independent-to-chain consolidations. At the same time, independent-to-independent mergers exhibit a pronounced long-term decline in all-cause mortality, a pattern that was obscured in the aggregated all-merger sample.

[76] The analysis of mergers by type involves the estimation of equation 4 where the treated units are from markets experiencing a merger of the reference type, and the control units are from markets experiencing no merger.

[77] This is the merger type under consideration in Eliason et al. (2020).

[78] For independent-to-chain mergers, the effect is estimated with greater variance.

G Appendix: Proofs

This appendix proves the results in Section 6, states and proves the aggregation and sparsity results referred to in Section 6.5, and then presents the policy extensions discussed in Section 8. The subgames other than closure with entry are characterized in Appendix H. Throughout, $c_A^0 \equiv c_A(n_A^0, \underline{q})$ and $c_T^0 \equiv c_T(n_T^0, \underline{q}_T)$ denote pre-merger variable costs per patient, and $N \equiv n_A^0 + n_T^0$.

Two bounds used throughout. First, every feasible quality lies in $[\underline{q}_T, \bar{q}]$, since $\underline{q}_T < \underline{q}$ (Assumption 1); because u is increasing, for any two feasible qualities

$$u(q_j) - u(q_i) \;\le\; u(\bar{q}) - u(\underline{q}_T) \;=\; \Delta. \tag{1}$$

Second, $b + \bar{\varepsilon} \ge b + \underline{\varepsilon} > \Delta > 0$ by condition (6), so Assumption 2 implies

$$V_0 \;<\; u(\underline{q}_T) - b - \bar{\varepsilon} \;<\; u(\underline{q}_T). \tag{2}$$

Hence a patient who stays at an open facility (utility at least $u(\underline{q}_T)$) or moves to an active facility of quality at least $\underline{q}_T$ (utility at least $u(\underline{q}_T) - b - \bar{\varepsilon}$) is always better off than if she leaves the market.

Proof of Proposition 1

Proof. **Step 1 (No patient switches or leaves).** Consider a patient k at open facility $i \in \{A, T\}$ and the other facility j. By (3) and (1), her net gain from switching satisfies

$$\big[u(q_j) - u(q_i)\big] - s_k^{ij} \;\le\; \Delta - \big(b + \varepsilon_k^{ij}\big) \;\le\; \Delta - \big(b + \underline{\varepsilon}\big) \;<\; 0,$$

where the last inequality is condition (6). So she does not switch, for any feasible qualities. She does not leave the market either: staying yields $u(q_i) \ge u(\underline{q}_T) > V_0$ by (2). Each facility's volume is therefore n_i^0, independent of (q_A, q_T).

Step 2 (Each facility chooses its floor). Given fixed volume, facility i's profit is $\pi_i(q_i) = p\, n_i^0 - F_i - c_i(n_i^0, q_i)\, n_i^0$, so

$$\frac{\partial \pi_i}{\partial q_i} = -\frac{\partial c_i(n_i^0, q_i)}{\partial q_i}\, n_i^0 \;<\; 0$$

by (1). Profit is strictly decreasing on $[\underline{q}_i, \bar{q}]$, so the unique maximizer is $q_i = \underline{q}_i$: $q_A^0 = \underline{q}$ and $q_T^0 = \underline{q}_T < \underline{q}$. □

Derivation of the First-Order Conditions (9) and (10)

Write $x \equiv u(q_A) - u(q_E)$, so that $r = G(x)$, $n_A^C = n_A^0 + G(x)\, n_T^0$, and $n_E^C = \big[1 - G(x)\big]\, n_T^0$.

Acquirer. Its operating profit is $\Pi_A^C = p\, n_A^C - c_A(n_A^C, q_A)\, n_A^C$. Since q_A affects Π_A^C both directly and through n_A^C, the chain rule gives

$$\frac{\partial \Pi_A^C}{\partial q_A} = p\, \frac{\partial n_A^C}{\partial q_A} - \left(c_A\, \frac{\partial n_A^C}{\partial q_A} + \frac{\partial c_A}{\partial n}\, \frac{\partial n_A^C}{\partial q_A}\, n_A^C + \frac{\partial c_A}{\partial q_A}\, n_A^C \right).$$

Collecting the terms in $\partial n_A^C / \partial q_A$ and using $\mathrm{MC}_A = c_A + (\partial c_A / \partial n)\, n_A^C$,

$$\frac{\partial \Pi_A^C}{\partial q_A} = \big(p - \mathrm{MC}_A\big)\, \frac{\partial n_A^C}{\partial q_A} - \frac{\partial c_A}{\partial q_A}\, n_A^C. \tag{3}$$

Because $\partial x / \partial q_A = u'(q_A)$,

$$\frac{\partial n_A^C}{\partial q_A} = n_T^0\, g(x)\, u'(q_A).$$

Substituting into (3) and setting it to zero yields (9).

Entrant. By the same steps, with $\mathrm{MC}_E = c_E + (\partial c_E / \partial n)\, n_E^C$,

$$\frac{\partial \Pi_E^C}{\partial q_E} = \big(p - \mathrm{MC}_E\big)\, \frac{\partial n_E^C}{\partial q_E} - \frac{\partial c_E}{\partial q_E}\, n_E^C. \tag{4}$$

Now $\partial x / \partial q_E = -u'(q_E)$, so

$$\frac{\partial n_E^C}{\partial q_E} = -\, n_T^0\, g(x)\, \frac{\partial x}{\partial q_E} = -\, n_T^0\, g(x)\, \big[-\, u'(q_E) \big] = n_T^0\, g\big(u(q_A) - u(q_E)\big)\, u'(q_E) \;>\; 0.$$

The two negative signs—one because the entrant's volume is $1 - G$, one because q_E enters x negatively—cancel. The density is therefore evaluated at $x = u(q_A) - u(q_E)$ for both firms; it is not reversed for the entrant. Substituting into (4) and setting it to zero yields (10).[79]

Proof of Lemma 1

Proof. **The floor is not a best response to the floor.** Evaluate (3) at $q_A = q_E = \underline{q}$. Then $x = 0$, $n_A^C = n_A^0 + G(0)\, n_T^0$, and

$$\left.\frac{\partial \Pi_A^C}{\partial q_A}\right|_{q_A = q_E = \underline{q}} = \big(p - \mathrm{MC}_A\big)\, n_T^0\, g(0)\, u'(\underline{q}) - \left.\frac{\partial c_A}{\partial q_A}\right|_{\underline{q}} \big[n_A^0 + G(0)\, n_T^0 \big] \;>\; 0$$

by the first inequality in (11). Since the right-derivative of Π_A^C at the lower boundary is strictly positive, some $q_A > \underline{q}$ yields strictly higher profit, so $q_A = \underline{q}$ is not a best response to $q_E = \underline{q}$. Evaluating (4) at the same point, where $n_E^C = \big[1 - G(0)\big]\, n_T^0$, the second inequality in (11) gives $\partial \Pi_E^C / \partial q_E > 0$, so $q_E = \underline{q}$ is not a best response to $q_A = \underline{q}$. Hence $(\underline{q}, \underline{q})$ is not a Nash equilibrium.

No equilibrium with one firm at the floor. Suppose $q_E^* = \underline{q}$. Then q_A^* is a best response to $q_E = \underline{q}$, which by the argument above lies strictly above $\underline{q}$. For $q_E^* = \underline{q}$ to be a

[79] If D_k is symmetric around zero—as it is when ε_k^{TA} and ε_k^{TE} are i.i.d.—then $g(x) = g(-x)$, and writing $g\big(u(q_E) - u(q_A)\big)$ in (10) is equivalent. Without symmetry, the argument must be $u(q_A) - u(q_E)$, given the definition $D_k \equiv \varepsilon_k^{TA} - \varepsilon_k^{TE}$.

best response to q_A^*, the right-derivative of Π_E^C at $q_E = \underline{q}$ must be non-positive. We require the second inequality in (11) to hold at $q_E = \underline{q}$ for every $q_A \in [\underline{q}, \bar{q}]$—that is, with $g(0)$ and $G(0)$ replaced by $g\big(u(q_A) - u(\underline{q})\big)$ and $G\big(u(q_A) - u(\underline{q})\big)$—and likewise for the acquirer; at $q_A = q_E$, these reduce to (11).[80] Then the right-derivative is strictly positive, a contradiction. The same argument rules out $q_A^* = \underline{q}$. Therefore $q_A^*(1,1) > \underline{q}$ and $q_E^*(1,1) > \underline{q}$. □

Proof of Proposition 2

We first record the acquirer's operating profit in each subgame (characterized in Appendix H):

$$\pi_A^*(0,0) = \big(p - c_A^0\big)\, n_A^0 + \big(p - c_T^0\big)\, n_T^0, \tag{5}$$

$$\pi_A^*(1,0) = \big[p - c_A(N, \underline{q})\big]\, N, \tag{6}$$

$$\pi_A^*(1,1) = \big[p - c_A(n_A^C, q_A^*)\big]\, n_A^C, \qquad n_A^C = n_A^0 + r(q_A^*, q_E^*)\, n_T^0. \tag{7}$$

Lemma 4 (Closure pays when entry is blocked)**.** *Under Assumption 1, $\pi_A^*(1,0) > \pi_A^*(0,0)$.*

Proof. Using $N = n_A^0 + n_T^0$, write $p\,N = p\,n_A^0 + p\,n_T^0$ and subtract (5) from (6):

$$\pi_A^*(1,0) - \pi_A^*(0,0) = \big[c_A^0 - c_A(N, \underline{q})\big]\, n_A^0 + \big[c_T^0 - c_A(N, \underline{q})\big]\, n_T^0.$$

Because $N > n_A^0$, scale economies give $c_A(N, \underline{q}) < c_A(n_A^0, \underline{q}) = c_A^0$, so the first bracket is strictly positive. By Assumption 1, $c_A^0 \le c_T^0$, so $c_A(N, \underline{q}) < c_A^0 \le c_T^0$ and the second bracket is strictly positive as well. Since $n_A^0, n_T^0 > 0$, the difference is strictly positive. □

The difference in Lemma 4 is the variable-cost saving from treating all N patients in one larger facility rather than in two separate ones: the acquirer's original patients become cheaper to treat because volume rises, and the target's patients are treated at the acquirer's lower cost.

Lemma 5 (The acquirer closes T rather than A)**.** *Suppose entry is blocked. Under Assumption 1, closing T yields the acquirer at least as much profit as closing A.*

Proof. If the acquirer closes T, all N patients are treated at A at quality $\underline{q}$ (Lemma 3), and its profit is $\big[p - c_A(N, \underline{q})\big]N - F_A$. If it instead closes A, the same argument applies with the roles reversed. T's patients stay, since $u(q_T) \ge u(\underline{q}_T) > V_0$ by (2). Each of A's displaced patients obtains $u(q_T) - b - \varepsilon_k^{AT} \ge u(\underline{q}_T) - b - \bar{\varepsilon} > V_0$ at T, so she moves to T for every draw of ε_k^{AT}. T's volume is then fixed at N, its profit is strictly decreasing in q_T, so T sets $q_T = \underline{q}_T$, and the acquirer's profit is $\big[p - c_T(N, \underline{q}_T)\big]N - F_T$. Closing T is therefore at least as profitable if and only if

$$\big[p - c_A(N, \underline{q})\big]N - F_A \;\ge\; \big[p - c_T(N, \underline{q}_T)\big]N - F_T \quad\Longleftrightarrow\quad F_A + c_A(N, \underline{q})\, N \;\le\; F_T + c_T(N, \underline{q}_T)\, N,$$

which is the third condition in (2). □

[80] This uniform version is what "at the floor" means in Lemma 1: each firm's margin on contested patients exceeds its cost of quality when it sits at the floor, whatever its rival does.

Proof of Proposition 2. We use the payoff table in Section 6.4.3 and the tie-breaking conventions in Section 6.3.

Step 1 (Entrant's best response). If $c = 1$, entering yields $\pi_E^*(1) - F_E = \overline{F}_E - F_E$ and staying out yields 0, so E enters if and only if $F_E \leq \overline{F}_E$. If $c = 0$, entering yields $-F_E < 0$ by Lemma 2, so E stays out. That $\overline{F}_E > 0$ follows from Lemma 1: at its equilibrium quality, E serves $n_E^C > 0$ patients at a positive margin.[81]

Step 2 (Acquirer's best response).

Case A: $F_E \leq \overline{F}_E$. Closing is followed by entry and yields $\pi_A^*(1,1) - F_A$; keeping T open is followed by no entry and yields $\pi_A^*(0,0) - F_A - F_T$. The acquirer closes if and only if

$$\pi_A^*(1,1) - F_A > \pi_A^*(0,0) - F_A - F_T \quad \Longleftrightarrow \quad F_T > \pi_A^*(0,0) - \pi_A^*(1,1),$$

with ties resolved in favor of keeping T open. This gives $(1,1)$ if $F_T > \pi_A^*(0,0) - \pi_A^*(1,1)$ and $(0,0)$ otherwise. Note that the relevant alternative to closing is the status quo $(0,0)$, not $(0,1)$: by Step 1, keeping T open always deters entry.

Case B: $F_E > \overline{F}_E$. The entrant stays out whatever the acquirer does. Closing yields $\pi_A^*(1,0) - F_A$ and keeping T open yields $\pi_A^*(0,0) - F_A - F_T$. The acquirer closes if and only if $F_T > \pi_A^*(0,0) - \pi_A^*(1,0)$. By Lemma 4 the right-hand side is negative, and $F_T > 0$, so the acquirer always closes: the outcome is $(1,0)$.

Step 3 (Exhaustiveness and uniqueness). Cases A and B partition $F_E > 0$, and the two sub-cases of Case A partition $F_T > 0$. Every (F_E, F_T) thus falls in exactly one region of (12). Within each region, the entrant's and acquirer's best responses are unique given the tie-breaking conventions, and the Stage-2 equilibrium of each subgame is unique, so the subgame-perfect equilibrium is unique. The subgame $(0,1)$ is never reached (Corollary 2).

Step 4 (Quality). In $(1,1)$, $q_A^*, q_E^* > \underline{q}$ by Lemma 1. In $(1,0)$, $q_A^* = \underline{q}$ by Lemma 3. In $(0,0)$, $(q_A^*, q_T^*) = (\underline{q}, \underline{q}_T)$ by Proposition 1. □

Aggregation Across Markets

Consider a population of merger markets ω, each satisfying the conditions of Proposition 2, with primitives—including $F_{E,\omega}$ and $F_{T,\omega}$—drawn from a joint distribution. Let c_ω and e_ω denote equilibrium closure and entry in market ω.

Proposition 3 (Net capacity)**.** *The expected changes in the number of acquired, non-acquired, and total facilities are*

$$\begin{gathered} \mathbb{E}[-c_\omega] = -\Pr(c_\omega = 1), \qquad \mathbb{E}[e_\omega] = \Pr\big((c_\omega, e_\omega) = (1,1)\big), \\ \mathbb{E}[e_\omega - c_\omega] = -\Pr\big(F_{E,\omega} > \overline{F}_{E,\omega}\big) \;\leq\; 0, \end{gathered} \tag{8}$$

with strict inequality whenever entry is blocked in a positive share of markets.

To measure market quality, let $\bar{Q} \equiv \frac{1}{N}\sum_i n_i q_i$ denote patient-weighted quality, summing over active facilities. By Proposition 1, pre-merger quality is $\bar{Q}^0 = (n_A^0 \underline{q} + n_T^0 \underline{q}_T)/N$.

[81] Strictly, the entrant earns at least the operating profit from choosing $q_E = \underline{q}$, which is positive whenever $1 - G\big(u(q_A^*) - u(\underline{q})\big) > 0$ and $p > c_E$; we assume the entrant's average margin is positive at the relevant volumes.

Corollary 1 (Market quality across equilibria)**.** *Relative to the pre-merger market,*

$$(1,1): \quad \bar{Q} - \bar{Q}^0 = \underbrace{\frac{n_T^0}{N}\big(\underline{q} - \underline{q}_T\big)}_{\text{exit of the target}} + \underbrace{\frac{n_A^C}{N}\big(q_A^* - \underline{q}\big)}_{\text{acquirer upgrade}} + \underbrace{\frac{n_E^C}{N}\big(q_E^* - \underline{q}\big)}_{\text{entrant above floor}} \; > \; 0, \tag{9}$$

$$(1,0): \quad \bar{Q} - \bar{Q}^0 = \frac{n_T^0}{N}\big(\underline{q} - \underline{q}_T\big) \; > \; 0, \tag{10}$$

$$(0,0): \quad \bar{Q} - \bar{Q}^0 = 0. \tag{11}$$

Hence, across markets, expected quality rises whenever the target closes with positive probability, even when $\mathbb{E}[e_\omega - c_\omega] < 0$*.*

In markets where closure is not followed by entry, the gain in (10) is purely compositional: no facility changes its quality, but the target's patients move from a facility providing $\underline{q}_T$ to one providing $\underline{q}$. Within-firm improvement by the acquirer and the entrant's higher quality, the last two terms of (9), arise only where entry follows closure, because only then must the acquirer compete for displaced patients.

Proof of Proposition 3

Proof. In market ω, closure removes one acquired facility and entry adds one non-acquired facility, so the change in the total number of facilities is $e_\omega - c_\omega$. By Proposition 2, each market ends in one of three equilibria:

$$(1,1): \; e_\omega - c_\omega = 0; \qquad (0,0): \; e_\omega - c_\omega = 0; \qquad (1,0): \; e_\omega - c_\omega = -1.$$

The first two expectations follow from the definitions of c_ω and e_ω, using that $e_\omega = 1$ only in $(1,1)$. For the third,

$$\begin{aligned}\mathbb{E}[e_\omega - c_\omega] &= 0 \cdot \Pr\big((1,1)\big) + 0 \cdot \Pr\big((0,0)\big) + (-1) \cdot \Pr\big((1,0)\big) \\ &= -\Pr\big((1,0)\big) = -\Pr\big(F_{E,\omega} > \overline{F}_{E,\omega}\big) \; \leq \; 0,\end{aligned}$$

where the last equality uses that $(1,0)$ arises if and only if $F_{E,\omega} > \overline{F}_{E,\omega}$. The inequality is strict whenever entry is blocked in a positive share of markets. □

Proof of Corollary 1

Proof. **Market size is fixed.** In every subgame, every patient is treated at an active facility in the market. Patients at open facilities stay (Proposition 1, Step 1). A displaced patient who moves to an active facility m obtains $u(q_m) - b - \varepsilon_k^{Tm} \geq u(\underline{q}_T) - b - \bar{\varepsilon} > V_0$, since $q_m \geq \underline{q} > \underline{q}_T$; she prefers this to leaving. Hence the number of patients is N in every subgame, and in $(1,1)$, $n_A^C + n_E^C = n_A^0 + r\, n_T^0 + (1-r)\, n_T^0 = N$.

Status quo $(0,0)$**.** Volumes and qualities are unchanged, so $\bar{Q} = \bar{Q}^0$.

Closure without entry $(1,0)$**.** All N patients are treated at A at quality $\underline{q}$ (Lemma 3), so $\bar{Q} = \underline{q}$ and

$$\bar{Q} - \bar{Q}^0 = \frac{N\underline{q} - n_A^0\underline{q} - n_T^0\underline{q}_T}{N} = \frac{n_T^0}{N}\big(\underline{q} - \underline{q}_T\big) \;>\; 0,$$

since $n_T^0 > 0$ and $\underline{q} > \underline{q}_T$.

Closure with entry $(1,1)$**.** Market quality is $\bar{Q} = \big(n_A^C q_A^* + n_E^C q_E^*\big)/N$. Adding and subtracting $\underline{q}$ at each facility and using $n_A^C + n_E^C = N$,

$$n_A^C q_A^* + n_E^C q_E^* = N\underline{q} + n_A^C\big(q_A^* - \underline{q}\big) + n_E^C\big(q_E^* - \underline{q}\big).$$

Dividing by N, subtracting $\bar{Q}^0$, and using the $(1,0)$ computation for $\underline{q} - \bar{Q}^0$ gives (9). The first term is strictly positive as above. The second and third are strictly positive because $q_A^*, q_E^* > \underline{q}$ (Lemma 1), $n_A^C \geq n_A^0 > 0$, and $n_E^C > 0$; the last holds because an entrant with no patients would earn no operating profit, whereas entry in equilibrium requires $\pi_E^*(1) \geq F_E > 0$.

Across markets. Each market's change is zero in $(0,0)$ and strictly positive in $(1,1)$ and $(1,0)$, so the expected change across markets is strictly positive whenever $\Pr(c_\omega = 1) > 0$, regardless of how often entry follows closure. □

Market Sparsity

Closures should be more likely where there are fewer facilities in the market before the merger. In less dense markets, the acquirer is, on average, the natural destination for the target's patients relative to other options, so it captures a larger share of the displaced pool. We formalize this with a sparsity parameter θ: higher sparsity (fewer competing facilities) strictly increases the acquirer's capture rate of displaced patients, $\partial r(q_A, q_E; \theta)/\partial\theta > 0$.

Proposition 4 (Density and catchment overlap)**.** *As sparsity θ increases, the operating profit the acquirer forgoes by closing the target when entry follows, $\pi_A^*(0,0) - \pi_A^*(1,1;\theta)$, strictly decreases, expanding the set of parameters under which closure occurs.*

Intuitively, when density is low, the acquirer acts as the dominant catchment for the target's displaced patients. Because it recaptures a larger share of them, the operating profit it forgoes by closing the target shrinks, and the fixed-cost saving F_T is more likely to justify closure. We delineate markets at three facilities in the data because markets with two or fewer facilities plausibly contain only the target and the acquirer, an extreme case of full catchment capture.

Proof of Proposition 4

Proof. Let $\Phi(\theta) \equiv \pi_A^*(0,0) - \pi_A^*(1,1;\theta)$ denote the operating profit forgone by closing when entry follows, where

$$\pi_A^*(1,1;\theta) = \Big[p - c_A\big(n_A^C(\theta), q_A^*\big)\Big]\, n_A^C(\theta), \qquad n_A^C(\theta) = n_A^0 + r(q_A^*, q_E^*; \theta)\, n_T^0.$$

Taking the total derivative with respect to θ,

$$\frac{d\pi_A^*(1,1;\theta)}{d\theta} = \frac{\partial \pi_A^*}{\partial \theta} + \frac{\partial \pi_A^*}{\partial q_A}\frac{\partial q_A^*}{\partial \theta}.$$

By the Envelope Theorem, $\partial \pi_A^*/\partial q_A = 0$ at the acquirer's optimum, so re-optimization of its own quality has no first-order effect. Treating the rival's strategic response as second-order, we evaluate the direct effect holding qualities fixed. Using $\partial n_A^C/\partial \theta = n_T^0\, \partial r/\partial \theta$ and the steps leading to (3),

$$\frac{\partial \pi_A^*(1,1;\theta)}{\partial \theta} = \big(p - \mathrm{MC}_A\big)\, n_T^0\, \frac{\partial r}{\partial \theta} \;>\; 0,$$

since $p > \mathrm{MC}_A$ and $\partial r/\partial \theta > 0$. The status-quo profit $\pi_A^*(0,0)$ does not depend on θ: while T is open, all patients are captive (Proposition 1). Therefore

$$\Phi'(\theta) = -\big(p - \mathrm{MC}_A\big)\, n_T^0\, \frac{\partial r}{\partial \theta} \;<\; 0,$$

and the set $\{F_T > \Phi(\theta)\}$ on which closure occurs strictly expands as θ rises.[82] □

Policy Extensions

Proposition 5 (Neutrality of transition support)**.** *Let τ index a transition-support policy that applies only to patients whose facility closes: it lowers the switching cost each displaced patient bears from b to $b_D(\tau)$, with $b_D(0) = b$ and $b_D'(\tau) < 0$. Patients whose facility remains open continue to face b, and F (hence G and g) and $u(\cdot)$ are unchanged. Then:*

(i) The Stage-2 equilibrium qualities $q_A^*(1,1)$ and $q_E^*(1,1)$ are invariant to τ.

(ii) The operating profits $\pi_A^*(c,e)$ and $\pi_E^*(c)$, the entry threshold $\overline{F}_E$, and hence the equilibrium in Proposition 2, are invariant to τ.

(iii) In every market in which the target closes, the aggregate switching cost borne by displaced patients, $b_D(\tau)\, n_T^0$, is strictly decreasing in τ.

Proof. **(i)** With the policy, a displaced patient pays $b_D(\tau) + \varepsilon_k^{Tm}$ at destination $m \in \{A, E\}$, so she prefers A to E if and only if

$$u(q_A) - b_D(\tau) - \varepsilon_k^{TA} \;\geq\; u(q_E) - b_D(\tau) - \varepsilon_k^{TE} \quad\Longleftrightarrow\quad D_k \leq u(q_A) - u(q_E).$$

The term $b_D(\tau)$ cancels, so $r(q_A, q_E)$, its derivatives, and the volumes n_A^C, n_E^C do not depend on τ. The first-order conditions (9)–(10) depend only on p, MC_i, n_T^0, g, u', $\partial c_i/\partial q_i$, and n_i^C, none of which depends on τ, so their unique solution is invariant to τ.

(ii) Operating profits depend only on volumes and qualities. In $(0,0)$, no facility closes, so the policy is never triggered. In $(1,0)$, every displaced patient is treated at A for any τ:

[82] When entry is blocked, the acquirer absorbs all displaced patients regardless of θ, and closure occurs for every θ (Proposition 2).

because $b_D(\tau) \le b$, $u(\underline{q}_T) - b_D(\tau) - \bar{\varepsilon} \ge u(\underline{q}_T) - b - \bar{\varepsilon} > V_0$, so Assumption 2 continues to hold. In $(1,1)$, volumes and qualities are invariant by part (i). Because the policy applies only to displaced patients, every patient at an open facility still faces b, so Proposition 1 and Lemma 2 are unaffected. Hence $\pi_A^*(c,e)$, $\pi_E^*(c)$, and $\overline{F}_E = \pi_E^*(1)$ are invariant to τ, and so is the partition (12), which depends only on these objects and on (F_E, F_T).

(iii) Each of the n_T^0 displaced patients bears $b_D(\tau)$, so $\frac{d}{d\tau}\big[b_D(\tau)\, n_T^0\big] = b_D'(\tau)\, n_T^0 < 0$. □

Remarks. Three features deliver neutrality. First, the switching cost a displaced patient bears is common to every destination available to her, so it does not affect *which* facility she chooses. Second, Assumption 2 ensures it does not affect *whether* she is treated in the market: if condition (5) failed, patients with poor fit at every active facility would leave, lowering b_D would bring more of them back into the contestable pool, and the entry threshold and the profit forgone by closing would change. Third, the policy is targeted at closure: a policy that lowers b for all patients changes condition (6) and therefore the pre-merger equilibrium, and one that lowers b for the acquirer's own patients after the merger makes them responsive to the entrant's quality and so changes the Stage-2 game.

Proposition 6 (Direct facility assignment)**.** *Suppose a fraction $\sigma \in [0,1]$ of displaced patients is directly assigned to the acquirer regardless of quality, while the remaining $1-\sigma$ choose between A and E as in Section 6.4.1. Assume g, u', MC_i, and $\partial c_i/\partial q_i$ are continuous, with $\partial c_A/\partial q_A > 0$ on the relevant compact domain. Then, in the $(1,1)$ subgame:*

(i) There exists $\sigma_A^* < 1$ such that for all $\sigma > \sigma_A^*$, the acquirer's quality reverts to the floor, $q_A^* = \underline{q}$, for any q_E.

(ii) For $\sigma \in (\sigma_A^*, 1)$, the entry threshold $\overline{F}_E(\sigma)$ is strictly decreasing in σ, so the $(1,1)$ region shrinks and the $(1,0)$ region expands.

(iii) At $\sigma = 1$, $\overline{F}_E(1) = 0$: entry never occurs and the equilibrium is $(1,0)$ with $q_A^* = \underline{q}$.

Proof. **Volumes.** Under assignment,

$$n_A^C(\sigma) = n_A^0 + \sigma n_T^0 + (1-\sigma)\, r(q_A, q_E)\, n_T^0, \qquad n_E^C(\sigma) = (1-\sigma)\big[1 - r(q_A, q_E)\big]\, n_T^0, \tag{12}$$

which nest the baseline at $\sigma = 0$, and $\partial n_A^C(\sigma)/\partial q_A = (1-\sigma)\, n_T^0\, g\big(u(q_A) - u(q_E)\big)\, u'(q_A)$.

(i) Substituting into (3),

$$\frac{\partial \Pi_A^C}{\partial q_A} = (1-\sigma)\big(p - \mathrm{MC}_A\big)\, n_T^0\, g\big(u(q_A) - u(q_E)\big)\, u'(q_A) - \frac{\partial c_A}{\partial q_A}\, n_A^C(\sigma). \tag{13}$$

Let $\mathcal{D}$ denote the compact set of (q_A, q_E, n) with $q_A, q_E \in [\underline{q}, \bar{q}]$ and $n \in [n_A^0, N]$. By continuity,

$$\bar{B} \equiv \max_{\mathcal{D}}\ \big(p - \mathrm{MC}_A\big)\, n_T^0\, g\big(u(q_A) - u(q_E)\big)\, u'(q_A), \qquad \underline{C} \equiv \min_{\mathcal{D}}\ \frac{\partial c_A}{\partial q_A}\, n$$

exist, with $\bar{B} > 0$ and $\underline{C} > 0$. Since $n_A^C(\sigma) \in [n_A^0, N]$, (13) satisfies $\partial \Pi_A^C/\partial q_A \le (1-\sigma)\bar{B} - \underline{C}$. Define $\sigma_A^* \equiv \max\{0,\, 1 - \underline{C}/\bar{B}\} < 1$. For $\sigma > \sigma_A^*$, $(1-\sigma)\bar{B} < \underline{C}$, so $\partial \Pi_A^C/\partial q_A < 0$ for all q_A and q_E, and the acquirer's unique best response is $q_A^* = \underline{q}$.

(ii) Fix $\sigma \in (\sigma_A^*, 1)$, so $q_A^* = \underline{q}$. The entrant's operating profit is

$$\overline{F}_E(\sigma) = \max_{q_E \in [\underline{q}, \bar{q}]}\ \Big[p - c_E\big(n_E^C, q_E\big)\Big]\, n_E^C, \qquad n_E^C = (1-\sigma)\Big[1 - G\big(u(\underline{q}) - u(q_E)\big)\Big] n_T^0.$$

By the Envelope Theorem,

$$\frac{d\overline{F}_E(\sigma)}{d\sigma} = \Big[p - \mathrm{MC}_E\Big]\frac{\partial n_E^C}{\partial \sigma} = -\Big[1 - G\big(u(\underline{q}) - u(q_E^*)\big)\Big]\, n_T^0 \Big[p - \mathrm{MC}_E\Big] \;<\; 0,$$

using $p > \mathrm{MC}_E$ and $1 - G(\cdot) > 0$, which holds because $u(\underline{q}) - u(q_E^*) \leq 0$ and $G(0) < 1$ when the support of D_k contains 0 in its interior. Hence $\overline{F}_E(\sigma)$ is strictly decreasing, so the set $\{F_E \leq \overline{F}_E(\sigma)\}$ shrinks and $\{F_E > \overline{F}_E(\sigma)\}$ expands (Proposition 2).

(iii) At $\sigma = 1$, $n_E^C = 0$ for any qualities, so $\overline{F}_E(1) = 0 < F_E$. The entrant never enters, and the equilibrium is $(1, 0)$ with $q_A^* = \underline{q}$ (Lemma 3). □

Remark. Conditional on entering, the entrant continues to upgrade: its entire volume is contestable, so the factor $1 - \sigma$ multiplies both sides of its first-order condition (10) and cancels. Direct assignment therefore erodes quality gains through two channels—it removes the acquirer's incentive to compete, and it deters entry by shrinking the contestable pool—rather than by lowering the entrant's quality.

H Appendix: Alternative Post-Merger Equilibria

This appendix characterizes each post-merger subgame. Throughout, condition (6) and Assumptions 1–3 hold, and $F_E, F_T > 0$. Table H.1 summarizes the resulting equilibria.

Subgame $(0,1)$: Entry Without Closure

Proof of Lemma 2. Fix $c = 0$, so T remains open and Assumption 3 does not apply. Every patient is at an open facility $i \in \{A, T\}$ and can avoid b by staying. A patient k at i switches to E only if $u(q_E) - b - \varepsilon_k^{iE} \geq u(q_i)$. By (1) and condition (6),

$$\big[u(q_E) - u(q_i)\big] - b - \varepsilon_k^{iE} \;\leq\; \Delta - \big(b + \underline{\varepsilon}\big) \;<\; 0,$$

so no patient at A or T switches to E, for any qualities (q_A, q_T, q_E). The entrant's demand is zero and its operating profit is $\pi_E^*(0) = 0$; entering yields $-F_E < 0$. □

Corollary 2 (The subgame $(0,1)$ is never reached)**.** *Because $\pi_E^*(0) = 0 < F_E$, an entrant facing an open target always stays out. Hence $(c, e) = (0, 1)$ is not an equilibrium outcome for any admissible parameters, and keeping T open always deters entry.*

Subgame $(1,0)$: Closure Without Entry

Proof of Lemma 3. With $c = 1$ and $e = 0$, A is the only active facility. Its original patients stay, since $u(q_A) \geq u(\underline{q}) > V_0$ by (2). A displaced patient who moves to A obtains

$$u(q_A) - b - \varepsilon_k^{TA} \;\geq\; u(\underline{q}_T) - b - \bar{\varepsilon} \;>\; V_0,$$

using $q_A \geq \underline{q} > \underline{q}_T$, $\varepsilon_k^{TA} \leq \bar{\varepsilon}$, and Assumption 2. She therefore prefers A to leaving the market for every q_A and every draw of ε_k^{TA}, so all n_T^0 displaced patients choose A. The acquirer's volume is fixed at N, independent of q_A, and

$$\Pi_A(q_A) = \big[p - c_A(N, q_A)\big]\, N, \qquad \frac{\partial \Pi_A}{\partial q_A} = -\frac{\partial c_A}{\partial q_A}\, N \;<\; 0.$$

Profit is strictly decreasing in q_A, so $q_A^*(1,0) = \underline{q}$ and $\pi_A^*(1,0) = \big[p - c_A(N, \underline{q})\big] N$. □

Subgame $(0,0)$: Status Quo

With $c = 0$ and $e = 0$, the open facilities and switching costs are those of the pre-merger market, so Proposition 1 applies unchanged: no patient switches, $q_A^* = \underline{q}$, $q_T^* = \underline{q}_T$, and $\pi_A^*(0,0) = (p - c_A^0)\, n_A^0 + (p - c_T^0)\, n_T^0$.

Subgame $(1,1)$: Closure With Entry

This subgame is characterized in Section 6.4.1: the Stage-2 equilibrium (q_A^*, q_E^*) solves (9)–(10), with $q_A^*, q_E^* > \underline{q}$ under the conditions of Lemma 1. The acquirer serves $n_A^C = n_A^0 +$

$r(q_A^*, q_E^*)\, n_T^0$ patients and earns $\pi_A^*(1,1)$ in (7); the entrant serves the rest and earns $\overline{F}_E = \pi_E^*(1)$.

Table H.1: Post-merger equilibria: parameter regions and outcomes

Equilibrium	**Entry cost**	**Target fixed cost**	**Quality**	**Δ Facil.**
$(1,1)$ Closure + entry	$F_E \leq \overline{F}_E$	$F_T > \pi_A^*(0,0) - \pi_A^*(1,1)$	$q_A^*, q_E^* > \underline{q}$	0
$(0,0)$ Status quo	$F_E \leq \overline{F}_E$	$F_T \leq \pi_A^*(0,0) - \pi_A^*(1,1)$	$\underline{q}$ at A, $\underline{q}_T$ at T	0
$(1,0)$ Closure, no entry	$F_E > \overline{F}_E$	any $F_T > 0$	$\underline{q}$ at A	-1
$(0,1)$ Entry, no closure	never reached (Corollary 2)			

I Appendix: Sensitivity of Cost-Benefit Analysis to Alternative Assumptions

This appendix examines the sensitivity of the cost-benefit calculations in Section 7 to alternative assumptions. To recap, the baseline analysis uses facility-level event-study estimates over a five-year post-merger evaluation window, without discounting, and incorporates the central value of a statistical life and life-expectancy estimates for individuals with ESRD. We calculate benefits by multiplying the estimated affected population of $109,295.21$ by the cumulative effect of mergers on hospitalizations (-0.248) and the cost per hospitalization ($27,055), yielding aggregate benefits of $688.22 million. We calculate costs by multiplying the affected population by the merger-year mortality coefficient (0.0065) to estimate the number of excess deaths attributable to mergers. We then multiply these excess deaths by the life expectancy at death of 4.3 years—the average remaining life expectancy for dialysis patients at the age of death in our sample—to estimate years of life lost, and value these years using the central value of a statistical life year. This yields aggregate costs of $848.46 million.

We next vary the baseline specification by introducing one methodological change at a time while holding all other parameters fixed at their baseline values. First, we make an adjustment to the valuation of benefits to reflect private insurance reimbursement rates. Private insurers paid, on average, a little more than twice the Medicare rate for comparable services in 2020 (Whaley et al. 2022).[83] Inflating all hospitalization benefits by a factor of two would therefore provide an upper bound on potential benefits—since in practice, only a minority of dialysis patients are covered by private insurance. This increases aggregate benefits to $1,376.44 million and, consequently, yields a net positive value of $527.98 million from mergers.

Second, we incorporate alternative valuations of a statistical life year (VSLY)—a low value of $130,000 and a high value of $423,000, obtained by dividing the low and high values of a statistical life reported in U.S. Department of Health and Human Services (2021) by average U.S. life expectancy. Because the VSLY affects only the valuation of mortality, these alternatives change estimated mortality costs to $396.76 million (low VSLY) and $1,291.00 million (high VSLY), while leaving hospitalization benefits unchanged. Under the high VSLY, therefore, costs exceed benefits by $602.78 million, while benefits exceed costs by frm-e91.00 million under the low VSLY.

Third, we utilize the general U.S. population average life expectancy of 16.30 years in place of the dialysis-specific figure to estimate years of life lost.[84] This adjustment scales mortality costs to $3,216.25 million, generating a large negative net value of mergers of $2.528 billion.

[83] Whaley et al. (2022) compare private insurance and Medicare prices for individual services within a claim and therefore do not directly report prices per hospitalization. We make the simplifying assumption that the ratio of service prices extends to overall hospitalization costs—an assumption that would hold if privately insured and Medicare patients receive comparable bundles of services during hospitalizations for similar health events.

[84] See life expectancy for ages 68–69 in Table 1 of National Vital Statistics Reports (2022).

Fourth, we include temporal discounting at a standard annual rate of 3%, applied simultaneously to both costs and benefits. For mortality costs, discounting the annual aggregate cost of lost life years (709.76 deaths $\times$ \$278,000 = \$197,313,280) over the 4.30-year life expectancy horizon, beginning at period 0, reduces total costs to \$808.01 million.[85] For hospitalization benefits, discounting annual annual aggregate hospitalization savings from year 0 to 5 reduces benefits to \$624.84 million.[86]

Fifth, we substitute market-level coefficients from panels B and C of Appendix Figure A.10 for facility-level coefficients. Doing so reduces mortality costs decrease to \$420.13 million, while hospitalization benefits decrease to \$500.78 million, so that net value is positive by \$80.64 million.

Finally, we calculate break-even time horizons under baseline methods to calculate how long we would have to extend our time horizon to for cumulative hospitalization benefits to surpass mortality costs. Assuming post-year 5 hospitalization effects remain constant at the year 5 coefficient level, the exact break-even point is reached at 5.77 years.[87] Under an alternative specification where post-year 5 hospitalization coefficients extrapolate linearly based on the observed trajectory from year 0 to 5, the break-even horizon occurs slightly earlier at 5.72 years.[88]

As detailed in the lower panel of Table I.1, isolating single-parameter shifts clarifies how alternative analytical assumptions drive overall welfare outcomes. Assumptions that affect only benefits (e.g., insurer prices) or only costs (e.g., VSLY and life expectancy) generate larger changes in the estimated net value of mergers, while assumptions that affect both sides of the valuation (e.g., the unit of analysis and discounting) generate smaller changes because their effects partially offset. Variation in life expectancy is the largest driver of the change in net value. Importantly, however, the more realistic assumptions in our sensitivity analysis—discounting, modestly upweighting prices to account for a small share of privately insured patients, and using market-level coefficients—all shift the baseline toward higher net benefits rather than higher costs.

[85] Discounted mortality costs are derived by calculating the present value of lost life years across a 4.3-year expectancy horizon: $PV = \sum_{t=0}^{3} \frac{\$197.31}{(1.03)^t} + \frac{\$197.31 \times 0.3}{(1.03)^4} \approx \$808.01$ million.

[86] We get this by calculating $\sum_{t=0}^{t=5} \text{Cost Per Hospital Discharge} \times \text{Exposed Population} \times \beta_t = \sum_{t=0}^{t=5} 25,433.38 \times 109,295.21 \times \beta_t$ where β_t are the coefficients from panel A of Figure 13.

[87] Adding 109,295.21 (exposed patients) $\times$ 25,433.38 (cost per hospital discharge) $\times$ 0.0756 (the hospitalization coefficient at event time 5) $\times$ 0.77 to the five-year aggregate benefits of \$688.22 million yields total benefits of \$849.96 million, just surpassing total costs.

[88] Linearly extrapolating the hospitalization coefficients from Figure 13 to a sixth post-merger year yields a coefficient of approximately -0.0811. Adding 109,295.21 (exposed patients) $\times$ 25,433.38 (cost per hospital discharge) $\times$ 0.0811 (the extrapolated hospitalization coefficient at event time 6) $\times$ 0.72 to the five-year aggregate benefits of \$688.22 million yields total benefits of \$850.49 million.

Table I.1: Cost-Benefit Estimates and Sensitivity Spread of Dialysis Facility Mergers ($ Millions)

(A): Cost-Benefit Comparatives			
Specification & Assumptions	(1) **Costs**	(2) **Benefits**	(3) **Net Value**
Baseline Estimates (Facility-Level)			
Assumptions: No Discounting, Central VSLY, Dialysis LE=4.3			
Baseline Level	848.46	688.22	–160.24
Comparatives: Private Insurance [Benefit-Only Shift]			
Private Insurance Prices Nearly Twice Medicare Prices	848.46	1,376.44	+527.98
Comparatives: Valuation of a Statistical Life Year (VSLY) [Cost-Only Shift]			
Low VSLY ($130k)	396.76	688.22	+291.46
High VSLY ($423k)	1,291.00	688.22	–602.78
Comparatives: Life Expectancy (LE) Assumption [Cost-Only Shift]			
U.S. Average LE (16.3, No Disc.)	3,216.25	688.22	–2,528.03
Comparatives: Discounting [Dual-Sided Shift]			
3% Discount Rate ($r = 0.03$)	808.01	624.84	–183.17
Comparatives: Market-Level Specification [Dual-Sided Shift]			
Market-Level Coefficients	420.13	500.78	+80.64
Magnitude Drivers			
Pricing Spread (Private vs. Medicare)	0.00	688.22	+688.22
VSLY Valuation Spread (Low vs. Central)	-451.7	0.00	+451.70
VSLY Valuation Spread (High vs. Central)	+442.54	0.00	–442.54
Life Expectancy Spread (U.S. vs. Dialysis LE)	+2,367.79	0.00	–2,367.79
Specification Spread (Market vs. Facility)	–428.33	–187.44	+240.88
Discounting Spread (3% vs. No Disc.)	–63.42	–63.38	–22.93
(B): Break-Even Analysis			
Exact years until Cumulative Benefit Exceeds Mortality Cost			
Constant post-Year 5 extrapolation[a]		5.77 Years	
Linear post-Year 5 extrapolation[b]		5.72 Years	

[a] Post-period coefficient held steady at -0.0755 per year. [b] Post-period coefficient assumes a continued linear trend based on the Year 0 to Year 5 effects.

J Appendix: Data

J.1 Determining Mergers and Markets

We obtain the Centers for Medicare and Medicaid Services (CMS) Provider of Service (POS) Files for the fourth quarter of each year from 1991 to 2021 from the NBER.[89] In cleaning the data, we assign each facility the modal state and county codes across its observations (or the lowest code if there are several modal values) when it appears in more than one state or county due to data errors. When a state and county code is available for a provider in some years, but missing in others, we fill out the missing values using the available information in other years. When the state and county code is missing in all years, we identify the state and county manually by web-searching the address. To address cases where the same address shows multiple provider numbers over the years, we identify situations where changes appear to reflect the same facility rather than distinct entities, and make appropriate corrections.[90]

The POS indicates a facility's chain ownership status. Missing values for this variable are assumed to be zero. For facilities with a missing year within a sequence, the panel is completed by carrying backward the values for chain ownership, state, and county from the next available non-missing year. We emphasize that we do not construct a balanced panel of facilities from 1991 to 2021; years preceding a facility's first appearance or following its last appearance are treated as the facility not existing, which is crucial for the analysis of entry and closures.

The POS file provides dates of facility ownership changes. We extract the year from this variable and tag each facility-year as a merger year when applicable. For facilities with multiple ownership changes, we use the most recent change, allowing identification of multiple mergers within a facility. We classify merger type using chain ownership information from the year prior to the merger for the target and from two years after the merger for the acquirer.[91]

Markets are defined at the county level. We divide three large counties, namely Cook County (IL), Los Angeles County (CA) and San Diego County (CA), into smaller markets. This classification of markets is adopted from the approach by Wollmann (2020).[92]

[89] NBER hosts the cleaned POS files here: `https://www.nber.org/research/data/provider-services-files`.

[90] Specifically, we compare the number of provider IDs at each state–county–address across years and within each year, and we merge IDs only when there is never more than one facility at an address within a year, so as not to combine genuinely separate providers. When the algorithm flags an apparent duplication of codes, we replace the provider number with the number recorded for that provider in the most recent year of the data.

[91] We use the value from two years after the merger, rather than the year of or immediately after, to account for delays in updates. Because our algorithm restricts treated units to those without merger activity for five years following the merger, the two-year post-merger value is unlikely to reflect subsequent mergers.

[92] Cook County (IL) is split into three markets where the first includes Evanston, the second includes Orland Park, and the third includes Elk Grove Village. San Diego County (CA) is split into two markets where the first consists of Encinitas, Escondido, Oceanside, San Marcos and Vista, and the second includes the rest of San Diego. Los Angeles County (CA) is subdivided into five markets where the first comprises Bellflower, Cerritos, Compton, Cudahy, Downey, Hawaiian Gardens, Huntington Park, Lakewood, Long Beach, Lynwood, Norwalk, Paramount, Santa Fe Springs, South Gate, Whittier, Rancho Dominguez, and San Pedro. The second comprises Alhambra, Arcadia, Azusa, Burbank, El Monte, Encino, Glendale, Glendora, Mission Hills, Monrovia, Montebello, Monterey Park, North Hollywood, Pacoima, Panorama City, Pasadena, San Gabriel, Sherman Oaks, South El Monte, Sun Valley, Temple City, Van Nuys, Boyle Heights. The third comprises Beverly Hills, Carson, Culver

J.2 Creating the Stacked Dataset

To estimate equation 1, we first construct a separate dataset for each cohort (sub-experiment) d, where d denotes the calendar year in which a merger occurred. For each cohort, we restrict the data to the periods between $\mathcal{L}'$ years before and $\mathcal{L}''$ years after year d, and include units that were treated in year d and units that were not treated in year d. Appendix Table B.1 shows the values of $\mathcal{L}'$ and $\mathcal{L}''$ for each outcome.

In each cohort d, treated units belong to markets that experience a merger in year d but no mergers in the 10 years prior to d or in the 5 years following d. This ensures that any post-merger effects observed during the 5 years after the merger can be attributed to the merger in year d and are not confounded by other consecutive mergers. Similarly, the 10-year pre-event window ensures not only that no merger occurs in the unit's market in the 5 years preceding the event of interest, but also that any effects from earlier mergers have had at most five years to dissipate. In all analyses, control units belong to markets that were not exposed to any merger in year d, for 10 years prior to year d and for 5 year after d. The procedure for identifying treated and control markets in the analysis of merger types (e.g., chain-to-chain mergers) is analogous, except that treated markets are those experiencing the reference-type merger in the year of interest but no merger of any type in the surrounding windows.

For outcomes measured from 1991 to 2021, we study mergers occurring between 1996 and 2018, allowing for complete data on all five pre-merger periods and at least four post-merger periods for each unit. Further, for mergers occurring between 2001 and 2016, we construct treated and control markets using the procedure described above, imposing restrictions on merger activity in the market during the 10 years prior to and 5 years after the merger. For mergers from 1991 to 2000, we allow for a shorter clean pre-merger period, calculated as the number of years between 1991 and the reference year (e.g., a 9-year clean pre-event window for mergers in 2000). Similarly, for mergers in 2017 and 2018, we allow for a shorter clean post-merger period, determined by the number of years between the reference year and 2021 (e.g., a 3-year clean post-event window for mergers in 2018). For outcomes measured from 2012 to 2021, the estimating equation includes fewer pre- and post-merger periods, but the restrictions on merger activity in markets in years surrounding the merger remain the same as previously described.

We define event time e as the calendar year minus the merger year d, with event time 0 corresponding to the merger year. This creates a cohort-specific dataset where the cohort is indexed by the merger year. We then append the cohort–specific datasets to create a stacked dataset aligned on event time. Markets may appear multiple times in the stacked dataset as they can serve as controls for more than one merger, and the same market may have multiple eligible mergers, separated sufficiently in time to be included as treated observations in the stacked dataset.

City, Gardena, Harbor City, Hawthorne, Inglewood Lomita, Redondo Beach, Santa Monica, Torrance. The fourth consists of Baldwin Park, Canoga Park, Canyon Country, Covina, Granada Hills, Hacienda Heights, La Puente, Lancaster, Newhall, Northridge, Palmdale, Pomona, San Dimas, San Fernando, Sylmar, Tarzana, Valencia, West Covina, Woodland Hills, City Of Industry, Irwindale, Santa Clarita, West Hills. The fifth includes the rest of Los Angeles County.

J.3 Market-Level Covariates

We compile a comprehensive set of county-level controls for inclusion in all regressions. For almost all markets, the covariates describe the market uniquely but for three large counties that were further split into smaller markets, covariate values are common within the county in which these markets lie. County median income is drawn from the Census SAIPE program. Because county estimates are missing for some early years (1991, 1992, 1994, 1996), we impute the 1993 value for 1991–1994 and otherwise carry forward the most recent available year. County population totals by age, sex, and race are compiled from the Census' July intercensal estimates. Dialysis certificate-of-need (CON) laws by state are taken from the AHPA matrix, available for 2004, 2008, 2012, and 2016. Only Missouri changed status during this period. We code CON law status for missing states as 0, use the 2004 value for all years up to 2007, and carry forward the last available value thereafter.

J.4 Variable Construction in the Analysis of Closures, Patient Volume, Congestion and Staffing

The USRDS facility data provide information on total dialysis patients and stations at each facility, the number of patients with Medicare (pending or approved), and staff counts by category and by full- or part-time status. These data are first merged with a crosswalk of USRDS and CMS facility IDs to enable linkage with the CMS POS data. The resulting market-level dataset retains all POS facilities. Facilities present in POS but absent from USRDS contribute to market-level counts of facilities but not to totals for staff, patients, or stations.

Each facility is classified as "acquired" or "non-acquired" within a cohort depending upon whether located in a treated market and is the facility undergoing a merger in the year indexed by the cohort. Both categories are classified as treated. The same classification is applied to staff, patients, and stations in treated markets. Facility-level outcomes are then used directly in the analyses of staffing and patients-to-stations.

We next construct a balanced panel covering the full study period from 1991 to 2021. For any years in which a market is not observed, either before its first appearance or after its last, outcome variables are set to zero. This yields a consistent panel of all U.S. markets that contained at least one facility at some point between 1991 and 2021. We then aggregate the numbers of facilities, stations, and patients at the market level, both overall and by acquired status. These aggregates serve as the outcomes in the market-level analyses. By definition, control markets contain no acquired or non-acquired facilities. Consequently, their outcomes are identical across categories (e.g., in control markets, acquired and non-acquired facility counts are set equal to total facility count).

J.5 Patient Insurance Information

For each patient and treatment spell, the USRDS provides payer history. We construct a patient–year indicator for whether Medicare was the primary payer. In each calendar year, we record a patient's modal payer, preferring the latest spell in any year when there is a tie. We then record whether the patient's modal primary payer was Medicare FFS. Patients

for whom Medicare was not the primary payer are excluded whenever claims data are used to define outcomes (e.g., hospitalizations) or to assign modal facilities, ensuring that counts accurately reflect primarily Medicare-covered care.

J.6 Linking Patients to Facilities

In the analysis of ESA prescription, fistula usage and all biomarkers and clinical endpoints, we utilize patient level data and aggregate it up to the facility level. To enable this, we first create unique annual linkages between patients and facilities. Following is the linkage algorithm we employ.

1. **Identify patient-facility pairs using outpatient claim revenue center codes (2002-2021)**.
    - Extract calendar year from the claim beginning dates. All claims end and begin in the same calendar year so this algorithm would work similarly when using claim end dates.
    - Count the number of dialysis sessions in each claim period for a patient and record the facility associated with the claim. Utilize the HCPCS revenue center code 90999 to identify dialysis.
    - Retain the patient-facility pair where the patient received the maximal dialysis sessions during a year. If there is a tie, prefer the facility with the later claim end date. If still tied prefer facility with earlier claim start date. If still tied, pick randomly (this happens rarely).
    - Merge with patient insurance information and keep only patients for whom Medicare FFS is the primary payer, since any information from claims is most accurate only for these patients.
2. **Identify patient-facility pairs using pre-cleaned outpatient claims (2011-2021)**.
    - In addition to raw outpatient claims, the USRDS also provides pre-cleaned claims files for years 2011 to 2021.
    - Like before, retain the patient-facility pair where the patient received the maximal dialysis sessions during a year. If there is a tie, prefer observations where the treatment modality was in-center hemodialysis.
    - Merge with patient insurance information and keep only patients for whom Medicare FFS was the primary payer.
3. **Identify patient-facility pairs using treatment spells (1991-2021)**.
    - For all patients, USRDS provides a detailed treatment history including precise dates of change in treatment modality and the facility where treatment was incurred during each such period.
    - Retain patients receiving hemodialysis.

- For each treatment spell, calculate the number of treatment days in each calendar year covered by the spell and record the facility where the patient received treatment.
- For each patient-year keep the facility with the maximum treatment days in any calendar year.

4. **Create the final patient-facility linkage using the following heirarchy**.
    - Begin with the linkages from Step 3.
    - When the facility ID is missing, utilize the linkage determined in Step 1.
    - When the facility ID is missing, utilize the linkage determined in Step 2.

This produces the complete annual patient-facility linkages which include all patients on dialysis, regardless of insurance type.

J.7 Construction of Patient-Level Outcomes

J.7.1 Dialysis Sessions (2002-2021)

The enumeration of total annual dialysis sessions by patients mirrors step 1 in Section J.6 above. We retain all patient-year-facility records.

J.7.2 CROWNWeb Outcomes (2012-2021)

We derive most biomarkers, ESA prescriptions and fistula use, from CROWNWeb data. CROWNWeb (now called the End Stage Renal Disease Quality Reporting System, or EQRS) is an electronic reporting system through which dialysis facilities submit patient information and lab results to comply with CMS coverage requirements and, in particular, the QIP. These data are reported monthly and include both Medicare and non-Medicare patients who receive treatment on the day of data collection.

For these patients, we compute the annual modal ESA receipt status and vascular access type at each facility where the patient receives treatment. We then construct a variable indicating whether the annual modal vascular access type was a fistula. Next, for each patient-year-facility, we calculate the mean values of single-pool Kt/V, corrected serum calcium, ESA dosage (if reported in units), pre-dialysis weight, post-dialysis weight, dialysis session length (minutes), serum ferritin, serum hemoglobin, iron saturation, serum albumin, and serum phosphorus.

We create dummy variables for quality thresholds, namely good Kt/V (mean $\geq$ 1.2), hypercalcemia (mean serum calcium $>$ 10.2 mg/dL), hypocalcemia (mean serum calcium $<$ 8.5 mg/dL), high hemoglobin (mean $>$ 12 g/dL), and low hemoglobin (mean $<$ 10 g/dL). These thresholds are derived from Medicare QIP measures (Kt/V, hypercalcemia) or the clinical literature. We also calculate each patient's mean annual ultrafiltration rate.

J.7.3 Mortality (1991-2021)

The USRDS provides a patient-level overview file contain demographic, death, transplant and high-level treatment information for each dialysis patient. The date of death of each patient provides information for the year of death. We use this to classify a death event in the year. Our mortality robustness analyses also exclude from the death-rate denominator patients who died before June, as determined by their date of death.

Upon merging with the patient panel data described in detail in Section J.8, the year corresponding to a patient's death is assigned a value of 1 for the death indicator, and all prior years are assigned a value of 0. By definition, the patient is excluded from the sample in all subsequent years.

J.7.4 Hospitalization, ICU Days, Transfusions (1991-2021)

Hospitalizations are identified from Medicare Part A claims. For each patient-year-facility, the number of hospitalizations equals the total count of unique claims in that year. Each claim also reports days spent in the ICU, which are summed to construct the ICU-days variable. Medicare claims include up to 25 ICD-9 or ICD-10 procedure codes. A red blood cell transfusion is flagged if any of these codes indicate transfusion of packed cells, of nonautologous red blood cells, of nonautologous frozen red cells, or of nonautologous whole blood into a central or peripheral vein.

J.8 Constructing Facility-Level Health Outcomes

Facility-level outcomes are constructed by aggregating patient-level outcomes up to the facility level. We begin with the unique patient–year–facility combinations generated by the algorithm described in Section J.6. At this step, we calculate each facility's annual counts of all patients. Next, we merge in death data using patient IDs and CROWNWeb data using patient ID, year, and facility ID so that each merged record retains its patient–facility linkage.[93] For each facility–year, we sum total deaths, we sum the values of dummy variables (e.g., fistula use, ESA prescription, hypercalcemia, good Kt/V), and compute means of continuous variables (e.g., ESA dose, ultrafiltration rate, serum albumin). For each facility–year, we also count the number of patients with non-missing data for each CROWNWeb outcome and take the maximum of these counts as the facility's total CROWNWeb patient count for that year.

Collapsing to the facility–year level yields annual averages for continuous outcomes. For dummy variables, we calculate facility-level shares by dividing the total count for that outcome (e.g., deaths, patients with a fistula, patients with good Kt/V) by the relevant patient denominator. For mortality, the denominator is total facility patients; for CROWNWeb outcomes, it is the facility's CROWNWeb patient count.

In a separate step, we construct facility-level health outcomes for Medicare FFS patients only. As before, we start from the unique patient–year–facility combinations described in Section J.6. We merge in patient insurance information from Section J.5 and retain only

[93] This means that a patient's information associated with their non-modal facility is dropped.

observations where Medicare FFS is the primary payer. We then merge in dialysis session counts and URR from Part B claims. If a patient–year appears multiple times in Part B, the merge retains the observation matching the patient's unique facility linkage. If a Part B claim has no matching patient–facility link, we retain the record but treat session and URR values as missing.

We next merge in hospitalization, ICU, and transfusion data at the patient–year level. Because these Part A data are not facility-linked, we attribute hospitalizations to the patient's unique facility by construction. If a patient–year does not merge, we assume the patient had no Part A claims (and therefore no hospitalizations) and set all hospitalization-related variables to zero. Finally, all patients in the resulting dataset form the denominator for hospitalization, ICU-days, and transfusions per patient, while only patients with non-missing session information form the denominator for sessions per patient and for the share of patients with good URR.

J.9 Facility-Level Covariates

The USRDS patient-level overview file provides demographic information. During the construction of the facility-level health outcome data, we merge these demographics using the unique patient ID. We create indicator variables for sex, two race categories, Hispanic origin, and 10-year age groups. Before collapsing to the facility level, we sum these indicators so that, in the all-patient dataset, we capture the total number of patients in each demographic category and, in the Medicare-only dataset, we capture the total number of Medicare FFS patients in each demographic category. In the collapsed dataset, we then calculate the share of the relevant patient population in each demographic category, which we use as covariates.

Separately, from the USRDS facility data we compute the share of patients with Medicare pending or approved status and include this measure among the facility-level covariates. Finally, using the CMS POS data, we identify each facility's first year of operation and calculate its age by subtracting this year from the calendar year of observation. We also use facility age as a covariate.