ESSAYS ON PRIVATE EQUITY AND HEALTHCARE

Tong Liu

A DISSERTATION

in

Finance

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2022

Supervisor of Dissertation

Itay Goldstein

Joel S. Ehrenkranz Family Professor
in Finance, Professor of Economics

Co-Supervisor of Dissertation

Lucian Taylor

John B. Neff Professor in Finance

Graduate Group Chairperson

Nancy Zhang, Ge Li and Ning Zhao Professor, Professor of Statistics and Data Science

Dissertation Committee

Guy David, Alan B. Miller Professor in Health Care Management, Professor of Medical Ethics and Health Policy
Michael Schwert, Assistant Professor of Finance

To my beloved wife and family

ACKNOWLEDGEMENT

I am deeply indebted to my advisors Guy David, Itay Goldstein, Michael Schwert, and Lucian Taylor for their invaulable guidance and mentorship. I am grateful to Sylvain Catherine, Vincent Glode, Joao Gomes, Yao Zeng and the rest of the Wharton finance faculty for their help and support throughout my time at Wharton.

I am fortunate to have met incredible colleagues and friends at Wharton. I thank my friends Aymeric Bellon, Roberto Gomez Cram, Paul Decaire, Mehran Ebrahimian, Chuck Fang, Marco Giometti, Alexandr Kopytov, Jianan Liu, Mauricio Medeiros Junior, Dominik Supera, Ruslan Sverchkov, Haotian Xiang, Ipek Yavuz, Yicheng Zhu.

Part of this dissertation was made possible by Decision Resources Group (now part of Clarivate) through its Real World Data (RWD) Product. Special thanks go to Ashwin Anand, Remy Corvese, Daniel Neilson, and Wouter van der Pluijm for assistance in providing and interpreting the data.

Finally, I owe immense gratitude to my wife, Peiyuan Zhang, for her unconditional support and belief in me; to my parents, for being an endless source of patience and encouragement. I could not have done this without you.

ABSTRACT

ESSAYS ON PRIVATE EQUITY AND HEALTHCARE

Tong Liu

Itay Goldstein

Lucian Taylor

This dissertation consists of two chapters that relate private equity and health care. In the first chapter, coauthored with Xuelin Li and Lucian Taylor, we study how common ownership affects innovation. We explore this question using project-level data on pharmaceutical startups and their venture capital (VC) investors. We find that common ownership leads VCs to shut down lagging drug projects, withhold funding from lagging startups, and redirect those startups' innovation. These results support theories dating back to Loury (1979): By coordinating R&D efforts across competing firms, a common owner can reduce duplication of R&D. Consistent with common ownership improving innovation efficiency, common ownership rates are positively correlated with the ratio of R&D output to funding. Though we highlight gains in innovation efficiency, common VC ownership can also impose social costs.

In the second chapter, I use proprietary health insurance claims data covering over 60% of privately insured individuals in the United States to study the impact of private equity (PE) hospital buyouts on hospital–insurer price negotiations, health spending, and patient welfare. I apply a novel structural approach that exploits state-level regulation changes as PE entry shocks. I find that PE buyouts lead to an 11% increase in total healthcare spending for the privately insured in affected markets, driven mostly by higher bargained prices at PE-backed hospitals and price spillovers to local rivals. PE investors' superior bargaining skills account for 43% of the price and spending increases, while financial engineering and bankruptcy threats contribute 40%, changes in patient demand contribute 10%, and reduced focus on social objectives contributes 8%. Operational efficiency gains reduce spending, but only by 1%. A counterfactual ban on PE hospital buyouts would increase patient surplus by an amount equivalent to 10.7% of health expenses. If antitrust regulators who conduct merger reviews ignore PE-backed acquirers' unique features, they risk greatly underestimating the impact of hospital mergers.

TABLE OF CONTENTS

CHAPTER 1

Common Ownership and Innovation Efficiency

1.1. Introduction

Competing firms sometimes share a large investor in common. Common ownership has grown significantly in recent years, leading to debates about its welfare implications. Recent papers suggest a welfare loss due to reduced product-market competition (e.g., Azar et al., 2018). Common ownership can also impact welfare, both positively and negatively, through its effects on innovation.

In this paper, we study how common ownership affects innovation. We find that common ownership leads investors to shut down lagging R&D projects, restrict their funding, and encourage their firms to pivot to new R&D projects. In this manner, common owners help reduce duplication of R&D efforts across competing firms. Theories dating back to Loury (1979) show that if firms compete, they invest more in R&D than is socially optimal. A common owner helps solve this problem by coordinating firms and reducing their excess R&D. Consistent with common ownership improving innovation efficiency, we show that common ownership is associated with more R&D output per dollar of R&D funding. This behavior, however, can also impose social costs: By shutting down projects, a common owner can reduce future product-market competition and consumer choice.

We study these issues in the venture capital (VC) setting. This setting is important, as VC-backed startups generate a large share of the innovation in our economy (e.g., Gompers and Lerner, 2004, Kaplan and Lerner, 2010). Also, the effects of common ownership should be especially strong for VC investors. VCs are sophisticated, active investors who hold concentrated portfolios, so they are likely to be attentive to spillovers across portfolio companies. VCs own large equity stakes and have strong control rights in their startups, so they have the power to influence those companies. We also find that common ownership by VCs is widespread: 39% of startups in our sample have a close competitor with a shared

VC investor. For all these reasons, one might expect the effects of common ownership to be even stronger in the VC setting than, for example, in the setting of passive index funds holding modest stakes in public companies.

Our main testable prediction is from Grossman and Shapiro (1987), an extension of Loury (1979). They model firms engaged in a multi-stage patent race, and they ask how a firm responds after seeing its competitor reach an intermediate milestone. The theory implies that if both firms share the same owner, then the lagging firm's project is likely to be shut down, because the common owner wishes to avoid the inefficient duplication of R&D efforts. If the firms instead have different owners, they fail to internalize the negative spillovers they impose on each other. The lagging project is therefore likely to continue, even if it is socially suboptimal.

Motivated by this theory, we predict that common ownership increases the likelihood that a lagging R&D project is held back. Our ideal experiment would feature two pairs of competing startups. We would randomly assign one pair to share a common VC investor and the other pair to not. One "pioneer" startup in each pair would reach an intermediate milestone, and we would compare the outcomes of the two pairs' "lagging" startups. Following Grossman and Shapiro (1987), we predict that the lagging startup is less likely to continue progressing in the pair that shares a common VC.

To apply this ideal experiment to real data, we need to identify competing startups, and we need publicly observable intermediate milestones. We overcome this challenge by using project-level data on pharmaceutical startups. Our data cover 1,045 Phase I drug projects conducted by 481 U.S. startups financed by 764 VC firms, from 2005 to 2018. We work at the level of VC firms, so "VC" stands for a VC firm rather than a VC fund or partner. Our data partition the pharmaceutical industry into 78 detailed product markets, so we can compare, for example, a pair of startups with competing tuberculosis projects to a pair of startups with competing arthritis projects. Regulation by the U.S. Food and Drug Administration (FDA) provides a clear, publicly observed intermediate milestone: seeing a drug project progress from Phase I to Phase II clinical trials. Besides having these useful properties, our data cover a sector of the economy that is both highly valuable and important for social welfare.

Another challenge is finding quasi-random variation in common ownership. We apply an instrumental variable (IV) approach that exploits the local nature of VC investing. VCs tend to invest in nearby companies in order to reduce the costs of search and monitoring. Our IV for whether two startups share a common VC is based on the startups' geographic proximity. The main identification assumption is that geographic proximity affects our dependent variable—the outcome at a lagging startup after a competing pioneer startup makes progress—only through the effect of proximity on whether the two startups share a common VC.

Our first tests examine the probability that a Phase I drug project progresses to Phase II after seeing a closely competing project—a pioneer—progress to Phase II. Consistent with the prediction above, we find that sharing a VC in common with the pioneer makes the lagging project less likely to progress to Phase II. More simply, common ownership leads VCs to hold back lagging projects. Economic significance is high: a shared VC reduces the probability of progressing by 0.53 of one standard deviation. The result obtains even if we include fixed effects for time by drug category, which amounts to comparing how two lagging tuberculosis projects (for example) react differently to seeing a third tuberculosis project reach Phase II, depending on whether the two lagging projects share a VC with the pioneer. By comparing how those two lagging projects react to the same pioneer, this exercise addresses potential concerns about the pioneer's existence being endogenous to common ownership.

How does sharing a common VC affect drug projects' outcomes? We find evidence of a VC financing mechanism. A VC can hold back a drug project by withholding follow-on funding. Doing so can make sense even though Phase I projects are very young, because progressing to Phase II is costly. We predict that after a VC sees a startup make progress, the VC is less likely to extend funding to a competing startup if both startups are in its portfolio. The data strongly support this prediction, with high levels of statistical and economic significance. We find evidence of a VC financing mechanism even if we compare across different VCs invested in the same startup and quarter, which effectively controls for a startup's demand for funding. Further, we show that when a common VC abandons a startup, other VCs do not step in to fill the financing hole, so the startup is less likely to raise financing from *any* VC.

We find interesting heterogeneity in these results. Supporting the patent-race interpretation from Grossman and Shapiro (1987), our main effects are at least twice as large when the pioneer and lagging projects are technologically similar, indicating they are closer competitors. Results are also stronger for VCs with larger ownership stakes in the startup. Such VCs have stronger control rights, making it more feasible for them to hold back projects. Results are stronger for less-diversified VCs as well. This result is consistent with Gilje et al. (2020), who find that common owners must be highly attentive to internalize spillovers within their portfolios.

We also show that common owners redirect lagging firms' innovation activities. When a startup sees a commonly owned competitor make progress in one drug category, the startup pivots away from that category toward non-overlapping categories. Specifically, the lagging startup is more likely to initiate new projects in non-overlapping drug categories, repurpose existing projects into those categories, and form financing alliances with large pharma companies in those categories. By redirecting the lagging firms' innovation efforts in this way, the common owner again helps avoid the duplication of R&D efforts.

To the extent that common ownership reduces unnecessary duplication of R&D costs, common ownership can improve innovation efficiency. Consistent with this story, we find that common ownership rates are positively correlated with a measure of innovation efficiency in the cross section of drug categories. We measure innovation efficiency as the total number of drugs reaching FDA approval within a drug category, scaled by the total amount of VC funding provided to all startups active in that category. This efficiency proxy is strongly correlated with the drug category's common ownership rate, with high levels of statistical and economic significance. This is a simple, descriptive result that does not have a causal interpretation. The result is consistent, however, with common ownership helping to avoid excess duplication of R&D, producing more approved drugs per dollar of aggregate R&D.

This result suggests that common ownership has a social benefit. Common ownership, however, may also impose social costs. By holding back lagging drug projects, common owners can reduce future product-market competition and consumer choice. Since the pairs of drugs we study are not always perfect substitutes, allowing the the lagging project to continue can improve welfare in some cases. Also, our tests examine a specific counterfactual in which we hold the set of projects or startups fixed while changing the ownership structure from non-common to common. Common ownership can influence which projects or startups exist in the first place, which affects welfare. Section 1.4 discusses welfare implications in detail.

We address various challenges to our identification strategy. One potential concern is that geographic proximity reflects the ease of poaching employees from a competing startup, which in theory could explain why a project is more likely to fail after a nearby competing project gains an edge. Inconsistent with that story, however, our results are not weaker when the startup's state strongly enforces employee non-compete agreements. We also analyze potential concerns related to technological proximity, information sharing, and spray-and-pray strategies.

Our main contribution is to document a new, important way in which common ownership affects innovation. Many papers study how common ownership affects product-market competition,[1] but relatively few study its effects on innovation. He and Huang (2017) show that common ownership promotes resource sharing across firms, which improves firm-level innovation productivity. Similarly, Kostovetsky and Manconi (2020) find that common institutional ownership facilitates innovation diffusion among portfolio firms. The mechanism in this paper is quite different: Common owners can improve innovation efficiency by redirecting innovation and avoiding duplication of R&D.

Cunningham et al. (2021) study "killer" acquisitions of competing drug companies. Similar to us, they show that an acquired drug project is less likely to progress if the acquirer has a similar drug project. There are important differences between our studies, though. Cunningham et al. study direct ownership, which is quite different from common ownership by a financial intermediary. Unlike "killer" acquirers, the VCs we study are minority shareholders with only partial control, and they presumably do not invest in startups with the intent to kill them. Therefore, it is somewhat surprisng that VCs hold back certain projects in their portfolios. Another distinction is that Cunningham et al. do not study the redirection of innovation, nor do they emphasize social benefits from reducing duplication of effort.

A few papers study spillovers within VC portfolios. Lindsey (2008) finds that startups sharing a common VC investor are more likely to form alliances, and González-Uribe (2020) finds they are more likely to share innovation resources. The mechanism we study is quite different, especially its focus on close competitors within a VC's portfolio. Closer to our paper, Eldar et al. (2020) ask whether common VC ownership affects startup performance. They show that common ownership improves information sharing among startups, which increases their growth and odds of success. Like us, they find that common VC ownership has social benefits, but the mechanism is clearly different. One advantage of Eldar et al. (2020) is that their data span all industries, making it easier to generalize results. An

[1] He and Huang (2017), Azar et al. (2018), and Azar et al. (2022) find the common ownership reduces product-market competition. Opposing evidence comes from Gramlich and Grundl (2017), Kennedy et al. (2017), Lewellen and Lewellen (2021), Kini et al. (2022), Dennis et al. (2021), Lewellen and Lowry (2021), Gilje et al. (2020), and Koch et al. (2020). Gutiérrez and Philippon (2017) show a negative relation between common ownership and corporate investment, which relates to our finding that common ownership leads VCs to hold back projects and funding.

advantage of using only pharmaceutical industry data, however, is that it lets us better identify firms and projects that are close competitors. For example, we assume a tuberculosis drug competes only with other tuberculosis drugs. Eldar et al. (2020) instead use a coarse industry classification that, for instance, combines all pharmaceutical companies into one category.

Our results also shed light on theories of common ownership and innovation. López and Vives (2019) show that a common owner internalizes firms' positive R&D spillovers, which can increase R&D and improve welfare. They predict that common ownership increases R&D if and only if R&D spillovers are sufficiently high. Viewed through the lens of their theory, our finding that common owners hold back lagging projects suggests that R&D spillovers are small in our setting. This result makes sense. If two drug startups are competing for share in a fixed market, then one startup's progress can impose negative spillovers on the other firm. Similarly, Benmelech et al. (2021) find that pharmaceutical R&D has small productivity spillovers. Of course, R&D spillovers can be large and positive in settings different from ours (e.g., Bloom et al., 2013, Matray, 2021). Related to López and Vives (2019), Antón et al. (2021) show theoretically that common ownership increases R&D when technological spillovers are large relative to product-market spillovers, and they provide supporting evidence using panel data on public firms.

We also contribute by documenting a "horse race" investment strategy used by some VCs. In this strategy, a VC invests in competing startups, waits for one to gain an edge, and then reduces funding to the lagging startup while redirecting its innovation. The Online Appendix provides a case study about one prominent VC firm following this strategy.[2] Related to this strategy, Ewens et al. (2018) show that VCs increasingly follow a "spray and pray" approach, investing small amounts in many startups with the expectation that many will be abandoned. Distinct from that strategy, VCs in our paper set up horse races between closely competing startups, and VCs help lagging startups to pivot. Stories of pivots are

[2]The Online Appendix is available on the authors' websites.

common among practitioners, but we are not aware of any other research on VC-enabled pivots.

1.2. Empirical approach

1.2.1. Data and institutional details

Drug development in the U.S. is strictly regulated by the FDA and follows several stages: discovery and Phase I, II, and III clinical trials. The discovery stage involves pre-clinical research for drug candidates in the laboratory and animal testing for basic safety. Phase I tests the safety and determines dosage of a drug candidate on a small group of humans. Phase II tests efficacy and side effects at a larger scale. Phase III, which involves thousands of participants, tests whether there is a treatment benefit. Favorable results are required to move from one phase to the next. All three phases feature high amounts of scientific uncertainty. We interpret clinical-trial expenses as R&D, following U.S. accounting standards, the U.S. government, Benmelech et al. (2021), and many others.[3]

We construct a sample of pharmaceutical projects initiated by U.S.-based companies that are funded by VCs. Detailed information on drug development and clinical trials comes from the Cortellis Life Sciences Healthcare Database. Cortellis obtains its data from public records, such as clinical trial registries, FDA submissions, patent filings, company press releases, and financial filings. We provide additional information on the Cortellis database in the Online Appendix. Cortellis records drug development history at the project level. A project is a sequence of trials for testing the safety and efficacy of a drug targeting a specific indication. An indication is a specific disease or medical condition. From Cortellis we collect the dates when projects progress through the various clinical phases; each project's suspension date (in the event of ultimate failure) or FDA drug approval date (in the event of success); detailed information on the company running the project, including the company's name, funding status, organization type, headquarters location, and major shareholders; information on alliance formation and drug repurposing; and detailed information on the portfolio of patents

[3] See https://www.cbo.gov/publication/57126 for the CBO's interpretation of clinical-trial expenses. For a summary of the relevant accounting rules, see http://www.pwc.com/us/pharma.

associated with the drug. We also scrape each patent's citation information from Google Patent.

We obtain VC investment records from SDC Platinum VentureXpert. We use a fuzzy matching algorithm to link funding records in SDC to company names and locations in Cortellis. For each matched pair, we manually check the matching accuracy. Our sample includes startups funded by corporate VCs (CVCs), but they account for only 2% of our panel dataset.

We apply several filters to obtain our final sample. First, we require both the startup and its VC investors to be in the U.S. We apply this filter because of data availability. Fortunately, the U.S. accounts for a very large share of global pharmaceutical R&D.[4] Second, we drop projects initiated before the first quarter of 2005, because Cortellis's coverage of development histories is less reliable before then. Third, we drop projects initiated after the first quarter of 2016 if no progress or suspension is observed. We do so because insufficient time has passed for these projects to reach any outcome. Next, given our focus on VC, we only include quarters when a given project is held within a VC's portfolio.[5]

We focus on projects in Phase I trials. A project enters our sample in the first quarter of its Phase I trials, and it exits our sample either when it progresses to Phase II, reaches the end of 2018, or is classified as abandoned. We classify a project as abandoned once it is suspended, discontinued, withdrawn, or coded as "no development" for four years. We focus on projects in Phase I for several reasons. Our startups typically raise their first VC money when their projects are in either the discovery or Phase I trials. We exclude the discovery stage, because it requires significantly less capital, tends to last a short period of time, and is often not publicly reported. Also, VCs have greater control over the startup during

[4]In 2016, the U.S. accounted for 58% of pharmaceutical R&D among 11 leading countries (ABPI, 2019).

[5]This period lasts from the first quarter when the startup receives its initial VC funding until either the company exits successfully (e.g., IPO or being acquired) or is written off. Following González-Uribe (2020), absent any exit, we say a startup is written off five years after its last VC financing round. If that date exceeds 2018Q4, we keep the project in our sample but record its final outcome as missing. Our main results are robust to using other cutoffs (e.g., 3, 7, and 10 years) to define the write-off date.

projects' early stages. By the time a project reaches Phase II or III, the startup has often gone public or been acquired, removing the VCs' influence. Since we want to study how VCs influence startups, we focus on the stage when that influence is most relevant. Finally, Phase I is inherently important, sometimes being called the "valley of death" in drug development (Seyhan, 2019). In the full Cortellis database, 55% of project failures occur during or prior to Phase I. Of the Phase I projects in our sample, only 21% progress to Phase II by the end of 2018.[6] Progressing from Phase I to Phase II increases a drug's probability of FDA approval from 10% to 31% (Thomas et al., 2016). While Phase I is traditionally viewed as testing safety and dosage, in VC-backed companies it often has more ambitious goals, such as identifying target patient populations, drug response rates, and biomarker evidence (Ivy et al., 2010; Roberts et al., 2004). The median duration of Phase I is 1.6 years, and for some projects it extends well beyond three years (Wong et al., 2019).

Our analysis requires identifying drug projects that are close competitors. We do so by partitioning the overall pharmaceutical industry into 78 sub-markets. We map Cortellis indications to the second chapter level of International Classification of Diseases (9^{th} Revision, "ICD"), and we refer to each chapter as an ICD category. Examples of ICD categories include tuberculosis, urinary system diseases, and Kaposi's sarcoma. Each ICD category contains one or more three-digit ICD codes showing more granular disease descriptions. We drop three-digit codes whose names contain "unspecified" or "ill-defined," since these diseases can be quite different from each other. After we apply this filter, indications in the same ICD category have a high pathological correlation, so drug projects in the same ICD category are plausibly substitutes and hence competitors. The Online Appendix tabulates our 78 ICD categories and their three-digit constituents.

We follow Krieger (2021) and Byrski et al. (2021) in intepretting each ICD category as a

[6]For comparison, the FDA website reports that 70% of Phase I projects progress to Phase II, and Wong et al. (2019) report a 66% probability of transitioning from Phase I to II. The transition rate is lower in our sample because many projects are still in progress at the end of our sample period, and VC-backed drug ventures tend to be riskier and more innovative. For example, 32% of our sample projects are related to oncology, a notoriously risky area of drug development.

distinct sub-market. In Section 1.6 we show that our main results are robust to working at the more granular three-digit ICD level. We prefer working at the ICD-category level, though, because it is common for a drug to target multiple three-digit diseases within the same ICD category. For example, the drug Posiphen explicitly targets Alzheimer's, Huntington's, and Parkinson's diseases, each of which belongs to a different three-digit code within a single ICD category. Thirty-three drugs in Cortellis explicitly target both Alzheimer's and Parkinson's diseases. Many drugs' indications are not even available at the three-digit level. Instead, their indications are at the broader ICD-category level, meaning these drugs target multiple related diseases.[7] Of course, it is possible that some drugs within an ICD category are not close competitors. To the extent that our sample contains such classification errors, our analysis contains noise that biases us against finding support for our predictions.

Finally, we drop a project if none of its VC investors are disclosed. These filters together produce a final sample of 57,316 VC-project-quarter observations from 1,045 projects across 78 drug categories. The projects are run by 481 startup companies funded by 764 VC firms. The sample covers 2005 to 2018.

1.2.2. Predictions and identification strategy

Our first prediction is that common ownership increases the likelihood that a lagging competitor project is held back. We provide additional background on the theory of Grossman and Shapiro (1987), which motivates our prediction. Grossman and Shapiro (1987) model firms competing in a winner-take-all patent race. They extend Loury (1979) by adding an intermediate milestone that allows one firm to gain an observable lead in the race. The theory maps well to drug development, as patent protection results in negative externalities from leading to lagging drug projects, and progressing to Phase II trials is an observable intermediate milestone. Grossman and Shapiro (1987) analyze whether the lagging firm continues its R&D project after seeing its competitor reach the milestone. The theory implies

[7] For example, 189 drugs in Cortellis list their targeted indication as the ICD category "chronic obstructive pulmonary disease and allied conditions" (codes 490-496), so these drugs claim to treat chronic bronchitis (491), emphysema (492), asthma (493), and the other constituent three-digit diseases.

that if the two firms share a common owner with strong control rights, the lagging project is more likely to be shut down, because the common owner wishes to avoid duplicating R&D costs.[8] Absent a common owner, the lagging project is less likely to shut down, because neither firm internalizes the other firm's R&D costs. Competition, in other words, results in over-investment in R&D, as in Loury (1979). By coordinating firms' R&D, a common owner reduces excess R&D and thereby helps solve a market failure associated with fierce competition.

While this story implies social benefits from common ownership, there is an alternative story implying social costs: Common owners hold back lagging projects in hopes of creating market power for the surviving projects. We explore these alternative motives and their welfare implications in Section 1.4.

The patent-race story clearly applies to situations where a patent blocks subsequent projects, but the story also applies when there is no such blocking patent. Even without a blocking patent, drugs in the same ICD category typically compete for share of a fixed market, and there are first-to-market advantages (Regnier and Ridley, 2015). Progress by one project therefore imposes a negative externality on the lagging project, which is the key ingredient for the predictions in Grossman and Shapiro (1987).

Why does shutting down projects at such an early stage make sense? After all, Phase I projects still have a chance of catching up. Shutting down a project after Phase I can make sense because progressing to Phase II is not free. According to DiMasi et al. (2016), Phase II costs $59 million on average. At the portfolio level, it can be optimal for a VC to avoid this cost and shut down a Phase I project whose stand-alone value is not too high.

[8]Specifically, in their application to intermediate patents, Grossman and Shapiro (1987) compare the ex post value if only the leading firm is allowed to continue (V_m in their notation) to the total ex post value if both firms continue ($V_{10} + V_{01}$ in their notation). A common owner with strong control rights can choose between those two options. They show that $V_m > V_{10} + V_{01}$, implying a common owner would shut down the lagging project, as long as discount rates and elasticities of marginal costs are not too high. Fulghieri and Sevilir (2009) generate a related prediction, which is that VCs have an incentive to invest in technologically similar companies, because it lets the VC shift resources from one portfolio company to another in case the first company fails.

Figure A.1 illustrates the ideal experiment for testing our prediction. Suppose there are two pairs of drug projects, A-B and C-D. The projects within each pair are in the same ICD category, so each pair contains close competitors. One pair (A-B) is randomly assigned to share a common VC investor, and the other (C-D) does not. One project in each pair achieves an observable intermediate success, defined as progressing from Phase I to Phase II. We call these projects, labeled A and C in the figure, our "pioneers." Our tests compare the outcomes of the two "lagging" projects, B and D. Our tests do *not* compare project A to B, C to D, or A to C, so we do not require those projects to have similar quality. Instead, our main prediction is that project B is less likely than D to progress to Phase II. Having a common VC makes it more likely that the lagging project gets held back.[9]

We can express this experiment in regression form as follows:

$$Progress_{it} = \gamma \, Lagging_{it} + \beta \, Lagging_{it} \times SharedVC_{it} + \Gamma' Controls_{it} + FEs + \eta_{it}. \quad (1.1)$$

$Progress_{it}$ is an indicator for whether project i progresses to Phase II in quarter t. The sample includes all projects i that are in Phase I at the beginning of quarter t. We include both lagging and non-lagging projects to obtain more precise estimates of Γ and fixed effects (FEs). We must therefore include $Lagging_{it}$, an indicator for whether drug project i has fallen behind a pioneering competitor. Specifically, $Lagging_{it}$ equals one if there exists a different project in the same ICD category (but different startup) that has progressed to Phase II between project i's initiation quarter and quarter $t-1$. The indicator $SharedVC_{it}$ equals one if project i and the pioneering project share a common VC in quarter t.[10] The coefficient of interest is β, which measures how lagging projects' outcomes depend on whether there is a shared VC. We predict $\beta < 0$, meaning a lagging project is less likely to progress

[9] Similar to us, Krieger (2021) studies how drug projects react to the failure of a competing project. We add an extra "diff" to the analysis by comparing the reaction between pairs with and without a common VC.

[10] We cannot include $SharedVC_{it}$ as an extra regressor, because it is only defined over project pairs containing a pioneer, so $SharedVC_{it}$ is only defined when $Lagging_{it} = 1$. We show in the Online Appendix that our main results are robust to controlling for a related variable, $CommonOwn_{it}$, an indicator for whether project i shares a VC with *any* other project in the same ICD in quarter t.

to Phase II in the presence of common ownership.

We use an OLS panel specification because our prediction is inherently dynamic: *after* falling behind the pioneer, the project is more likely to be shut down. The Online Appendix shows that our results also hold using the Cox hazard model. We prefer the linear probability model because the hazard model does not easily accommodate a large number of fixed effects.

At first glance, Eq. (1.1) looks like a "triple-diff" regression. It examines projects over time (the first diff), depending on whether the project is lagging (the second diff), and also depending on whether the projects share a VC (the third diff). Unlike in a triple-diff, though, $SharedVC_{it}$ is not randomly assigned, so OLS estimation of Eq. (1.1) does not have a causal interpretation.

We take an IV approach, which requires finding quasi-random variation in $SharedVC$, the indicator of common VC ownership. Our instrument exploits the local nature of VC investing. VCs prefer investing in nearby companies to reduce the costs of screening and monitoring (Lerner, 1995). For example, Gompers et al. (2020) show that VCs rely heavily on their networks to source deals, and networks are often local. We use the geographic proximity of two startups to create an instrument for whether they share a common VC investor. Specifically, we create an instrument for the endogenous interaction term $Lagging_{it} \times SharedVC_{it}$ as follows. For project i and quarter t, we collect the set S_{it} of projects j that are in the same ICD category as i, and that progress from Phase I to Phase II between the birth of project i and quarter $t-1$. These projects j are the pioneers that make project i "lagging." Let d_{ij} denote the distance, in miles, between the Metropolitan Statistical Areas (MSAs) of startups' headquarters for those startups running projects i and j. We scale the distance by a constant of 2,600 miles, which is roughly the air travel distance from Boston to San Francisco. This scaling helps produce a strong first-stage regression. Given the scaled distance d'_{ij}, we compute $f\left(d'_{ij}\right) = \exp(-d'_{ij})$ as a proxy for the probability that projects i and j

share a common VC. Our IV for the variable $Lagging_{it} \times SharedVC_{it}$ is then

$$Lagging_{it} \times Proximity_{it} = \sum_{j \in S_{it}} \exp(-d'_{ij}). \tag{1.2}$$

We set the instrument to zero if the set S_{it} is empty, meaning $Lagging_{it} = 0$. We compute a sum in Eq. (1.2) because common ownership is more likely if there are multiple nearby pioneer projects. Summation is not driving our results, however; our results are actually stronger if we control for the number of pioneer projects in set S_{it}. We show later that geographic proximity strongly explains common ownership, and the first stage of our IV regressions is quite strong. The relevance condition holds for this instrument, in other words.

Our identifying assumption is that the instrument affects the outcome variable $Progress_{it}$ only through its effect on $Lagging_{it} \times SharedVC_{it}$. Why does this exclusion restriction plausibly hold? Our IV test boils down to checking whether, after falling behind a pioneer, a project is more likely to fail if it is geographically closer to the pioneer. More simply, the IV's reduced form tests whether geographic proximity predicts opposite outcomes for the lagging and pioneer projects. Reverse causation is not an issue here, because the startups' locations are determined well before the projects' outcomes. The bigger concern is omitted variable bias, meaning there exists some omitted variable W that is correlated with both geographic proximity and the projects' achieving opposite outcomes. No such W is apparent to us. To reduce portfolio risk, a VC may prefer investing in projects with negatively correlated outcomes, and unobserved negative correlation could explain projects' opposite outcomes. This is exactly why we need an instrument for $SharedVC$. We find no evidence that our IV, based on geographic proximity, coincides with negatively correlated project fundamentals. For example, we show in the Online Appendix that there is no significant relation between proximity and shared patent citations, a proxy for technological similarity. Section 1.5 addresses other possible identification challenges (e.g., employee poaching, information sharing), and Section 1.3.2 addresses potential concerns about the endogeneity of *Lagging*.

An implicit assumption is that projects B and D, i.e., the lagging projects with and without common ownership, are of similar quality. In the Online Appendix, we show that projects B and D do not differ significantly on observable proxies for quality. Proxies we consider include project age, startup age, whether the startup has advanced previous projects to Phase II or Phase III, and whether the project is located in Boston or San Francisco. The lack of correlation holds whether we compare lagging projects based on the endogenous variable *SharedVC* or based on the instrument *Proximity*. The exclusion restriction cannot be tested, but it is comforting that neither the endogenous variable nor the instrument correlates significantly with proxies for project quality. Also, our distance-based instrument is not significantly related to outcome variables in the subsample of projects without common ownership, which further supports the exclusion restriction.

1.3. Empirical results

1.3.1. Summary statistics and frequency of common ownership

Figure A.2 shows the geographic distribution of drug projects in our sample. As expected, there are many projects in California, Massachusetts, and New York. Approximately 46% of projects, however, are located outside these states.

Table A.1 shows summary statistics for our main variables in the full sample as well as the "Never Treated" and "Ever Treated" subsamples. The "Never Treated" group contains projects or startups for which $Lagging \times SharedVC$ equals zero across their entire lifespan. Remaining projects or companies are in "Ever Treated."

Panel A shows project-level variables. The full-sample mean of *Lagging* is 0.504, meaning in roughly half of our sample a project has seen a close competitor progress to Phase II. *Progress* is roughly three times higher in the "Never Treated" group, which foreshadows our result that sharing a VC with a competing, successful project significantly reduces the likelihood of progressing to Phase II.

Panel B contains startup-level variables. In a typical quarter, the average startup has 1.66 projects covering 1.42 ICD categories. A typical startup will run several projects over its lifespan, so the total number of projects per startup is much higher than 1.66. The typical startup has 4.7 VC firms invested in it at a given point in time, although there is wide variation in this number.

Next, we show that common ownership by VCs is quite common. Table A.2 shows that a sizeable 39% of our startups have at least one close competitor that shares at least one VC in common, in a typical quarter. At first glance, this common ownership rate seems high given that our sample contains 764 distinct VC firms. The rate is high for a few reasons. First, the average startup has 4.7 VCs, so there is a decent chance that at least one VC is shared with a competitor's (multiple) VCs. Also, Table A.2 shows that VCs specialize in specific ICD categories. While a VC could potentially invest across all 78 drug categories, the typical VC holds stakes in just 2.5 categories in the typical quarter, making it more likely that the VC holds multiple startups in the category. Put differently, while there are 764 VCs in our full sample, the typical ICD category has only 16.4 VCs invested in it at once, increasing the chances of VC overlap. This high degree of specialization makes sense given the high level of scientific expertise needed to evaluate drug startups. Finally, some ICD categories contain many startups, making it easier to find at least one other competing startup that shares a VC. The rate of common ownership appears lower, 8.6%, if we measure it instead as the probability of a shared VC between two startups chosen at random within an ICD category.

1.3.2. Evidence from project outcomes

Our main results begin in Table A.3, which contains results from the panel regression in Eq. (1.1). Except in columns 2 and 4, the dependent variable is $Progress_{it}$, the indicator for whether project i progresses to Phase II in quarter t. We include fixed effects (FEs) for the startup, ICD category, and quarter to control for unobserved heterogeneity.[11] We include three controls: the project's age ($\ln(Age)$), number of Phase I projects being developed by the startup ($NProjects$), and number of VC firms owning a stake in the startup company ($NVCs$). We find that $\ln(Age)$ is a particularly important control, as it takes time to obtain Phase I results. Standard errors are computed by two-way clustering at the ICD category and startup levels.

The first column contains OLS estimates. This regression does not have a causal interpretation, but it provides a useful description of the data. Since the dependent and independent variables of interest are indicators, the slope coefficients can be interpreted as average changes in probability. The OLS results point in the predicted direction. The coefficient on $Lagging \times SharedVC$ is negative and highly statistically significant ($t = -2.95$). Sharing a common VC is associated with a 0.019 lower probability that, if a similar project makes progress, the project progresses to Phase II during the quarter. The standard deviation of $Progress$ is 0.132, so common ownership is associated with a 0.14 standard deviation lower probability of progressing. The insignificant slope on $Lagging$ indicates that, absent a shared VC, falling behind a related project makes it neither more nor less likely that the lagging project progresses. Combining both coefficients, we see that a project becomes less likely to progress after falling behind a commonly owned competitor.

Our first IV test is in columns 2–3.[12] As expected, in column 2 we see a strong, positive relation between the instrument and $Lagging \times SharedVC$, consistent with VCs preferring

[11] The ICD category FEs are useful because it may be more difficult to reach Phase II for certain diseases. The quarter FEs controls for aggregate changes in the FDA process, conditions in the macroeconomy or pharmaceutical industry, etc. The startup company FEs are useful, because the typical startup runs several projects over its lifespan, and some startups may be systematically better at getting projects through Phase I.

[12] Since the focal regressor, $Lagging \times SharedVC$, is binary, we follow the advice of Wooldridge (2010) and Angrist and Pischke (2008) and use a probit-2SLS procedure. In an initial step, we estimate a probit regression of the endogenous regressor, $Lagging \times SharedVC$, on $Lagging \times Proximity$ and exogenous controls. The probit model's predicted probability, denoted $P(Lagging \times Proximity)$, becomes the instrument in the first stage of 2SLS. The Online Appendix tabulates the probit results. We find very similar results if we instead follow a 2SLS approach without the probit step (Online Appendix). The probit-2SLS approach has been used recently in the finance literature by Saretto and Tookes (2013) and Ewens and Marx (2018). For additional details on this approach, see page 191 of Angrist and Pischke (2008).

to invest in nearby companies. The first-stage F-statistic is 57, indicating no problems with weak instruments. The IV estimate in column 3 echoes the OLS results. Consistent with our prediction, after seeing a similar "pioneer" project progress to Phase II, a project is less likely to progress to Phase II if it shares a VC with the pioneer project. Compared to OLS, the IV results have higher statistical and economic significance. A shared VC reduces a lagging project's probability of progressing to Phase II by 0.070, which is 53% of the dependent variable's standard deviation.

Jiang (2017) shows that it is common for IV estimates to be much larger than their OLS counterparts. It makes sense that our IV slopes exceed the OLS slopes. Our IV picks up commonly owned startups that are geographically near each other. Our IV-compliers are therefore VCs whose investment decisions are more influenced by the monitoring costs associated with geographic proximity. This subgroup of VCs is likely to exhibit larger local average treatments effects (LATEs), for at least three reasons. A VC who intends to run a horse race between two startups knows it will have to monitor them closely, so the VC will want them closer together. Second, our IV-compliers arguably care more about monitoring in general, making them more aware of spillovers within the portfolio, which is necessary for our prediction to hold. Finally, the IV-compliers are arguably more cost-conscious, making them more eager to avoid duplicating R&D costs, which in turn makes them more willing to suppress a lagging project.[13]

Columns 4 and 5 show that our result is robust to using a bivariate probit model instead of 2SLS. Bivariate probit models are suited to IV settings like ours, in which the dependent and endogenous independent variables are both binary.[14] The bivariate probit estimates again support our prediction. Statistical significance is somewhat weaker ($t = 1.90$), but the estimate implies a sizeable 0.038 lower probability of progressing if the lagging project has a common VC. An advantage of the bivariate probit model over 2SLS is that it recognizes

[13]Supporting evidence is in the Online Appendix. Consistent with the IV-compliers being more attentive to spillovers, IV-compliers hold less-diversified portfolios. IV-compliers are also slightly smaller, which arguably correlates with being more cost-conscious.

[14]For more details, see Angrist and Pischke (2008), pp. 197–205, and Wooldridge (2010), pp. 594–599.

that our dependent and independent variables are binary. A disadvantage of the bivariate probit model is that it imposes a strong distributional assumption on the data, namely, that error terms are normally distributed. A bigger disadvantage is that bivariate probit does not easily accommodate the large number of FEs in our 2SLS specification, so we omit all FEs in columns 4 and 5. Nevertheless, we continue to report bivariate probit results alongside 2SLS results when possible.

One potential concern is that *Lagging*, the indicator for falling behind a close competitor, is somehow endogenous to common ownership, making *Lagging* a "bad control." We address this concern by adding ICD-by-quarter FEs in column 6, so that *Lagging* (not interacted) is dropped from the regression. Our test now asks how lagging projects react to exactly the same pioneering project. The ICD-by-quarter FEs isolate, for example, variation across different Phase I tuberculosis drug projects in a given quarter. If a tuberculosis project progresses to Phase II in quarter $t-1$, then all remaining Phase I tuberculosis projects in quarter t would have $Lagging = 1$. (Note the project that progresses to Phase II during quarter $t-1$ is no longer in our sample during quarter t.) The variable $Lagging \times SharedVC$ would therefore only pick up variation in $SharedVC$ across remaining Phase I tuberculosis projects in quarter t, and the regression would test how those remaining, lagging projects respond differently to the same pioneer depending on whether they share a VC with the pioneer. Column 6 shows that including ICD-by-quarter FEs has a negligible effect on the OLS magnitude and decreases its statistical significance only slightly. We show in the Online Appendix (Table OA4.2) that our IV results also continue to hold if we include ICD-by-quarter FEs.

Some projects fail to reach Phase II not because their owners actively suppress them, but simply because they have bad Phase I outcomes. If common owners play an active role in holding back lagging projects, then our results should hold even in the subset of projects with good Phase I outcomes. To test this idea, we re-estimate the previous models after excluding projects that, according to Cortellis, fail to reach the primary Phase I end point or experience an adverse Phase I event, such as a dangerous side effect or death. Results are in Table A.4. Economic significance is even stronger compared to Table A.3. The negative coefficients on $Lagging \times SharedVC$ remain statistically significant, even though the sample has shrunk by more than half. These results support the idea that common owners play an active role in holding back lagging projects.

Figure A.3 Panel A describes these effects' dynamics. We estimate an OLS regression similar to Eq. (1.1), except we include multiple leads and lags of $Lagging$ and $Lagging \times SharedVC$. The figure plots the OLS coefficients on $Lagging \times SharedVC$ at the various leads and lags. The horizontal axis denotes the number of quarters after the "event" quarter, i.e., the quarter when $Lagging \times SharedVC$ first equals one for each project. The category $\tau = 12$, for example, contains projects that existed and became "lagging" 12 quarters earlier. We see no significant coefficients in the six quarters leading up to the event, indicating no anticipation effects or pre-trends. After the event, all coefficients have negative point estimates, and several are statistically significant. There is no indication of a delayed response or reversion back to zero. Overall, the figure implies that when a commonly owned competitor makes progress, on average there is a prompt, persistent drop in the project's odds of progressing. To address concerns that our panel structure spuriously drives our result, we show in Section 1.6 that our result continues to hold in a purely cross-sectional test.

To summarize, these project-level results support the prediction that common ownership leads VCs to hold back lagging projects. Specifically, we find that common ownership makes it less likely that a lagging project progresses to Phase II after a competing project does so. The result holds in OLS, IV, and bivariate probit specifications. This result goes through whether we examine how projects respond to different pioneers or the same pioneer. The result holds even if we study only those projects that have the potential to progress to Phase II. The average effects begin soon after projects fall behind the pioneer, and they are persistent.

Our findings support the patent-race theory of Grossman and Shapiro (1987). According to their theory, a common owner internalizes both firms' R&D costs and avoids duplicating costs. As a result, common ownership makes it more likely that a lagging project is held back, which improves innovation efficiency.

1.3.3. A VC financing mechanism

We provide evidence of a VC financing mechanism. A common VC can hold back a lagging project by refusing to provide follow-on funding to its startup. Startups typically have negative cash flows, so without follow-on funding they are unable to continue their projects.

For this mechanism to work, a startup that loses funding from one VC cannot easily raise funding from another source, including other VC firms. VCs' private information makes this condition likely to hold, at least in some situations. The logic comes from the relationship-lending literature (e.g., Rajan, 1992; Petersen and Rajan, 1994). A startup will typically raise multiple rounds of financing from one or more VC firms. After investing in and working with a startup, a VC gains private information about the startup. A VC's choice to stop funding a startup can send a negative signal, leading other VCs to either avoid funding the startup or offer unfavorable terms.[15] This private information channel is not necessary, though. VCs often have contractual protections such as veto power over the startup's future financing rounds. A VC with strong enough veto power can effectively prevent the startup from fundraising.[16]

To test this mechanism, we consider a variation on our ideal experiment. As before, we consider two pairs of similar drug projects, and one project (the "pioneer") within each pair progresses to Phase II. In the ideal experiment, we would randomly assign one pair of projects to share a VC and the other pair to not share a VC. Whereas we previously studied whether the lagging projects progress to Phase II, we now study whether the startup companies running those lagging projects receive VC funding. We predict that, if there is a shared

[15]This mechanism might not work if outsiders perfectly infer that funding is withheld due to common ownership rather than a negative private signal. But as long as outsiders cannot perfectly distinguish between a common ownership motive and a negative private signal, the mechanism should continue to work.

[16]In theory, a lagging firm can also turn to banks or government for funding. While those sources are broadly important for funding innovation (e.g., Hombert and Matray, 2017, Hombert and Matray, 2018), they are much less relevant for pharma startups (Wessner, 2009, Harrington, 2012).

VC, then that VC is less likely to fund the lagging project's company. We can express this experiment via the following regression:

$$ExtendFunds_{ijt} = \gamma\, Lagging_{ijt} + \beta\, Lagging_{ijt} \times SharedVC_{ijt} + \Gamma' Controls_{ijt} + FEs + \eta_{ij}, \tag{1.3}$$

where $ExtendFunds_{ijt}$ equals one if VC j extends funding to startup i in quarter t. To create a company- rather than project-level variable, we redefine $Lagging_{ijt}$ to equal one if at least one project in startup i has seen another project (a "pioneer") in the same ICD category (but different startup) progress to Phase II between the quarter when VC j invests in startup i and quarter $t-1$. $SharedVC_{ijt}$ equals one if VC j holds a stake in both startup i and at least one of the startups owning the pioneering projects that produced $Lagging_{ijt} = 1$. To create a company-level version of our instrument, $Lagging_{ijt} \times Proximity_{ijt}$, we modify Eq. (1.2) so that we sum our proximity measure over all the pioneering projects that make any of startup i's projects "lagging." Similar to before, our main identifying assumption is that geographic proximity between startups affects our dependent variable—whether a VC funds a startup—only through the effects of geographic proximity on common ownership. The defense of the exclusion restriction is the same as in our previous tests. We also include controls for the number of startups and ICD categories in the VC's portfolio, the size of and time elapsed since the startup's previous financing round, the number of VC firms invested in the startup, and the number of projects owned by the startup. The first columns include VC firm FEs (because some VCs invest more frequently than others), startup FEs (because some startups raise more VC rounds than others for unrelated reasons), and year-quarter FEs (to soak up aggregate unobservables). We predict $\beta < 0$, meaning a VC is less likely to extend funding to a startup, after the startup sees a competitor make progress, if the VC is invested in both companies.

Support for this prediction is in Table A.5. The OLS slope coefficient on $Lagging \times SharedVC$ is indeed negative and statistically significant at the 5% level. The coefficient's magnitude, -0.030, implies that common ownership is associated with a 0.030 lower probability that the lagging firm raises money from a VC. Column 2 shows a very strong first-stage regression, and column 3's 2SLS slope on $Lagging \times SharedVC$ is negative and statistically significant at the 5% level. The 2SLS estimate, -0.133, implies a 0.133 lower probability of raising VC money if the lagging firm shares a VC with the pioneer. For comparison, the standard deviation of the dependent variable, $ExtendFunds$, is 0.247, so we find that common ownership reduces a VC's funding probability by roughly half of one standard deviation. Results from the bivariate probit model also strongly support the prediction, with a t-statistic of -9.2 and similarly high economic significance.

These results support the idea that common ownership makes a VC less willing to fund a startup that has fallen behind a close competitor. We can reject an alternative interpretation of these results, namely, that investors are discouraged from investing in lagging startups, regardless of common ownership. That story is about the coefficient on $Lagging$, whereas our predictions and main results are about the coefficient on $Lagging \times SharedVC$. In other words, $Lagging$ (not interacted) controls for any effects that falling behind might have on VC financing, and our tests focus on variation in common ownership across lagging startups. Interestingly, $Lagging$ has a large positive coefficient, which might result from VCs funding lagging startups to help them catch up with the pioneer.

The variable $Duration_{ijt}$ deserves comment. $Duration_{ijt}$ equals the log of one plus the number of years since VC j last funded startup i. VCs—especially early-stage ones—will often fund a company for a few rounds, then stop once the round sizes get too large. The large negative slope on $Duration$ likely reflects that, once a VC has stopped funding a company, the VC is unlikely to start funding it again.

One potential concern with the previous test is that a startup's demand for funding, an omitted variable, may differ depending on whether there is a shared VC or on geographic considerations. We address this concern by adding startup-by-quarter FEs in column 6 of Table A.5. These FEs isolate variation across VCs within the same startup-quarter, and hence they effectively control for the startup's funding demand. For example, suppose startup i has two existing VC investors, A and B, and startup i sees a competing startup i' make progress. VC A is invested in both startups, but VC B is only invested in startup i. The regression tests whether VC A is less likely than B to extend funding to startup i in the given quarter. As we predict, column 6 shows that the coefficient on $Lagging \times SharedVC$ remains negative, indicating that VC A (in our example) is indeed less likely than B to extend funding. Economic significance is similar to column 1, and statistical significance is slightly higher. The Online Appendix shows that our IV results also continue to hold after including startup-by-quarter FEs.

Panel B of Figure A.3 illustrates these effects' dynamics. As in the project-level results, we see no significant pre-trends, the effects begin to appear quite soon after the pioneer appears, and the effects are highly persistent.

One potential challenge to the financing mechanism is that other VCs could step in and fill the hole left by a VC who abandons the startup. Table A.6 shows this not to be the case. When a startup with common ownership falls behind a pioneer, the startup is significantly less likely to raise money from *any* VC (column 1), and the amount of new capital raised from all VCs in a given quarter is significantly lower (column 2). These results indicate that new VCs, or existing VCs without common ownership, do not fill the hole left by the existing VC with common ownership. In fact, column 3 shows that new VCs are significantly less likely to invest in a lagging startup with common ownership. As explained earlier, this behavior can result from VCs' contractual protections or negative signals from abandonment by an earlier investor. Column 4 tests whether the existing non-common VCs offset the effects of the common VC. The regressor of interest is $Lagging \times SharedVC \times NonCommonVC$, an indicator for whether startup i is lagging behind a pioneer, at least one of startup i's VCs is shared with the pioneer, and VC j is not the shared VC. We find no significant relation between this regressor and $ExtendFunds$, an indicator for whether VC j extends funds to startup i in quarter t. This result implies that existing non-common VCs do not significantly offset the common VC's reduction in investment.

We find interesting heterogeneity in the VC-funding results. Table A.7 shows results from estimating the baseline VC-funding regression by OLS in various subsamples. (We estimate by OLS because the subsamples are too small to deliver significant first-stage IV regressions.) The first two columns indicate a much stronger VC-funding mechanism for lead VCs, defined as the VC with the highest total amount invested across the given startup's VCs. For non-lead VCs, the coefficient on $Lagging \times SharedVC$ is close to zero. This result makes sense, because a lead VC arguably has stronger control rights over the startup. Also, lead VCs are expected to provide greater follow-on funding, so they would suffer more from duplication of R&D. The last two columns indicate a stronger VC-funding mechanism for less-diversified VCs. We create a simple diversification proxy equal to the number of unique startups, including non-pharmaceutical startups, the VC invests in during the previous five years.[17] The subsample with below-median diversification has a coefficient on $Lagging \times SharedVC$ that is much larger in magnitude (-0.049 versus -0.009). This result supports the finding in Gilje et al. (2020) that, in order for common ownership to have any effects, the common owners must be attentive to spillovers within the portfolio. Since less-diversified VCs can be more attentive to each portfolio company, it makes sense that common ownership has stronger effects among such VCs.

To summarize, the previous three tables are consistent with a VC financing mechanism by which common ownership affects project outcomes. An existing VC with common ownership is less likely to fund a startup that is lagging behind a (commonly owned) competitor. Making matters worse, new VCs are deterred when the existing VC abandons the firm, and other existing VCs do not offset these effects. By starving the lagging startup of capital, the common VC owner can hold back the startup's projects. The financing channel is especially strong for lead VCs and less-diversified VCs.

Besides the financing mechanism, there may be other ways that common ownership results

[17] Computing a more sophisticated diversification proxy is challenging, because portfolio holdings are not marked to market, we do not know when VCs write off failed startups, and our data do not allocate a VC syndicate's investment across the syndicate members.

in lagging projects being held back. A common VC owner can use their board seats, voting rights, option to replace the CEO, or other control rights. A VC can simply exert less effort helping a lagging startup. These mechanisms are not mutually exclusive. Our earlier results on project outcomes reflect these mechanisms' combined effects.

1.3.4. Redirecting innovation

Next, we explore how common owners redirect innovation at lagging firms. Figure A.4 illustrates our prediction.

Suppose VC 1 is a common investor in Startups A and B, and Startup B has a project in ICD category b that progresses to Phase II. That project corresponds to the pioneering project that makes the other startup, denoted A, fall behind. Our previous tests focus on Startup A's projects in the overlapping ICD category, b. We now examine Startup A's projects in non-overlapping categories, such as ICD a. To the extent that Startup A continues innovating, the common VC would like to redirect A's innovation efforts so they do not overlap with Startup B's pioneering project in ICD b. We therefore predict that Startup A's innovation activity increases in ICD a relative to ICD b. The common VC may redirect innovation because it wishes to avoid duplicating R&D costs in ICD b, or because it wants to reduce competition faced by Startup B.

The Cortellis data allow us to measure the redirection of innovation in four ways. Our first measure, $Initiation_{ijt}$, is an indicator for whether startup i initiates a new project in ICD category j in quarter t. $RepurposeInto_{ijt}$ is an indicator for whether the startup repurposes an existing project into ICD category j from a different ICD category j' during t.[18] In the example above, we predict that Startup A repurposes its project from ICD b into ICD a. Conversely, $RepurposeAway_{ijt}$ is an indicator for whether startup i repurposes an existing project from ICD category j into a different ICD category j' during t. In the example above, we predict that Startup A is less likely to repurpose away from ICD a.

[18] Drug companies sometime repurpose a chemical compound from treating one disease to another. For example, Pfizer created Viagra by repurposing a PDE5 inhibitor from coronary artery disease to erectile disfunction.

Finally, $Alliance_{ijt}$ is an indicator for whether the startup i forms an alliance with another pharma company—typically a large, public company—in ICD j and quarter t. Encouraging an alliance in ICD a is one way that VC 1 can redirect innovation toward ICD a without funding the innovation itself.

We use a different specification than in previous sections, because we must now test how innovation activity shifts across ICD categories within a given startup. We regress the outcome variables above on $Treated \times NonSharedICD_{ijt}$, an indicator for whether (1) startup i has a project in a different ICD category j' that is lagging, meaning a different project (at a different startup) in ICD j' progressed to Phase II; and (2) startup i and the startup with the pioneering project share a common VC. In the figure above, $Treated \times NonSharedICD$ would equal one for ICD a of Startup A, and it would equal zero for ICD b of Startup A. We include startup-by-quarter FEs and controls related to those in previous tables.

Results are in Table A.8. Column 1 shows that the lagging startup is 3.4 percentage points (pp) more likely to initiate projects in the non-common ICD. This magnitude is roughly 12% of *Initiation*'s standard deviation (28 pp). Column 2 shows that the lagging startup is 2.6 pp more likely to repurpose existing projects into the non-common ICD, and column 3 shows it is 3.6 pp less likely to repurpose projects away from the non-common ICD. In column 4, we see that the lagging project is 0.9 pp more likely to form an alliance in the non-common ICD. That result is significant only at the 10% level, whereas the previous results' t-statistics range from 2.2 to 3.8 in magnitude. All four results are consistent with the lagging startup redirecting innovation toward ICD categories without common ownership. Using the example from the previous figure, we find that Startup A's innovation efforts increase in ICD a relative to ICD b. This redirection occurs through starting new projects, repurposing existing projects, and creating alliances with large pharma companies.

How can we square these results with the earlier finding that lagging startups are likely to lose their common VCs' funding? Recall that this analysis includes startup-by-quarter

FEs. Those FEs are important, because a lagging startup is likely to lose the VC's funding and hence is likely to reduce innovation activities in all ICD categories. The interesting comparison is therefore across ICDs within a startup-quarter. Even if the startup's overall level of innovation decreases, we predict more activity in its non-overlapping ICD category relative to the overlapping one. Our results support that prediction. Even if the common VC stops funding the lagging startup, the VC's existing control rights plausibly give it enough power to redirect innovation.

1.4. Welfare implications

We show above that common ownership leads investors to hold back lagging projects and startups, and to redirect their innovation. By reducing the duplication of R&D costs in this way, common ownership can improve innovation efficiency. Specifically, common ownership can reduce the total R&D cost, aggregated across all startups and projects, per drug that makes it all the way to FDA approval. This subsection provides descriptive evidence that common ownership indeed improves innovation efficiency, which has a welfare benefit. We then discuss other ways that common ownership can affect welfare, positively and negatively. Since we do not quantify all these channels, we make no conclusion about whether the total effect of common ownership on welfare is positive or negative.

We show that common ownership rates are positively correlated to a proxy for innovation efficiency. We perform a simple cross-sectional comparison of drug categories. We measure an ICD category's common ownership rate as the probability that a randomly selected pair of startups active in the ICD share at least one VC in common, averaged over startup pairs and the sample period. Our efficiency proxy seeks to measure the number of approved drugs per dollar of R&D spending, aggregated across all projects. The proxy equals the number of drugs within the ICD category that reach FDA approval during our sample period, divided by the total dollar amount of funding provided by VCs to all startups (successful and unsuccessful) that are active in the given ICD during the sample period. If a startup has projects in multiple ICD categories, we allocate the startup's VC funding to

ICD categories in proportion to the startup's number of projects in each category. Using VC funding as a proxy for R&D spending is reasonable given that most pharma VC dollars go toward drug development as opposed to, say, marketing.

Figure A.5 plots innovation efficiency versus the common ownership rate, with each circle representing an ICD category. Common ownership rates vary significantly across ICD categories. Several ICD categories see no drug approvals and therefore have efficiency measures equal to zero. Other ICD categories have efficiency measures as high as 0.005, which corresponds to 1/0.005=$200 million of aggregate VC funding per approved drug. The dashed best-fit line shows a positive relation between the efficiency measure and the common ownership rate. Column 1 of Table A.9 shows the underlying regression and confirms that the positive relation is highly statistically significant ($t = 3.6$). Economic significance is also quite high. A one standard deviation increase in the common ownership rate is associated with a 0.34 standard deviation increase in the efficiency measure.[19]

We perform a number of additional checks. In column 2 we control for potentially important omitted variables. To capture the idea that drug development is costlier or more prone to failure in certain ICDs, we control for the average duration between project initiation and reaching Phase III and the probability of reaching Phase III. We compute these measures for Phase III rather than FDA approval, since many ICDs experience no FDA approvals. One limitation of our efficiency measure is that it excludes funding raised after the VCs exit, e.g., in an IPO or after an acquisition. Although it is not clear why those exclusions would relate to common ownership rates, we control for VCs' average holding period to proxy for how early or late VCs exit their investments. None of these controls enter significantly, and the positive relation between efficiency and common ownership weakens only slightly. In the Online Appendix, we show that results are robust to measuring common ownership rates differently or measuring efficiency based on the number of projects reaching Phase III rather than FDA approval.

[19]The standard deviations of *Approval Efficiency* and *Common Ownership Rate* are 0.00111 and 0.0123, respectively, and $0.34 = 0.031 \times 0.0123 / 0.00111$.

Our efficiency proxy ignores that some drugs have larger impacts than others. We find no evidence, however, that commonly owned drugs have significantly higher or lower impact. We proxy for impact using drugs' patent citation counts, and we find the relation between common ownership and impact to be statistically and economically insignificant (Online Appendix).

Even with these checks, this remains a simple, descriptive analysis. Common ownership rates are not randomly assigned to ICDs, so these results do not have a causal interpretation. Nevertheless, the results are consistent with common ownership improving innovation efficiency by avoiding the duplication of R&D costs.

There are at least three other ways that common VC ownership can improve welfare. Commonly owned startups are more likely to share resources and information (e.g., González-Uribe, 2020, Eldar et al., 2020). Common ownership can create tournaments within VC portfolios, strengthening incentives (e.g., Inderst et al., 2007). Common ownership can also improve welfare by promoting entry. Since our analysis holds the set of projects or startups fixed, we take an ex post view that ignores entry. In theory, common ownership can promote entry, because the option to abandon lagging startups ex post can encourage VCs to invest in more startups ex ante.

Common VC ownership can also impose social costs, in two related ways. First, there is a view that R&D is below its socially optimal value, due to positive R&D spillovers (see Hall et al. (2010) for a survey). If this is true, then our paper implies that common VC ownership curtails innovation further by shutting down competing projects. Mitigating this cost, though, pharmaceutical R&D arguably has low spillovers (e.g., Benmelech et al., 2021) and often goes toward developing less-novel "me-too" drugs with limited social value (DiMasi and Faden, 2011, Krieger et al., 2022).

Second, by holding back lagging Phase I projects, common owners can potentially reduce the number of drugs that make it all the way to FDA approval. Limiting consumer choice in this way can have negative health consequences in cases where the lagging and pioneering drugs are not close substitutes, because the suppressed drug project may be a better treatment for certain patients. Also, if fewer drugs reach the market, there can be less competition and higher drug prices. To explore these social costs, we ask whether common ownership is indeed associated with a lower fraction of drugs reaching FDA approval. We compute each drug category's ratio of number of approved drugs to number of initiated drug projects, aggregating years. We regress this ratio on the drug category's common ownership rate. We find a positive relation, inconsistent with common ownership reducing future product-market competition (Online Appendix). This is just a descriptive analysis, however, so we cannot rule out that common ownership imposes these social costs.

1.5. Identification challenges

Next, we address four potential identification challenges. The first is that our proximity instrument simply picks up the ease of poaching employees from a competing startup (e.g., Hombert and Matray, 2017). For example, if startup A experiences success, it could poach employees from a nearby, competing startup B, causing B to fail. This story, in theory, can explain why nearby startups are more likely to have opposite outcomes. If this story were true, then our results should be weaker in states that strongly enforce employee non-compete agreements, because poaching is less feasible there.[20] The evidence in Table A.10 does not support that prediction, however. We create a dummy variable for whether non-compete agreements are strongly enforced in startup i's state in quarter t. We then interact that non-compete dummy variable with $Lagging \times SharedVC$ in our main project-outcome regressions. If the competing story were true, we should find a significantly positive slope on the triple-interaction term in column 1. Instead, it has a slightly negative coefficient with a t-statistic of -0.3. This result indicates that employee poaching is not spuriously generating our main results.

[20]Supporting this idea, Ewens and Marx (2018), Jeffers (2021), and others confirm empirically that non-compete agreements restrict worker mobility, especially in knowledge-intensive sectors. Our analysis does not require emlpoyee non-competes to be homogenously applied within each state. We simply need differences in enforceability to generate differences in average poaching propensity across states, which is plausible.

What about other forms of poaching? Poaching nearby customers or suppliers is irrelevant, because the demand for novel drugs is not local, and neither are suppliers. Poaching clinical-trial sites is also irrelevant, since trials are often performed nationally and, in many cases, internationally. Poaching nearby investors also does not explain our results. The competing story here is that nearby lagging projects shut down not because of common ownership, but simply because the pioneering startup poaches nearby investors. If this concern were valid, then we should see a startup struggle to raise financing after a nearby competitor makes progress, even if the two startups do not share a common owner. We do not see that pattern in the data, however, so this competing story seems unlikely (see Online Appendix).

Contrary to all these poaching stories, our main results continue to hold after controlling for whether the lagging and pioneering projects are in the same MSA (see Online Appendix). All forms of poaching are more likely for startups that are very nearby, e.g., in the same MSA. It is reassuring that our results are not confined to pairs of startups in the same MSA.

There is also an opposing force to these poaching stories. Suppose nearby drug projects A and B are using a similar technology to cure a given disease. If drug A experiences success, that sends a positive signal about the scientific underpinnings of drug B.[21] This positive signal makes B's startup and its employees (or investors) less willing to separate from each other, so poaching is less likely to occur. If instead the drugs use unrelated technologies, the poaching motive would arguably be lower because B's employees do not have expertise with A's technology.

That logic also helps address a second identification challenge, which is that our geographic-proximity instrument simply picks up the effects of easier information sharing between nearby startups. Rather than explaining our results, though, this story works against our results. Success of a nearby, related drug A sends a signal about B's scientific underpinnings that is either positive (if technologically similar) or neutral (if technologically unrelated),

[21] Supporting this logic, Krieger (2021) finds that a drug project is more likely to be terminated after a close competitor terminates its own project, implying that technological learning dominates competition effects.

which on average makes it more likely that B progresses to Phase II. We instead find that progress by A is *less* likely to predict progress by B if the two startups are nearby. An even simpler argument against the information-sharing story is that the type of information we study (i.e, whether a project progresses to Phase II) is not more likely to flow between nearby startups. Instead, progressing to Phase II is public information revealed through clinical trial registries, FDA submissions, patent filings, and company press releases. Further, information sharing is more likely to occur between pairs of startups in the same MSA, but our main results hold even after controlling for a same-MSA indicator (Online Appendix).

A third potential challenge is that we find opposite outcomes for nearby projects because they are more technologically similar, e.g., because scientists share similar backgrounds. Fearing it will not be first to market, a startup may abandon a project after seeing a technologically similar peer gain an edge. We address this concern by controlling for overlapping patent citations, a proxy for technological similarity.[22] We repeat our main panel regressions while controlling for $Lagging \times SharedCite$, an indicator for whether the lagging startup's drug shares at least one citation in common with the progressing drug. OLS regressions still produce a significantly negative slope on $Lagging \times SharedVC$ when we control for $Lagging \times SharedCite$, for both project and funding outcomes (Table A.11, columns 1 and 3). In the Online Appendix, we show that our IV results continue to hold after controlling for $Lagging \times SharedCite$. Also, if the competing story were true, then we should find a negative correlation between $Progress$ and $Lagging \times SharedCite$. Instead, $Progress$ is positively related to $Lagging \times SharedCite$, consistent with the argument that seeing a close competitor make progress, even absent a shared VC, is a positive signal, not a negative signal. Overall, these results indicate that our instrument is not simply picking up the omitted effects of technological similarity.

[22]Cunningham et al. (2021) use an alternative proxy based on drugs' mechanism of action. We do not use that proxy because the data on mechanism of action are missing much more often for small, private companies than large, public companies. Yet another proxy is based on chemical similarity of small-molecule drugs (e.g., Krieger et al., 2022). That proxy, however, is unavailable for biological products, which make up a large portion of our sample.

A final concern is that our results stem from VCs following a "spray and pray" investing approach in certain markets (Ewens et al., 2018). Suppose in some product markets exactly one project will ultimately succeed, which makes projects' outcomes negatively correlated. If VCs "spray and pray" in such markets, then these markets will feature more common ownership, so common ownership will be associated with projects' reaching opposite outcomes. This story cannot explain our results, however. Even our simplest OLS regressions include ICD FEs, which difference out any effects of a negative correlation at the product-market level. The story is even less likely in our test with ICD-quarter FEs, which difference out potentially time-varying correlations within product markets. Also, this story cannot explain our IV results, because it is not clear why nearby drug startups would be more likely to belong to a winner-take-all market.

1.6. Robustness

For our main predictions to hold, the pioneering and lagging projects should be close competitors. Next, we show that our results are stronger when that condition is more likely to hold. Our proxy for "close competitor" is $Lagging \times SharedCite$, our indicator for whether the pioneering and lagging projects have an overlapping patent citation. We interact that term with $SharedVC$, creating a triple-interaction term, and we include it in our project-outcome and VC-funding regressions. We predict a negative coefficient on this term, meaning the negative relation between the dependent variable and $Lagging \times SharedVC$ is even more negative if the projects' technologies are more similar. Results are in columns 2 and 4 of Table A.11. The estimated coefficients are indeed negative. When the dependent variable is project-level success ($Progress$, column 2), the slope on the triple-interaction term has only marginal statistical significance. It is very large in magnitude, however, implying the effects of $Lagging \times SharedVC$ are roughly twice as large ($[0.017+0.020]/0.017]$) if $SharedCite = 1$. When the dependent variable is company-level VC funding ($ExtendFunds$, column 4), the slope on the triple-interaction term is large in magnitude and statistical significance.

This result implies that common owners are especially likely to hold back lagging projects that compete closely with the pioneering projects. These results support the patent-race story of Grossman and Shapiro (1987). They also help rule out a simpler story, which is that competition is irrelevant and VCs simply have limited capital. In that story, success by any portfolio company results in less funding for all other portfolio companies, regardless of how closely the companies compete. As a result, the triple-interaction terms would not enter significantly. Since they do enter significantly, our interpretation based on competition appears more plausible.

A potential concern is that drugs in the same ICD category are not always competitors. Those classification errors would inject noise into our analysis. To address this concern, we impose the more restrictive assumption that only drugs sharing the same three-digit ICD code are close competitors. Results are in the Online Appendix. This assumption increases the number of product markets from 78 to 200. It reduces the common-ownership rate (i.e., the percent of startups with a close competitor held by the same VC) from 39% to 25%, which is still quite high. More important, our main results on project and funding outcomes continue to hold. Levels of economic and statistical significance are similar in most cases and even higher in some cases.

Our main project-level tests take the form of panel regressions. For robustness, we show that results are similar when we collapse the panel into a single observation per project. We redefine $Progress_i$ to be an indicator for whether project i ultimately progresses to Phase II. $Progress_i$ equals zero if the project is eventually suspended or remains "in progress" indefinitely. We redefine $Lagging_i$ to be an indicator for whether project i falls behind a pioneer at any time between i's birth and the end of its Phase I. The results, shown in the Online Appendix, are quite similar to those in the full panel regressions. Whether we estimate by OLS, 2SLS, or bivariate probit, we find a significantly negative loading of $Progress$ on $Lagging \times SharedVC$, consistent with our prediction.

We also find support for our project-level predictions if we replace the dependent variable $Progress_{it}$ with its opposite, $Suspend_{it}$, an indicator for whether project i is suspended in quarter t. We focus on $Progress$ in our main tests because it is difficult to accurately measure suspensions. We predict a positive relation between $Suspend$ and $Lagging \times SharedVC$, meaning common ownership makes a lagging project more likely to be suspended. The data support this prediction (Online Appendix).

We also show in the Online Appendix that our main results are largely robust to dropping either Boston or San Francisco, our sample's two largest MSAs. Dropping either MSA has little effect on economic significance in the OLS and IV specifications. The OLS results remain highly statistically significant, but the IV results become significant at only the 5% confidence level, due largely to a weaker first-stage regression.

1.7. Conclusions

We study the effects of common ownership on innovation. Our evidence comes from detailed project-level data on pharmaceutical startups and their VC investors. We find that common ownership leads VCs to shut down lagging drug projects and withhold funding from lagging startups. Also, startups redirect their innovation away from areas where they have fallen behind commonly owned peers. Taken together, these results are consistent with common owners reducing duplication of R&D costs, which can help solve a market failure among competing firms (Loury, 1979). Consistent with common ownership improving innovation efficiency, common ownership is positively correlated with the number of approved drugs per dollar of aggregate VC funding.

We do not claim that common VC ownership improves welfare. Instead, we document a specific social benefit of common ownership. At the same time, common ownership can entail social costs by reducing product-market competition and consumer choice. Weighing these and other channels to determine common ownership's total effect on welfare is an important area for future research.

CHAPTER 2

Bargaining with Private Equity: Implications for Hospital Prices and Patient Welfare

2.1. Introduction

The steep rise in U.S. healthcare prices has become a source of grave concern. In 2019, total U.S. health expenditure reached 18% of GDP, making the healthcare sector larger than the manufacturing sector (11% of GDP) and the energy sector (6% of GDP) combined (Panel A of Figure A.6). Healthcare costs are projected to grow at an average rate of 5.5% per year in the next decade, primarily driven by the surge in healthcare prices. At the same time, private equity (PE) investors have shown a keen interest in the healthcare sector, with a 10-fold increase in deal values over the last decade (Panel B of Figure A.6). Active PE involvement has the potential to greatly change this sector, which previously was dominated by public and not-for-profit players. But we still know very little about how PE works and through what mechanisms PE actually affects real outcomes. Proponents of PE argue that it can improve operational efficiency, lowering health costs. Opponents claim that PE investors are tough negotiators who care only about profits and load up hospitals with debt, hence PE deals push up health expenses and make patients worse off. Because of the unknown consequences for patients, providers, and healthcare prices, PE involvement in the healthcare sector has raised concerns among regulators.[23] For all these reasons, it is important to understand the effects of PE on healthcare pricing and patient welfare.

This paper studies the rise of PE in healthcare using the hospital setting. This setting is important as hospitals constitute the largest component of healthcare spending ($1.2 trillion dollars in 2019) and play a central role in meeting population health needs. Focusing on the setting, I address the following questions: How do PE buyouts of hospitals affect healthcare pricing and spending? What are the specific channels by which PE investors

[23] For example, Congress has recently passed legislation to curb the "surprise medical billing" crisis, in which PE-backed physician staffing firms predominate (Appelbaum and Batt, 2020).

create and redistribute value among stakeholders? What are the implications for patient welfare and antitrust regulations? There are several key novelties of this paper in exploring these understudied questions. First, I use new, nationwide large-scale insurance claims data to capture detailed hospitals' transaction prices. It overcomes the data challenge since hospital prices are typically negotiated between hospitals and private insurers and thus difficult to observe. Second, I apply a novel structural approach that exploits state-level regulation changes as PE entry shocks to quantify PE's impacts on negotiated prices and patient welfare. Importantly, it overcomes the empirical challenge in the existing literature to quantify the channels by which PE creates and redistributes value and to decompose PE's effects into operational versus financial effects. The model also demonstrates new economic mechanisms in bilateral price bargaining, which apply to other business-to-business settings beyond the hospital sector. Third, this paper sheds light on antitrust regulations. It highlights that traditional antitrust regulations fail to account for the role of PE-backed acquirers in merger reviews, potentially leading to an underestimation of proposed mergers' impacts on consumer welfare.

Despite its importance, empirical research on how PE affects healthcare pricing and price bargaining in a business-to-business setting is scarce. One challenge is to assemble a dataset of sufficient detail and scope to credibly identify the impact of PE buyouts. Typically, the outcomes of hospital–insurer price negotiations are treated as commercially sensitive and thus are largely unavailable to researchers. I introduce new insurance claims data covering over 60% of individuals with private health insurance in the United States from 2013 to 2019. My final sample includes more than 600 million claims sourced from around 5,000 hospitals. The data include hospitals' detailed transaction prices charged to private insurers for each medical procedure provided during patients' visits. For PE-ownership information, I manually match claims data with a comprehensive list of PE investments in hospitals over the last two decades. The novel combination of datasets enables me to investigate PE's impacts on hospital pricing.

I begin by presenting descriptive evidence that PE investment in hospitals leads to sharp increases in their negotiated prices with private insurers. These increases cannot be attributed to market consolidation alone. After accounting for a rich set of observables and fixed effects, I find that negotiated prices rise about 69% within a hospital–insurer pair after PE interventions compared to a matched control sample. The results hold even when I compare prices of a subset of homogeneous medical procedures—specific medical imaging procedures—contracted between a hospital and an insurer. In contrast, hospital service utilization does not change significantly; I find that the units of medical resources used to treat a patient with a given diagnosis do not change after PE buyouts. This result suggests that PE investors mainly affect prices, not quantities.

PE intervention also imposes heterogeneous spillovers on the negotiated prices of rivals in the local market. After a PE investor acquires any hospital in a region, local non-PE-backed rivals that share common insurers with the PE-backed hospital experience an average 8.1% increase in negotiated prices. No such pattern is observed among other local rivals. This contrast shows that the existence of spillover effects closely hinges on whether a rival is connected with a PE-backed hospital through a common insurer in the "bargaining network."

Motivated by these patterns in the data, I develop and structurally estimate a model of price bargaining between hospitals and insurers. The structural approach helps overcome several empirical challenges. First, the existence of heterogeneous spillovers in price bargaining implies that the stable unit treatment value assumption (SUTVA) is violated, possibly introducing bias in reduced-form potential outcomes. Second, reduced-form techniques allow measuring the direction of various channels for changes in negotiated prices. However, evaluating their magnitudes requires a model. Third, quantifying patient welfare and exploring counterfactual experiments also requires the structural model.

The model presented here features PE buyouts, leverage choices of target hospitals, bilateral hospital–insurer bargaining over prices, and hospital choices of patients. It extends the price bargaining model in Gowrisankaran et al. (2015) and Ho and Lee (2017) to include PE investments as well as the associated financial and operational changes. On the demand side, I model patients' hospital choices as dependent on a rich set of patient and hospital characteristics, including patients' health conditions, age, gender, and travel distance to the hospital; hospitals' size and availability of specialty services; and other factors. I also incorporate the impact of PE buyouts on patients' choices by including a PE-ownership indicator and its interaction terms with patient demographics in the demand model. These variables capture any impact on patients' choices due to changes in service quality after buyouts.

On the supply side, private insurers bilaterally negotiate a benchmark price with hospitals in the local market. Price negotiations are modeled as a modified Nash-in-Nash bargaining game in which all bilateral negotiations occur separately and simultaneously; renegotiations occur when the local market structure shifts due to hospital bankruptcies. In the bargaining process, PE-backed hospitals can credibly threaten bankruptcy if their expected revenues are unable to meet the debt burden. PE investors optimally choose whether to invest in a hospital. They also choose financial leverage to maximize the expected payoffs of target hospitals.

PE investors can affect bargained prices through five channels in the model. These channels cover key mechanisms of PE creating and redistributing value through financial, operational, and governance engineering. First, financial engineering of PE has direct implications: High debt loads render credible threats of bankruptcy from PE firms. On the one hand, bankruptcies of PE-backed hospitals make the outside option of insurers worse since they will have to face a more concentrated hospital sector and higher service prices charged in the local market. Therefore, insurers have incentives to "subsidize" hospitals and keep them competing with each other. On the other hand, bankruptcy threats make the outside option of equityholders more attractive as they can walk away if net revenues are under water. This mechanism spotlights a close and unique relation between financial characteristics of PE buyouts and price bargaining outcomes. It is also consistent with the descriptive evidence that PE-target hospitals with higher leverage have higher prices, whereas leverage is not significantly related to prices among non-PE-backed facilities.

Second, PE firms, as veteran negotiators in financial markets, could potentially bring in new negotiating strategies and contracting expertise that improve hospitals' bargaining power. The model is flexible enough to separate bargaining power (the inherent ability of negotiators to secure a fixed share from the bargain) from the general concept of bargaining position (the outside option of negotiators if the bargain fails). This decomposition allows me to examine how PE buyouts affect hospitals' bargaining power.

Third, the model recognizes that PE intervention affects hospitals' operational efficiency. As shown in the descriptive evidence, a hospital's adjusted average costs per patient discharge indeed decrease after PE investments. In the model, operational efficiency is reflected by the marginal costs of hospitals and thus is factored into price negotiations.

Fourth, the model also incorporates a non-pecuniary motive of hospitals. After PE buyouts, hospitals are more profit-driven and have less focus on their social objectives. Providing medical services becomes less appealing for them unless it is lucrative enough, which presses insurers to offer a higher price. Supporting this mechanism, I show in the descriptive evidence that price increases are more pronounced among hospitals that were not-for-profit prior to buyouts.

Last, the model allows PE investors to alter service quality and hence change patients' demand for the hospital. This demand shift has implications for bargaining outcomes because the redistribution of patient flows in the local region affects both insurers' and hospitals' outside options.

The model also captures spillover effects of price bargaining in the local market. If a PE-backed hospital increases its negotiated price with an insurer in equilibrium, then reallocating patients from local rivals to the PE-backed hospital becomes more costly. Thus, when negotiating with these rivals, the value of the insurer's outside option would diminish if the bargaining breaks down. As a result, the negotiated prices between the insurer and local rivals are predicted to rise as well. Note that only insurers that have contracts with PE-backed hospitals experience disruptive changes in outside options. Stronger spillover effects are expected among hospitals in these insurers' networks, consistent with the heterogeneity described earlier.

The main identification challenges in estimating the model are that PE investors endogenously select target hospitals, and PE-backed hospitals endogenously choose their financial leverage. I tackle the endogenous selection of target hospitals by exploiting state-level regulatory changes regarding the corporate practice of medicine doctrine (CPOM), which prohibits nonprofessional business corporations (e.g., PE firms) from practicing medicine or employing physicians to perform medical services. The content of CPOM is generally shaped by state laws, court rulings, and medical board opinions, and has evolved over time in different states. These changes provide exogenous variation in PE's entry costs into the hospital sector. I manually collect a series of legislative events and construct a CPOM regulation index as an instrumental variable. For the endogenous selection of leverage, I use the annual ICE BofA U.S. high yield index option-adjusted spread as the instrumental variable. The rationale for this approach is that the market price of credit risk affects PE's debt financing costs but is uncorrelated with idiosyncratic shocks to hospitals.

Estimates of demand parameters reveal preferences of patients and their expected utility for each hospital. The granularity of insurance claims data enables observation of a patient's choice set in the local market. Estimates from the multinomial logit choice model reveal that enrollees respond to PE intervention. Typical patients prefer visiting PE-owned hospitals, particularly for female enrollees, yet more senior patients and those with more severe diseases are more likely to go to a non-PE-backed hospital. Consistent with previous literature, I find that enrollees prefer hospitals nearby, hospitals they have visited before, and hospitals that provide specific services accommodating their needs.

Estimates of supply parameters (the bargaining model) come from moment conditions based on orthogonality restrictions on negotiated prices, marginal costs, and leverage choices.I find that PE significantly increases hospitals' bargaining power, consistent with PE sponsors providing superior negotiating expertise. On average, PE-backed hospitals can extract 50% more of the surplus generated by a hospital–insurer relationship than non-PE-backed hospitals. I also find that a higher level of financial leverage renders a more credible bankruptcy threat, improving PE-backed hospitals' bargaining positions. In addition, parameter estimates imply not-for-profit hospitals enjoy non-pecuniary benefits, equivalent to \$102 of pecuniary rewards for every unit of service provided to local communities. These non-pecuniary benefits vanish among PE-backed hospitals because they focus less on social objectives. The estimates also show that PE improves hospitals' operational efficiency by reducing marginal costs by 8% on average.

I use the estimated model as a laboratory to conduct a series of counterfactual analyses. First, I consider a counterfactual PE ban in the hospital sector and its implications for negotiated prices and spending, spillover effects, channel decomposition, and patient surplus. Similar policies were recently considered in several state legislatures.[24] The results show that banning PE-based hospital takeovers would lead to an 11% reduction of health spending in affected regions, with savings mainly coming from a decrease of negotiated prices. The quantity-weighted average prices would drop by almost 11%, while the quantities of provided services would barely change in the counterfactual. The savings also exhibit heterogeneity across regions. A local market having a higher percentage of PE-backed hospitals would save more in the counterfactual.

Looking at the spillover effects in the PE-ban counterfactual, the results show that PE-backed hospitals would contribute 86% of the total savings if PE ownership were banned. Another 14% would be contributed by local rivals that share common insurers with the PE-backed hospital, while the remaining hospitals (local rivals that do not share insurers with

[24]For instance, the California legislature in 2020 introduced a bill (SB-977) that aimed to rein in PE healthcare buyouts. The bill gave the state attorney general power to review and potentially block a broad range of merger deals in the healthcare sector involving private equity groups.

the PE-backed one) would contribute almost nothing to the total savings. These results are largely consistent with the reduced-form results regarding heterogeneity in spillover effects. More specifically, I find that banning PE ownership results in a 49% decrease of the quantity-weighted average prices among PE-backed hospitals and a 3.2% decrease among local rivals that share insurers with the PE-backed one. Both groups experience only a modest change in service quantities in the counterfactual: PE-backed hospitals see a 5.5% increase in service provisions while local rivals that share insurers experience a 3.1% reduction.

Next, I use the model to quantify the relative contribution of various channels to the total savings. I do so by adding one channel at a time into the counterfactual and computing total spending changes in each scenario. I find that superior bargaining skills of PE sponsors contribute about 43% of the savings, making it the largest contributor. There are many anecdotes of PE investors bring better negotiating expertise to their portfolio companies. I find strong empirical support for those stories in the hospital–insurer setting. Financial engineering and bankruptcy threats contribute an additional 40% of the savings, implying that the financing structure of PE buyouts can generate significant real effects in markets. Further savings come from the other two channels: changes in patient demand and the focus on social objectives, which contribute 10% and 8%, respectively. In contrast, changes in operational efficiency contribute about −1% because banning PE would eliminate a small gain in operational efficiency.

Finally, I quantify how patient surplus would be affected in the PE-ban counterfactual. Under the assumption that the savings on medical expenses would benefit patients through reductions in insurance premiums, the model implies that banning PE would bring patient-surplus gains equivalent to 10.7% of the documented health expenses in PE-target regions. The gains largely come from savings in outpatient expenses, even though small reductions in service quality would negatively impact patient surplus if PE ownership were banned.

In the second counterfactual, I use the model to evaluate policy implications for merger reviews. Specifically, I consider 100 hypothetical mergers in 2013 wherein a PE-backed hospital system is randomly assigned to acquire one rival system locally. I assess the impacts of those mergers using the full model estimated in this paper (PE model) and a re-estimated plain model after removing PE-related features (No-PE model). Note that the No-PE model is essentially the same model used by regulators to evaluate proposed mergers and only takes into account the market concentration effect induced by mergers. I compare the two models' predictions regarding total spending, quantity-weighted average prices, and patient surplus before and after mergers.

The results imply that regulators risk greatly underestimating the impacts of proposed mergers if they disregard PE acquirers' unique features. Not surprisingly, both the full and the plain models predict an increase in total spending and quantity-weighted prices in those regions after hypothetical mergers. The magnitudes are positively correlated with the change in the Herfindahl-Hirschman index (HHI) of the hospital sector after mergers. However, the models produce very different magnitudes: Percentage changes in total spending and quantity-weighted prices after mergers predicted by the No-PE model can be 10 points lower than the PE model. The prediction gap between both models still remains when we examine the changes in patient surplus after mergers. The underestimation mainly comes from merging hospitals, among which the prediction gap in quantity-weighted average prices, for example, could be as high as 25%. For non-merging rival hospitals, the underestimation is much less severe. These results highlight the importance of taking PE acquirers' traits into account when reviewing proposed mergers.

Contribution to the Literature

This work contributes to the empirical literature on the effects of PE buyouts.[25] A number of papers have examined the effects of PE buyouts on firm productivity and growth (e.g., Kaplan, 1989, Lichtenberg and Siegel, 1990, Boucly et al., 2011, and Cohn et al., 2014), firm employment (e.g., Davis et al., 2014 and Davis et al., 2019), industry performance (e.g.,

[25] Kaplan and Strömberg (2009) provide a nice survey on the economics of leverage buyout and PE. More recently, Gulliver and Jiang (2020) give a comprehensive review on studies of PE's impacts on target firms' productivity, revenue per employee, return on assets, and innovation.

Bernstein et al., 2017), financial stability (e.g., Bernstein et al., 2019 and Johnston Ross et al., 2021), workplace safety (e.g., Cohn et al., 2020), restaurant quality (e.g., Bernstein and Sheen, 2016), student outcomes (e.g., Eaton et al., 2020), and product prices (e.g., Fracassi et al., 2021). This paper adds to this literature by exploring how PE creates and redistributes value among relevant stakeholders in a business-to-business price-bargaining setting. Prominently, it uses a novel structural approach to decompose PE's real effects into operational versus financial channels, and to quantify its implications for hospital pricing, patient welfare, and antitrust regulations. It is the first to introduce the insurance claims "big" data from the DRG Real World Data product and a novel identification strategy by exploiting state-level regulation changes to address the endogenous selection concern of PE investments.

More closely related is a strand of literature in health economics studying the impact of PE investment (e.g., Stevenson and Grabowski, 2008, Pradhan et al., 2014, Huang and Bowblis, 2019, Gondi and Song, 2019, Bruch et al., 2020, Gandhi et al., 2020, Bruch et al., 2021, Cerullo et al., 2021, and Offodile et al., 2021). In recent work, Gao et al. (2021) examine employment and patient outcomes at hospitals acquired by PE and Gupta et al. (2021) study whether PE buyouts affect the mortality outcomes of patients in nursing homes. My paper differs in several important ways. First, the research question is different. Besides patient outcomes, this paper investigates how PE ownership impacts bargained healthcare prices, healthcare costs, and consumer surplus in the market of private insurance. Second, this paper provides a new structural IO-finance approach combined with a novel identification strategy to quantify the effects of PE on consumer surplus. It enables a decomposition of PE's impacts into various channels, which brings new insights on why and how PE creates and redistributes value in the market. Third, this paper complements the findings in Gao et al. (2021) and Gupta et al. (2021). In additional to the new pricing implications, I find that PE on average improves hospital service quality. Furthermore, I document that PE buyouts have heterogeneous spillovers to the price negotiation of local hospitals.

More broadly, this paper contributes to the literature on firm ownership by unveiling how changes in ownership type from not-for-profit to PE-owned affect firms' pricing strategies and consumer welfare. Early discussions include Hansmann (1980) and Lichtenberg et al. (1987). The hospital sector is a perfect setting to study related questions given its diverse ownership types.[26] The seminal work of Duggan (2000) uses a plausibly exogenous change in hospital financing to test three competing theories of firm ownership. He finds that different types of ownership respond to financial incentives differently. More recently, Capps et al. (2020) use detailed California data and find no evidence indicating that not-for-profit hospitals are more likely to provide charity care compared to for-profit ones. Different from these papers, I focus on PE ownership transitions and their implications on price negotiations and patient welfare.

It also contributes to the industrial organization literature on hospital transaction prices and hospital competition. Unlike most of the literature focusing on variation in Medicare hospital prices (e.g., Finkelstein et al., 2016),[27] this paper is the first, to my knowledge, to use a national multiyear dataset of transaction prices from the DRG Real World Data in the finance or economics literature. The newly introduced data contribute to the discussion of hospital prices for the privately insured and shed light on how price variations arise from PE buyouts. Regarding hospital competition,[28] this paper extends the models of bargaining and competition between hospitals and insurers from Gowrisankaran et al. (2015) and Ho and Lee (2017). Unlike their papers, I examine a different research question and uncover new channels, in particular the financing side of PE buyouts, that impact hospital price bargaining. Through my second counterfactual experiment, I discuss new policy implications in antitrust regulations.

[26]There is a literature in finance using the hospital context to study how firms of different ownership types respond to financial shocks, including Adelino et al. (2015), Adelino et al. (2020), and Aghamolla et al. (2021).

[27]One prominent exception is Cooper et al. (2019). They quantify price variations within hospitals by using insurance claims data from the Health Care Cost Institute (HCCI) that cover the privately insured nationwide from 2007 to 2011. Lopez et al. (2020) provide a literature review on the differences of hospital transaction prices for private insurers and Medicare.

[28]In survey papers, Gaynor and Town (2011) and Gaynor et al. (2015) introduce the general framework to study the impact of competition on price and quality.

Finally, my analysis relates to the empirical literature on bargaining. Besides those papers on the healthcare sector mentioned above,[29] there is a strand of literature discussing bargaining problems in the financial market. Matsa (2010) provides evidence that firms use debt financing as a strategic tool to improve their bargaining position with organized labor. Robles-Garcia (2020) analyzes the welfare implications of regulating brokers' services and compensation when brokers and lenders bargain over commission rates. Dou et al. (2020) study sources of inefficiency in the bankruptcy process in which senior and junior creditors bargain with two-sided incomplete information. Brown et al. (2009) investigate if leveraged buyouts strengthen the bargaining positions of portfolio companies when facing their suppliers. Unlike these papers, I study a different question, namely, how PE buyouts affect portfolio companies' bargaining power and position in a business-to-business setting.

2.2. Institutional Details and Data

2.2.1. Institutional Details

Hospitals are major providers of medical, diagnostic, and treatment services in the United States. In 2019, these facilities accounted for over one-third of total health spending (Panel A of Figure A.7), and accommodated over 145.6 million patient visits in emergency rooms. According to the most recent American Hospital Association (AHA) annual survey, there are over 6,000 hospitals in the United States; among them, about 3,000 hospitals are not-for-profit, 1,200 are investor-owned for-profit, and the rest are government-owned (federal, state, or local government). Typically, hospitals provide both inpatient and outpatient services: inpatient includes any services for patients who are formally admitted to a hospital by doctor's order, while outpatient includes emergency department services, observation services, outpatient surgery, lab tests, or any other services when doctors do not admit a patient.

Hospital spending is mainly determined by patient volume and payment rates per service

[29]Other empirical papers studying hospital price bargaining include Grennan (2013), Lewis and Pflum (2015), Ho and Lee (2019), Dafny et al. (2019), and Barrette et al. (2021). A survey by Grennan and Swanson (2021) discusses the most recent developments in the literature.

or diagnosis, most of which are paid by insurers. Besides public health insurance programs like Medicare and Medicaid, commercial insurers are major payers and cover over 50% of the total population.[30] For public health insurance, the reimbursement rates are set unilaterally by the government (e.g., the Centers for Medicare & Medicaid Services [CMS]) and identically applied to providers across the nation, except for adjustments to account for local price and wage variations and other cost factors. Regarding commercial insurance, insurers must periodically negotiate local-market hospital contracts on service prices, which could be tremendously different from the charged prices in hospitals' chargemaster (listing prices on their menus). Such negotiations are commonly bilateral bargaining and could take months to reach an agreement. The finalized contract will not specify negotiated prices for each item as millions of prices would have to be listed. In most cases, a benchmark price is negotiated and used to multiply the service weights provided to calculate the final prices, which is similar to the reimbursement rules in Medicare's Hospital Inpatient/Outpatient Prospective Payment System. Given these features, hospitals in a local market compete with each other by deciding which private insurers' networks to join and negotiating prices bilaterally, unlike the standard Bertrand price competition. Patient demand for hospital service is very inelastic because patients pay little out-of-pocket costs.

Though reimbursement rates diverge between public and private insurance, hospitals are able to sustain their profitability by negotiating much higher private payment rates. Previous research has shown that hospitals have quite stable profit margins, which have risen to their highest levels in the past decade thanks to the expansion of health insurance coverage after the Affordable Care Act. A 2016 AHA report reveals that the profit margin of the hospital industry was estimated to be 7.8%, outperforming many other sub-industries in the healthcare sector, including health insurers, pharmacies, and pharmacy benefit managers.

[30] More statistics are shown at the Kaiser Family Foundation website https://www.kff.org/other/state-indicator/total-population/. Specifically, Medicare and Medicaid cover 14.2% and 19.8% of the total U.S. population, respectively. Among privately insured individuals, 89.3% are covered by employer-based private insurance, which is typically purchased by employers for their employees and financed through employer or joint employer–employee contributions. The rest are covered by non-group insurance, which is purchased directly from an insurance company either as policy holder or as dependent.

Both the huge cash flow and stable profitability of the industry greatly intrigue PE investors. The last decade has witnessed tremendous PE money and deals in the hospital arena. Panel B of Figure A.7 exhibits the number of PE deals and the number of hospitals that were involved in buyouts between 2006 and 2019.

On the one hand, PE can bring many turnarounds to the table. As deep-pocket investors, PE firms are able to provide capital for necessary investments such as IT system upgrades. As professional investors, PE can add value by improving operational efficiency and recruiting the very best management for hospitals. When dealing with hospitals' suppliers and customers, they can also provide useful skills and experience to negotiate better terms. As veteran financial advisers, they can advise hospitals to adopt a better financial structure and more flexible financial tools. On the other hand, many argue that PE leads to adverse consequences. PE demands a higher rate of return than other traditional financing sources and such a high-profit-driven motive conflicts with hospitals' missions. High debt loads resulting from financial engineering potentially press hospitals to take more aggressive cost-cutting measures and staff reductions, hence compromising their service quality. PE has a structured investment timeline, and such a short horizon renders it impossible to focus on long-term success.

2.2.2. Data

The main data used in the paper are from the proprietary database DRG Real World Data Product (RWD Product). The RWD product was developed by Decision Resources Group,[31] a global information and technology services company that provides proprietary data and solutions to the healthcare industry. A more detailed discussion about the data is provided in the Online Appendix. The raw RWD data contain over 26 billion claims, track more than 300 million longitudinal patient lives in the United States, and cover over 98% of U.S. healthcare plans (private insurance, Medicare, and Medicaid, but not uninsured patients). The database includes claim rejection rates, adjusticated payments, and out-of-pocket costs

[31] The company was acquired by Clarivate Analytics from Piramal Enterprises Limited in May 2020.

for detailed cost analysis. The database is built on the DRG Repository that sources claims from various originators, including hospitals, physician offices, pharmacies, long-term care facilities, and nursing facilities.

One contribution of this paper is that it is the first, to my knowledge, to introduce and use DRG RWD data in healthcare finance and economics. The prominent alternative insurance claims data include the IBM MarketScan and Health Care Cost Institute (HCCI) databases, which have been used in several previous papers. Both are closed-network sources in the sense that MarketScan's claims are sourced primarily from commercial and/or employee-sponsored health plans and HCCI's claims come from three insurance companies in the United States: Aetna, Humana, and UnitedHealth.[32] In contrast to those closed-network sources, the DRG RWD Product contains open-network claims sourced from billing software used by healthcare service providers and pharmacies. So, it has an overall larger capture of patients with claims data and is more representative across multiple payer categories compared to closed-network sources. Another advantage is its ability to track patients even after they change payers, which is usually infeasible in closed-network databases. Most importantly, DRG RWD claims can be updated daily, granting access to the most up-to-date medical claim information, while others, such as the HCCI database, can only access claims data with a delay.

The RWD Product includes a unique provider identifier (the National Plan and Provider Enumeration System Identifiers [NPI]),[33] a unique patient identifier, a unique payer/insurer identifier, the date services were provided, the date payments were processed, the hospital's charges and negotiated prices (sum of the payer paid amounts and patient paid amount), and a detailed decomposition of the patient's payments such as deductibles, coinsurance payments, and copayments. In addition, the data contain full diagnosis information in

[32] HCCI suspended their data services in 2019 when UnitedHealth terminated their agreement of claim data sharing. HCCI's dataset "2.0" is currently under construction with Blue Health Intelligence as a new raw data provider.

[33] As described in the sample construction procedure in the Online Appendix, I leverage the primary taxonomy code in the CMS NPI Registry to restrict the main sample to hospital facilities. Nursing homes, physician offices, and other facilities are excluded in the analysis.

terms of International Classification of Disease (ICD) codes; medical procedures in terms of Healthcare Common Procedure Coding System (HCPCS) codes; patients' demographic information such as 3-digit zip code prefix of their residency, gender, and age; insurer's identity if they are commercial, Medicare, or Medicaid; and types of patients' insurance plans. One advantage of the RWD data is the ability to identify payers, which enables comparison of price variations within hospital–insurer pairs and makes it possible to study the impacts of PE on price bargaining outcomes given a hospital–insurer pair.

For the purpose of this analysis, I focus on a sample of outpatient hospital claims of private insurers in the RWD data. For a given set of provided health services, final payments are generally broken down by institutional payments that paid for facility services, and professional payments that go to physicians. Institutional and professional payments are processed separately. As the first filter, this paper restricts analysis to insurance claims with institutional payments in the database. The second filter limits the sample to the outpatient setting, though previous work mainly focuses on the inpatient one.[34] I do this for two primary reasons. First, in the sample, PE-backed hospitals have better coverage of outpatient visits, which provides more power to detect the impact of PE intervention. Second, to construct the relative service-mix weights and benchmark prices in the inpatient setting requires the use of diagnosis-related group (DRG) codes, which are not reliably reported in the RWD product. Reconstructing DRG classifications from the data without proprietary software would introduce additional noise into the analysis. The detailed method for identifying outpatient claims in the sample is provided in the Online Appendix. The last filter restricts the sample exclusively to claims sourced from private insurers, which means that Medicaid, Medicare, and Medicare Advantage claims are excluded from the analysis. In the end, I obtain a sample of over 600 million commercial insurance claims in outpatient settings.

In the main analysis, I group claim lines to each patient visit and calculate the payment and

[34] Some exceptions, for example, include Liebman (2016) and Prager and Tilipman (2020).

medical information up to the patient-visit level. To identify the *Disease* category for each patient visit, I follow previous literature (e.g., Shepard, 2016) to group ICD-10 (or ICD-9) diagnosis codes in medical claims into 285 mutually exclusive Clinical Classification Software (CCS) single-level categories. The sample contains 78,742,811 patient visits. I drop any observations of patient visits to federal government-owned hospitals (e.g., army, veterans) and the Kaiser Foundation's hospital system. Any patient visits with missing information on patient age, gender, CCS category, payer identification, and total paid amounts are dropped. Patient visits with negative paid amounts or negative service-mix weights are excluded as well. Lastly, I require that a hospital has at least 10 patient visits a year in the sample. Applying all these filters leads to a final sample of 72,663,354 patient visits between 2013 and 2019.

The data on PE deals are from four proprietary sources: PitchBook, Preqin, Capital IQ, and SDC Platinum. I synthesize them to construct a comprehensive database on PE investments in hospitals. The primary source is a list of deals from PitchBook. Specifically, I restrict the industry to healthcare services with "Clinics/Outpatient Services" and "Hospitals/Inpatient Services" and select all PE deals that are already completed. I focus on deals that took place between 1994 and 2019 and that invested in companies located in the United States. In the end, I obtain a sample of 2,324 companies involved in 2,827 deals initiated by 1,339 investors. As a complement, I also download all PE deals in Preqin and restrict those deals to the healthcare sector in the United States within the same time frame. In addition, I restrict deal types to buyouts, public to private, merger, add-on investment, and growth capital. In a similar fashion, I download PE investment from the Capital IQ and SDC Platinum databases to supplement PitchBook and Preqin for a few missing deals.

For characteristics of hospitals, I tap the AHA annual survey sample from 2005 to 2019. It provides detailed information on operational, staffing, technology, finances, and IT characteristics. There are many categories of hospitals in the sample. A typical example is a general acute care hospital, but in my analysis, I also include specialty hospitals for rehabilitation, mental healthcare, transplants, and long-term care. To match AHA survey data to RWD claims, I follow Cooper et al. (2019) and use the NPI codes in both datasets to link them using the detailed procedure provided in the Online Appendix.

I also use the Healthcare Cost Report Information System (HCRIS) to complement the AHA survey data. HCRIS contains information from cost reports submitted annually to CMS by all Medicare-certified hospitals. It supplements the AHA data with detailed financial and cost information. To evaluate the quality information of hospitals, I leverage the quality scores from various CMS-managed quality programs: outpatient imaging efficiency scores, 30-day mortality rates, 30-day readmission rates, patient safety indicators, and surveyed patients' satisfaction scores. Other supplemental data used in the paper include annual ambulatory payment classifications (APC) code files and annual relative value unit (RVU) files from CMS to compute the relative service-mix weights, HRR and HSA files from the Dartmouth Atlas, and hospital mergers and acquisitions (M&As) data from the Irving Levin Associates' healthcare services acquisition reports.

2.2.3. Summary Statistics

For the sample of PE hospital buyouts from 2006 to 2019, Table A.12 reports key features and statistics of the deals. It records a total of 243 deals. Panel A classifies them into six categories based on deal types. Over half are add-on deals: transactions where an existing PE-backed hospital system acquired other facilities; these deals involve 4.9 hospitals on average. The second-largest group is private-to-private buyouts, including transactions where PE investors directly bought out a private hospital system. Compared to other deal types, private-to-private buyouts typically have large transaction sizes, with 13.7 hospitals per deal on average. The third group is PE growth deals, which occur when hospitals invite PE money and use the financing to upgrade technologies or expand facilities. A total of 32 transactions belong to this group. A typical deal involves around 4.4 hospitals, the smallest transaction size across all categories. The remaining deals are 18 secondary buyouts, 6 public-to-private buyouts, and 5 management buyouts. Secondary buyouts are transactions where a PE bought hospital systems from another PE firm. In public-to-private buyouts, PE firms target hospital systems that are publicly traded companies. Note that the public-to-private buyouts entail the largest transaction size with 26 hospitals per deal on average.

A total of 838 hospital facilities were ever involved in these PE deals.[35] Panel B of Table A.12 tabulates descriptive statistics of these hospitals' characteristics. The majority of target hospitals are located in urban areas and do not have teaching or critical access statuses. Panel C reports hospital ownership types prior to PE buyouts. About two-thirds of them were for-profit: there are 466 facilities previously owned by for-profit corporations, 101 facilities owned by for-profit partnerships, and a handful of facilities owned by individuals. On the other hand, about one-third of hospitals were not-for-profit prior to PE investments, including 125 community-owned or other not-for-profit facilities, 105 church-owned facilities, and 34 facilities owned by local government. These not-for-profit hospitals must abruptly transform their ownership types from not-for-profit to for-profit after buyouts because only for-profit entities are allowed to distribute profits to their shareholders.

To explore how PE-target hospitals are geographically spread, Figure A.8 demonstrates their location distribution across hospital referral regions (HRRs), which is a geographic concept commonly used to delineate distinct local markets of hospital services in the United States.[36] Panels A and B exhibit the number and the relative ratio of hospitals ever invested in by PE in each HRR. As observed, PE-target hospitals are widely spread across the United States. However, they tend to be more concentrated in the south and west, particularly in states like Texas, Florida, and Tennessee. In the Online Appendix, I provide figures of the geographic

[35] Note that a hospital might experience multiple PE deals across the sample period. For example, some hospitals would be double counted in Panel A of Table A.12 if they experienced a secondary buyout—the first-time buyout and the secondary buyout. Therefore, the net count of hospitals reported in Panel B is smaller than that in Panel A.

[36] Hospital referral regions (HRRs) are developed by the Dartmouth Atlas. It divides the United States into 306 distinct HRRs. Each represents regional healthcare markets for tertiary medical care that generally requires the services of a major referral center. An HRR is determined at the zip code level using an algorithm reflecting commuting patterns and is required to have at least one city where both major cardiovascular surgical procedures and neurosurgery are performed. Many previous studies, including Cutler and Sheiner (1999), Chandra and Staiger (2007), and Ericson and Starc (2015), use HRRs to define boundaries of local hospital markets and argue that it is preferable to conduct analyses at this level.

distribution of PE-backed hospitals at the county level.

For the insurance-claim sample of my main analysis, Table A.13 reports summary statistics in the full sample as well as the "Never Treated" and "Ever Treated" subsamples. The "Ever Treated" group contains hospitals that were ever targeted by PE or received PE investment during the sample period. The remaining hospitals are in the "Never Treated" group.

Panel A of Table A.13 shows variables at the patient-visit level. The mean age of patients is 46 and about 62% of them are female. The gender and age differences are insignificant between the "Never Treated" and "Ever Treated" groups. A typical visit requires a 42-minute drive, which is calculated between the centroid of a patient's 3-digit zip code prefix and the location of the hospital under normal traffic conditions using the *HERE API*.[37] Also, patients exhibit some degree of inertia, as they prefer hospitals that they have visited before. The full sample mean of *relative service-mix weights*, a variable measuring how much medical services or resources are used to treat a disease during a patient's visit, is 6.5.[38] The relative service-mix weights are roughly one-third higher in the "Ever Treated" group, either because patients of more severe diseases prefer to visit them, or those hospitals overutilize resources to treat patients. These two explanations will be distinguished later in Section 2.3.2. In terms of hospital prices, the mean listing price on hospitals' chargemaster for a typical visit is around $2,675 (dollars in 2019 adjusted by GDP deflators). In contrast, the average negotiated payment per visit is only $814, of which insurers cover about 82%. The remaining balance has to be paid by patients themselves, although there is wide variation in

[37] I rely on the Stata command "georoute" developed by Weber and Péclat (2020) to retrieve travel distance and travel time between two geographical coordinates. The coordinates of a patient's zip code prefix centroid are calculated by the AgGIS Pro software.

[38] The detailed procedure to assign the relative service-mix weights to each patient visit is depicted in the Online Appendix. The weights play an important role in pricing services for both public (e.g., Medicare) and private insurance. For example, in the Medicare Hospital Outpatient/Inpatient Prospective Payment System (PPS), hospitals are reimbursed based on a benchmark price set by CMS and the relative weights of incurred services. The benchmark price is equivalent to the price per unit of service weights. Of course, the final payments for individual services will also adjust by a conversion factor and other factors to take into account the geographic differences in input prices. So, a lot of private insurers follow similar pricing schemes by negotiating a benchmark price with providers and setting the payments using the relative service-mix weights. The structural model would adopt this assumption, which has been extensively used in previous literature (e.g., Gowrisankaran et al., 2015 and Ho and Lee, 2017).

this number as the median patient pays zero out-of-pocket costs. The total paid amounts do not exhibit significant differences between the "Never Treated" and "Ever Treated" groups, which by no means implies that negotiated prices are unchanged after PE intervention. A closer scrutiny should look into price variations within the hospital–insurer pair.

Panel B of Table A.13 reports variables at the hospital-year level. It shows that PE-backed hospitals tend to have less personnel, though the number of hospital beds is similar to other hospitals. They are less likely to have teaching status and tend to be located in rural areas. They accommodate fewer outpatient visits and Medicaid patients, but more Medicare patients.

Panels C and D of Table A.13 summarize variables at the local region level. Median household income is higher in counties where PE-target hospitals are located. Their residents have a slightly lower coverage ratio of private insurance, but slightly higher ratio of Medicaid. At the HRR level, PE tends to target those more competitive regions with more hospitals and lower HHIs in hospital beds and inpatient days. One conjectured reason is that PE firms would like to reduce pre-merger scrutiny from antitrust regulators by picking targets in more competitive markets (e.g., Wollmann, 2020).

2.3. Descriptive Evidence

2.3.1. Empirical Specification

To explore the impacts of PE ownership, I adopt the following specification:

$$Y_{i(m)jdt} = \alpha_1 \mathrm{PE}_{jt} + \alpha_2 \mathbf{x}_{it} + \alpha_3 \mathbf{y}_{jt} + \mathrm{FEs} + \varepsilon_{i(m)jdt} \tag{2.1}$$

in which $Y_{i(m)jdt}$ is an outcome variable in a visit for patient i to hospital j with disease d at time t, and $i(m)$ represents that patient i enrolled in a health plan provided by carrier m. The dependent variable includes the natural logarithm of patients' out-of-pocket costs, payer paid amounts, and total paid amounts (negotiated prices between hospitals and insurers) defined as the sum of the previous two variables.[39] All dependent variables are winsorized at the 1st and 99th percentiles within each CCS category. PE_{jt} is a dummy equal to one if facility j is owned by PE firms at time t. $\mathbf{x}_{it}$ is a vector of patient characteristics, including their gender, ages, plan types (HMO, PPO, etc.), and relative service-mix weights. $\mathbf{y}_{jt}$ is a vector of facility characteristics, including total number of beds, for-profit status, teaching school status, fraction of total admission days of Medicare patients, and fraction of total admission days of Medicaid patients. FEs are a set of fixed effects to control for unobserved time-invariant characteristics, which include $Disease \times Year$ FEs, $\gamma_d \times \tau_t$, to isolate variations within a given disease and a year, and most importantly, $Hospital \times Insurer$ FEs, $\mu_j \times \eta_m$, to compare negotiated prices within a hospital–insurer pair before and after PE investments.[40] Standard errors are clustered at the hospital level.

In order to determine the exit time of PE deals, I conduct extensive web searches on each hospital and infer the ultimate outcomes of these transactions based on news reports. For those deals I am still unable to determine the exit time for, I assume the holding period of PE investors to be ten years, which is consistent with the evidence documented by Strömberg (2008) that the median firm stays in leveraged buyout (LBO) ownership for more than nine years in a sample of worldwide LBO transactions from 1970 to 2007.

The coefficient of interest, α_1, essentially captures the difference of negotiated service prices between PE- and non-PE-backed hospitals for patients of identical disease complications (and hence who required the same treatment) in a given period. In addition, these patients are supposed to have insurance plans from an identical insurer. In a separate set of tests, I

[39]I use the logarithm of the outcomes plus one dollar to ensure there are no zeros in the dependent variable, which are otherwise quite frequent in, for example, patients' out-of-pocket payments (Cuesta et al., 2019). The results are robust when using the inverse hyperbolic sine transformation of the dependent variable, $\tilde{y} = \ln(y + \sqrt{y^2+1})$, which is recommended by Burbidge et al. (1988) and has been widely adopted in the literature (e.g., Browning et al., 1994 and Kale et al., 2009). The results are displayed in the Online Appendix.

[40]Note that it is important to include $Hospital \times Insurer$ FEs in the specification given the nature of price setting in this sector, i.e., through bilateral bargaining between each pair of hospitals and insurers in a local market. As argued in Craig et al. (2021), *Hospital* FEs alone cannot distinguish whether the price changes are due to renegotiation between hospitals and insurers, or from provider switching on the insurer's end.

use the natural logarithm of the relative service-mix weights of each visit as the dependent variable to investigate how PE buyouts affect the medical resource utilization, given patients stricken by the same disease. I discuss regression results in the following section.

2.3.2. Impacts of PE Buyouts on Negotiated Prices

The OLS regression results are exhibited in Table A.14. Panel A of Table A.14 uses the log of total paid amounts as the dependent variable. The coefficient on PE is positive and statistically significant at the 5% level, and its magnitude is quite stable across varying inclusions of controls and fixed effects.[41] It suggests that after PE buyouts, negotiated prices between a hospital and an insurer increase by 32%. To split the payment increase between insurers and patients, Panels B and C of Table A.14 report the regression results using the log of patient/insurer paid amounts. Most of the price increase is covered by insurers: the payer paid amounts increase by 28% after PE buyouts, while patient paid amounts slightly increase by only about 4% and the coefficient is insignificant for most specifications. The last panel explores whether PE-backed hospitals change the utilization of medical resources for patient treatment. Instead of paid amounts, I use the log of relative service-mix weights of each visit as the dependent variable and re-estimate the model. It shows that PE ownership has no impacts on service utilization given the same diagnosed disease, which contrasts with the significant increase in negotiated prices observed in the data.

In addition to the full-sample OLS regressions, I run another model in which I match the PE-target hospitals to a subset of non-PE-backed ones that are similar along observable dimensions. Specifically, I match each PE-backed hospital to three control hospitals using the optimal Mahalanobis matching, wherein the "distance" between hospitals is defined using five variables in the year before PE buyouts: ratio of Medicare patients, ratio of Medicaid patients, number of hospital beds, for-profit status, and the average benchmark price. I also require that the control hospital share the same metro status as the target one, be in the same census division, and have the same service code. The results of the matched-sample estimates

[41] The main results are robust after excluding PE growth/expansion deals in the sample, as exhibited in Table OA.3 in the Online Appendix.

are reported in Table A.15. Similar to the full-sample OLS regressions, negotiated prices between hospitals and insurers increase by about 69% after PE buyouts. This increase is mostly shouldered by insurers rather than patients, while the relative service-mix weights of patient visits do not change significantly after PE intervention. The increased magnitude of price estimates in the matched sample deserves comment: though OLS and matched-sample estimates are unable to fully resolve the endogenous selection concern of PE investments, the matched sample provides more precise estimates of the treatment effect as it creates a balanced panel and eliminates potential influence of confounding factors.

To explore the dynamics of these effects, I conduct an event study based on the matched sample.[42] I implement it at the quarterly frequency in order to better visualize how PE buyouts affect prices and whether the pre-trend assumption is violated. Specifically, I look into a time window of four years before and after PE investments and bin distant relative periods (Abraham and Sun, 2020). The estimated model is similar to Regression (2.1), except I include 16-quarter leads and lags of PE. Following the literature, I exclude the quarter right before the event time as a benchmark. Results are displayed in Figure A.9. For all dependent variables, there are no significant coefficients in the 16 quarters leading up to the event, indicating no anticipation effects or pre-trends. After completion of PE deals, the top-left and bottom-left panels of Figure A.9 witness a spike in the total paid amounts and insurer paid amounts. The effects are significantly positive and persistent for the following quarters. Such pattern is not observed for the patient paid amounts[43] or the relative service-mix weights of each visit as shown in the top-right and bottom-right panels. Overall, the figure implies that when a hospital is bought out by PE, there is a prompt, persistent rise in the negotiated prices between PE-backed hospitals and insurers, and such price increases are mostly borne by insurers. There is no indication that PE-backed hospitals

[42]Recent econometric literature points out several potential concerns regarding the conventional dynamic diff-in-diff design with staggered treatment timing, such as under-identification, negative weighting, and so on (e.g., Ivanov et al., 2020 and Borusyak et al., 2021). However, as discussed in Borusyak et al. (2021), these concerns largely disappear with a large never-treated group, which is the case in this paper.

[43]I further decompose patient paid amounts into copayments and coinsurance, and explore the dynamic effects of PE buyouts on them. Results are shown in Figure OA.3 of the Online Appendix.

increase utilization of medical resources or quantities of services to boost revenues.

One potential concern of pooling all outpatient visits in the above regressions is that PE intervention might lead to a change in service quality and hence impact negotiated prices. I address this concern by focusing on a subset of medical procedures—specific medical imaging procedures, which are widely considered as the least differentiated healthcare service with respect to clinical quality among millions of available health services items (Chernew et al., 2021). I follow Brown (2019) to pool X-ray and MRI scan procedures together and regress negotiated prices of one procedure (rather than prices aggregated to patient visit level) on the PE-ownership indicator. Results are demonstrated in Table A.16. The first column contains OLS estimates by only including hospital–insurer pair fixed effects. Consistent with Panel A of Table A.14, it implies that after PE investments, prices of imaging procedures increase by approximately 10%, significant at 5% level. Other columns vary the inclusions of controls and other fixed effects and deliver very similar estimates. For example, the last column indicates that prices will increase about 11% at 1% significance level. Furthermore, I run the main regression within each specific procedure of the top 35 X-ray and MRI scans ranked by total usage in the sample. Estimates are collected in Figure A.10. Each cross represents the point estimate for a specific imaging procedure and the capped spikes denote the 95% confidence intervals. For most imaging procedures, the figure supports the conclusion that after PE investments, negotiated prices between hospitals and insurers increase significantly.

2.3.3. Spillover Effects

PE investment in hospitals not only impacts their own negotiated prices with insurers, but also imposes spillover effects on negotiated prices of local rivals via insurers' "bargaining network." I use a variant of Regression (2.1) and test the spillover effects by examining two separate subsamples of hospitals: (1) a subset of non-PE-backed hospitals that share at least one insurer with the PE-backed facilities in the same local markets and (2) a subset of non-PE-backed hospitals that do not share any. In both subsets, I check how their negotiated prices respond to PE buyouts of other hospitals in the local market.[44] In these tests, the PE dummy is turned on if there is any PE-backed hospital in the local market in a given year. Results are reported in Table A.17.

Column 1 documents strong positive spillover effects for the first sample. It implies that after PE buyouts, local rivals that share an insurer with the PE-backed one increase their negotiated prices by 8.1% on average, which is statistically significant at 5% level. For local rivals that do not share any insurer with the PE-backed one, column 2 reports insignificant effects on prices and the estimated magnitude is close to zero. The results reveal that PE investment imposes spillover effects heterogeneously across the "bargaining network" in the local market. The magnitude of spillover effects closely hinges on the hospital's position in the "bargaining network" of insurers. This, in fact, poses a threat to the traditional reduced-form approach to evaluate the impacts of PE ownership on healthcare prices. I will discuss its implications in detail in Section 2.3.5.

2.3.4. Potential Channels

The hospital sector has experienced phenomenal consolidation in recent decades. Some might attribute the increase in negotiated prices after PE buyouts to the market consolidation effect as a result of M&As. In this section, I compare PE deals to typical M&As and argue that the observed effects of PE buyouts are unique. I then explore several potential channels associated with PE intervention.

Table A.18 compares the impact of PE to that of M&As. The sample only includes hospitals that ever experienced M&As or received PE investments. Key regressors are $M\&A$, an indicator equal one if a hospital was acquired or merged with another hospital before a given year, and $M\&A \times PE$, an indicator equal one if the deal involved any PE investors. The coefficient of $M\&A \times PE$ would pick up any effects of PE buyouts on negotiated prices that cannot be ascribed to M&As. Columns 1 to 4 show that all estimates of $M\&A \times PE$

[44] Note that this is equivalent to a "diff" specification by comparing negotiated prices before and after a "shock" occurs. The results are robust to a standard diff-in-diff model by including hospitals in regions that never had PE-backed hospitals as a control group.

are statistically significant and the economic magnitude is comparable to that in Panel A of Table A.14, indicating that the impacts of PE are unique. I explore several novel channels that might explain the observed price increases after PE intervention.

Financial Engineering

One of the most prominent features involved in PE buyouts is the heavy use of debt (so-called financial engineering). Investors typically inject large amounts of debt and jack up the target firm's leverage (e.g., Kaplan and Strömberg, 2009). Higher leverage dictates higher probability of bankruptcy, and as discussed in previous literature (e.g., Brown et al., 2009 and Matsa, 2010), could enhance the firm's bargaining position. Panel A of Table OA.4 in the Online Appendix examines how a hospital's financial leverage changes after PE investment, in which leverage is defined by the ratio of total long-term liabilities to total assets. I winsorize it at the 1st and 99th percentiles.[45] Column 4 shows that after PE buyouts, the leverage of a hospital increases by 10% on average, statistically significant at the 1% level. This is a nontrivial increase in debt levels given that the mean leverage in the sample is 13.7%.

To explore how leverage affects negotiated prices of a PE-backed hospital, in Panel B of Table OA.4 I include the leverage and its interaction term with PE as regressors. Two things are worth noting: First, estimates of $PE \times Leverage$ suggest a significantly positive effect: A 10 percentage point increase of leverage of a PE-backed hospital is associated with a 65 basis point increase of negotiated prices. Second, the coefficient of $Leverage$ itself is insignificant and has a negligible effect on prices, which suggests that an increase in leverage for non-PE-backed hospitals does not affect price negotiations. This supports the idea that

[45]Both long-term liabilities and total assets are book values. The total long-term liabilities are the sum of debt with a maturity longer than 12 months, including mortgages payable (line 46 in the HCRIS database), notes payable (line 47), unsecured loans (line 48), and other long-term liabilities (line 49); and the total assets are equal to the sum of total current assets (line 11), total fixed assets (line 30), and total other assets (line 35). I focus on this definition for two reasons: (1) leveraged buyouts usually take debt with relatively longer maturities; and (2) in HCRIS data, short-term liabilities include items such as accounts payable due suppliers, or salaries and wages payable due employees, which might contaminate the analysis. To avoid the pitfall of excluding relevant short-term debt, as a robustness test I introduce a variant of the leverage measure, defined as total long-term liabilities plus notes and loans payable (line 40), and then divided by total assets. All results are robust under the new definition.

the financial engineering channel is unique to PE-backed hospitals, consistent with anecdotal evidence that PE investors have a reputation for closing down distressed facilities and are able to credibly follow through on the threat of bankrupting hospitals by taking up debt loads.

Non-pecuniary Motive

Another prominent feature of PE firms is that they are profit-centered due to their fiduciary duty to investors (e.g., Shleifer and Summers, 1988 and Lerner et al., 2011). This contrasts with the ownership type of most not-for-profit hospitals in the sector. After a not-for-profit hospital is bought out by PE, it must transition to a for-profit entity with less focus on social objectives. This could potentially boost the negotiated prices of PE-backed hospitals. To test this channel, I compare price changes of PE-target hospitals which were previously for-profit to those that were previously not-for-profit. In the estimation, I extend Regression (2.1) by including an interaction term $PE\times$Previously For-profit, an indicator whether the PE-backed hospital was previously for-profit. Note that the Previously For-profit dummy itself is absorbed by hospital fixed effects. The prediction is that they will experience a smaller price increase after PE investment, and the coefficient of $PE\times$Previously For-profit is expected to be negative.

Estimation results of Table OA.6 in the Online Appendix support the prediction: The estimate of $PE\times$Previously For-profit is indeed significantly negative, indicating that previously for-profit hospitals, compared to other types of entities, experience smaller price increases after PE investment. Summing estimates of $PE\times$Previously For-profit and PE, it suggests that their negotiated prices on average increase by only 8.5%, which is smaller than the average price increases in Table A.14.

Operational Efficiency

Previous literature documents that PE improves operational efficiency of its portfolio companies, commonly translating into lower marginal costs (e.g., Davis et al., 2014, Bernstein and Sheen, 2016, and Bircan et al., 2020). Theoretically, lower marginal costs are associated with lower prices, which can be another channel affecting negotiated prices of PE-backed hospitals. I therefore check how PE buyouts affect the operational costs of hospitals by exploiting the measure of total cost per adjusted discharge proposed by Schmitt (2017). Specifically, it captures the average operational cost of a hospital to take care of a patient, defined as

$$\mathrm{AC}_{it} = \frac{\mathrm{TC}_{it}}{D_{it} \times (1 + \frac{R_{it}^O}{R_{it}^I})},$$

in which AC_{it} is the cost per adjusted discharge for hospital i in fiscal year t, TC_{it} is the total costs of hospital i during the year, D_{it} is the number of inpatient discharges, R_{it}^O is the outpatient charges, and R_{it}^I is the inpatient charges.[46]

I depict the average cost dynamics before and after PE investment in Figure OA.4 in the Online Appendix. It shows that the average costs significantly trend down two years after PE buyouts.

Service Quality and Patients' WTP

Patients' willingness-to-pay (WTP) is another determinant of a hospital's bargaining position, and it is typically correlated with the service quality provided by the hospital. For example, if one hospital is believed to provide good services and hence patients assign high WTP to it, the hospital is able to bargain a relatively high price compared to other local providers. So as a potential channel, PE might affect hospitals' bargaining positions by adjusting service quality.

To explore this channel, I use 42 different quality measures from CMS to evaluate how they change after PE intervention, including 30-day mortality rates, 30-day readmission rates, patient safety indicators, outpatient imaging efficiency, and consumer assessment scores. Results are exhibited in Figure OA.5 in the Online Appendix. Each cross denotes the point

[46] Essentially, AC_{it} measures the average cost for an effective inpatient discharge by converting outpatient activities into the equivalent number of inpatient ones. Given that our main analysis focuses on the outpatient prices, the underlying assumption here is that PE buyouts have an overall effect on hospitals' operational efficiency for both outpatient and inpatient settings, which allows us to infer the operational efficiency change of a hospital for an outpatient visit.

estimate for one specific quality measure and the capped spikes denote its 95% confidence interval. In contrast to Gupta et al. (2021), quality implications after PE buyouts are mixed: In general, there is no significant difference in service quality before and after PE buyouts. For some measures, hospitals perform better after PE buyouts. However, patient satisfaction significantly worsens after PE investments, indicating that ancillary-service quality worsens.

2.3.5. Challenges of the Reduced-form Approach

Though the reduced-form regression is a transparent way to examine the impacts of PE, there are a couple of challenges to that approach that support a structural approach. First, as shown in Section 2.3.3, PE investment exhibits heterogeneous spillover effects. This violates the SUTVA assumption embedded in the reduced-form approach (Berg et al., 2020). So, more structural assumptions are needed to accommodate the spillover effects. Second, it is challenging to use the reduced-form approach to quantify the relative contribution and importance of various channels on price bargaining outcomes, though it does provide evidence on the directional effects. In addition, some potential channels are difficult to find proxies for and to isolate using the reduced-form techniques. For example, PE firms are believed to be professional investors in financial markets and have rich experience in business negotiations. It is possible that after PE investment, hospitals learn better negotiation skills and increase their bargaining power. But finding a proxy for negotiation skills in reduced form is difficult. On the contrary, the structural approach can serve the purpose with the help of bargaining theory to isolate the associated impacts. Third, a structural model is an indispensable tool if we would like to quantify patient surplus and explore counterfactuals. It provides guidance to empiricists so that one can estimate parameters linked, either directly or indirectly, to the underlying theory of PE investment and hospital price determination. As a laboratory, the counterfactual analyses also shed light on relevant policy debates.

2.4. Structural Approach

2.4.1. Model Specification

I extend a bargaining model in the style of Capps et al. (2003), Gowrisankaran et al. (2015) and Ho and Lee (2017) wherein I detail the interaction between PE firms, hospitals, insurance companies, and patients in three stages, as depicted in Figure 2.4.1. In the first stage, PE firms pick target hospitals to invest in and choose an optimal debt leverage to finance the deal. Next, hospitals in the local market engage in bilateral bargaining to determine benchmark prices with insurers. Finally, patients select hospitals to visit when they become sick. I discuss each stage in reverse order.

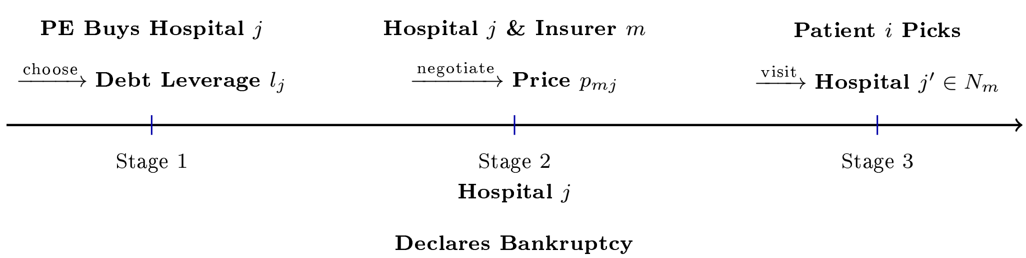

Patient Hospital Choice

Within an HRR, which is deemed as a local market, there is a set of hospitals indexed by $j = 1, \cdots, J$, and a set of insurance companies indexed by $m = 1, \cdots, M$. The hospitals are partitioned into $S \leq J$ systems. Let J_s denote the set of hospitals in system s.[47] There is a set of enrollees denoted by $i = 1, \cdots, I$, each of which has a health insurance plan issued by a particular insurer. Let $m(i)$ denote enrollee i of insurer m. The subset of hospitals that insurer m agrees to contract with in their network is denoted by N_m. Each insurer m and hospital system s negotiate a benchmark price p_{ms} in the price bargaining stage. $\mathbf{p}_m$

[47] If a single hospital in the local market does not share a system with any other local hospitals, I treat it as a single-hospital system. Hereafter, hospital system s is sometimes abbreviated to hospital s in the main text.

denotes the vector of all negotiated prices for insurer m.

Enrollee $m(i)$ who is stricken by illness $d = 0, 1, \cdots, D$, where $d = 0$ represents the status of no illness, picks a hospital in the network to visit. Following the common pricing practice with the notion of severity weights, w_d represents the relative service-mix weights of illness d, which measures the intensity of resources used to treat the disease, and $w_0 = 0$. So, the total price paid for treatment of disease d at hospital $j \in J_s$ by insurer m is $w_d p_{ms}$, that is, the base price multiplied by the relative service-mix weights. For each illness $d = 1, \cdots, D$, the patient seeks hospital care at the hospital that gives them the highest utility. The ex post utility of patient i insured by insurer m receiving care from hospital $j \in N_m$ is given by

$$
\begin{aligned}
U_{ijmd} = {} & \underbrace{\beta_1(\mathbf{X_{id}}) \cdot d_{ij} + \beta_2(\mathbf{Y_j}) \cdot d_{ij} + \beta_3 \cdot d_{ij}^2}_{\text{Distance}} + \underbrace{\gamma_1(\mathbf{X_{id}}) \cdot \mathbf{Y_j} + \gamma_2 \text{PastPatient}_{ij} + \eta_j w_d}_{\text{Hospital Characteristics} \times \text{Patient Observables}} \\
& + \underbrace{\gamma_3(\mathbf{X_{id}}) \cdot \text{PE}_j}_{\text{PE Owned} \times \text{Patient Observables}} + \eta_j + e_{ij}. \quad (2.2)
\end{aligned}
$$

This specification encompasses different models used in the literature (e.g., Town and Vistnes, 2001, Capps et al., 2003, Gowrisankaran et al., 2015, Shepard, 2016, and Ho and Lee, 2017) to account for preference heterogeneity. The main covariates include travel time (in minutes) between patient i's residence and hospital j's location, d_{ij}, and travel time squared, d_{ij}^2, as well as interacted terms with a vector of patient observables $\mathbf{X_{id}}$ (e.g., patient age, gender, and relative service-mix weights), and various hospital characteristics $\mathbf{Y_j}$ (e.g., number of hospital beds, for-profit status, and teaching status). In addition to travel time, the indirect utility depends on other hospital characteristics interacted with patient observables such as major diagnoses dummies interacted with hospital-provided services[48] and

[48]These interaction terms potentially capture the impact on patients' utility from changes in hospitals' service lines after buyouts. For example, as shown by Eliason et al. (2020) in the dialysis industry, the change in procedures provided to patients is an important margin for patient welfare. However, I examine 95 categories of hospital service to provide descriptive evidence that PE investors do not dramatically change lines of service within a hospital after buyouts, as exhibited in Figure OA.6 in the Online Appendix. So, it should be less of a concern that the changing composition of procedures after buyouts would affect the welfare implications.

relative service-mix weights interacted with hospital dummies. Following Shepard (2016), the indirect utility specification also includes the past outpatient status, PastPatient_{ij}, to capture relationships between patients and a hospital's physicians, which is observed to be a key source of heterogeneity in hospital choice. Note that, following previous literature (e.g., Capps et al., 2003 and Ho and Lee, 2017), I exclude out-of-pocket costs in the utility function. This simplifying assumption is consistent with (1) limited cost-sharing faced by patients and (2) empirical evidence of Section 2.3.4 indicating negligible impacts of PE intervention on patient paid amounts. I will add out-of-pocket costs back when inferring implications for patient surplus.

I include the PE-owned indicator to account for the effects of PE buyouts on patients' hospital choices as a result of changes in service quality and patient experience, as documented in Section 2.3.4. The specification includes $\text{PE}_j \in \{0, 1\}$, an indicator if hospital j is backed by PE firms as well as its interaction terms with patient observables. The goal is to provide a flexible enough choice model to incorporate possible heterogeneity of patients in responding to the quality changes after PE buyouts.

Finally, hospital fixed effects, η_j, are added to control for unobserved time-invariant characteristics of hospitals. e_{ij} is an idiosyncratic error with i.i.d. type 1 extreme value distribution that is known by the patients at the time of choosing providers. Patients may visit a hospital in their network, $j \in N_m$, within an HRR. The outside option is modeled as choice 0, which corresponds to patients going to an out-of-network hospital or not going at all, and the delivered utility is normalized as $u_{i0d} = e_{i0}$.

Define $\delta_{ijmd} = u_{ijmd} - e_{ij}$ as the observed expected utility. The logit model implies that the choice probability for patient i with disease d as a function of patient and hospital characteristics is

$$s_{ijmd}(N_m) = \frac{\exp(\delta_{ijmd})}{\sum_{k \in \{0, N_{m(i)}\}} \exp(\delta_{ikmd})}.$$

Then, the expected utility for a patient of disease d in need of outpatient services is

$$CS_{imd}(N_m) = \ln\left(\sum_{k\in\{0,N_{m(i)}\}} \exp(\delta_{ikmd})\right).$$

Insurance Company and Hospital Bargaining

In the local market, insurer m will negotiate a price with hospital system s with hospitals $J_s \subseteq N_m$, wherein N_m denotes the set of hospitals in m's network. Correspondingly, hospital system s will bargain over a price with every insurer $m' \in M_s$, wherein M_s denotes the set of insurers that include hospital system s in their networks. In the baseline model, N_m and M_s are taken as given and matched with the network structure observed in the data.[49]

The bargaining process occurs between hospital system s and insurer m over a benchmark price p_{ms}. Before discussing the details of bargaining protocol, I cover one new feature emphasized in this paper: the structure of the local market impacts the bargaining positions of hospitals and insurers. Specifically, if any hospital system s in N_m declares bankruptcy, the local market of hospitals is changed (and becomes less competitive), and the rest of the hospitals within $N_m - J_s$ will restart negotiations with insurer m. Otherwise, the bargaining process follows the standard protocol of the Nash-in-Nash bargaining model widely used in the applied literature (e.g., Horn and Wolinsky, 1988, Crawford and Yurukoglu, 2012, Grennan, 2013, Gowrisankaran et al., 2015, Ho and Lee, 2017, and Collard-Wexler et al., 2019). This new feature resulting from hospital bankruptcy introduces an extra consequence to insurance companies: Letting a hospital fail may bring additional costs. It lowers its disagreement point during the bargaining process. I illustrate the idea by a simple motivating model in the Online Appendix.

Given N_m and M_s, the payoff functions to insurers and hospitals are derived as follows.

[49]Previous studies, including Gowrisankaran et al., 2015, Lewis and Pflum (2015), Ho and Lee, 2017, Brown (2019), and Raval et al. (2020), assume that the contracted networks of hospitals for insurers are fixed. In Section 2.6, I extend the baseline model to incorporate the network-formation process following Ghili (2016) and Prager and Tilipman (2020).

Notice that the normalized quantity of patients from insurer m to hospital system s is represented by

$$q_{ms}(N_m) = \sum_{j \in J_s} \sum_{i=1}^{I} \mathbf{1}\{m(i) = m\} w_d s_{ijmd}(N_m).$$

It is assumed that insurers maximize a weighted difference of their enrollees' expected utility and costs of paying for health care. Denote the sum of enrollees' expected utility as $W_m(N_m) = \sum_{i=1}^{I} \mathbf{1}\{m(i) = m\} CS_{imd}(N_m)$. Then, the objective function of insurer m can be represented by

$$V_m(N_m, \mathbf{p}_m) = \alpha W_m(N_m) - \sum_{J_s \subseteq N_m} p_{ms} q_{ms}(N_m),$$

wherein α is the insurer's weight on enrollees' expected utility and measures how much insurers care about enrollees' welfare relative to the costs.

If hospital system s declares bankruptcy, the set of intended-to-negotiate hospitals becomes $N_m - J_s$, and they will be renegotiating with insurer m. Rather than sticking to the price $\mathbf{p}_{-s}$, let the renegotiated prices be $\mathbf{p}'_{-s}$ and the payoffs to insurer m become $V_m(N_m - J_s, \mathbf{p}'_{-s})$. Notice this is different from the payoffs in which hospital system s stays alive while not reaching an agreement with insurer m. In this alternative case, the payoffs to m are $V_m(N_m \backslash J_s, \mathbf{p}_{-s})$ following the Nash-in-Nash bargaining protocol.

Next, I turn to depicting the behavior of hospitals. In a similar fashion to the service-weight pricing scheme, let mc_{ms} denote the marginal cost of hospital $j \in J_s$ for treating a patient from insurer m with disease weight $w_d = 1$. Then, the treatment costs for an illness with relative weights w_d is $w_d \text{mc}_{ms}$.

Therefore, the profit that hospital system s expects to earn from insurer m is

$$\pi_s(\mathbf{p}_m, N_m, \text{PE}_s) = q_{ms}(N_m) \times (p_{ms} - \text{mc}_{ms}).$$

And the total expected profits of hospital system s are $\sum_{m \in M_s} \pi_s(\mathbf{p}_m, N_m, \text{PE}_s)$, wherein

M_s is the set of insurers that include hospital system s in their networks. To incorporate the non-pecuniary motive of hospitals, it is assumed that hospitals care about the weighted quantity of patients they serve (e.g., Gaynor and Vogt, 2003 and Lakdawalla and Philipson, 2006) besides its profits. I denote it by τ_{NP}, measuring the non-pecuniary benefits of a not-for-profit hospital for every unit of service it provides. So, the total benefits of hospital s from being included in insurer m's network are

$$\begin{aligned} \varpi_s(\mathbf{p}_m, N_m, \text{PE}_s) &= \pi_s(\mathbf{p}_m, N_m, \text{PE}_s) + \text{NP}_s \cdot (1 - \text{PE}_s) \cdot \tau_{\text{NP}} q_{ms}(N_m) \\ &= q_{ms}(N_m) \cdot [p_{ms} - \text{mc}_{ms} + \text{NP}_s \cdot (1 - \text{PE}_s) \cdot \tau_{\text{NP}}], \end{aligned} \tag{2.3}$$

wherein NP_s is an indicator of whether hospital s has not-for-profit status and PE_s is an indicator of whether hospital s is owned by PE investors. Notice $\text{NP}_s \cdot (1 - \text{PE}_s)$ captures the impact of PE buyouts on the non-pecuniary motive of hospitals: τ_{NP} would vanish after PE investment. This is because PE firms have to turn hospitals into for-profit entities if they would like to distribute profits. To incorporate the idea that PE investors can potentially change the operational efficiency of target hospitals, mc_{ms} is assumed to be a function of PE_s, which will be specified later in the estimation step. Summing across all contracted insurers, the objective function of hospital system s (without taking into account the possibility of going bankrupt) is

$$\Pi_s(M_s, \{\mathbf{p}_m\}_{m \in M_s}, \{N_m\}_{m \in M_s}, \text{PE}_s) = \sum_{m \in M_s} \varpi_s(\mathbf{p}_m, N_m, \text{PE}_s). \tag{2.4}$$

Debt and Threat of Bankruptcy

Leveraged buyouts typically involve massive financial engineering that would lead to a large debt burden for target hospitals. Given PE's choice of debt level D_s for hospital system s, I assume the interest payments and due debt repayments are a function of the debt level, denoted by $C(D_s)$. PE firms will declare bankruptcy of invested hospitals when their net "income" fails to meet the debt burden, that is, $C(D_s) > \widetilde{\Pi_s}(\cdot)$, wherein $\widetilde{\Pi_s}(\cdot)$ is a function of $\Pi_s(\cdot)$ described in Equation (2.4). By parameterizing $C(D_s) = \tilde{\theta} D_s$, with $\tilde{\theta}$ reflecting the magnitude of hospitals' debt burden, the bankruptcy rule is

$$\underbrace{\tilde{\theta} D_s}_{\text{Debt Burden}} > \underbrace{\widetilde{\Pi_s}(\cdot)}_{\text{Net Revenue}} . \tag{2.5}$$

To characterize the disagreement points for both parties in the bargaining game, the bankruptcy scenario as well as its probability has to be considered. Two assumptions are made in order: (1) PE-backed hospitals are able to credibly threaten bankruptcy by following the above bankruptcy rule, while other hospitals cannot. It is argued that PE investors are profit-driven with hard budget constraints while other hospitals tend to face soft budget constraints (Kornai et al., 2003 and Kornai, 2009) because of government subsidies, private grants, and donations, which can potentially help them overcome most financial difficulties. Often, local governments tend to bail out distressed non-PE-backed hospitals to prevent a loss of access to medical services in local communities. In addition, PE investors have the reputation of closing facilities due to financial considerations,[50] which makes the bankruptcy threat more credible. Most importantly, this assumption is consistent with the empirical results in Section 2.3.4 that financial leverage of non-PE-backed hospitals does not affect service pricing. (2) The second assumption is about the specification of $\widetilde{\Pi_s}(\cdot)$. If hospital system s successfully reaches agreements with all insurers in M_s, we have $\widetilde{\Pi_s}(\cdot) = \Pi_s(M_s, \cdot)$. Otherwise, say the negotiation between hospital system s and insurer m, $m \in M_s$, was halted; then, $\widetilde{\Pi_s}(M_s \backslash m) = \Pi_s(M_s \backslash m, \cdot) - v_{st}$, wherein v_{st} is a random variable to capture the unexpected losses to hospital s due to failed negotiation. To sum up, $\widetilde{\Pi_s}(\cdot)$ is characterized as

$$\widetilde{\Pi_s}(\cdot) = \begin{cases} \Pi_s(M_s, \cdot) & \text{if negotiation is successful with all } m \in M_s \\ \Pi_s(M_s \backslash m, \cdot) - v_{st} & \text{if negotiation between } s \text{ and } m \text{ fails.} \end{cases} .$$

[50] Critics have accused buyout firms of preferring to close invested hospitals rather than keep them running. Stories prevail on media. For instance, the *Wall Street Journal* reports that PE investors threatened to shut down the Easton Hospital in Pennsylvania in the midst of the COVID-19 outbreak, alluding to the reputation of PE closing up unprofitable hospitals.

Scaling Equation (2.5) by hospitals' total assets, the bankruptcy rule can be rewritten as

$$\tilde{\theta} l_s > \widetilde{g_s}(\cdot) = \begin{cases} h_s(M_s) & \text{if negotiation is successful with all } m \in M_s \\ h_s(M_s \backslash m) - \tilde{v}_{st} & \text{if negotiation between } s \text{ and } m \text{ fails} \end{cases},$$

in which $l_s = \frac{D_s}{\mathrm{TA}_s}$ is hospital s's leverage, $\widetilde{g_s}(\cdot) = \frac{\widetilde{\Pi_s}(\cdot)}{\mathrm{TA}_s}$ is the return on assets (ROA), $h_s(\cdot)$ is equal to $\frac{\Pi_s(\cdot,\cdot)}{\mathrm{TA}_s}$, and $\tilde{v}_{st} = \frac{v_{st}}{\mathrm{TA}_s}$ is assumed to be i.i.d. distributed and to follow a logistic distribution with the location parameter $\tilde{\mu}$ and the scale parameter $\tilde{s} > 0$.[51]

One direct implication is that in equilibrium, PE-backed hospitals will not choose an extreme level of D_s such that $\tilde{\theta} D_s > \Pi_s(M_s, \cdot)$ because it is not rational for PE firms to let hospitals go bankrupt directly on the equilibrium path. This implicitly imposes a constraint on the choice of D_s in the first stage. Therefore, the bankruptcy risk emerges in the model only if disagreements between the hospitals and the insurer occur, which reflects the idea emphasized in the illustrative model of the Online Appendix: bankruptcy risk arises because of potential losses of income when hospitals do not reach an agreement with insurers, which makes it harder for them to meet their debt obligations.

For PE-backed hospital s, the expected bankruptcy probability is calculated by taking the expectation of $\mathbf{1}\{\widetilde{\Pi_s}(\cdot) < C(D_s)\}$, denoted by $\rho(l_s, \cdot) = \mathbf{E}\left(\mathbf{1}\{\widetilde{\Pi_s}(\cdot) < C(D_s)\}\right)$. For example, on the equilibrium path where hospital s successfully reaches agreements with all the insurers in M_s, the previous discussion implies that the expected bankruptcy probability satisfies $\rho(l_s, M_s) = 0$.

Bargaining Stage

Consider the case where hospital s is backed by PE investors. The payoffs of hospital s and insurer m under agreement and disagreement scenarios are specified as follows. With

[51] The logistic distribution can be rationalized by adding error terms to both sides of Equation (2.5) and assuming those terms follow i.i.d. type I extreme value distribution, which is a typical assumption in the discrete choice literature.

successful negotiations, the expected payoffs to s and m are

$$\Omega_s^A(M_s, \mathbf{p}_s) = \underbrace{\widetilde{\Pi_s}(M_s) - C(D_s)}_{\text{Payoffs of } s \text{ after debt}}$$
$$\Omega_m^A(N_m, \mathbf{p}_m) = V_m(N_m, \mathbf{p}_m),$$

and with negotiations broken down, the expected payoffs become

$$\Omega_s^{NA}(M_s \backslash m, \mathbf{p}_{s-ms}) = \underbrace{\mathbf{E}\left(\left[1 - \mathbf{1}\{\widetilde{\Pi_s}(M_s \backslash m) < C(D_s)\}\right] \times \left[\widetilde{\Pi_s}(M_s \backslash m) - C(D_s)\right]\right)}_{\text{Expected payoffs of } s \text{ if not bankrupt}}$$
$$\Omega_m^{NA}(N_m \backslash J_s, \mathbf{p}_{m-s}) = \underbrace{[1 - \rho(l_s, M_s \backslash m)] \times V_m(N_m \setminus J_s, \mathbf{p}_{m-s})}_{\text{Expected payoffs of } m \text{ if } s \text{ survives while excluded}} + \underbrace{\rho(l_s, M_s \backslash m) \times V_m(N_m - J_s, \mathbf{p}'_{m-s})}_{\text{Expected payoffs of } m \text{ if } s \text{ goes bankrupt}},$$

in which $\rho(l_s, M_s \backslash m) = \frac{1}{1+\exp(\varrho h_s(M_s \backslash m) - \theta l_s - \mu)}$ represents the expected bankruptcy probability if hospital system s is excluded by insurer m, wherein $\varrho = \frac{1}{\tilde{s}} > 0$, $\theta = \frac{\tilde{\theta}}{\tilde{s}}$, and $\mu = \frac{\tilde{\mu}}{\tilde{s}}$. The payoffs to hospital equityholders are zero after bankruptcy. Note that the outside option for insurer m consists of two parts: The first corresponds to the scenario wherein hospital s survives when bargaining breaks down. Since the market structure does not change, other hospitals will not renegotiate the contract immediately. The second part embodies the scenario in which hospital s going bankrupt leads to a more concentrated hospital sector. As a result, other hospitals start to renegotiate with insurer m, changing the payoffs of m when facing a less competitive hospital sector. The price renegotiation resembles the bargaining protocol wherein negotiations are non-binding and contingent on the network of hospitals never reaching disagreement (Stole and Zwiebel, 1996). However, I emphasize a distinct mechanism: In a repeated setting, insurers have to take into account the adverse impact of a more concentrated hospital market if a local hospital fails. Therefore, insurers have the incentive to "subsidize" hospitals to keep them alive and competing with each other.[52] This intuition shares a similar spirit with the idea of Raskovich (2003) in which pivotal buyers

[52]Practitioners are found to use this strategy quite frequently. For example, in the article "The Death of Hahnemann Hospital" (June 7, 2021, issue of the *New Yorker*), the management team of Hahnemann explicitly describes: "...The insurance companies had an incentive to compromise. If Hahnemann closed, the privately insured patients treated there would go to other city hospitals, where the cost of their care would rise. You go into Blue Cross and you say, we need some help, and it's in your best interest to help us...Give us ten million dollars more per year—versus losing fifty million per year...".

tend to have worse bargaining positions since they have to cover a larger portion of fixed costs incurred by the seller. But for the applied purpose of this paper, the dynamic feature is abstracted away and the implication of market structure changes is reflected in the hybrid bargaining protocol.[53]

Therefore, the Nash bargaining problem is characterized by the exponential product of net values:

$$NB^{m,s}(p_{ms} \mid \mathbf{p}_{m-s}, \mathrm{PE}_s) = \left(\Omega_s^A - \Omega_s^{NA}\right)^{B_{sm}} \times \left(\Omega_m^A - \Omega_m^{NA}\right)^{1-B_{sm}},$$

in which $B_{sm} = b_0 + \mathbf{b} \times \mathbf{Y_s^b} + \mathrm{PE}_s \cdot g_b$ is the bargaining power of hospital s facing insurer m. In the specification of B_{sm}, b_0 is a constant term representing the base bargaining power of a typical hospital system. $\mathbf{Y_s^b}$ is a vector of hospital characteristics, including whether s is a multisystem hospital, its teaching status and for-profit status, whether it is a rural hospital, log of the number of hospital beds, physician arrangement, shares of patient days in local markets, and number of insurers in the local market. PE_s equals one if the hospital is backed by PE firms, and g_b measures the change of hospitals' bargaining power after PE intervention.

The above Nash bargaining problem can be further simplified as

$$\begin{aligned} & NB^{m,s}(p_{ms} \mid \mathbf{p}_{m-s}, \mathrm{PE}_s) \\ & \quad = \left(\varpi_s(\mathbf{p}_m, N_m, \mathrm{PE}_s) - \underbrace{\mathbf{E}\left(\mathbf{1}\{\widetilde{\Pi}_s(M_s \backslash m) < C(D_s)\} \times \left[C(D_s) - \widetilde{\Pi}_s(M_s \backslash m)\right]\right)}_{\text{Denoted by } \Delta\varpi_s} \right)^{B_{sm}} \\ & \times \left(V_m(N_m, \mathbf{p}_m) - V_m(N_m \setminus J_s, \mathbf{p}_{m-s}) + \underbrace{\rho(l_s, M_s \backslash m) \times \left[\sum_{k \in N_m \backslash J_s} q_{mk}(N_m \backslash J_s) \cdot (p'_{mk} - p_{mk})\right]}_{\text{Denoted by } \Delta V_m} \right)^{1-B_{sm}}. \end{aligned} \tag{2.6}$$

In contrast, for non-PE-backed hospitals, as they are unable to credibly threaten bankruptcy,

[53] Consistent with the prediction of this mechanism, I find that if a PE-backed hospital accommodates a larger share of patient flow from an insurer, this hospital is able to negotiate a higher price with the insurer, as shown in Table OA.7 in the Online Appendix.

the bargaining problem boils down to

$$NB^{m,s}(p_{mj} \mid \mathbf{p}_{m-s}, \mathrm{PE}_s) = \left(\varpi_j(\mathbf{p}_m, N_m, \mathrm{PE}_s)\right)^{B_{sm}} \times \left(V_m(N_m, \mathbf{p}_m) - V_m(N_m \setminus J_s, \mathbf{p}_{m-s})\right)^{1-B_{sm}}, \tag{2.7}$$

which is a special case of Equation (2.6) when $\rho(l_s, M_s \backslash m) = 0$. Note that this goes back to a typical bargaining problem discussed in the literature (e.g., Gowrisankaran et al., 2015, Lewis and Pflum, 2015, and Ho and Lee, 2017).

First-Order Condition of the Bargaining Problem with PE-backed Hospitals

The Nash bargaining solution is a negotiated price p^*_{ms} that maximizes Equation (2.6), and it should satisfy the first-order condition (FOC) as follows:

$$\begin{aligned} p^*_{ms} = {} & (1 - B_{sm}) \cdot \left(\underbrace{\mathrm{mc}_{ms}(\mathrm{PE}_s)}_{\text{Operational Efficiency}} - \underbrace{\mathrm{NP}_s \cdot (1 - \mathrm{PE}_s) \cdot \tau_{\mathrm{NP}}}_{\text{Non-pecuniary Objective}} \right) + \\ & \frac{B_{sm}}{q_{ms}(N_m)} \left(\underbrace{\alpha \left(W_m(N_m) - W_m(N_m \backslash J_s)\right)}_{\text{Willingness to Pay (WTP)}} + \underbrace{\sum_{k \in N_m \backslash J_s} p_{mk} \left(q_{mk}(N_m \backslash J_s) - q_{mk}(N_m)\right)}_{\text{Insurer Cost Change due to Exclusion}} \right) + \\ & \mathrm{PE}_s \times \frac{1}{q_{ms}(N_m)} \left(\underbrace{B_{sm} \Delta V_m + (1 - B_{sm}) \cdot \Delta \varpi_s}_{\text{Financial Engineering}} \right), \end{aligned} \tag{2.8}$$

wherein ΔV_m is described in Equation (2.6) and $\Delta\varpi_s = \frac{1}{\varrho}\mathrm{TA}_s \ln\left(1 + \exp\left(\theta l_s + \mu - \varrho h_s(M_s \backslash m)\right)\right)$. Detailed derivation of the FOC is presented in the Online Appendix. To see how various channels of PE buyouts affect negotiated prices, several predictions are discussed based on Equation (2.8).

Bargaining power: Under general conditions, the FOC predicts that a hospital's prices increase with its bargaining power. If PE is able to bring in new bargaining skills and expertise to the hospital, reflected by an increase in B_{sm}, higher negotiated prices would emerge in equilibrium.

Bankruptcy threats and financial engineering: The FOC predicts that bankruptcy threats potentially boost negotiated prices of PE-backed hospitals in both extensive and intensive margins. Compared to non-PE-backed hospitals, the ability to credibly threaten bankruptcy grants the PE-backed hospital a better bargaining position. Moreover, among PE-backed hospitals, the FOC predicts that higher leverage would result in an increase in negotiated prices, as shown in the Online Appendix.

Change of demand: The FOC gives ambiguous predictions regarding how changes of consumer demand would affect negotiated prices. On the one hand, PE alters the attractiveness of hospitals relative to others. This directly changes the distribution of patient quantities q_{mk}. On the other hand, hospital quality improvement or deterioration after PE intervention affects the WTP of patients, which leads to changes of negotiated prices.

Non-pecuniary objectives: If PE-backed hospitals care less about their social objectives and more about their fiduciary duties, transition to PE ownership would lead to a decline of non-pecuniary motives for previously not-for-profit hospitals. Equation (2.8) predicts an increase in negotiated prices.

Operational efficiency: Suppose PE improves operational efficiency of target hospitals. The FOC predicts that, quite intuitively, a decrease in mc_s would lead to a decrease in hospitals' negotiated prices with insurers.

Spillover effects: The spillover effects are also captured by the FOC. The intuition is that as PE buyouts induce higher negotiated prices at PE-target hospitals, the insurer's cost of excluding any other hospitals in the network also changes, which directly impacts the bargaining outcomes of these local rivals.

PE Investment

I model the decision of PE investment as a binary choice. Denote it by $\mathrm{PE}_s \in \{0, 1\}$, in which $\mathrm{PE}_s = 1$ means that PE firms decide to invest in hospital system s. Specifically, PE firms weigh investment costs against potential gains. After agreeing to invest, an optimal debt level is chosen. The cost of debt financing at the buyout stage has to be factored in by

PE. It is different from the debt burden modeled in the bargaining stage. Potential risks of using larger amounts of debt include higher costs resulting from debt overhang (Hennessy and Livdan, 2009 and Hege and Hennessy, 2010), larger risk of financial distress (Myers, 2001), and more severe agency conflict between general and limited partners (Axelson et al., 2009). I model the average cost of debt for PE firms as $C_d(l_s)$ and assume it is increasing in the leverage level of the portfolio company.

Given leverage level l_s, the potential gains for PE investors are derived from its intervention, calculated as

$$\triangle\Pi_s(l_s^*) = \max_{l_s} \mathrm{TA}_s \times \left(\underbrace{\frac{\Pi_s(M_s, \{\mathbf{p}_m^*\}, \{N_m\} \mid \mathrm{PE}_s = 1) - \Pi_s(M_s, \{\mathbf{p}_m\}, \{N_m\} \mid \mathrm{PE}_s = 0)}{\mathrm{TA}_s}}_{\text{ROA of PE Deal}} - C_d(l_s) \right),$$

wherein p_{ms}^* with $m \in M_s$ is a function of l_s. Therefore, the optimal level l_s^* satisfies $\partial \triangle\Pi_s(l_s^*)/\partial l_s = 0$, which implies that

$$\underbrace{\sum_{m \in M_s} \left[B_{sm} \frac{\partial \Delta V_m}{\partial l_s} + (1 - B_{sm}) \frac{\partial \Delta \varpi_s}{\partial l_s} \right] / \mathrm{TA}_s}_{\text{Marginal Benefits of Debt Leverage}} = \underbrace{C_d'(l_s^*)}_{\text{Marginal Costs of Debt Leverage}}. \tag{2.9}$$

However, as a (nonprofessional) corporate investor in medicine, various regulatory requirements are imposed across states, which would lead to variations in entry costs. I assume the entry costs of investing in hospital s in state w in year t to be EC_{swt}. Therefore, the entry/investment decision of PE firms is that $\mathrm{PE}_s = 1$ if and only if

$$\triangle\Pi_s(l_s^*) \geq \mathrm{EC}_{swt}.$$

In reality, the potential gains are split between the PE firm and the original hospital owners under a fixed sharing rule. I follow SØrensen (2007) to assume that PE expects to receive a fraction, λ, of the total gains. This assumption rules out transfers. Since PE's investment decision does not quite depend on the constant λ, I abstract the sharing rule away by assuming $\lambda = 1$ for simplicity.

2.4.2. Estimation

Estimation and Identification of Patient Demand

The patient demand model is estimated by maximum likelihood using the patient-visit data aggregated from the RWD Product. Identification of the demand parameters in Equation (2.2) follows the standard arguments for the multinomial logit model (e.g., Train, 2009).

The independence of irrelevant alternatives (IIA) property of demand implied by the multinomial logit structure raises a potential concern. To alleviate this, patient choice is parameterized by including rich data at the patient–hospital level to capture the heterogeneity, such as travel time and its interaction terms between diagnosis and hospital facilities. The model also includes hospital fixed effects and interactions of those with disease weights to capture the unobserved quality of hospitals in treating different diseases. The identification comes from the cross-sectional and longitudinal variations in observed hospital choices when characteristics of hospitals, patients, or the choice set vary. For example, the coefficient of travel time between a patient's residence to a hospital's location is identified by the variation of choices of a given hospital among patients who live close relative to patients who live further away in the same year within the same HRR area.

Estimation and Identification of Bargaining and PE Investment Model

In this section, I describe the estimation strategy for parameters of the supply side of the model. I exploit the generalized method of moments (GMM) with the TikTak-Melder-Mead algorithm to search for the global optimization solution.[54]

I parameterize marginal costs and use the FOC in Equation (2.8) to obtain the first two sets of moment conditions. The marginal cost of an outpatient visit is assumed to vary across providers and years. In addition, it is potentially different before and after PE buyouts as

[54] Arnoud et al. (2019) describe the TikTak algorithm and present its applications. In their paper, they benchmark seven global optimization algorithms by comparing their performance on multidimensional test functions as well as a method of simulated moments estimation of a panel model of earnings dynamics. The authors conclude that the TikTak method overall outperforms other algorithms. More details can be found in the Online Appendix.

PE might change a hospital's operational efficiency. Specifically, I parameterize it in an additively separable form multiplying potential operational efficiency gain or loss due to PE buyouts:

$$\mathrm{mc}_{ms} = (1 + \mathrm{PE}_s \cdot g_c) \exp\left(\eta_0 + \eta_t + \eta_r + \boldsymbol{\eta} \times \mathbf{Y}_{\mathbf{s}}^{\mathrm{mc}}\right) + \varepsilon_{ms},$$

wherein η_t and η_r are year and census region fixed effects, η_0 is a constant term, $\mathbf{Y}_{\mathbf{s}}^{\mathrm{mc}}$ is a vector of hospital characteristics including average Medicare outpatient costs per user and average HCC risk score at the HRR where hospital s is located, teaching status, for-profit status, rural status, log of the number of hospital beds, Medicare patient ratio, and Medicaid patient ratio. In addition, PE_s is an indicator of whether hospital s is backed by PE, g_{c} measures the proportional change of operational efficiency for hospitals after PE buyouts, and ε_{ms} is the component of cost that is not observable to econometricians.

Notice that the current specification of marginal costs does not include insurer fixed effects, which implies that there are no systematic factors making the marginal cost of treating patients vary across insurers for a given hospital. It corresponds to the second specification proposed in Gowrisankaran et al. (2015), which enables me to identify the variation of bargaining power across hospital–insurer pairs. There are two reasons why this assumption is appropriate for the current setting. First, hospitals are generally believed not to discriminate against patients based on their insurance carriers (among private insurers) when choosing treatment and medical resources for them. Second, for non-treatment costs in the outpatient setting, such as administrative costs, hospitals face similar burdens regardless of patients' insurance carriers. Plugging the parameterized marginal cost into the FOC given by Equation (2.8), the error term now becomes

$$\begin{aligned}\varepsilon_{ms} = &-(1 + \mathrm{PE}_s \cdot g_c) \cdot \exp\left(\eta_0 + \eta_t + \eta_r + \boldsymbol{\eta} \times \mathbf{Y}_{\mathbf{s}}^{\mathrm{mc}}\right) + \mathrm{NP}_s \cdot (1 - \mathrm{PE}_s) \cdot \tau_{\mathrm{NP}} + \frac{p_{ms}}{1 - B_{sm}} - \\ &\frac{B_{sm}}{(1 - B_{sm}) \cdot q_{ms}(N_m)} \left(\alpha \left[W_m(N_m) - W_m(N_m \backslash J_s) \right] + \sum_{k \in N_m \backslash J_s} p_{mk} \left[q_{mk}(N_m \backslash J_s) - q_{mk}(N_m) \right] \right) - \\ &\mathrm{PE}_s \times \frac{1}{q_{ms}(N_m)} \left(\frac{B_{sm}}{1 - B_{sm}} \Delta V_m + \Delta \varpi_s \right). \qquad (2.10)\end{aligned}$$

ΔV_m and $\Delta \varpi_s$ are represented in Equation (2.6). Equation (2.10) forms the basis for the price moment conditions used in the estimation. In the Online Appendix, I describe how to derive the equilibrium benchmark prices from claims data in detail.

The second set of moment conditions uses cost information filed by hospitals. I adopt a strategy akin to Crawford and Yurukoglu (2012), Byrne (2015), and Hackmann (2019) to construct an additional "cost" moment using the average cost data from HCRIS. Specifically, it is assumed that the average marginal cost $\bar{\text{mc}}_s$ in the outpatient setting of hospital system s is linear to the cost per adjusted discharge in a year, AC_s, computed in Section 2.3. It is expressed as $\bar{\text{mc}}_s = \lambda_1 \cdot \text{AC}_s + \lambda_2$.[55] The underlying argument is that any factor that impacts operational efficiency is hospital-wide and affects the marginal costs in both outpatient and inpatient services proportionally. Previous work in medical research (e.g., Alexander et al., 1996 and Spang et al., 2001) commonly use AC_s to measure hospital operational efficiency. More closely, Ho and Lee (2017) use an average cost measure as the approximation for the marginal cost of inpatient care. This is a restrictive parameterization; however, it does capture the marginal cost changes due to PE intervention. By taking the mean of mc_{ms} across insurers during a year, the hospital-wide average marginal cost is $\bar{\text{mc}}_s = \sum_{m \in M} \frac{q_{ms}}{\sum_{m \in M} q_{ms}} \text{mc}_{ms}$. By first-order differencing $\bar{\text{mc}}_s$ and AC_s across years, the following equation holds:

$$\triangle \bar{\text{mc}}_s = \lambda_1 \triangle \text{AC}_s \Longrightarrow \triangle \left(\sum_{m \in M} \frac{q_{ms}}{\sum_{m \in M} q_{ms}} \text{mc}_{ms} \right) / \triangle \left(\text{AC}_s \right) = \lambda_1,$$

which generates the cost moment condition as follows:

$$\triangle \bar{\varepsilon}_s = \triangle \frac{\triangle \left\{ (1 + \text{PE}_s \cdot g_c) \exp\left(\eta_0 + \eta_t + \eta_r + \boldsymbol{\eta} \times \mathbf{Y}_{\mathbf{s}}^{\text{mc}} \right) \right\}}{\triangle \left(\text{AC}_s \right)}. \tag{2.11}$$

The last set of moment conditions comes from the optimal leverage choice. I parameterize the costs of debt raising as $C_d'(l_s) = \mu_1 + \mu_2 l_s + \xi_s$, which is a linear function of hospitals' leverage. ξ_s is the component of debt costs not observable to econometricians, and it introduces

[55] This assumption can be relaxed by allowing hospital-specific coefficients, λ_1^j and λ_2^j, which yields the same moment conditions as shown afterwards.

another source of endogeneity to the model. Plugging $C_d'(l_s)$ into Equation (2.9) yields

$$\xi_s = \sum_{m \in M_s} \left[B_{sm} \frac{\partial \Delta V_m}{\partial l_s} + (1 - B_{sm}) \frac{\partial \Delta \varpi_s}{\partial l_s} \right] / \mathrm{TA}_s - \mu_1 - \mu_2 l_s^*, \tag{2.12}$$

wherein expressions of $\frac{\partial \Delta V_m}{\partial l_s}$ and $\frac{\partial \Delta \varpi_s}{\partial l_s}$ are provided in the Online Appendix.

Parameters in the supply side of the model, namely τ_{NP}, the non-pecuniary motive for not-for-profit hospitals; $(b_0, \boldsymbol{b})$, the bargaining-power coefficients; $(\eta_0, \eta_t, \eta_r, \boldsymbol{\eta})$, the marginal-cost coefficients; (ϱ, θ, μ), the debt burden and bankruptcy probability parameters; (μ_1, μ_2), the ex ante debt-raising costs; and (g_b, g_c), the changes to the bargaining power and marginal costs after PE buyouts, are estimated by the following moment conditions:

$$\mathbf{E} \left(\begin{array}{c} \varepsilon_{ms} \\ \triangle \bar{\varepsilon}_s \\ \xi_s \end{array} \middle| \boldsymbol{Z}_{ms} \right) = 0,$$

wherein $\boldsymbol{Z}_{ms}$ is a vector of exogenous variables.

Identification

Parameter identification relies on the variation of data and the way it affects constructed moments. The identifying assumption underlying this research design is that exogenous variables and instrumental variables are orthogonal to unobserved shocks to hospitals' marginal costs and debt raising costs. Utilizing price moment condition (2.10), the bargaining-power parameters $(b_0, \boldsymbol{b})$ are informed by variations of the relative importance of an insurer's patient-redistribution costs if hospitals are excluded on respective negotiated prices within the insurer. For example, suppose we observe very different negotiated outcomes between a particular insurer and two hospitals, both of which impose similar patient-reallocation costs for the insurer if either hospital is excluded. Such price variations can translate into hospitals' differential bargaining power. The parameter on bargaining power changes due to PE buyouts (g_b) is identified by variations of bargaining power changes before and after

PE investment for the same hospital.

Given estimates of the bargaining-power parameters, the insurer's weight on enrollee surplus (α) is identified by variations of the relative importance of enrollees' expected utility on negotiated prices within an insurer. One implicit assumption is that α is identical across insurers, akin to Gowrisankaran et al. (2015). Though the model is flexible enough to accommodate the case with α varying across insurers (e.g., α_m being insurer-specific as in Prager and Tilipman, 2020), this enormously increases the dimension of parameters to be estimated, making computation extremely challenging. Furthermore, the simplifying assumption of constant α seems to be innocuous given that I exclusively focus on the commercial health insurance market, wherein payers share quite similar behaviors. Lastly, the debt burden and bankruptcy probability parameters (ϱ, θ, μ) are identified by price variations across PE-backed hospitals that have different levels of debt leverage and revenues from various insurers.

Unlike Gowrisankaran et al. (2015), marginal-cost parameters $(\eta_0, \eta_t, \eta_r, \boldsymbol{\eta})$ are not solely identified from price variations across hospitals in the price moment conditions. Introducing cost moment condition (2.11) also helps to pin down cost parameters. Under certain circumstances, price moment conditions alone cannot distinguish the parameter on the change of marginal costs after PE intervention (g_c) from the non-pecuniary motive parameter (τ_{NP}). For example, if all ownership transition events from not-for-profit to for-profit in the sample coincide with PE buyouts, PE would impact negotiated prices by changing the marginal costs and the not-for-profit status simultaneously, making g_c and τ_{NP} indistinguishable. Therefore, it is important to introduce the cost moment conditions to uniquely identify g_c by comparing the average costs of target hospitals before and after PE intervention. τ_{NP} can then be disentangled by the price differences between the for-profit and the not-for-profit hospitals.

Leverage moment condition (2.12) is helpful to identify the ex ante debt-raising cost parameters (μ_1, μ_2). Given the expected patient quantities from demand estimation and observed prices, μ_1 and μ_2 are essentially identified by variations of expected revenues and chosen leverage among PE-backed hospitals.

Instrumental Variables

In this subsection, I describe a set of instrumental variables used in the estimation. The source of endogeneity is that the model assumes bargaining participants know mc_{ms} and $C'_d(l_s)$ as well as their error terms ε_{ms} and ξ_s. This assumption implies that the price p_{ms}, the leverage choice l_s, and the PE investment choice $\mathrm{PE}_s \in \{0, 1\}$ are potentially correlated with the error terms and hence are endogenous. To identify the effect of prices, I follow the previous literature by including three instruments: enrollees' expected utility for the hospital system, the expected utility per enrollee for each insurer, and the predicted hospital quantity. I follow a common assumption in the industrial organization literature that marginal-cost shifters do not affect preferences directly.

For PE investment decisions, I construct a novel instrumental variable by exploiting the exogenous changes of state regulations—the corporate practice of medicine (CPOM) doctrine. The regulation has been adopted in one form or another by many state legislatures and refined by state courts, state professional licensing boards, and other health regulatory agencies. If a state prohibits the doctrine, only licensed professionals, such as physicians and dentists, may own or control the provision of health care. It essentially bans unlicensed entities (e.g., PE firms) from engaging in the professional practice of medicine. However, it is worth noting that imposing the CPOM prohibition does not mean that corporate or other nonprofessional entities are completely barred from operating in the healthcare sector. Certain types of entities are exempt from the CPOM prohibition in a few states. There is also some leeway for PE firms to circumvent the regulation by, for example, forming management services organizations (MSOs) with target hospitals, rather than directly owning them. Nevertheless, more stringent prohibition of CPOM can hugely increase deal costs since PE must cautiously structure the deal and face potential risks in the future. More discussions about CPOM are provided in the Online Appendix.

In a similar vein as Rice and Strahan (2010), Cain et al. (2017), and Karpoff and Wittry (2018), I construct a CPOM regulation index to measure the regulation's strictness across states based on three aspects: (1) state statutes and regulations prohibiting/allowing CPOM; (2) legal precedent/case law prohibiting or allowing CPOM; and (3) attorney general opinions or state medical board opinions prohibiting or allowing CPOM. The index construction starts from a summary by Michal et al. (2006) of the CPOM statuses across 50 states in 2006.

As the first step, three panels of indicators are created: State statute indicator in 2006 equals (negative) one if the statute of a state explicitly prohibits (allows) CPOM, and zero otherwise. Legal precedent indicator in 2006 equals (negative) one if any previous state supreme court rulings were against (in favor of) CPOM, and zero otherwise. Board opinion indicator in 2006 equals (negative) one if opinions from the attorney general or medical board of a state are against (in favor of) CPOM, and zero otherwise. For years after 2006, I track changes of these indicators by manually searching state statutes, legislation files, court cases, administrative materials, and other secondary materials in the Lexis Advance Research Database. If any record was to relax (strengthen) the CPOM prohibition in a state of year t, a new value will be assigned as the previous value minus (plus) one. For example, Texas expanded the exemptions from the CPOM prohibition by passing a legislative bill (S.B. 894) in 2011, so I change the state statute indicator of Texas from one to zero after 2011. Detailed descriptions of relevant legal events are provided in the Online Appendix.

In the second step, I run a regression of PE-ownership dummies for hospitals on the indicators collected in the first step, after controlling hospital fixed effects, year fixed effects, and time-varying characteristics of hospitals. Results are exhibited in Table OA.10 of the Online Appendix. The CPOM regulation index is then calculated as the sum of predicted values of the three indicators for state s in year t multiplied by one hundred. A lower score implies stronger CPOM prohibition. Figure A.11 exhibits time series of the CPOM regulation indices for a subset of 16 states from 2006 to 2019.

Though not a formal validation of the first stage, Table A.19 provides a heuristic first-stage regression by examining the correlation between PE investment decisions and the CPOM regulation index. Standard errors are clustered at the hospital level. As expected, we see a strong, positive relation between PE and the CPOM regulation index, indicating that PE firms are more likely to target hospitals in states with more lenient CPOM regulations. The Kleibergen-Paap Wald F statistic is 10.4 in column 3, strongly rejecting weak instruments at the 0.1% significance level. Results are robust when (1) clustering standard errors at the state by year level, (2) controlling for Medicaid expansion of states after ACA, (3) dropping hospitals in Texas or Tennessee, and (4) dropping observations in the year of 2012, as shown in the Online Appendix.

To alleviate concerns about the exclusion restriction, I take several steps: (1) Compare hospitals' observables. Specifically, I explore whether significant differences in observables occur between hospitals facing more and less stringent CPOM regulation. A hospital is classified in facing more (less) stringent CPOM regulation if the CPOM index of the state where the hospital is located in a given year is within the bottom (top) 20% of all indices across the state's history. Table OA.12 in the Online Appendix describes the results. I find no significant differences between the two sets of hospitals across a list of observables such as hospital size, ratio of Medicare patients, ratio of Medicaid patients, total outpatient visits, total inpatient days, etc. Only the share of for-profit hospitals is significantly higher when in more lenient states, consistent with the first-stage results that PE is more likely to target and convert hospitals facing less stringent CPOM regulation. (2) Conduct a placebo test. It is expected that the CPOM only affects PE's acquisition of hospitals. It should not impact mergers and acquisitions between hospitals. In the placebo test, I regress an indicator if the hospital involved in any M&As without PE backing on the CPOM regulation index. Table OA.13 reports the results, and indeed the CPOM regulation does not significantly affect the M&A decisions of non-PE-backed hospitals.

The annual ICE BofA U.S. high yield index option-adjusted spread is another instrumental variable used in the estimation and is helpful in identifying the effect of leverage. It is a valid IV since the market prices of credit risk presumably do not correlate with ε_{ms} and ξ_s but do correlate with the leverage choices of PE-backed hospitals (e.g., Axelson et al., 2013 and Davis et al. 2019). Estimates are robust by using an alternative measure—credit spread between high-yield U.S. corporate bonds and the U.S. LIBOR. In addition, the instrument set includes all fixed effects and other exogenous time-varying variables appearing in the marginal-cost and bargaining-power specifications.

2.4.3. Results

Demand Estimates

I estimate the patient demand separately for each HRR, so Table A.20 summarizes the estimates by reporting the visit-number-weighted coefficients and standard errors of all HRRs. The first set of coefficients reports PE-ownership impacts on utility. The positive and statistically significant estimate of the PE dummy indicates that, on average, patients respond positively to the PE ownership of hospitals and quality changes after PE intervention, though different groups respond differently: females are more willing to visit PE-backed hospitals, while senior patients and patients with more severe diseases (reported by the PE $\times$ Weight interaction) are less likely. This result echoes the reduced-form evidence that some hospital quality measures improve after PE buyouts while others are unchanged and consumer satisfaction scores deteriorate. Hence, patients respond heterogeneously to the operational adjustments, and their expected utility towards PE-backed hospitals changes.

Consistent with prior literature, the coefficient of travel time is negative and statistically significant, indicating that patients prefer nearby hospitals. The willingness to travel is increasing in the size of hospitals, teaching status, severity weights of diseases, and patient age, and decreasing if patients are female. The coefficient of past use of the hospital is significantly positive, implying that patients are inertial and more likely to stick to the hospital they have visited before. In addition, patients in need of specific services are more likely to choose hospitals that are able to accommodate their needs. For instance, patients with a psychological or cancer diagnosis are more likely to visit a hospital providing psychological and oncological services, respectively.

Supply Estimates

Table A.21 reports the coefficients and standard errors of parameters related to hospitals' bargaining power and marginal costs as well as impacts of PE intervention. Panel A exhibits estimates of the bargaining-power parameters, which are largely consistent with previous studies (e.g., Lewis and Pflum, 2015). The model assumes that bargaining weights are determined linearly by hospitals' characteristics. To have a better sense, the mean and standard deviation of hospitals' bargaining weights in the sample are 0.374 and 0.269. Several hospital characteristics are associated with larger bargaining power. The estimates imply that multi-hospital systems, for-profit hospitals, teaching hospitals, and hospitals affiliated with any type of physician organization have, on average, higher bargaining power, corresponding to an increase of 0.29, 0.18, 0.28, and 0.11, respectively, in their bargaining weights at a significance level of 0.1%. Other features turn out to affect bargaining power negatively. Hospitals in rural areas have smaller bargaining power. Hospitals with larger bed numbers have smaller bargaining power. At first glance, it seems counterintuitive that larger hospitals have smaller bargaining power. But, as argued in Lewis and Pflum (2015), this does not mean that larger hospitals are unable to exercise market power via their better bargaining positions: I also find that the local market structure impacts hospitals' bargaining power. The estimates suggest that hospitals with larger market shares measured by inpatient days in the local market would enjoy larger bargaining power, while hospitals in markets with more insurers tend to have less bargaining power.

Panel B demonstrates estimates for the marginal-costs specification, which is assumed to be determined exponentially by hospitals' characteristics. The estimated parameters imply that the marginal costs of for-profit hospitals are 6% lower on average; for teaching hospitals, they are 63% higher; and for rural hospitals, they are 31% lower. The results also suggest economies of scale for hospitals: A 1% increase in the number of hospital beds is associated with a 0.2% decrease in marginal costs. In addition, the estimates indicate that hospitals accepting a larger portion of Medicare and Medicaid patients have higher marginal costs. They increase by about 0.1% if the Medicare patient ratio increases by 1%, though the estimate is insignificant. The marginal costs would rise by 1.9% if the Medicaid patient ratio increases by 1%. This is consistent with previous evidence that Medicaid reimbursements are usually lower than those of Medicare, and much lower than private payers. Providers typically lose money by treating Medicaid patients. When a hospital accepts a higher portion of Medicaid patients, subsidizing effects from private payers to Medicaid would be reflected in higher marginal costs in the model. Lastly, average costs in the local market are incorporated into the marginal-cost specification. The result intuitively suggests a positive correlation between the marginal costs of hospitals and the Medicare average outpatient costs of the HRR where these hospitals are located.

Panel C highlights estimates of PE's impacts on price bargaining and other parameters in the model. The estimate for insurers' preference, α, is approximately 1,237. The magnitude is comparable to that in Prager and Tilipman (2020). One possible interpretation is that private insurers treat one unit of consumers' utils as equivalent to \$1,237 in their objective function. The second coefficient is the non-pecuniary motive for not-for-profit hospitals, estimated to be approximately 102 and significantly different from zero. This indicates that not-for-profit hospitals place strictly positive values on social objectives. Providing one unit of medical service brings an equivalent of \$102 in benefits besides profits. This supports prior discussions that not-for-profit hospitals have different objectives than for-profit ones (e.g., Lakdawalla and Philipson, 2006). The third set of coefficients includes the impacts of bankruptcy threats on price bargaining. The estimate of debt burden θ is significantly positive at 0.1% level, consistent with the descriptive evidence in Section 2.3 that higher leverage of PE-backed hospitals imposes greater impacts on price negotiations. The negative estimate of μ and relatively large estimate of ϱ, two parameters governing the distribution of bankruptcy shocks, indicate that typical bankruptcy risks to PE-backed hospitals are low (while increasing in debt leverage) but with large variances. The fourth set of estimates concerns the ex ante debt-raising costs for PE firms. Though statistically insignificant, it suggests that PE-backed hospitals face convex debt-raising costs: the marginal cost of debt is increasing in the amount of debt taken up by a PE-backed hospital. The last set of estimates directly examine PE's impacts on hospitals' operational efficiency and bargaining power. I find that after PE buyouts, hospitals' operational efficiency improves—marginal costs decrease by approximately 8%. At the same time, hospitals' bargaining power increases by 0.19. The magnitude is economically significant in contrast to the average bargaining weights in the sample.

2.4.4. Model Fit

To evaluate the model fit, I conduct two exercises of comparing model predictions with those observed in the data. First, I compare the predicted distribution of hospital–insurer negotiated prices to the distribution of negotiated prices in the data. Figure A.12 displays kernel density plots of the natural logarithm of predicted and observed negotiated price distributions, pooled across all hospital–insurer pairs and years in the sample. Panel A of Figure A.12 demonstrates the price distribution of the full sample, while Panels B and C assess fit of price distributions separately for hospitals ever bought out by PE firms and hospitals never acquired by PE firms. The vertical lines in each plot represent the arithmetic mean of the respective distributions. Recall that individual hospital fixed effects are not included in estimating the bargaining model due to the sizable computation burden. The predicted negotiated prices are purely driven by hospital, insurer, and local market characteristics; patients' demand; PE ownership; and year and region time-invariant factors. Given the relative inflexibility of the model with respect to price variations across hospitals, the overall fit is pretty good. For the subsample of PE-backed hospitals, the model-predicted price distribution does not perfectly align with the realized one. The model-predicted distribution has a left-biased spike and a fat right tail. But the distribution matches the first and second moments well. For the subsample of non-PE-backed hospitals, the whole price distribution matches well.

Second, I compare the predicted distributions of HRRs' outpatient spending to distributions of outpatient spending observed in the data. The total spending in an HRR is determined by the negotiated prices and the weight-adjusted patient quantities in a year. Panel A of Figure A.13 displays the spending distribution for the full sample, while Panels B and C exhibit spending distributions of subsamples of HRRs that ever had PE-backed hospitals and those that never had. The spending distributions fit well in the full sample and the two subsamples, which indicates that the structural model not only captures price variations nicely, but also matches the hospital choices of patients well.

As an additional exercise, I derive the local sensitivity measures following Andrews et al. (2017) to assess the asymptotic bias in the estimated parameters implied by violations of the exclusion restrictions. Figure OA.7 in the Online Appendix presents the results for two selective parameters, g_c and g_b, which measure the impacts of PE buyouts on hospitals' marginal costs and bargaining power. The figure delivers some valuable insights on the estimation. It shows that the asymptotic bias in both parameters is very sensitive to the CPOM regulation index, credit spreads, and expected utility per enrollee. It makes a lot of sense that the CPOM regulation index and credit spreads are vital to identify g_c and g_b, since these variables closely relate to PE investment decisions. It also suggests that the demand-side instruments, such as the expected utility per enrollee, really matter for the validity of the estimates. This is consistent with the practice in prior literature of using demand-side instruments as the main variation sources to identify cost and bargaining parameters.

2.5. Counterfactual Analysis

By using the baseline estimates from Section 2.4, I perform a series of counterfactual analyses. The first considers a scenario in which PE investments in the United States were banned by regulators. Through the lens of the experiment, I quantify the impacts of PE ownership on negotiated prices, total spending, and patient surplus. It also enables me to assess the spillover effects and decompose the relative contribution of various channels. In the second counterfactual, I demonstrate the importance of taking into account the PE ownership of acquirers when reviewing proposed mergers. I compare predictions with regard to post-merger outcomes of a plain model, which ignores the PE ownership and is typically used by Federal Trade Commission (FTC) and Department of Justice (DOJ), and my model.

2.5.1. Ban PE Investments

Whether to ban PE investments in the healthcare sector has been at the forefront of current policy discussions. For example, Congress has recently passed legislation to curb the "surprise medical billing" crisis, in which PE-backed physician staffing firms predominate (Appelbaum and Batt, 2020). So, it is important to understand the impacts of PE buyouts in the sector.

To this end, I examine a counterfactual scenario in which PE investments were prohibited. Specifically, I look into a sample of hospitals in 339 HRR-years that ever had PE-backed hospitals between 2013 and 2018.[56] For each HRR-year, I simulate negotiated prices between hospitals and insurers for the realized case and the case with a PE ban. All prices are adjusted by GDP deflators to dollars in 2019. As an interim step of the simulation, I re-estimate patients' choices of hospitals for the counterfactual ban scenario. Figure A.14 summarizes the aggregate outcomes. The gray bars represent variables in the realized case while the red bars denote the PE-banned case. The total realized outpatient spending in the sample is about \$27 billion. Banning PE ownership leads to a reduction of approximately \$3 billion in spending, which accounts for 11% of the total realized spending. The saving largely comes from decreases in negotiated prices: The quantity-weighted average price is about \$284 in the counterfactual ban and \$317 in the realized case, almost the same percentage reduction as that of the total spending. In contrast, the weight-adjusted quantities of patients barely change from the realized case to the counterfactual scenario.

It is important to note that the weight-adjusted quantities of patients change in the counterfactual because patients are allowed to choose an outside option in the logit model. So, aggregate shift in patient quantities does not capture the reallocation of patients across providers, which would directly impact patients' expected utility and outside options of hos-

[56]Hospitals from a subset of 40 HRR regions are dropped due to missing key variables (e.g., debt levels for PE-backed hospitals). A few other HRR areas are discarded because the model fails to converge for them in the simulation. So, the final sample in the counterfactual analysis only includes hospitals in those 339 HRR-years.

pitals/insurers in price bargaining. I will discuss the implications of this when decomposing the relative contribution of various channels and evaluating patient welfare.

To explore the heterogeneity, Figure A.15 demonstrates how the savings of hospital expenses in the counterfactual correlates with the degree of presence of PE-backed hospitals in the local market. The x-axis of Panel A is the ratio of PE-backed hospitals in an HRR during a year, and the y-axis is the percentage change of spending when banning PE ownership in the HRR. Observations are grouped for every 5% interval in the x-axis. The circle size represents the number of observations within each bin. The red dashed line is a linear fit of data, which exhibits a noticeable downward trend, indicating that hospital expense savings would grow with the number of PE-backed hospitals in a region if PE ownership were banned in the counterfactual. Panel B depicts an alternative specification with the x-axis being the ratio of PE-backed hospital beds in an HRR of a year. The conclusion of more PE presence leading to more savings under the counterfactual is robust.

Next, I investigate the spillover effects and quantify how they vary across rival hospitals. To do so, I categorize hospitals into three groups: (1) PE-backed hospitals (*PE-backed*), which would be directly impacted by PE intervention; (2) non-PE-backed hospitals sharing insurers with the PE-backed one in the local market (*Non-PE/Shared*), which could potentially be affected through changes in insurers' outside options; and (3) remaining hospitals (*Non-PE/Non-Shared*). Figure A.16 reports levels of total health spending, quantity-weighted average prices, and weight-adjusted patient quantities for each group in the realized scenario. The top panel focuses on the *PE-backed* group. The realized spending is about \$5.5 billion, accounting for about 20% of total spending documented in the sample. The quantity-weighted average price is \$588, and the weight-adjusted patient quantity is about 9.3 million. In contrast, the *Non-PE/Shared* group, depicted in the bottom-left panel, has slightly higher total spending with \$6.8 billion, though its quantity-weighted average price is only \$234, less than half of the price in the *PE-backed* group. The total quantity of patients accommodated by *Non-PE/Shared* hospitals is much larger, about 29 million after adjusting for the relative service-mix weights. The bottom-right panel presents the realized amounts for the last group. It shows that hospitals in this group treated the majority of patients in the sample, with total spending amounting to \$15 billion and total quantity of patients being about 48 million. Their quantity-weighted average prices lie in between that of the *PE-backed* and the *Non-PE/Shared* group at \$315.

Figure A.17 reports how these groups respond to the counterfactual ban. The top panel documents the relative contribution of different groups to the total savings after banning PE ownership. Not surprisingly, the *PE-backed* contributes the most among all groups, with about 86% of total saved expenses. Noticeably, hospitals in the *Non-PE/Shared* group also contribute a non-negligible share, reaching 14% of the total savings. This captures the spillover effects of PE intervention: non-PE-backed hospitals respond to the shock, though they are not directly targeted. In contrast, hospitals in the *Non-PE/Non-Shared* group are barely affected, as their contribution is almost negligible. Furthermore, I dissect those changes within the first two groups that have positive contribution to the total savings. Hospitals in the *PE-backed* group undergo dramatic changes, described in the bottom-left panel of Figure A.17. Their total healthcare spending gets cut by 46%, most of it coming from a 49% decrease in quantity-weighted average prices. Nevertheless, they experience an uptick in weight-adjusted patient quantities by 5.5% in the counterfactual ban. In the *Non-PE/Shared* group, hospitals experience modest changes in the counterfactual. The bottom-right panel of Figure A.17 shows that their own total spending decreases by 6.2%, along with a 3.2% reduction in quantity-weighted average prices and a 3.1% decrease in weight-adjusted patient quantities. Noticeably, the price decrease of *Non-PE/Shared* hospitals reflects the spillover effects in the price bargaining. By removing PE ownership, the treatment effects on negotiated prices between PE-backed hospitals and insurers are wiped out. This in turn changes the outside options of local rivals with which these insurers are simultaneously negotiating, resulting in a new level of price for the *Non-PE/Shared* group in equilibrium.

Besides quantifying the aggregate effect in the counterfactual ban, the structural model allows me to decompose it to measure the relative contribution of various channels. Figure A.18 reports the results. Recall that the model highlights five key channels of how PE buyouts affect price bargaining: First, as veteran investors, PE firms have superior experience in business negotiations. They are able to hand down this expertise to portfolio hospitals and strengthen their bargaining skills. To mute this channel in the counterfactual, I set $g_b = 0$. Second, PE has a reputation for closing down financially distressed facilities. By ratcheting up hospitals' debt loads in buyouts via financial engineering, credible bankruptcy threats give hospitals an edge in price negotiations. To evaluate the contribution of this channel, I make both ΔV_m and $\Delta \varpi_s$ equal to zero in the counterfactual. Third, PE intervention can lead to changes in service quality and consumer satisfaction, resulting in a reshuffling of patient demands in the local market. The relative contribution of this channel is reflected in the counterfactual by removing the PE dummy and its interaction terms in the multinomial logit model and then recomputing patients' demand. Fourth, PE would change portfolio hospitals' focus on social objectives, shifting hospitals' willingness to reach agreements in price bargaining. I mute this channel by setting $\tau_{\text{NP}} = 0$ in the counterfactual. Lastly, PE can improve operational efficiency of target hospitals and lower their marginal costs. This channel is muted by letting $g_c = 0$ in the counterfactual. In Figure A.18, the first and the last bars are copied from Figure A.14, corresponding to the total spending in the realized and the counterfactual scenarios. The waterfall bars in between represent the relative contribution of the above five channels. Prominently, PE investors' superior bargaining skills and bankruptcy threats are two major channels affecting bargaining outcomes. The change in bargaining skills explains about 43% of the price and spending decreases in the counterfactual, and the bankruptcy threats contribute around 40%. The other two channels, changes in patient demand and social objectives, contribute around 10% and 8%. The only channel that leads to an increase in spending in the counterfactual is the change of operational efficiency, which increases spending by only 1%. It suggests that PE investors indeed improve the operational efficiency of hospitals, though with relatively small magnitude. Thus, lower efficiency increases the total spending in the counterfactual if PE ownership were to be banned.

So far, I have explored the implications of PE ownership on hospital prices and spending. To push the model to the limit, Table A.22 reports how patient surplus is affected in the counterfactual. I calculate the change of patient surplus on a dollar basis in the following way:

$$\Delta Surplus = \Delta Spending + \Delta Quality.$$

Two assumptions are in order. First, I assume that changes in hospital spending for insurers will be passed to consumers in the form of insurance premium reductions. Due to lack of data, the model does not include the stage of insurers setting insurance premiums. But it is a reasonable assumption that insurers' interests are largely aligned with patients' and a reduction of medical spending would translate to a decrease in premiums. Moreover, markups in health insurance are regulated by the Medical Loss Ratio provision in the ACA (Cicala et al., 2019). Insurers' profit margins are capped at 15% (20% in the individual and small group market segments), thus insurance premiums typically go hand-in-hand with incurred medical expenses. Second, a mild assumption is imposed on the price-to-utility sensitivity. Recall that in the step of estimating patient demand, I exclude the out-of-pocket costs from the indirect utility to simplify the main analysis. In order to obtain the price sensitivity estimate, I add them back here and re-estimate the multinomial logit model of patient choice. I then leverage the sensitivity estimate to convert the change in patients' expected utility into dollar terms in the counterfactual. Results are collected in the first row of Table A.22. Changes in patients' expected utility are mainly due to the change in hospitals' service quality. The results indicate the counterfactual ban worsens service quality on average, leading to an equivalent reduction of $22 million in patients' expected utility. This implies that PE buyouts in general improve hospital services, though the impact is fairly modest compared to the $2.95 billion savings from hospital price changes. Summing both terms, the aggregate patient surplus in the counterfactual increases by $2.92 billion, accounting for 10.7% of the total spending documented in these regions.

It is worth noting that this counterfactual examines a specific scenario in which other sources of capital would have always filled the hole from PE not investing in those targeted hospitals. This is consistent with the majority of previous studies that examine the impacts of PE. However, it leaves out an important alternative scenario in which hospitals have to shut down without the backing of PE investors. This would be an interesting extension to explore in future work.

2.5.2. Evaluate Mergers and Acquisitions

It was not until recently that antitrust enforcers put a spotlight on the consolidation activities of PE firms. For example, FTC chairwoman Lina Khan has recently said she would investigate whether the private equity investment model encourages unfair business practices.[57] Given these trends, in this counterfactual, I underscore the importance of considering PE acquirers' unique traits when regulators reviewing proposed mergers. Otherwise, regulators could potentially underestimate the impacts of mergers on local markets.

Specifically, I conduct merger reviews by using two different sets of tools: One is the typical model used by regulators (e.g., the FTC) to evaluate mergers in the hospital sector (hereafter the *No-PE model*), which mainly focuses on the market consolidation effects arising from mergers. The other is the full-fledged model considered in this paper (hereafter the *PE model*).

To replicate the No-PE model, I re-estimate the patient choice by dropping all PE-related variables. Then, based on the new demand estimates, I re-estimate the bargaining model by removing all PE-related features (e.g., bankruptcy threats) and moment conditions based on PE-related instrumental variables. The estimation results are tabulated in the Online Appendix. Compared to the PE model, most estimates of bargaining weight coefficients are quite stable. For the marginal-cost estimates, the coefficients of for-profit and teaching status become smaller, while the estimates for rural area, Medicare and Medicaid patient

[57] More discussions are provided in the *Wall Street Journal* article (September 30th, 2021) titled "Antitrust Regulators Fix Their Sights on Private Equity" (https://www.wsj.com/articles/antitrust-regulators-fix-their-sights-on-private-equity-11632999600).

ratios, HCC scores, and the average outpatient costs in the local market are larger than those in the PE model. In addition, the No-PE model delivers a larger estimate for the non-pecuniary motive of not-for-profit hospitals and a smaller estimate for insurer's preference on enrollee surplus.

To investigate the discrepancy of predictions between the PE model and the No-PE model, I construct a sample of hypothetical merger cases in 2013 and randomly choose 100 cases to evaluate using both models. Each merger includes a hypothetical acquirer, which must be a PE-backed hospital system, and a target, which is a non-PE-backed hospital system in the same HRR. The hypothetical mergers include, for example, the Steward Health Care System backed by the Cerberus Capital Management in the Boston Hospital Referral Area (HRR code 227) acquiring the Anna Jaques Hospital in the same region. For each hypothetical merger, it is assumed that the resulting number of hospital beds, inpatient day shares in the local market, and total assets are equal to the combined values of the acquirer and the target. It is also assumed that the merging hospital's financial leverage stays the same as that of the PE-backed acquirer prior to mergers, and the network of insurers in the local market does not change after mergers. In the counterfactual, I simulate predictions of hospital–insurer negotiated prices and total spending in the local market before and after each hypothetical merger, and compute their percentage changes. I also look at predictions on patient-surplus changes from both models.

Figure A.19 displays the results. Panel A demonstrates how the total spending in an HRR changes after a merger. The x-axis denotes the change in the HHI of hospital beds in an HRR, and the y-axis denotes the percentage change in the total spending of an HRR. I bin together the hypothetical mergers for every 25-unit increase of HHI and compute the average percentage changes of spending within each bin. The size of a circle in the scatter plot represents the number of merger cases contained in an interval. The blue circles are predictions from the PE model, while the orange circles represent predictions from the No-PE model. Dashed lines in the figure are linear fits of data. As a sanity check, both models predict that changes in the total spending positively relate to changes in the HHI after mergers, which is in line with the theoretical predictions on market competition. However, there is a prominent gap of predictions between the PE model and the No-PE model, reaching as high as 10% in some circumstances. This shows that regulators could systematically underestimate the impacts of mergers on total spending if they follow the typical tool used for merger reviews. Panel B of Figure A.19 delivers a similar message. Instead of the total spending, Panel B examines how the quantity-weighted average prices change after mergers. The pattern resembles the one observed in Panel A. But again, there is a persistent gap of predicted negotiated price changes between the PE and No-PE models. In Panel C, I examine model predictions about the changes in patient surplus. Following a similar strategy as in Section 2.5.1, I translate the surplus changes into dollar terms and compute the change as a percentage of the HRR's total spending. For mergers inducing only an uptick in the HHI ($<$50), the PE model and No-PE model provide indistinguishable predictions about patient-surplus changes. But for mergers prompting larger disruptions in the market structure (increases in the HHI$>$50), the No-PE model clearly underestimates the impacts of mergers on patient surplus compared to the PE model.

To explore the prediction gaps in different groups, Figure A.20 looks into the price predictions for merging hospitals and non-merging rival hospitals after each hypothetical merger. Panel A exhibits percentage changes in the quantity-weighted average prices of the merging hospitals (acquirers plus targets), while Panel B exhibits that of the non-merging rival hospitals (other rival hospitals), evaluated by both models. The prediction gap of negotiated prices is asymmetric for these two groups. For the merging hospitals, as shown in Panel A, the prediction gap could reach as high as 25%. In contrast, Panel B implies that the No-PE model only underestimates the price changes to a modest level, 3% on average, among the non-merging rival hospitals. This, to some extent, underscores that the gaps in predictions largely come from the underestimation of merging parties in mergers.

This counterfactual exercise emphasizes the impacts of mergers resulting from the unique traits of PE acquirers, instead of the anti-competitive behaviors caused by typical market consolidations. Therefore, traditional antitrust laws might not have much bite from a regulatory perspective. However, for a benevolent regulator whose objective is to protect consumer welfare, the exercise conveys a clear message: Mergers involving PE firms systematically exhibit distinct attributes. Regulators could possibly risk underestimating these mergers' impacts on healthcare prices and patient surplus if they treat them as typical mergers.

2.6. Discussion and Robustness

2.6.1. Network Formation

In the baseline model, the networks between insurers and hospitals, namely, N_m and M_s, are assumed to be fixed exogenously. Here, I show that the main estimates are robust to estimating an extended model by adding an endogenous network-formation stage between hospitals and insurers. The new model timeline is exhibited below.

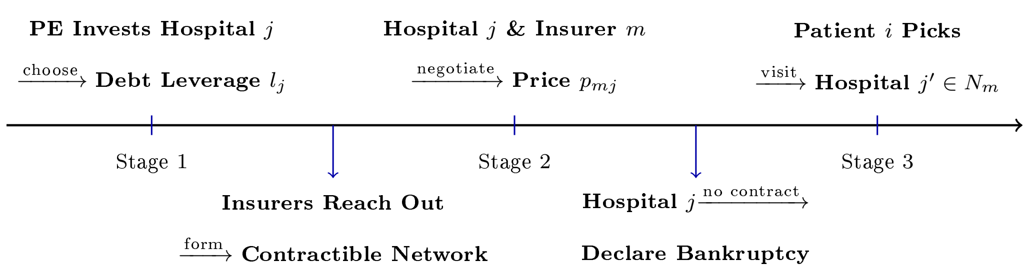

Specifically, in the added stage, any pair of hospitals and insurers must decide whether to form a contractible link when they formally negotiate prices. It is assumed that each insurer reaches out to a subset of hospitals with which it would like to form negotiation links (and hence maximize its expected payoffs). The locked-in bilateral relation results in a contractible network for insurer m, which is a subset of hospitals in the local market, denoted by N_m, and a contractible network for hospital system s, which is a subset of insurers in the local market, denoted by M_s. As a result, insurer m will negotiate prices with every hospital system s with $J_s \subseteq N_m$, and hospital system s will negotiate prices with every single insurer $m' \in M_s$.[58]

This leads to an optimization step for insurers and introduces a set of additional inequality moments from the network-determination decisions. In particular, a subset of necessary conditions are implied by the observed equilibrium network structure: adding hospitals from outside or dropping hospitals within N_m would generate lower expected payoffs to the insurer than what it earned in equilibrium. I construct these moments following previous literature (e.g., Crawford and Yurukoglu, 2012, Pakes et al., 2015, and Prager and Tilipman, 2020). More details about the estimation process are provided in the Online Appendix, and the new estimates are tabulated in Table OA.15 of the Online Appendix. Using the new estimates produces \$2.74 billion savings in the first counterfactual of banning PE ownership, comparable to the main estimates in the paper (\$2.95 billion).

2.6.2. New Entry

Though the structural model incorporates potential hospital closures through PE's bankruptcy threats, it does not explicitly consider entry of new hospitals. Here, I briefly discuss implications of new entry and argue that the structural model is flexible enough to tackle the issue. First, recall that in the model, the price bargaining process is analyzed separately for each local market given a set of hospitals. Any new hospital that enters the market and generates insurance claims that are captured by the database will be included in the set of hospitals and considered in the price bargaining process. So, the entry of new hospitals should not impact the estimation outcomes.

Second, I argue that new entry should have a very moderate—if any—impact on local markets, due to the role of Certificate of Need (CON) regulations across states. CON regulations, adopted by 35 states so far, require healthcare providers to obtain state approval before opening up a new facility or expanding existing facilities, which greatly increases entry

[58] One innocuous timing assumption, which simplifies the pairwise network-stability analysis later on, is that hospitals and insurers simultaneously implement network formation and price bargaining in the model.

costs and deters new entry (e.g., Cutler et al., 2010 and Ho, 2020). Therefore, it is relatively rare to observe entry of new hospitals in the data, and they would have very limited impact on the estimates.

2.7. Conclusion

PE investment in healthcare has ballooned over the past decade. Though it has drawn considerable policy interest among regulators, there is a lack of systematic studies on the impacts of PE buyouts on healthcare markets, regulations, and patient welfare. This paper addresses these questions by introducing novel insurance claims data, structurally estimating a new model, and quantifying PE investors' equilibrium impacts. The paper finds that PE buyouts lead to an 11% increase in bargained prices between PE-backed hospitals and insurers, meaning that insurers pay hospitals more for the same services after a PE buyout. Local rivals respond by also negotiating higher prices, though responses exhibit strong heterogeneity. The counterfactual simulations imply that banning PE ownership would bring gains in patient surplus, mainly by reducing hospital prices rather than changing the quantity or quality of service. It also shows that regulators potentially underestimate the impact of proposed mergers if they ignore PE acquirers' unique features.

Several prior studies show that PE buyouts negatively impact patient welfare mainly through changes in health service quality. Their findings can be reconciled with mine by noting that they study nursing homes, for which the major payer is Medicare. The Centers for Medicare & Medicaid Services (CMS) sets fixed reimbursement rates in Medicare. As a result, PE firms must take aggressive cost-cutting measures to increase revenues, which could negatively affect service quality. When it comes to the hospital sector—the focus of this paper—prices are widely negotiated between hospitals and private insurers. Downgrading service quality might be an unwise strategy for PE firms, because negotiated prices are responsive to service quality. So, as documented in this paper, PE firms make other changes to directly impact the price bargaining outcomes. These mechanisms are closely related to the unique features of PE. However, they can be applied beyond the healthcare sector studied in this paper.

Policymakers should therefore be aware of them when evaluating PE investments involving any business-to-business bargaining settings.

APPENDIX

Tables and Figures

A.1. Figures

Figure A.1: Ideal Experiment

This figure presents the ideal experiment for our main analysis.

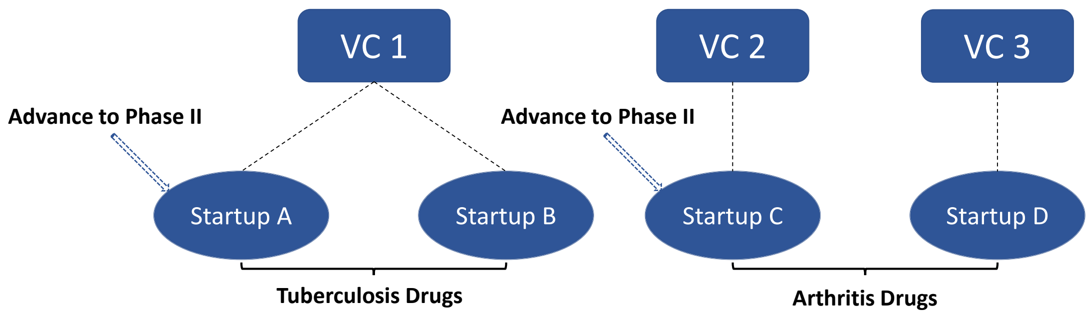

Figure A.2: Geographic Distribution of Drug Projects

This figure shows the number of drug projects in our sample in each state.

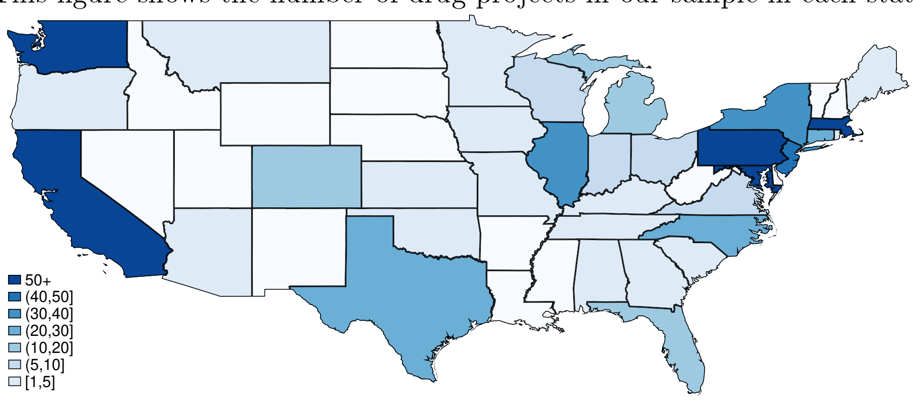

Figure A.3: Coefficient Dynamics

This figure presents the coefficients dynamics in an event-study framework. Panel A contains results for project outcomes. It plots the OLS coefficients β_τ from the following regression:

$$Progress_{i,t} = \sum_{\tau=-6}^{15+} \gamma_\tau Lagging_{i,\{t-\tau_0=\tau\}} + \sum_{\tau=-6}^{15+} \beta_\tau Lagging \times SharedVC_{i,\{t-\tau_0=\tau\}} + FEs + Controls + \varepsilon_{i,t}.$$

where τ_0 is the event quarter, defined as the first quarter when $Lagging \times SharedVC = 1$ for project i. The x-axis denotes τ, the number quarters after τ_0, with $\tau = 0$ as the excluded category. Periods more than 6 quarters before τ_0 are binned into the $-6th$ quarter and denote as $\tau <= -6$ in the figure. Periods more than 15 quarters after τ_0 are binned into the $15th$ quarter and denoted as $\tau >= 15$ in the figure. Panel B plots results for VC financing outcomes. To create Panel B, we replace $Progress_{it}$ by $ExtendFunding_{ijt}$ in the regression above and plots the resulting estimates β_τ. Remaining details are the same as in Tables A.3 and A.5.

Panel A: Project Outcomes

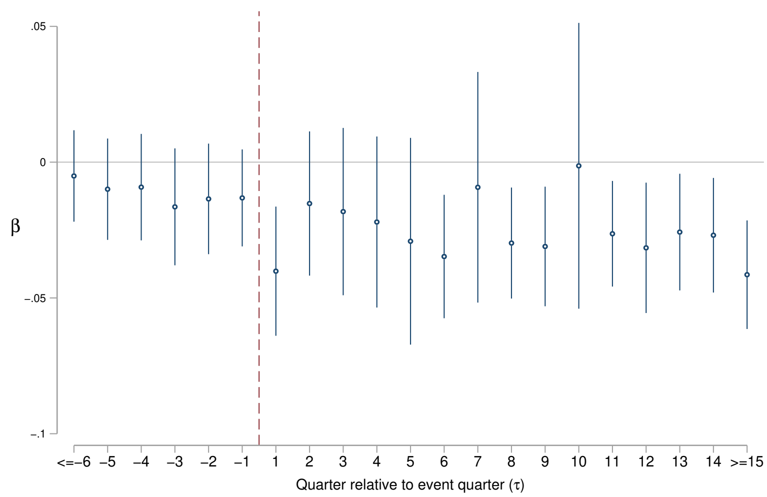

Panel B: VC Financing Outcomes

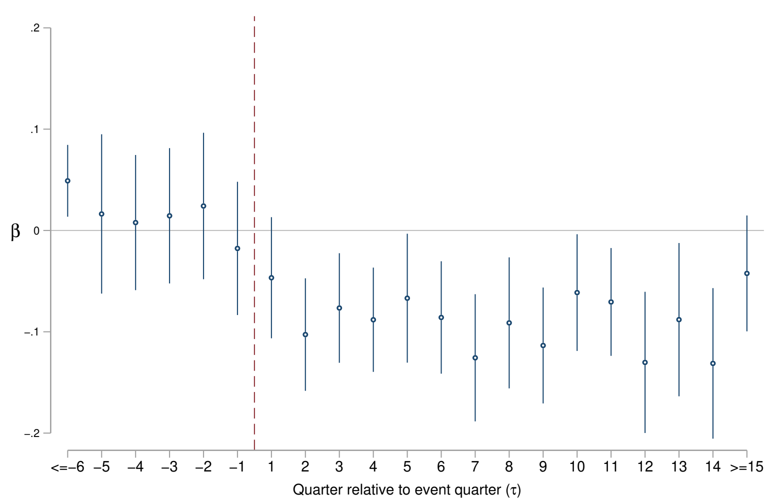

Figure A.4: Redirecting Innovation

This figure presents the ideal experiment for the pivoting analysis.

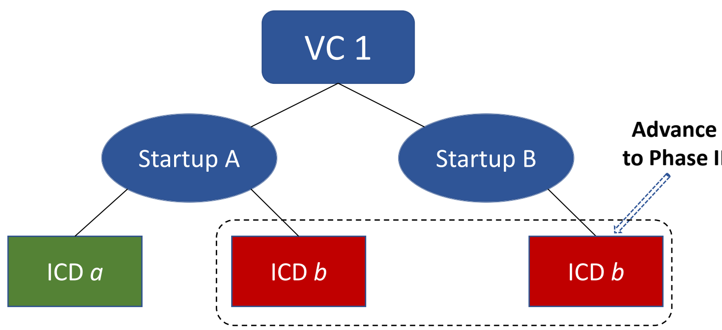

Figure A.5: Innovation Efficiency

This figure shows the relation between innovation efficiency and the common ownership rate. Each circle corresponds to an ICD category. An ICD category is included in the sample if it has at least one drug initiated later than 2005Q1 by a VC-funded startup. The x-axis denotes the ICD's *Common Ownership Rate*, defined as the probability that a randomly chosen pair of startups within the ICD shares at least one VC in common. The y-axis denotes the efficiency proxy, defined as number of FDA approved drugs divided by total VC funding (in millions of dollars), where the numerator and denominator are both aggregated across startups and over time within the ICD. We include VC funding through all drug development phases, not just Phase I, which extends the sample of ICDs to 94 (rather than 78 in our main regression sample), and which reduces the common ownership rate relative to Table A.2. If a startup is operating in multiple ICDs, we allocate the startup's VC funding to each ICD proportionally based on the number of its projects in that ICD. The dashed line denotes the best-fit line; the underlying regression is in column 1 of Table A.9.

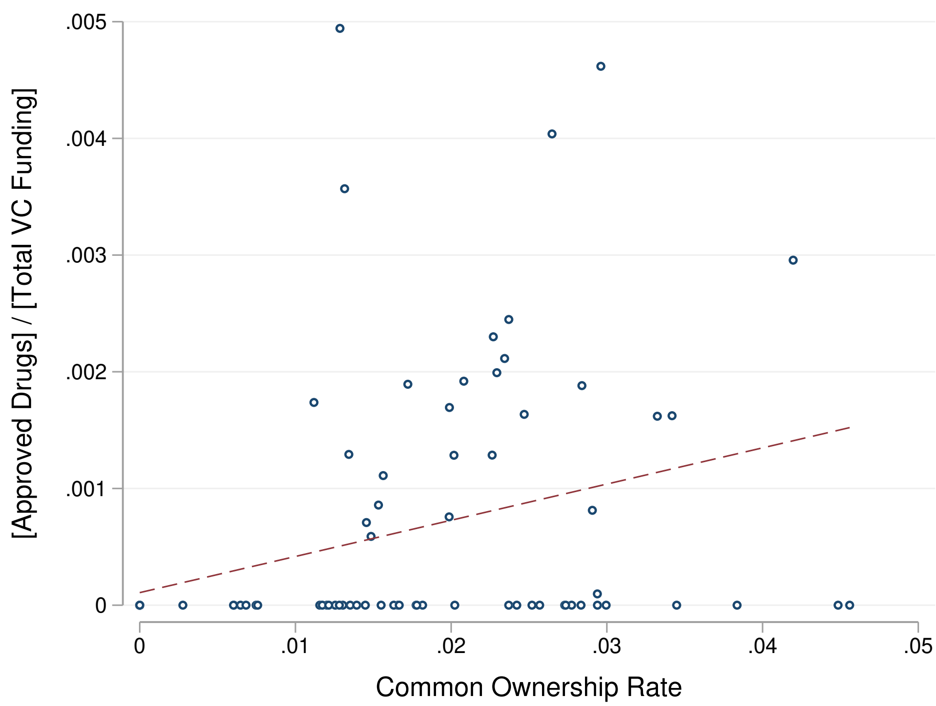

Figure A.6: Total Health Expenditure Trend and PE Investment in Healthcare

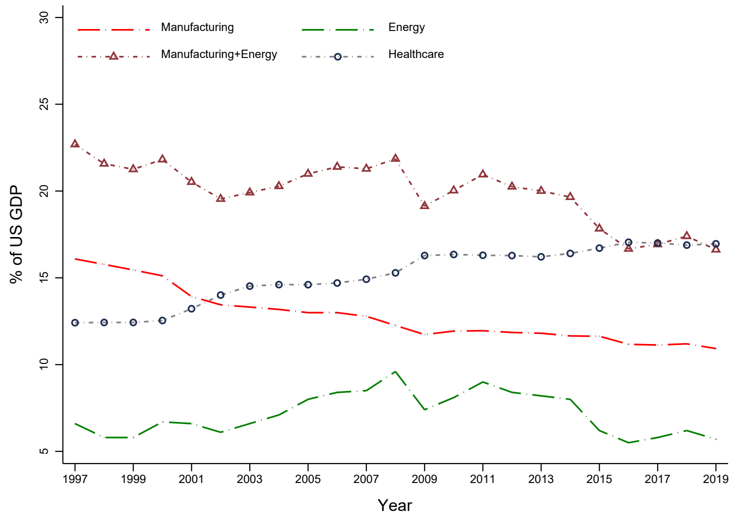

(a) Health Expenditure versus Other Sectors

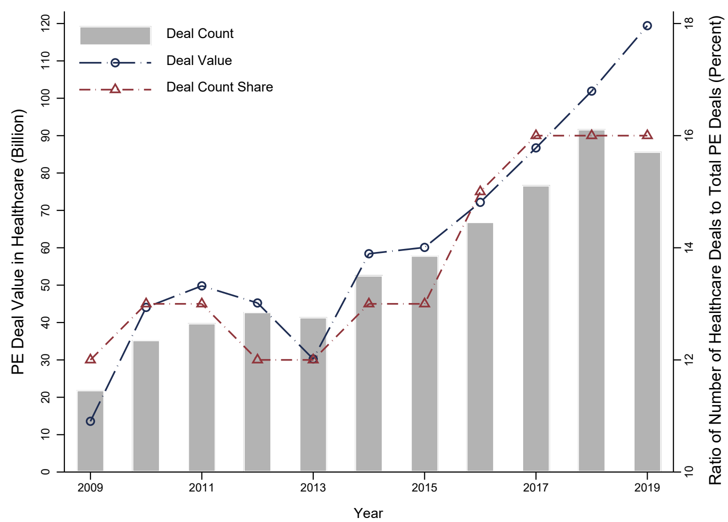

(b) PE Investments in Healthcare

This figure shows the aggregate trends in health spending and PE investments in the healthcare sector. Panel A plots the U.S. health expenditures as well as manufacturing sector and energy expenditures as a percentage of U.S. gross domestic product (GDP) from 1997 to 2019. The data on health-related expenditures within GDP is from the National Health Expenditure Accounts (NHEA). More details about the estimates of healthcare spending are provided here https://www.bea.gov/system/files/papers/BEA-WP2020-8.pdf. Panel B plots the PE healthcare deal values and deal counts as well as the ratio of healthcare deals to total PE deal counts in the United States from 2009 to 2019.

Figure A.7: Hospital Spending and PE Investment

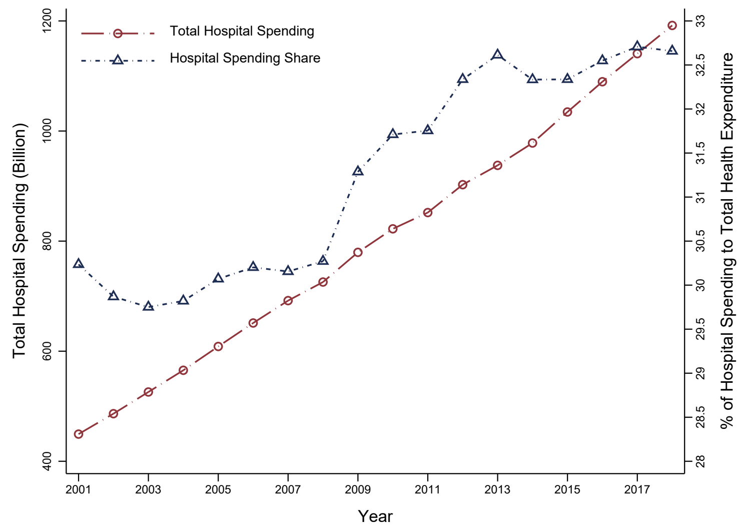

(a) Hospital Spending per Year

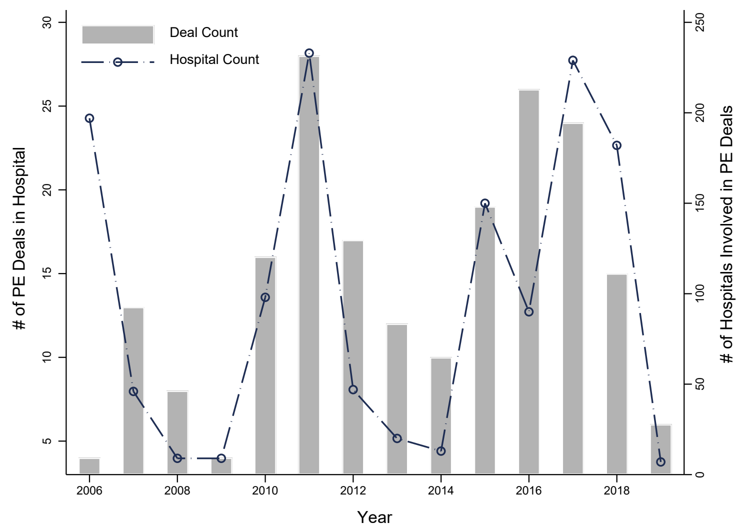

(b) PE Hospital Buyouts within Sample

This figure shows the aggregate hospital spending and PE investments in the hospital sector in the United States. Panel A plots the total hospital spending and the ratio of hospital spending to total health expenses in the United States from 2001 to 2018. Panel B plots the number of PE deals in the hospital sector and the number of hospital facilities involved in these deals in the United States from 2006 to 2019.

Figure A.8: Geographic Distribution of PE-backed Hospitals

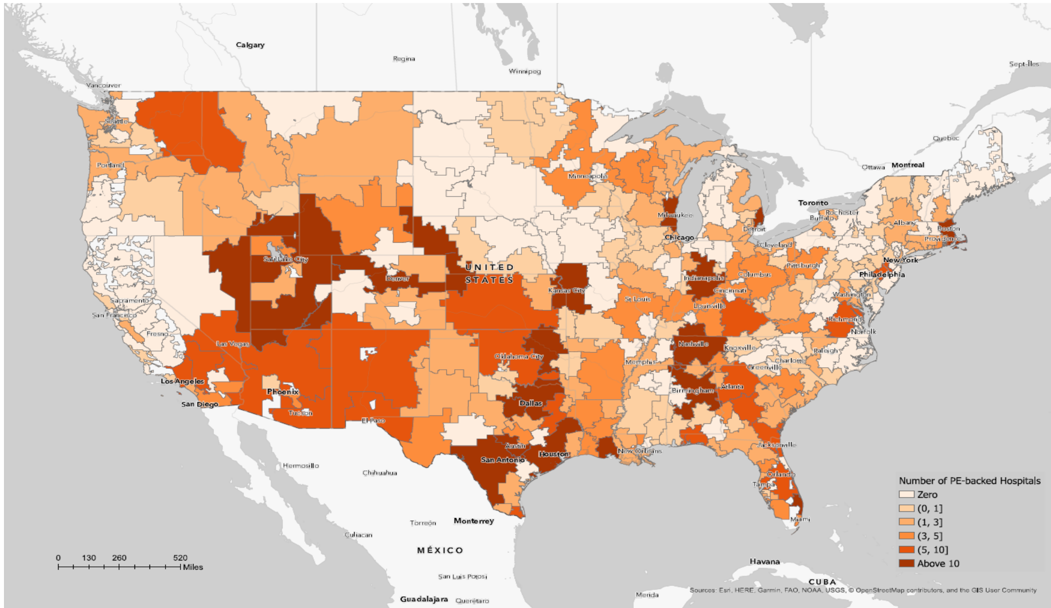

(a) Number of PE-backed Hospitals

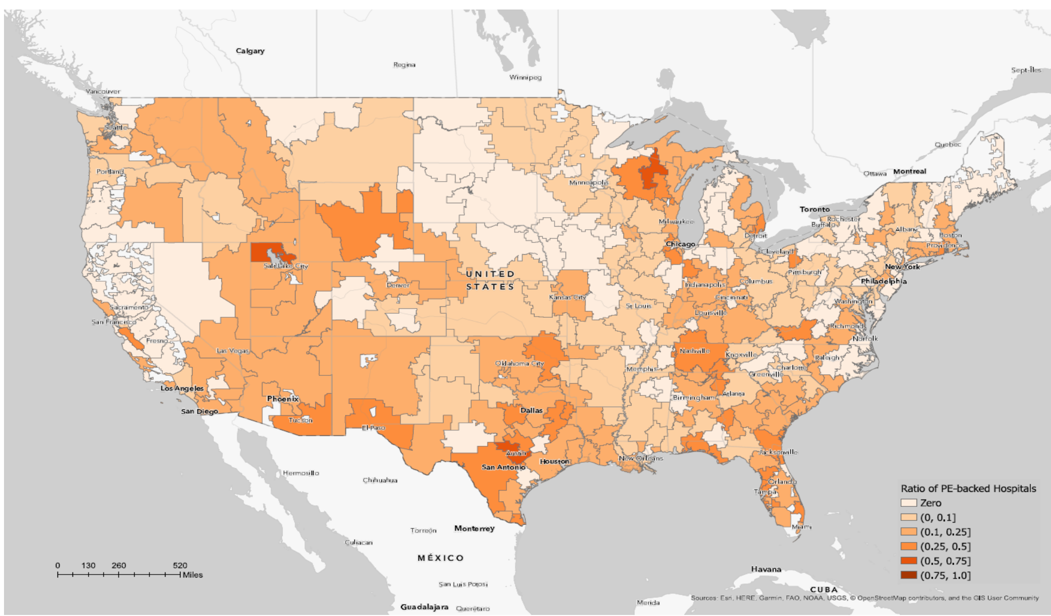

(b) Ratio of PE-backed Hospitals

This figure shows the geographic distribution of PE-backed hospitals at the hospital referral region (HRR) level between 2006 and 2019. Panel A counts the number of hospitals that were ever involved in any PE buyouts in each HRR. Panel B exhibits the ratio of the number of PE-target hospitals to the total number of hospitals in each HRR.

Figure A.9: Dynamic Effects of PE Intervention

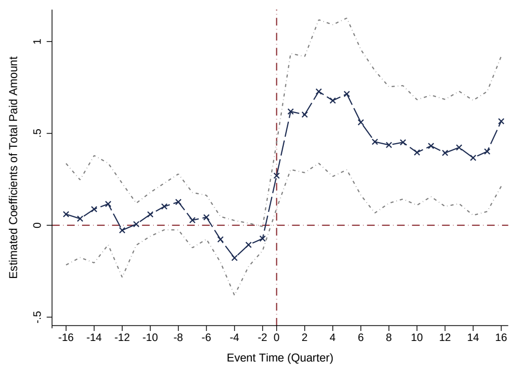

(a) Log of Total Paid Amounts

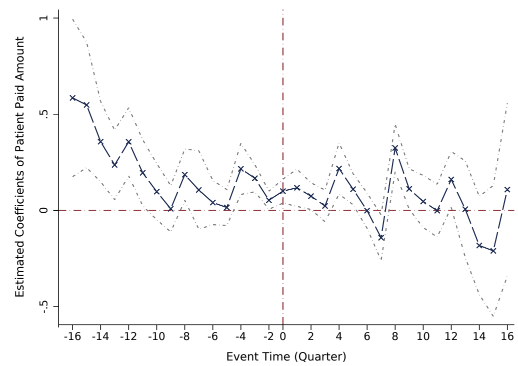

(b) Log of Patient Paid Amounts

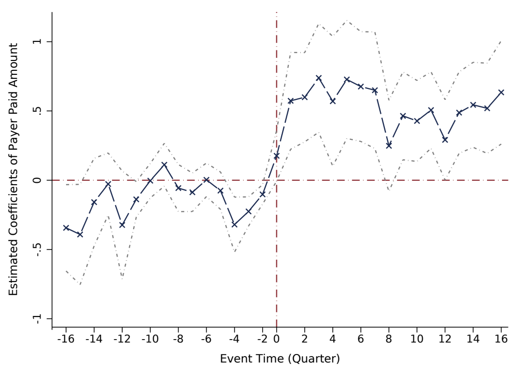

(c) Log of Payer Paid Amounts

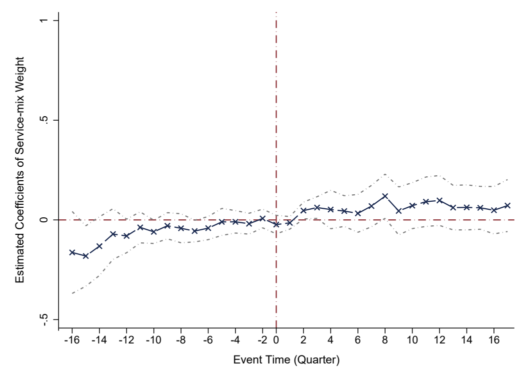

(d) Log of Service-mix Weights

This figure presents the dynamic treatment effects of PE buyouts in event studies. It plots the OLS coefficients α_τ from the following regression:

$$Y_{i(m)jdt} = \sum_{\tau=-16,\tau\neq-1}^{16} \alpha_\tau \mathrm{PE}_{j,\{t-t_0=\tau\}} + \mathrm{Controls} + \mathrm{FEs} + \varepsilon_{i(m)jdt},$$

wherein t_0 denotes the first quarter when $\mathrm{PE}_j = 1$ for hospital j. The coefficient with $\tau = -1$ is excluded as a benchmark category. Any quarters beyond 16 (-16) are binned into the 16th (-16th) quarter. Panel A exhibits the coefficient dynamics using the natural logarithm of the total paid amounts as the dependent variable. Panel B uses the natural logarithm of the patient paid amounts, and Panel C uses the natural logarithm of the payer paid amounts. Panel D uses the natural logarithm of the relative service-mix weights as the dependent variable. All standard errors are clustered at the hospital level. Gray dotted lines represent 95% confidence intervals.

Figure A.10: Regression Coefficients for Medical Imaging Procedures

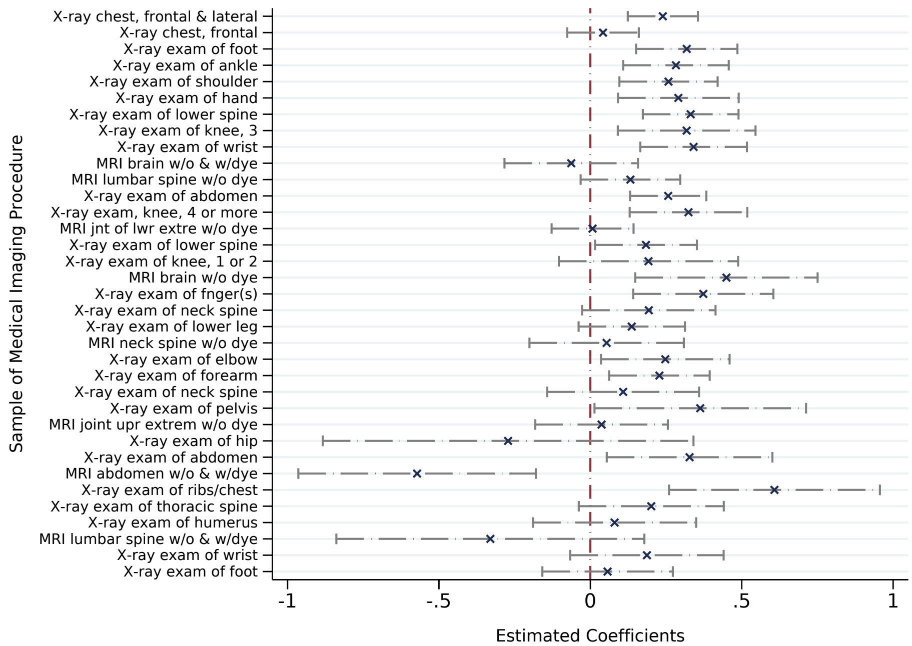

This figure exhibits the estimated coefficients of the PE dummy in Regression (2.1) for the subsamples of the top 35 medical imaging procedures. The y-axis denotes the names of medical imaging procedures. All standard errors are clustered at the hospital level. Capped spikes represent 95% confidence intervals.

Figure A.11: Constructed CPOM Regulation Index across States

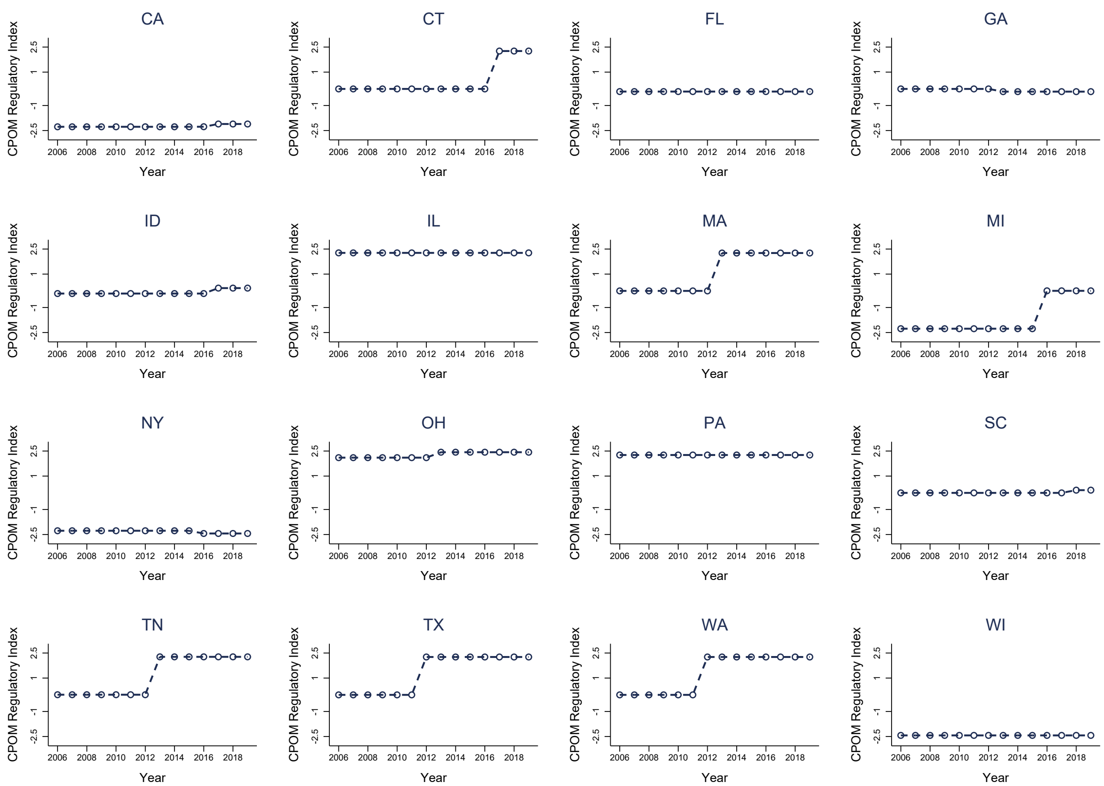

This figure shows the time series of the constructed CPOM regulation indices for 16 states between 2006 and 2019. The detailed construction procedure of the CPOM regulation index is collected in the Online Appendix.

Figure A.12: Model Fit: Negotiated Prices

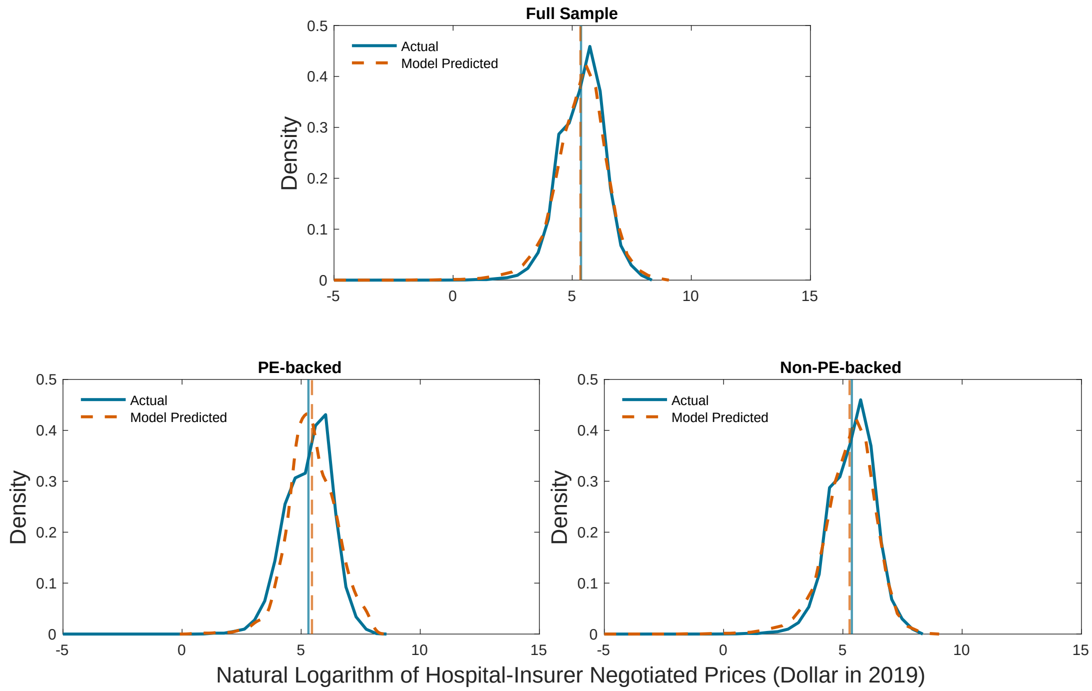

This figure demonstrates the kernel density plots of the predicted and observed distributions of negotiated prices between hospitals and insurers on a log scale for the full sample, the subsample of hospitals ever targeted by PE, and the subsample of non-PE-backed hospitals. All years are pooled, and prices are adjusted to dollars in 2019 by GDP deflators. An observation in the sample is a hospital–insurer–year. Vertical lines denote the mean of the respective distributions. Predicted prices are simulated based on the estimates from Tables A.20 and A.21.

Figure A.13: Model Fit: HRR Outpatient Spending

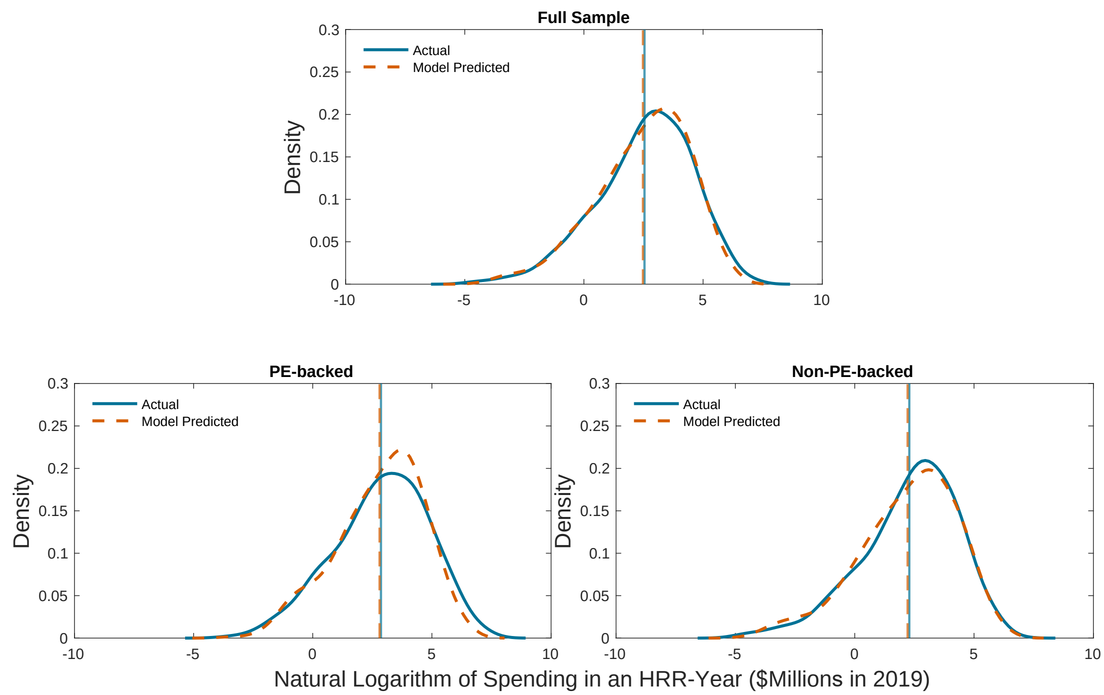

This figure demonstrates the kernel density plots of the predicted and observed distributions of total outpatient spending in local markets ($millions) on a log scale for the full sample, the subsample of HRRs that ever had PE-backed hospitals, and the subsample of HRRs that were never targeted by PE. All years are pooled, and prices are adjusted to dollars in 2019 by GDP deflators. An observation in the sample is an HRR-year. Vertical lines denote the mean of the respective distributions. Predicted spending is simulated based on the estimates from Tables A.20 and A.21.

Figure A.14: 1st Counterfactual: Aggregate Changes

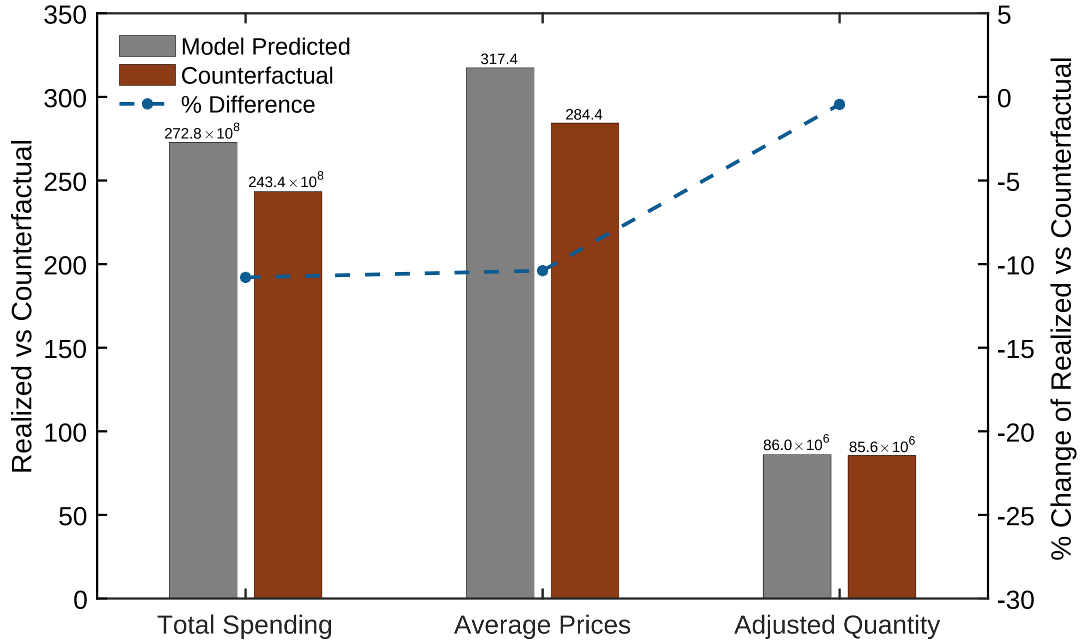

This figure presents outcomes of a sample of 339 HRR-years in the PE-ban counterfactual between 2013 and 2018. Total spending is the sum of the outpatient payments captured by the sample. Average prices are the quantity-weighted negotiated prices between hospitals and insurers. Adjusted quantity is the sum of the relative service-mix weights of all patient visits. Spending and prices are adjusted to dollars in 2019 by GDP deflators. The gray bars represent the model-predicted outcomes in the realized scenario. The orange bars represent the counterfactual outcomes if PE ownership was banned. Both are simulated based on the estimates from Table A.21. Numerical values are indicated on top of each bar. The dashed line represents the percentage differences in outcomes between the realized and the counterfactual scenarios.

Figure A.15: 1st Counterfactual: Spending Changes across Regions

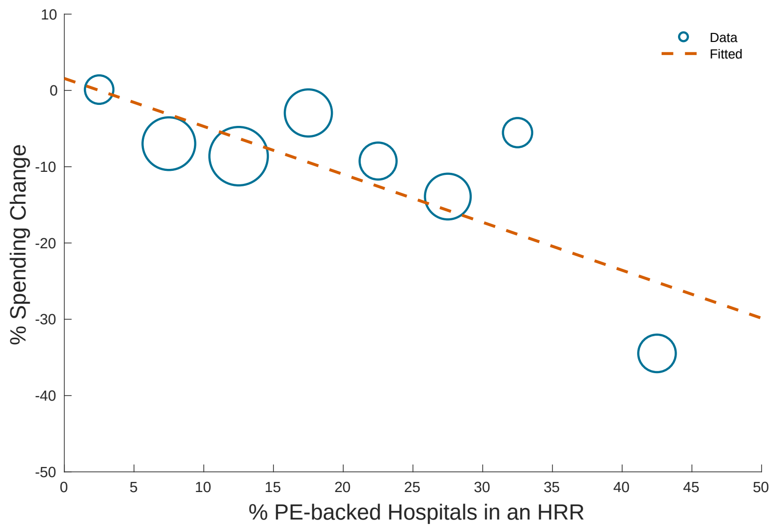

(a) %PE-backed Hospital of HRRs and Spending Change in Counterfactual

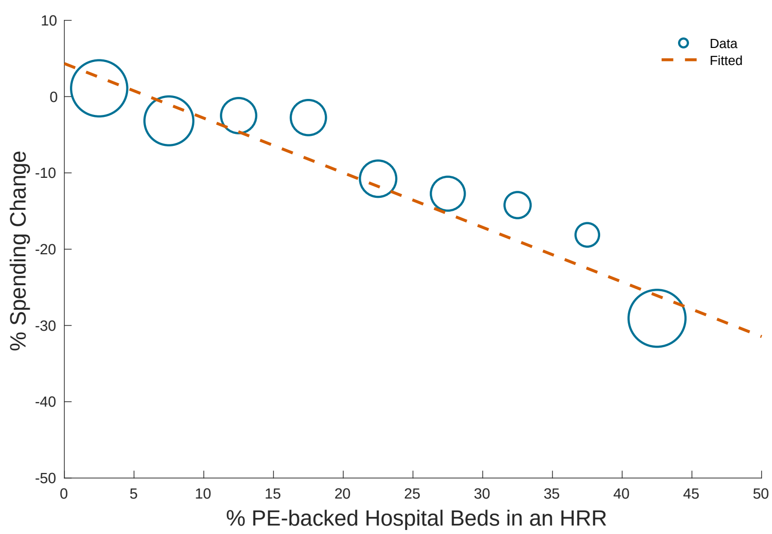

(b) %PE-backed Hospital Beds of HRRs and Spending Change in Counterfactual

This figure presents how the hospital spending varies across regions with different degrees of PE intervention in the PE-ban counterfactual. The unit of observation is an HRR-year. Observations are grouped for every five-percent interval on the x-axis. In Panel A, the x-axis denotes the percentage of PE-back hospitals in an HRR of a year. In Panel B, the x-axis denotes the percentage of PE-backed hospital beds in an HRR of a year. The y-axis denotes the percentage change in hospital spending. Each circle corresponds to the mean percentage change in a bin. Circle size represents the number of observations within each bin. The dashed line denotes the best-fit line.

Figure A.16: 1st Counterfactual: Model-predicted Amounts across Groups

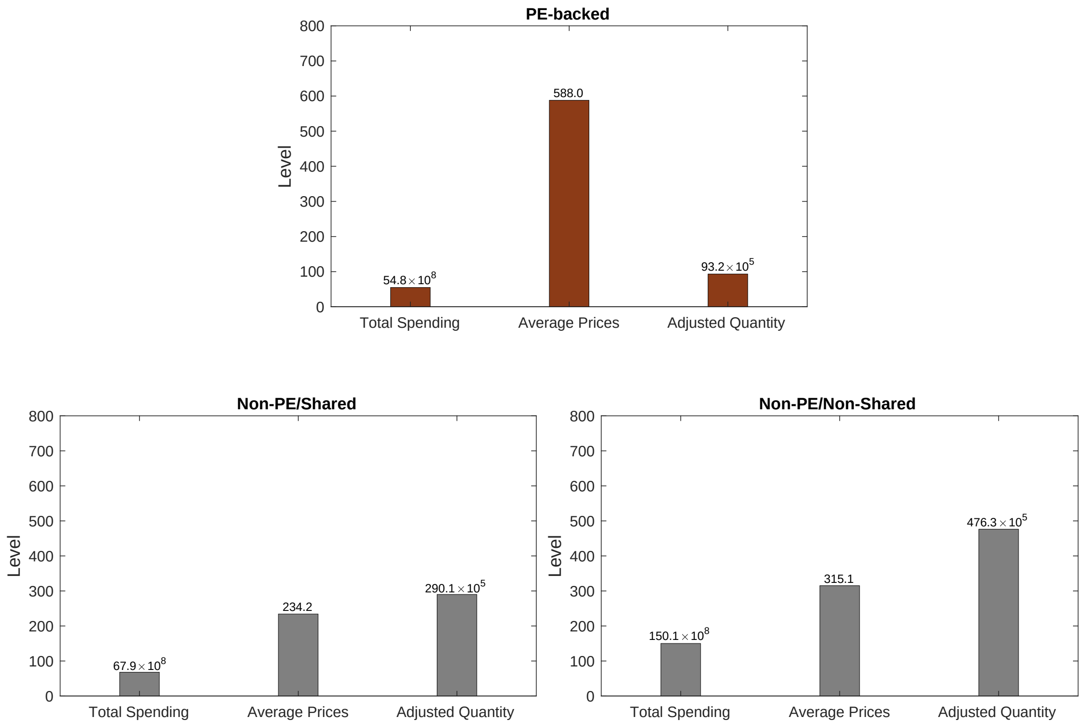

This figure presents the model-predicted outcomes across groups. The top panel exhibits the model-predicted outcomes of the *PE-backed* group, which includes hospitals under PE ownership. The bottom-left panel exhibits the model-predicted outcomes for the *Non-PE/Shared* group, which consists of non-PE-backed hospitals that share common insurers with the PE-backed one in an HRR. The bottom-right panel exhibits the model-predicted outcomes of the *Non-PE/Non-Shared* group, which includes other hospitals that do not belong to the previous two groups. Remaining details are the same as Figure A.14.

Figure A.17: 1st Counterfactual: Quantify Spillover Effects

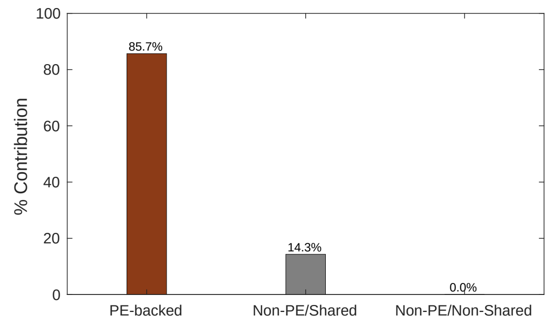

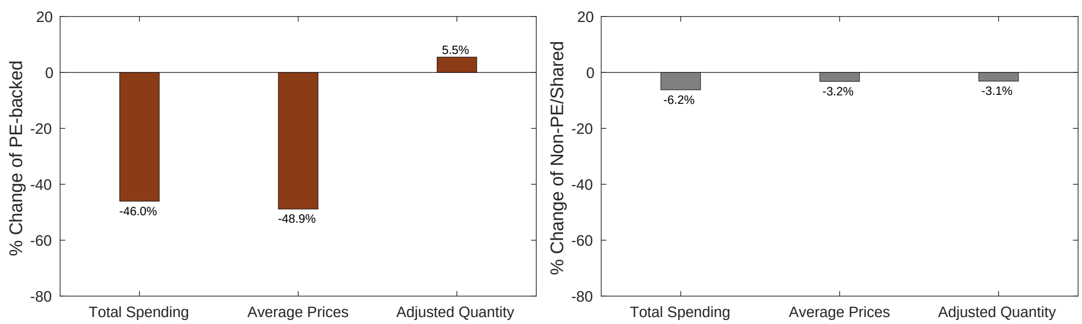

This figure quantifies the spillover effects of PE intervention in local markets. The top panel presents the relative contribution to the total savings in the counterfactual across different groups, including *PE-backed*, *Non-PE/Shared*, and *Non-PE/Non-Shared*. The bottom-left panel presents the percentage changes in the total spending, quantity-weighted prices, and total relative service-mix weights of hospitals within the *PE-backed* group in the counterfactual. The bottom-right panel presents the percentage changes in the total spending, quantity-weighted prices, and total relative service-mix weights of hospitals within the *Non-PE/Shared* group in the counterfactual. Remaining details are the same as Figure A.16.

Figure A.18: 1st Counterfactual: Decomposition by Channels

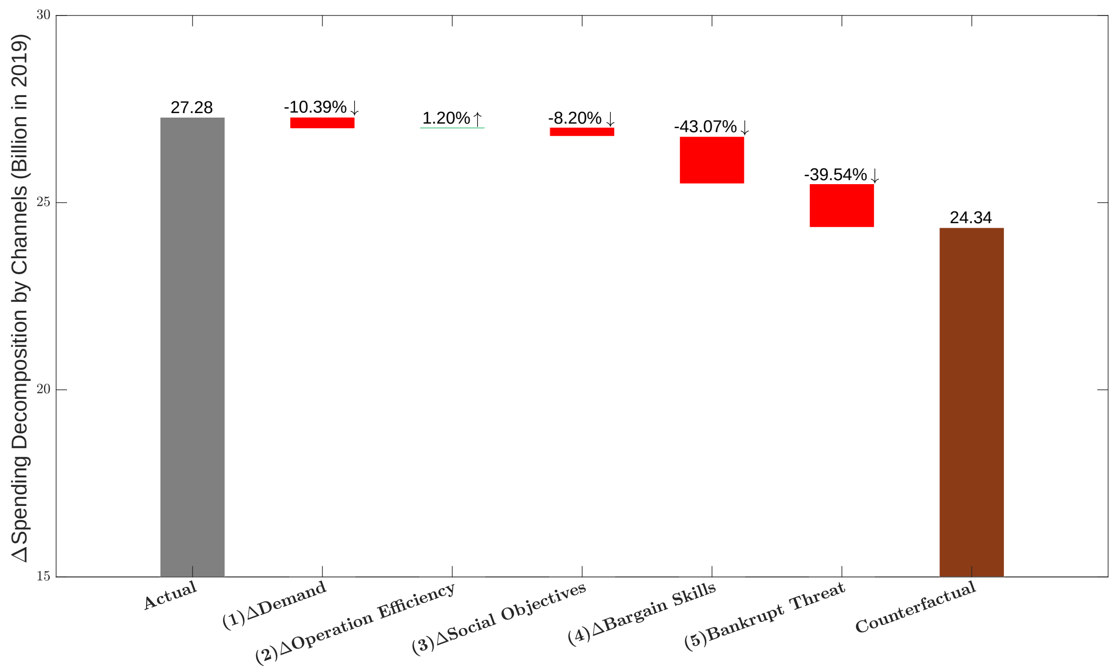

This figure presents the relative contribution of various channels to the total savings in the counterfactual. The first gray bar denotes the model-predicted spending for the realized scenario, and the last orange bar denotes the counterfactual spending if PE ownership were to be banned, which replicates results in Figure A.14. The waterfall charts in between denote the relative contribution of five key channels highlighted in the model: (1) changes in patients' demand; (2) changes in the operational efficiency; (3) changes in the focus on social objectives; (4) changes in hospitals' bargaining skills; and (5) bankruptcy threats and financial engineering. Numbers on top of each bar denote their respective contribution percentages. Red (green) bars indicate positive (negative) contribution in the counterfactual.

Figure A.19: 2nd Counterfactual: Changes after Mergers

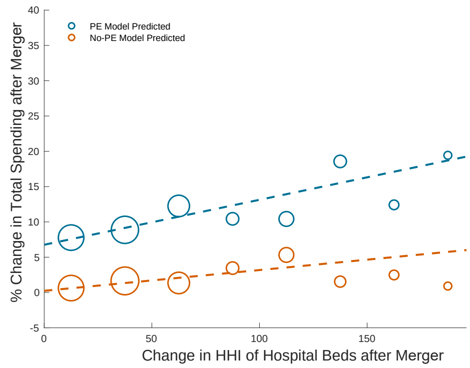

(a) %Spending Changes after Mergers

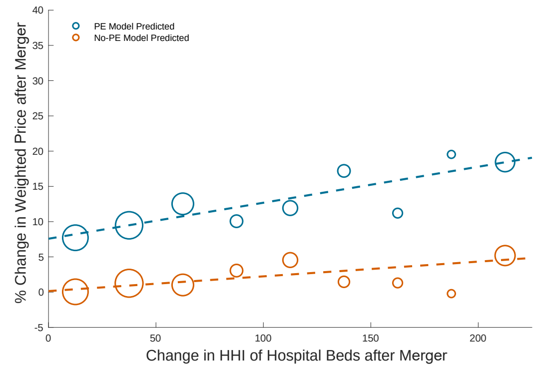

(b) %Price Changes after Mergers

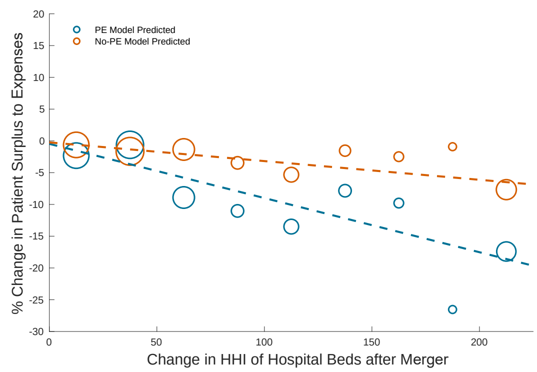

(c) %Surplus Changes after Mergers

This figure presents predictions about changes in the total spending and prices in a sample of 100 hypothetical mergers using the *PE model* and the *No-PE model*. The *PE model* refers to the full-fledged model estimated in Section 2.4.3. The *No-PE model* refers to a plain model after removing all PE-related features in the *PE model*. The sample of mergers is randomly chosen among hypothetical cases with PE-backed hospital systems acquiring local rivals in 2013. The unit of observation is an HRR. Observations are grouped for every 25-unit interval on the x-axis, which denotes the change in the Herfindahl-Hirschman index (HHI) of hospital beds after mergers. In Panel A, the y-axis denotes the percentage change of the total spending in an HRR after mergers. In Panel B, the y-axis denotes the percentage change of the quantity-weighted prices in an HRR after mergers. In Panel C, the y-axis denotes the changes in patient surplus (in dollar terms) after mergers as a percentage of the total spending in the HRR. Changes in patient surplus are calculated following the method described in Section 2.5.1. Each circle corresponds to the mean percentage change in a bin. Circle size represents the number of observations within each bin. The blue circles represent predictions from the *PE model* and the orange circles represent predictions from the *No-PE model*. The dashed line denotes the best-fit line.

Figure A.20: 2nd Counterfactual: Price Changes in Subgroups after Mergers

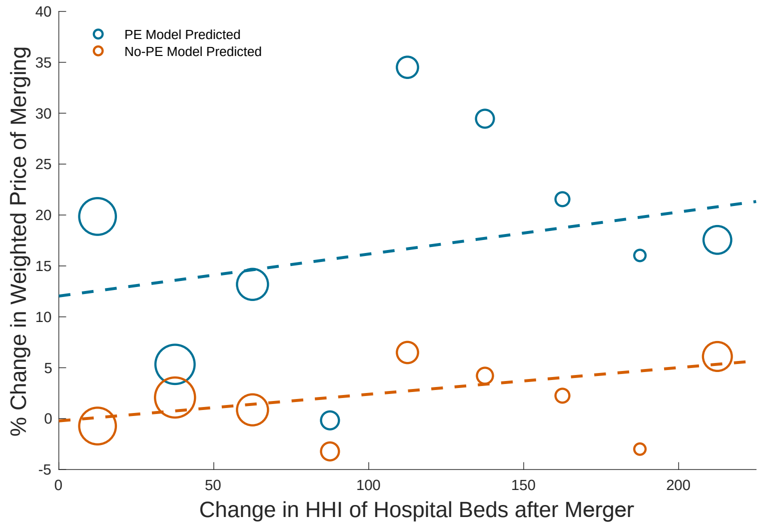

(a) %Price Changes of Merging Hospitals after Mergers

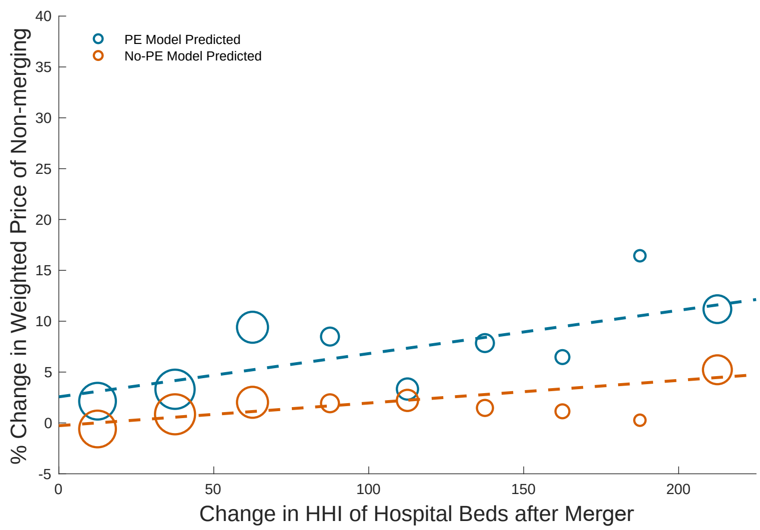

(b) %Price Changes of Non-merging Rival Hospitals after Mergers

This figure presents predictions about the price changes across subgroups in a sample of 100 hypothetical mergers using the *PE model* and the *No-PE model*. Panel A reports the percentage change of the quantity-weighted prices of merging hospitals (acquirers and targets) in an HRR after mergers (relative to the quantity-weighted prices of acquirers before mergers). Panel B reports the percentage change of the quantity-weighted prices of non-merging rival hospitals after mergers. Remaining details are the same as in Figure A.20.

A.2. Tables

Table A.1: Summary Statistics

This table contains summary statistics for our sample of Phase I drug projects in U.S. VC-backed startups. The unit of observation is the project×quarter in Panel A, and the startup company×quarter in Panel B. We show summary statistics for the full sample and the subsamples of observations that are "Never Treated" versus "Ever Treated." The last column reports the difference between those subsamples. Projects/startup companies are categorized as "Never Treated" if the variable $Lagging \times SharedVC$ is equal to zero across its lifespan, and "Ever Treated"' otherwise. $Lagging$ is an indicator for whether a different project (at a different startup) in the same ICD category progresses from phase I to II between the given project's initiation quarter and the previous quarter. $Progress$ is an indicator for whether a project progresses to Phase II in a given quarter. Age equals the number of quarters since the project's initiation. $Number\,of\,Projects$ is the number of projects being developed within the startup company at all clinical stages in the given quarter. $Number\,of\,ICD\,Categories$ is the number of distinct ICD categories covered by those projects. $Number\,of\,VCs$ is the number of VC firms invested in the startup. Ln($Size\,of\,Last\,VC\,Round$) is the log dollar amount that the startup raised in its most recent VC financing round. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively, two-way clustering at ICD category and startup company levels.

	Full Sample						Never Treated	Ever Treated	
Variable	*Mean*	*Std Dev*	*P25*	*Median*	*P75*	*Obs*	*Mean*	*Mean*	*Difference*
				Panel A: Project-Level Variables					
Lagging	0.504	0.500	0.000	1.000	1.000	12,481	0.460	0.762	−0.302***
Progress	0.018	0.132	0.000	0.000	0.000	12,481	0.020	0.005	0.015***
Age	13.444	9.266	6.000	12.000	19.000	12,481	13.305	14.267	−0.962
				Panel B: Startup-Level Variables					
Number of Projects	1.659	1.374	1.000	1.000	2.000	7,745	1.617	2.292	−0.675*
Number of ICD Categories	1.419	0.882	1.000	1.000	2.000	7,745	1.395	1.781	−0.386
Number of VCs	4.688	3.924	2.000	4.000	6.000	7,745	4.698	4.537	0.161
Ln(*Size of Last VC Round*)	8.992	1.522	8.161	9.313	10.096	6,690	9.007	8.780	0.227

Table A.2: Frequency of Common Ownership by VCs

Row 1 reports the percent of startup-quarter observations in which the startup has a close competitor. Startup i is defined as having a close competitor in a quarter if there is another startup in the same quarter developing a drug in Phase I in the same ICD category as one of i's Phase I drugs. Row 2 contains the fraction of startup-quarter observation in which the startup has a close competitor, and that close competitor shares at least one VC in common with the startup in question. Row 3 reports the average number of distinct VCs that own stakes in Phase I drug projects in a given ICD category, averaging across drug categories and quarters. Row 4 reports the average number of distinct ICD categories in which a typical VC owns a stake in a startup with a Phase I project, averaging across VCs and quarters. Row 5 contains the percent of competing startup pairs that share a common VC. Specifically, we count the unique pairs of startups in a given ICD category and quarter that share a VC, and we divide it by the number of possible unique pairs of startups in the ICD category and quarter; we then average this ratio across quarters and ICD categories.

Percent of startups with a close competitor	93.2%
Percent of startups with a close competitor held by same VC	38.6%
Average number of VCs per drug category	16.4
Average number of drug categories per VC	2.52
Percent of competing startup pairs with a common VC	8.6%

Table A.3: Project Outcomes

This table contains estimates of Eq. (1.1). The unit of observation is the project by quarter. The sample contains all project-quarters that are in Phase I trials. Dependent variables are indicated in the column titles. $Progress_{it}$ is an indicator for whether project i progresses to Phase II in quarter t. $Lagging_{it}$ is an indicator for whether another project (at a different startup) in the same ICD category as project i progresses from phase I to II between project i's initiation quarter and quarter $t-1$. $SharedVC_{it}$ is an indicator for whether that progressing project shares a VC with project i in quarter t. $\ln(Age_{it})$ is the log of one plus the number of quarters since the project's i initiation. $NProjects_{it}$ is the number of Phase I projects being developed in project's i's company during quarter t. $NVCs_{it}$ is the number of VC firms that own a stake in project i's startup company in quarter t. $NProjectsperICD_{it}$ is the number of Phase I projects being developed in project's i's ICD during quarter t. Column 1 reports results from OLS regressions. Column 2 reports results from the first stage of the 2SLS regression. $P\,(Lagging \times Proximity)$ is the predicted probability from a probit model with dependent variable $Lagging \times SharedVC$ and independent variable $Lagging \times Proximity$, which is defined in Eq. (1.2) and captures the distance between the lagging project i and the corresponding pioneer project. Footnote 12 describes the probit-2SLS procedure. The table reports the Kleibergen-Paap Wald rk F statistic, with its associated p-value in parenthesis. Columns 4 and 5 report results from the bivariate probit method, with standard errors clustered by ICD category. The coefficient of $Lagging$ is not identified in Column 4. Column 6 reports results from the OLS regression that includes $ICD \times Quarter$ FEs. Standard errors are computed by two-way clustering at the ICD category and startup company levels in the OLS and 2SLS regressions. t-statistics are in parentheses. FEs are noted in the bottom row. ***,**, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS *Progress*	**1st Stage** $Lagging\times$ $SharedVC$	**2SLS** *Progress*	**Biprobit** $Lagging\times$ $SharedVC$	**Biprobit** *Progress*	**OLS** *Progress*
$Lagging \times SharedVC$	-0.019***		-0.070***		-0.038*	-0.018**
	(-2.95)		(-3.82)		(-1.90)	(-2.37)
$P\,(Lagging \times Proximity)$		0.721***				
		(7.55)				
$Lagging \times Proximity$				0.445***		
				(7.05)		
Lagging	0.000	0.101***	0.007		0.000	
	(0.16)	(3.40)	(1.56)		(0.04)	
$\ln(Age)$	0.009***	-0.016**	0.010***	0.212**	-0.053	0.006*
	(2.89)	(-2.17)	(3.03)	(2.05)	(-1.28)	(1.74)
$NProjects$	0.002**	0.005	0.002**	0.000	-0.022	0.002**
	(2.20)	(1.47)	(2.48)	(0.02)	(-1.35)	(2.06)
$NVCs$	0.000	-0.012*	0.000	0.084***	0.023**	-0.001
	(0.10)	(-1.88)	(0.05)	(7.24)	(2.26)	(-0.43)
$NProjectsperICD$	-0.034**	-0.008	-0.034**	-0.060	-0.094**	
	(-2.53)	(-0.20)	(-2.49)	(-0.61)	(-2.47)	
1st stage F-stat		57.07				
		(0.000)				
Startup FE	Yes	Yes	Yes	No	No	Yes
Yr-Qtr FE	Yes	Yes	Yes	No	No	No
ICD FE	Yes	Yes	Yes	No	No	No
ICD × Qtr. FE	No	No	No	No	No	Yes
N	12,469	12,469	12,469	12,481	12,481	11,507
Adj. R^2	0.073	0.553				0.078

Table A.4: Project Outcomes After Excluding Bad Phase I Outcomes

This table repeats the estimation in Table A.3, except now we restrict the sample to drug projects that have no adverse readouts from Phase I clinical trials. Specifically, we match each drug project in the sample with detailed clinical trials information from the Clinical Trials Database of Cortellis. We exclude projects with adverse events (e.g., death of trial participants) in the process of clinical trials and those that failed to reach their primary endpoints. In total, 6,880 observations involving 525 drug projects are dropped. All other details are the same as the previous table.

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS *Progress*	***1st Stage*** *Lagging×SharedVC*	***2SLS*** *Progress*	**Biprobit** *Lagging×SharedVC*	**Biprobit** *Progress*	***OLS*** *Progress*
$Lagging \times SharedVC$	-0.039**		-0.207**		-0.135*	-0.042**
	(-2.36)		(-2.61)		(-1.76)	(-2.19)
$P\,(Lagging \times Proximity)$		0.838***				
		(3.34)				
$Lagging \times Proximity$				0.301***		
				(3.14)		
$Lagging$	0.004	0.114***	0.025		0.002	
	(0.57)	(3.22)	(1.66)		(0.03)	
$\ln(Age)$	0.029***	-0.013	0.030***	0.241**	0.003	0.025**
	(3.99)	(-1.13)	(4.41)	(2.12)	(0.04)	(2.50)
$NProjects$	0.003	-0.004	0.002	0.012	-0.015	0.002
	(1.47)	(-1.59)	(1.17)	(0.51)	(-1.08)	(0.97)
$NVCs$	0.008	-0.014**	0.008	0.092***	0.035**	0.005
	(1.21)	(-2.47)	(1.11)	(5.06)	(2.00)	(0.64)
$NProjectsperICD$	-0.082***	0.002	-0.079**	-0.002	-0.039	
	(-2.82)	(0.03)	(-2.58)	(-0.03)	(-0.95)	
1st stage F-stat		11.12				
		(0.001)				
Startup FE	Yes	Yes	Yes	No	No	Yes
Yr-Qtr FE	Yes	Yes	Yes	No	No	No
ICD FE	Yes	Yes	Yes	No	No	No
ICD × Qtr. FE	No	No	No	No	No	Yes
N	5,592	5,592	5,592	5,601	5,601	4,578
Adj. R^2	0.088	0.563				0.098

Table A.5: VC Funding Outcomes

This table reports results from estimating Eq. (1.3). The sample includes all pairs of VC firms j that invest in startups i. For each pair $\{i,j\}$, we include quarters t when the startup has at least one project in Phase I trials and the quarter is between j's first investment in i and i's exit (e.g., IPO, trade sale). The dependent variable, $ExtendFunds_{ijt}$, is an indicator for whether VC j invests in startup i in quarter t. $Lagging_{it}$ equals one if at least one project in startup i has seen another project (a "pioneer") in the same ICD category (but different startup) progress to Phase II between the former project's inception and quarter $t-1$. $Lagging \times SharedVC_{ijt}$ equals one if VC j holds a stake in both startup i and at least one of the startups owning the pioneering projects that produced $Lagging_{it} = 1$. $SelfProgress_{ijt}$ equals one if startup i has progressed at least one drug project to the next stage of clinical trials between the quarter when VC j first invested in i and quarter $t-1$; we include this variable as a proxy for startup performance. $NCats_{jt}$ is the number of ICD categories covered by all of VC j's portfolio companies during quarter t; this variable measures the VC's diversification. $Duration_{ijt}$ is the log of one plus the number of years since startup i last received funding from VC j. $NProjects_{it}$ is the number of projects under development at startup i in quarter t. $NVCs_{it}$ is the number of VC firms that own a stake in startup i at the beginning of quarter t. $PortfolioSize_{jt}$ is the number of startups in VC j's portfolio in quarter t. $PrevRoundSize_{it}$ is the log dollar amount that startup i raised in its most recent VC financing round before quarter t. Column 1 contains results from OLS. Columns 2–3 contain results from a probit-2SLS specification similar to that in Table A.3. Columns 4 and 5 report results from the bivariate probit model. Column 6 reports results from the OLS regression by controlling $Startup \times Quarter$ FEs. Standard errors are clustered at the startup company level. We report $t-$statistics in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS *ExtendFunds*	$1st$ **Stage** *Lagging*× *SharedVC*	**2SLS** *ExtendFunds*	**Biprobit** *Lagging*× *SharedVC*	**Biprobit** *ExtendFunds*	**OLS** *ExtendFunds*
$Lagging \times SharedVC$	-0.030**		-0.133**		-0.141***	-0.028**
	(-2.26)		(-2.38)		(-9.20)	(-2.32)
$P\left(Lagging \times Proximity\right)$		0.808***				
		(6.00)				
$Lagging \times Proximity$				0.269***		
				(8.83)		
$Lagging$	0.098***	0.018***	0.099***		0.235***	
	(9.91)	(3.21)	(9.86)		(15.55)	
$SelfProgress$	0.022*	0.001	0.022*	0.238	1.700***	0.211***
	(1.75)	(0.10)	(1.75)	(1.43)	(7.84)	(7.70)
$NCats$	-0.005***	0.006**	-0.005***	-0.004	-0.018	-0.005***
	(-3.03)	(2.00)	(-2.69)	(-0.13)	(-0.71)	(-3.78)
$NProjects$	0.001	-0.006	0.002	0.003	-0.061	
	(0.31)	(-1.48)	(0.38)	(0.09)	(-0.76)	
$NVCs$	0.002	-0.000	0.002	0.015	0.016	
	(0.93)	(-0.19)	(0.89)	(1.10)	(1.39)	
$PortfolioSize$	0.005**	-0.011**	0.005**	0.096***	0.075***	0.005***
	(2.22)	(-2.24)	(2.26)	(2.93)	(2.69)	(2.45)
$Duration$	-0.213***	-0.007	-0.214***	0.017	-20.008***	-0.262***
	(-27.00)	(-1.28)	(-26.82)	(0.31)	(-24.49)	(-26.96)
$PrevRoundSize$	-0.001	-0.007	-0.002	0.006	-0.105***	
	(-0.32)	(-1.54)	(-0.54)	(0.12)	(-2.90)	
$1st$ *stage F-stat*		36.04				
		(0.000)				
VC Firm FE	Yes	Yes	Yes	No	No	Yes
Startup FE	Yes	Yes	Yes	No	No	No
Yr-Qtr FE	Yes	Yes	Yes	No	No	No
Startup× Qtr. FE	No	No	No	No	No	Yes
N	32,537	32,537	32,537	32,552	32,552	34,414
Adj. R^2	0.344	0.342				0.578

Table A.6: Aggregate Funding Outcomes for Lagging Firms

The unit of observation in columns 1–3 is the startup (i) by quarter (t). The dependent variable in column 1 is $ExtendFundAgg$, an indicator for whether any VC invests in startup i in quarter t. The dependent variable in column 2 is $AmountRaised$, the percentage increase in cumulative VC funding during quarter t for startup i. Specifically, $AmountRaised$ equals $\log(1 + NewFunds_{it}/CumFunds_{i,t-1})$, where $NewFunds_{it}$ is the dollar amount raised by startup i from all VCs in quarter t, and $CumFunds_{i,t-1}$ is the cumulative amount of VC funding raised by firm i through the end of quarter $t-1$; we take the log of 1 plus the amount since $NewFunds_{it}$ can equal zero. The dependent variable in column 3 is $NewInvestors$, an indicator for whether startup i raises funding from any new VC during quarter t. In columns 1–3, $Lagging_{it}$ equals one if at least one project in startup i has seen another project (a "pioneer") in the same ICD category (but different startup) progress to Phase II between the former project's inception and quarter $t-1$, and $Lagging \times SharedVC_{it}$ equals one if there is any VC that holds a stake in startup i and at least one of the startups owning the pioneering projects that produced $Lagging_{it} = 1$. Control variables $SelfProgress_{it}$, $Duration_{it}$, $NProjects_{it}$, $NVCs_{it}$, and $PrevRoundSize_{it}$ are defined as in Table A.5, except we collapse them to the startup-quarter level in columns 1–3. Column 4 resembles the specification of column 1 in Table A.5, which is at VC-startup-quarter level. The dependent variable is $Extendfunds_{ijt}$, an indicator for whether VC j invests in startup i in quarter t. One distinction is that our independent variable $Lagging \times SharedVC_{it}$ is collapsed to the startup-quarter level, as in columns 1–3. Additionally, we interact our focal variables with $NonCommonVC_{ijt}$, a dummy for whether VC j is not a common owner of a pioneering project, if there is any, competing with startup i in quarter t. We estimate all models by OLS. Standard errors are clustered at the startup company level. FEs are noted in the bottom row. ***,**, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
	$ExtendFundsAgg$	$AmountRaised$	$NewInvestors$	$ExtendFunds$
$Lagging \times SharedVC$	-0.086***	-0.075***	-0.024*	-0.059***
	(-3.42)	(-4.60)	(-1.70)	(-3.09)
$Lagging \times SharedVC \times NonCommonVC$				0.018
				(0.96)
$Lagging$	0.112***	0.015	-0.018*	0.085***
	(6.62)	(0.99)	(-1.72)	(4.28)
$Lagging \times NonCommonVC$				-0.011
				(-0.60)
$SelfProgress$	0.008	0.019	-0.033**	0.017
	(0.35)	(1.06)	(-2.22)	(1.62)
$NProjects$	0.012	-0.003	0.001	0.002
	(1.51)	(-0.43)	(0.23)	(0.37)
$NVCs$	0.030***	0.039***	0.021***	0.014***
	(4.55)	(4.39)	(4.83)	(5.17)
$Duration$	-0.230***	-0.042***	-0.047***	-0.122***
	(-23.79)	(-4.87)	(-8.64)	(-19.28)
$PrevRoundSize$	-0.013**		0.014***	-0.009**
	(-2.01)		(4.10)	(-2.49)
$\log(CumFunds_{i,t-1})$		-0.276***		
		(-4.35)		
$NCats$				-0.003**
				(-2.05)
$PortfolioSize$				0.002
				(1.16)
VC Firm FE	No	No	No	Yes
Startup FE	Yes	Yes	Yes	Yes
Yr-Qtr FE	Yes	Yes	Yes	Yes
N	7,704	7,640	7,704	34,045
Adj. R^2	0.371	0.290	0.079	0.182

Table A.7: VC Influence and Diversification

This table presents results from estimating the VC-funding regression (Eq. 1.3) in subsamples based on either the VC's financial commitment to the startup or a measure of the VC's portfolio diversification. The dependent variable is an indicator for whether VC j extends funding to startup i in quarter t. Columns 1 and 2 compare results across subsamples of non-lead and lead VCs, where a lead VC is defined as the VC whose total amount invested to date is the highest across all the startup's VCs. Our data report the amount invested by each VC syndicate but not by each syndicate member. In cases where these missing data create ambiguity about the lead-VC measure, we assume all syndicate members invest equal amounts. Columns 3 and 4 compare results across subsamples of low- versus high-diversification VCs, where a high-diversification VC is defined as one whose average active portfolio size is above the sample median (19.6). Active portfolio size is the number of unique startups (all industries) invested by the VC in the past five years. This table reports OLS estimates. Remaining details and variable definitions are the same as in Table A.5.

	(1)	(2)	(3)	(4)
	Non-Lead	**Lead**	**Low Diversification**	**High Diversification**
$Lagging \times SharedVC$	-0.005	-0.083***	-0.049*	-0.009
	(-0.40)	(-3.35)	(-1.87)	(-0.86)
$Lagging$	0.052***	0.099***	0.039***	0.079***
	(8.23)	(8.49)	(6.17)	(10.19)
$SelfProgress$	0.012*	0.004	0.019***	0.011
	(1.76)	(0.22)	(2.83)	(0.93)
$NCats$	-0.002	-0.006	0.000	-0.003*
	(-1.28)	(-1.55)	(0.04)	(-1.70)
$NProjects$	0.000	0.006	0.001	0.001
	(0.22)	(0.88)	(0.26)	(0.22)
$NVCs$	0.011***	0.022***	0.012***	0.014***
	(6.58)	(2.97)	(4.18)	(6.45)
$PortfolioSize$	-0.001	0.010***	-0.002	0.002
	(-0.28)	(3.20)	(-0.53)	(0.95)
$Duration$	-0.099***	-0.175***	-0.072***	-0.156***
	(-17.86)	(-13.83)	(-12.99)	(-16.82)
$PrevRoundSize$	-0.009***	-0.006	-0.003	-0.012***
	(-2.74)	(-0.59)	(-0.67)	(-2.78)
VC Firm FE	Yes	Yes	Yes	Yes
Startup FE	Yes	Yes	Yes	Yes
Qtr FE	Yes	Yes	Yes	Yes
N	24,532	7,988	16,726	15,804
Adj. R^2	0.139	0.229	0.105	0.217

Table A.8: Redirecting Innovation

The unit of observation is the startup (i) by ICD category (j) by quarter (t). For each startup and quarter, the sample includes all ICD categories in which the startup has an active Phase I project at some point during the quarter. In column 1, the dependent variable $Initiation_{ijt}$ is an indicator for whether startup i initiates any new project in ICD category j during quarter t. In column 2, the dependent variable $RepurposeInto_{ijt}$ is an indicator for whether startup i repurposes one of its existing drugs into ICD category j during quarter t. In column 3, the dependent variable $RepursposeAway_{ijt}$ is an indicator for whether startup i repurposes a drug project from ICD category j into a different ICD category during quarter t. In column 4, the dependent variable $Alliance_{ijt}$ is an indicator for whether startup i forms an alliance with another pharmaceutical company in ICD area j during quarter t. The independent variable of interest is $Treated \times NonSharedICD_{ijt}$, an indicator for whether (1) startup i has a project X in a *different* ICD category $j' \neq j$ that has experienced $Lagging = 1$, meaning a different project (from a different firm) in ICD category j' progresses to Phase II between the project X's initiation and quarter $t-1$; and (2) startup i and the startup that caused $Lagging = 1$ share a VC investor in common. $ShockedICD_{ijt}$ is an indicator for whether another project (from a different startup) in ICD category j progressed to Phase II between startup i's earliest initiation in ICD j and $t-1$. $SelfICDProgress_{ijt}$ is an indicator for whether startup i has a project in ICD j progressing to Phase II in quarter t. $NICDProjects_{ijt}$ is the number of active projects that startup i has in ICD category j in quarter t. $\text{Ln}(ICDAge_{ijt})$ is the logarithm of one plus the number of quarters from startup i's earliest initiation in ICD j to t . All models are estimated by OLS, and standard errors are clustered at the startup level. FEs are noted in the bottom panel. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1) *Initiation*	(2) *RepurposeInto*	(3) *RepurposeAway*	(4) *Alliance*
$Treated \times NonSharedICD$	0.034**	0.026***	-0.036***	0.009*
	(2.19)	(3.10)	(-3.79)	(1.73)
$ShockedICD$	0.016	0.017**	-0.019**	-0.001
	(1.28)	(2.21)	(-2.02)	(-0.27)
$SelfICDProgress$	-0.016	-0.084***	0.091***	-0.003
	(-0.37)	(-3.00)	(2.83)	(-0.48)
$NICDProjects$	0.003***	0.002***	-0.001***	-0.000
	(3.99)	(4.52)	(-4.14)	(-0.22)
$\ln(ICDAge)$	-0.122***	-0.080***	0.069***	-0.002
	(-8.46)	(-7.65)	(6.33)	(-0.82)
Startup-Qtr FE	Yes	Yes	Yes	Yes
ICD FE	Yes	Yes	Yes	Yes
N	5,374	5,374	5,374	5,374
Adj. R^2	0.313	0.209	0.437	0.119

Table A.9: Innovation Efficiency

We estimate OLS regressions where the unit of observation is the ICD category. The dependent variable is our efficiency proxy, the number of approved drugs scaled by aggregate VC funding. $Common\,Ownership\,Rate$ is the probability that a randomly chosen pair of startups within an ICD shares at least one VC in common. In column 2, we add following variable as controls: *Duration to Phase III* is the logarithm of one plus the average number of quarters between project initiation date and reaching Phase III, averaging across all projects within the ICD that reach Phase III. We measure the duration through Phase III rather than FDA approval since some ICDs never see a project reach approval. *Prob. Reach Phase III* is the probability that a project initiated within this ICD eventually reaches Phase III. *Num. VCs per Startup* is the average number of VC firms per startups active in the ICD. *VC Holding Duration* is the logarithm of one plus average number of quarters between a VC firm's first investment and exit from the startup, averaging across all startups within the ICD. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
Common Ownership Rate	0.031***	0.026***	0.026***	0.025**	0.025**
	(3.621)	(2.701)	(2.703)	(2.387)	(2.345)
Duration to Phase III		-0.000	-0.000	-0.000	-0.000
		(-1.240)	(-1.010)	(-0.982)	(-0.976)
Prob. Reach Phase III			0.001	0.001	0.001
			(0.347)	(0.379)	(0.371)
Num. VCs per Startup				0.000	0.000
				(0.397)	(0.378)
VC Holding Duration					-0.000
					(-0.139)
N	94	94	94	94	94
R^2	0.125	0.139	0.140	0.142	0.142
Adj. R^2	0.115	0.120	0.112	0.103	0.093

Table A.10: Enforcement of Employee Non-Compete Agreements

This table examines whether our results differ across states with strong versus weak enforcement of employee non-compete agreements. Column 1 repeats our OLS regression of project outcomes, from column 1 of Table A.3, except we include interactions terms $Lagging \times SharedVC \times Noncompete$ and $Lagging \times Noncompete$. $Noncompete_{it}$ is an indicator for whether startup i's state in year t strongly enforces non-compete agreements. We obtains states' current non-compete statutes from Beck Reed Riden LLP (http://www.beckreedriden.com/50-state-noncompete-chart-2/). We record states' historical statute changes following Ewens and Marx (2018) and Jeffers (2021). We then use the U.S. Department of Treasury's report (https://www.treasury.gov/resource-center/economic-policy/Documents/UST%20Non-competes%20Report.pdf) to classify the enforceability of non-compete agreements into 5 levels: not enforced, undecided, red pencil, blue pencil, and reformation. We classify blue pencil and reformation as having a high degree of enforceability. Column 2 repeats our OLS regression of VC funding outcomes, from column 6 of Table A.5, with a similar triple-interaction term. Standard errors are computed by two-way clustering at the ICD category and startup company levels. FEs are noted in the bottom row. ***,**, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	(1) $Progress$	(2) $ExtendFunds$
$Lagging \times SharedVC$	-0.015**	-0.021**
	(-2.15)	(-2.04)
$Lagging \times SharedVC \times Noncompete$	-0.002	-0.008
	(-0.28)	(-0.48)
$Lagging$	0.001	0.265***
	(0.11)	(7.26)
$Lagging \times Noncompete$	0.000	0.013
	(0.06)	(0.30)
$Noncompete$	0.008	
	(1.01)	
$\ln(Age)$	0.010***	
	(2.98)	
$NProjects$	0.002**	
	(2.23)	
$NVCs$	0.000	
	(0.14)	
$SelfProgress$		0.091***
		(3.61)
$NCats$		-0.003**
		(-2.13)
$PortfolioSize$		0.002
		(1.36)
$Duration$		-0.172***
		(-18.96)
Startup FE	Yes	No
Startup× Qtr. FE	No	Yes
VC Firm FE	No	Yes
Yr-Qtr FE	Yes	No
ICD FE	Yes	No
N	12,469	34,414
Adj. R^2	0.071	0.465

Table A.11: Technological Similarity

This table reports how the results from regression Eqs. (1.1) and (1.3) vary depending on technological similarity between the lagging project and the pioneering project. Columns 1 and 2 report the project-level results. Columns 3 and 4 report the VC-funding results. The proxy for technological similarity is based on patent citations, which we collect from the Google Patents database. $Lagging \times SharedCite$ equals one if the lagging and pioneering drug projects have at least one outgoing citation in common. $Lagging \times SharedVC \times SharedCite$ equals one if the lagging drug shares at least one citation in common with at least one progressing drug of a different company, and both those companies share a common VC. To provide a reference point, Columns 1 and 3 estimate the regressions without a triple-interaction term. Remaining details and variable definitions are the same as in Table A.3 and A.5.

	(1)	(2)	(3)	(4)
	Progress	*Progress*	*ExtendFunds*	*ExtendFunds*
$Lagging \times SharedVC$	-0.020***	-0.017**	-0.021*	-0.009
	(-3.03)	(-2.50)	(-1.85)	(-0.89)
$Lagging \times SharedVC \times SharedCite$		-0.020*		-0.103***
		(-1.84)		(-3.60)
$Lagging$	-0.000	-0.000	0.069***	0.069***
	(-0.06)	(-0.08)	(6.79)	(6.77)
$Lagging \times SharedCite$	0.010	0.012	-0.003	0.001
	(1.00)	(1.12)	(-0.16)	(0.08)
$NProjects$	0.002**	0.002**	0.001	0.001
	(2.20)	(2.24)	(0.26)	(0.27)
$NVCs$	0.000	0.000	0.012***	0.012***
	(0.13)	(0.11)	(4.37)	(4.34)
$\ln(Age)$	0.009***	0.009***		
	(2.85)	(2.82)		
$NProjectsperICD$	-0.034**	-0.034**		
	(-2.58)	(-2.58)		
$SelfProgress$			0.019*	0.019*
			(1.70)	(1.67)
$NCats$			-0.003**	-0.004**
			(-2.16)	(-2.20)
$PortfolioSize$			0.003	0.003
			(1.32)	(1.25)
$Duration$			-0.121***	-0.122***
			(-19.31)	(-19.52)
$PrevRoundSize$			-0.009**	-0.009**
			(-2.46)	(-2.53)
VC Firm FE	No	No	Yes	Yes
Startup FE	Yes	Yes	Yes	Yes
Yr-Qtr. FE	Yes	Yes	Yes	Yes
ICD FE	Yes	Yes	No	No
N	12,469	12,469	32,537	32,537
Adj. R^2	0.073	0.073	0.179	0.180

Table A.12: Summary Statistics of PE Deals

This table reports summary statistics for the sample of PE hospital deals in the United States between 2006 and 2019. The unit of observation is the PE deal in Panel A, and the hospital in Panels B and C. Panel A focuses on the classification of deal types, and Panels B and C focus on the characteristics of PE-target hospitals and their ownership statuses prior to PE intervention.

Panel A: Classification of PE deals

Deal Type	# of Deal	Avg. # of Hospitals
Add-on	149	4.85
Private to Private	33	13.71
PE Growth/Expansion	32	4.42
Secondary Buyout	18	9.06
Public to Private	6	26
Management Buyout	5	9.40
Total	243	6.91

Panel B: Characteristics of PE-target hospitals

	Yes	No
Rural Area	153	685
Teaching School	12	826
Critical Access	66	525

Panel C: Previous ownership Status

Ownership Status	# of Hospitals
Local Government/Hospital Authority	34
Church Operated	105
Other Not-for-profit	125
For-profit (corporation)	466
For-profit (partnership)	101
For-profit (individual)	7

Table A.13: Summary Statistics

This table reports summary statistics for the sample of outpatient visits aggregated from the DRG insurance claims between 2013 and 2019. The unit of observation is the patient visit in Panel A, the hospital in Panel B, the county in Panel C, and the HRR at Panel D. Summary statistics are presented for the full sample and the subsamples of "Never Treated" and "Ever Treated." The last column reports the difference between these subsamples. Patient visits are categorized as "Never Treated" if the visiting hospital is non-PE-backed, and "Ever Treated" otherwise. Counties and HRRs are categorized as "Never Treated" if no hospitals in the region were ever owned by PE across the sample period, and "Ever Treated" otherwise. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively, clustering at the hospital level.

	Full Sample			Never Treated	Ever Treated	
Variable	Mean	SD	Median	Mean	Mean	Diff
Panel A: Patient Visit Level						
Female	0.620	0.485	1.000	0.620	0.622	-0.002
Patient Age	45.720	23.589	50.000	45.627	46.582	-0.955
Travel Time (Minutes)	41.677	48.447	31.517	41.932	39.025	2.907
Visit Before	0.472	0.499	0.000	0.478	0.378	0.100***
Service-mix Weight	6.508	49.604	1.435	6.357	9.017	-2.660***
Charge Amount ($)	2,675.287	8,508.866	726.308	2,625.482	3253.815	-628.333***
Total Paid Amount ($)	814.324	77,525.560	193.316	810.419	810.158	0.261
Patient Paid Amount ($)	143.182	77,484.770	0.000	142.408	148.095	-5.687
Payer Paid Amount ($)	671.422	2,440.494	141.614	668.283	662.385	5.898
Panel B: Hospital Level						
Num of Beds	150.103	184.304	82.000	149.929	152.566	-2.637
Total Personnel	841.584	1,423.151	346.000	855.721	642.016	213.705***
Teaching (%)	4.689	21.140	0.000	4.939	1.155	3.784***
Rural Area (%)	34.759	47.621	0.000	36.060	16.399	19.661***
Inpatient Days (k)	36.342	52.029	17.429	36.490	34.261	2.229
Outpatient Visits (k)	116.124	209.429	47.668	118.507	82.482	36.025***
Medicare Ratio (%)	48.676	22.775	51.893	48.417	52.328	-3.911***
Medicaid Ratio (%)	18.438	16.734	15.168	18.589	16.303	2.286***
Panel C: County Level						
Poverty Rate (%)	15.922	6.032	15.100	15.933	15.821	0.111
Median Household Income ($k)	46.487	12.366	44.312	46.016	50.901	-4.885***
Insurance Coverage (%)	88.291	5.647	89.115	89.063	86.252	2.811***
Private Insurance Coverage (%)	68.784	9.461	69.321	70.091	65.339	4.752***
Medicaid Coverage (%)	17.414	6.499	16.767	17.057	18.355	-1.298**
Medicare Coverage (%)	16.147	4.370	15.590	16.187	16.041	0.146
Panel D: HRR Level						
HHI in Hospital Beds	0.183	0.126	0.154	0.212	0.146	0.066***
HHI in Inpatient Days	0.209	0.143	0.171	0.241	0.170	0.071***
Num of Hospitals	19.560	18.314	14.000	14.735	25.671	-10.936***

Table A.14: Impacts of PE Ownership

This table reports the results of Regression (2.1) for the full sample of patient visits between 2013 and 2019. Panels A to D examine the impacts of PE ownership on the natural logarithm of the total paid amounts, patient paid amounts, payer paid amounts, and relative service-mix weights. All columns contain hospital×payer fixed effects. Columns (2) to (4) include diagnosis×year fixed effects. Patient controls, including *gender*, *age group*, *insurance type*, and *relative service-mix weights* (except in Panel D), are added in columns (3) and (4). Hospital controls, including *hospital bed numbers*, *teaching status*, *rural status*, *for-profit status*, *ratio of Medicare patients*, and *ratio of Medicaid patients*, are added in column (4). Standard errors are clustered at the hospital level. t-statistics are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Logarithm of total paid amount

PE	0.300**	0.304**	0.309**	0.319**
	(2.056)	(2.180)	(2.278)	(2.340)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.227	0.285	0.308	0.308
Observations	70,861,556	70,861,542	70,861,542	70,861,542

Panel B: Logarithm of patient paid amount

PE	0.009	0.016	0.037	0.046*
	(0.267)	(0.449)	(1.375)	(1.681)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.290	0.347	0.352	0.352
Observations	70,852,714	70,852,700	70,852,700	70,852,700

Panel C: Logarithm of payer paid amount

PE	0.295**	0.290**	0.279**	0.280**
	(2.083)	(2.132)	(2.065)	(2.053)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.124	0.168	0.183	0.183
Observations	70,861,127	70,861,113	70,861,113	70,861,113

Panel D: Logarithm of relative service-mix weight

PE	0.013	0.017	0.017	0.019
	(0.440)	(0.635)	(0.695)	(0.720)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.214	0.303	0.305	0.305
Observations	70,861,556	70,861,542	70,861,542	70,861,542

Table A.15: Impacts of PE Ownership in Matched Sample

This table reports the results of Regression (2.1) for the matched sample of patient visits between 2013 and 2019. The matched sample is constructed by matching each PE-backed hospital to three control hospitals using the optimal Mahalanobis method. Remaining details are the same as in Table A.14.

Panel A: Logarithm of total paid amount

PE	0.649***	0.644***	0.672***	0.689***
	(2.923)	(3.134)	(3.415)	(3.456)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.182	0.258	0.288	0.288
Observations	11,140,533	11,140,503	11,140,503	11,140,503

Panel B: Logarithm of patient paid amount

PE	0.045	-0.028	0.003	-0.005
	(0.938)	(-0.545)	(0.075)	(-0.126)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.233	0.314	0.322	0.322
Observations	11,139,158	11,139,128	11,139,128	11,139,128

Panel C: Logarithm of payer paid amount

PE	0.671***	0.725***	0.729***	0.746***
	(3.220)	(3.794)	(3.823)	(3.862)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.087	0.140	0.156	0.156
Observations	11,140,502	11,140,472	11,140,472	11,140,472

Panel D: Logarithm of relative service-mix weight

PE	0.076*	0.037	0.037	0.049
	(1.784)	(0.897)	(0.999)	(1.352)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.182	0.290	0.294	0.294
Observations	11,140,533	11,140,503	11,140,503	11,140,503

Table A.16: Outcomes on Medical Imaging Procedures

This table reports the results of Regression (2.1) for the sample of the top 35 medical imaging procedures. All procedures are pooled together in the estimation. The dependent variable is the natural logarithm of the total paid amounts. All columns contain hospital×payer fixed effects. Columns (2) to (5) include year fixed effects. Columns (3) to (5) include imaging procedure fixed effects. Patient controls, including *gender*, *age group*, *insurance type*, and *relative service-mix weights* (except in Panel D), are added in columns (4) and (5). Hospital controls, including *hospital bed numbers*, *teaching status*, *rural status*, *for-profit status*, *ratio of Medicare patients*, and *ratio of Medicaid patients*, are added in column (5). Standard errors are clustered at the hospital level. t-statistics are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:				
	Logarithm of total paid amount				
	(1)	(2)	(3)	(4)	(5)
PE	0.099**	0.121***	0.104***	0.098***	0.112***
	(2.225)	(3.003)	(2.689)	(2.617)	(3.031)
Hospital Controls	N	N	N	N	Y
Patient Controls	N	N	N	Y	Y
Procedure FE	N	N	Y	Y	Y
Year FE	N	Y	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y	Y
Adj. R^2	0.316	0.317	0.438	0.448	0.449
Observations	15,940,360	15,940,360	15,940,357	15,940,357	15,940,357

Table A.17: Evidence of Spillover Effects within HRRs

This table shows the heterogeneous spillover effects of PE intervention in the local market. Column (1) examines a subsample of non-PE-backed hospitals which share common insurers with the PE-backed one in an HRR. Column (2) examines a subsample of non-PE-backed hospitals which do not share any insurer with the PE-backed one. The dependent variable is the natural logarithm of the total paid amounts. The independent variable, PE_{jt}, is an indicator for whether the HRR where hospital j is located has any PE-backed hospitals in year t. All columns contain hospital×payer and diagnosis×year fixed effects. Patient controls include *gender*, *age group*, *insurance type*, and *relative service-mix weights*. Hospital controls include *hospital bed numbers*, *teaching status*, *rural status*, *for-profit status*, *ratio of Medicare patients*, and *ratio of Medicaid patients*. Standard errors are clustered at the hospital level. t-statistics are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Non-PE-backed Hospitals Share an Insurer with PE-backed	Non-PE-backed Hospitals Not Share Any Insurer with PE-backed
PE	0.081**	−0.033
	(2.338)	(−0.923)
Hospital Controls	Y	Y
Patient Controls	Y	Y
Diagnosis×Year FE	Y	Y
Hospital×Payer FE	Y	Y
Adj. R^2	0.283	0.352
Observations	22,608,229	21,195,762

Table A.18: Compare with Merger and Acquisition (M&A) Deals

This table compares the effect of PE intervention with that of M&As. The sample includes insurance claims of hospitals that ever experienced M&As or received PE investments between 2013 and 2019. The dependent variable is the natural logarithm of the total paid amounts. The independent variable, $M\&A_{jt}$, is an indicator for whether hospital j experienced any M&As by year t. $M\&A \times PE_{jt}$ equals one if hospital j is under PE ownership in year t. All columns contain hospital×payer fixed effects. Columns (2) to (4) include diagnosis×year fixed effects. Patient controls, including *gender*, *age group*, *insurance type*, and *relative service-mix weights*, are added in columns (3) and (4). Hospital controls, including *hospital bed numbers*, *teaching status*, *rural status*, *for-profit status*, *ratio of Medicare patients*, and *ratio of Medicaid patients*, are added in column (4). Standard errors are clustered at the hospital level. t-statistics are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:			
	Logarithm of total paid amount			
	(1)	(2)	(3)	(4)
M&A × PE	0.295**	0.314**	0.320**	0.330**
	(2.002)	(2.251)	(2.352)	(2.423)
M&A	0.030*	-0.036	-0.038	-0.036
	(1.680)	(-1.581)	(-1.633)	(-1.577)
Hospital Controls	N	N	N	Y
Patient Controls	N	N	Y	Y
Diagnosis×Year FE	N	Y	Y	Y
Hospital×Payer FE	Y	Y	Y	Y
Adj. R^2	0.232	0.293	0.318	0.318
Observations	30,357,358	30,357,338	30,357,338	30,357,338

Table A.19: Examine the First Stage of CPOM Regulation Index

This table examines the first-stage correlation between PE investment decisions and the CPOM regulation index. The sample includes U.S. hospitals in the AHA's Annual Survey data between 2006 and 2019. The dependent variable is an indicator of whether a hospital is under PE ownership in a year. The independent variable, *CPOM Regulation Index*, measures how lenient states are about the corporate practice of medicine regulations. The detailed procedure of the index construction is provided in the Online Appendix. All columns contain hospital fixed effects. Columns (2) to (3) include year fixed effects. Control variables, including *hospital bed numbers*, *teaching status*, *rural status*, *for-profit status*, *ratio of Medicare patients*, and *ratio of Medicaid patients*, are added in the last column. The 1st stage *F-stat* is the *Kleibergen-Paap Wald F statistic*, whose p-value is indicated in parentheses. Standard errors are clustered at the hospital level. t-statistics are in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable:		
	PE Indicator		
	(1)	(2)	(3)
CPOM Regulation Index	0.026***	0.011***	0.010***
	(8.514)	(3.488)	(3.220)
ln(# Hospital Beds)			-0.003
			(-0.618)
Teaching			-0.025*
			(-1.694)
Rural			-0.002
			(-0.262)
For-profit			0.117***
			(7.364)
Medicare Patient Ratio			-0.011
			(-1.406)
Medicaid Patient Ratio			-0.003
			(-0.317)
F-stat	72.489	12.164	10.368
	(0.000)	(0.001)	(0.001)
Year FE	N	Y	Y
Hospital FE	Y	Y	Y
Adj. R^2	0.635	0.642	0.646
Observations	83,684	83,684	83,673

Table A.20: Estimation of Multinomial Logit Model of Patient Choice

This table shows the estimates for the multinomial logit hospital choice model. Since the patient choice is estimated separately for each HRR, the table reports the visit-number-weighted coefficients and standard errors. *PE indicator* equals one if hospital j is under PE ownership in year t. Its interactions with patient gender, age, and relative service-mix weights are included. *Travel time* (minutes) is computed between the centroid of a patient's 3-digit zip code prefix and the location of a hospital under normal traffic conditions. The square term of *Travel Time* and its interactions with hospital and patient characteristics are included. *Visit Before* is an indicator of whether a patient has visited a specific hospital before. The Diagnosis×Hospital Services includes ten interactions: *Mental Illness×Psychiatric Care*; *Pregnancy×Obstetrics Services* and *×NICU*; *Injury Diagnosis×Level 1 Trauma Center*; *Nervous, Circulatory, and Muscle Diagnosis×Magnetic Resonance Imaging*; *Cardiac Diagnosis×Cath Lab*, *×Interventional Cardiology*, and *×Heart Surgery Services*; *Cancer×Oncology Services*; *Musculoskeletal Diagnosis×Arthritis Services*. Standard errors are in parentheses.

VARIABLE	Coeff.	Std. Error
PE Intervention		
PE Indicator	1.4527	(0.1349)
PE×Female	0.0276	(0.0601)
PE×Age/100	−0.1200	(0.1510)
PE×Weight	−0.0423	(0.0177)
Travel Time to Hospital		
Travel Time	−0.1057	(0.0021)
Travel Time Squared	3.7466×10^{-4}	(0.2066×10^{-4})
Travel Time Interactions		
×Beds/100	0.0011	(0.0003)
×Age/100	0.0064	(0.0007)
×For-profit	0.0021	(0.0886)
×Teaching	0.0162	(0.0010)
×Wgt/1000	0.0906	(0.0474)
×Female	−0.0008	(0.0003)
Past Use of this Hospital		
Visit Before	0.4770	(0.0292)
Hospital Characteristics		
Hospital Dummy	Yes	
Hospital Dummy×Weight	Yes	
Teaching×Weight	−0.0190	(0.0071)
Diagnoses×Hospital Services (largest coeffs)		
Pregnancy: Obstetrics Services	1.0878	(0.0646)
Mental: Psych. Services	0.7375	(0.0823)
Cancer: Oncology Services	0.4790	(0.0327)

Table A.21: Estimation of Bargaining Model Parameters

This table presents estimates for the bargaining model. Panel A shows the estimates related to hospitals' bargaining weights. Panel B shows the estimates related to hospitals' marginal costs. Panel C collects the estimates related to PE's impacts and other parameters in the model. Standard errors are in parentheses.

Panel A: Bargaining Weight Parameter

VARIABLE	Coeff.	Std. Error
Multi-hospital System	0.2908	(0.0028)
For-profit	0.1750	(0.0001)
Teaching Status	0.2775	(0.0002)
Physician Arrangement	0.1143	(0.0001)
Rural Hospital	−0.0938	(0.0005)
ln(#Hospital Beds)	−0.0852	(0.0001)
Market Share of Inpatient Days	1.4793	(0.0012)
# Insurer in HRR	−0.0172	(0.0000)
Constant	0.6421	(0.0028)

Panel B: Marginal Costs Parameter

For-profit	−0.0616	(0.0405)
Teaching Status	0.6333	(0.0278)
Rural Area	−0.3092	(0.0315)
ln(#Hospital Beds)	−0.2204	(0.0138)
Medicare Patient Ratio	0.1176	(0.1520)
Medicaid Patient Ratio	1.8653	(0.1825)
HCC Score	−0.6175	(0.0833)
HRR Medicare Avg. OP. Cost/1000	0.0426	(0.0306)
Census Region FEs	Yes	
Year FEs	Yes	

Panel C: Impacts of PE and Other Parameters

Insurer Preference		
Insurer Weight (α)	1,236.8354	(2.0389)
Social Objectives		
Non-pecuniary Motive (τ_{NP})	101.9772	(2.0051)
Bankruptcy Threat		
Debt Burden (θ)	0.0050	(0.0000)
Location of Logistic Dist. (μ)	−2.1088	(0.0025)
Scale of Logistic Dist. (ϱ)	70.4975	(0.1767)
Ex-ante Debt Raising Costs		
Linear Debt Costs (μ_1)	0.6792	(14.1961)
Quadratic Debt Costs (μ_2)	1.0190	(53.3584)
Impacts of PE Intervention		
Marginal Cost Change (g_c)	−0.0835	(0.0205)
Bargaining Power Change (g_b)	0.1878	(0.0002)

Table A.22: Implications for Patient Surplus in Counterfactual

This table shows how banning PE ownership affects patient surplus in the counterfactual. The first row represents patient surplus changes in dollar terms due to the alteration of service quality. The second row represents patient-surplus changes resulting from hospital expense savings. The third row represents the aggregate changes of patient surplus by adding the previous two rows. The last row computes the ratio of the aggregate patient-surplus changes in dollar terms to the total hospital spending documented in the sample. All dollar terms are adjusted to dollars in 2019 by GDP deflators.

Counterfactual: Banning PE Ownership	
ΔQuality (\$billion)	-0.022
ΔSpending (\$billion)	2.945
Summing up...	
= ΔSurplus (\$billion)	2.923
Equivalent to...	
% Total Spending	10.71%

BIBLIOGRAPHY

ABPI. 2019. Worldwide pharmaceutical company R&D expenditure by country. URL https://www.abpi.org.uk/facts-and-figures/science-and-innovation/worldwide-pharmaceutical-company-rd-expenditure-by-country/.

Abraham, S., and L. Sun. 2020. Estimating dynamic treatment effects in event studies with heterogeneous Treatment effects. *Journal of Econometrics* Forthcoming.

Adelino, M., K. Lewellen, and W. B. McCartney. 2020. Hospital financial health and clinical choices: Evidence from the financial crisis. *Management Science* Forthcoming.

Adelino, M., K. Lewellen, and A. Sundaram. 2015. Investment decisions of nonprofit firms: Evidence from hospitals. *Journal of Finance* 70:1583–1628.

Aghamolla, C., P. Karaca-Mandic, X. Li, and R. T. Thakor. 2021. Merchants of death: The effect of credit supply shocks on hospital outcomes. *Working Paper* .

Alexander, J. A., M. T. Halpern, and S. Y. Lee. 1996. The short-term effects of merger on hospital operations. *Health Services Research* 30:827–847.

Andrews, I., M. Gentzkow, and J. M. Shapiro. 2017. Measuring the sensitivity of parameter estimates to estimation moments. *Quarterly Journal of Economics* 132:1553–1592.

Angrist, J. D., and J.-S. Pischke. 2008. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Antón, M., F. Ederer, M. Giné, and M. C. Schmalz. 2021. Innovation: The bright side of common ownership? *Working paper* .

Appelbaum, E., and R. Batt. 2020. Private equity buyouts in healthcare: Who wins, who loses? *Working Paper* .

Arnoud, A., F. Guvenen, and T. Kleineberg. 2019. Benchmarking global optimizers. *Working Paper* .

Axelson, U., T. Jenkinson, P. Strömberg, and M. S. Weisbach. 2013. Borrow cheap, buy high? The determinants of leverage and pricing in buyouts. *Journal of Finance* 68:2223–2267.

Axelson, U., P. StrÖmberg, and M. S. Weisbach. 2009. Why are buyouts levered? The financial structure of private equity funds. *Journal of Finance* 64:1549–1582.

Azar, J., S. Raina, and M. C. Schmalz. 2022. Ultimate ownership and bank competition. *Financial Management* 51:227–269.

Azar, J., M. C. Schmalz, and I. Tecu. 2018. Anticompetitive effects of common ownership. *Journal of Finance* 73:1513–1565.

Barrette, E., Gautam Gowrisankaran, and R. Town. 2021. Countervailing market power and hospital competition. *Review of Economics and Statistics* Forthcoming.

Benmelech, E., J. Eberly, D. Papanikolaou, and J. Krieger. 2021. Private and Social Returns to R&D: Drug Development and Demographics. In *AEA Papers and Proceedings*, vol. 111, pp. 336–40.

Berg, T., M. Reisinger, and D. Streitz. 2020. Handling spillover effects in empirical research. *Journal of Financial Economics* Forthcoming.

Bernstein, S., J. Lerner, and F. Mezzanotti. 2019. Private equity and financial fragility during the crisis. *Review of Financial Studies* 32:1309–1373.

Bernstein, S., J. Lerner, M. Sorensen, and P. Strömberg. 2017. Private equity and industry performance. *Management Science* 63:1198–1213.

Bernstein, S., and A. Sheen. 2016. The operational consequences of private equity buyouts: Evidence from the restaurant industry. *Review of Financial Studies* 29:2387–2418.

Bircan, C., M. Biesinger, and A. Ljungqvist. 2020. Value creation in private equity. *Working Paper* .

Bloom, N., M. Schankerman, and J. Van Reenen. 2013. Identifying technology spillovers and product market rivalry. *Econometrica* 81:1347–1393.

Borusyak, K., X. Jaravel, and J. Spiess. 2021. Revisiting event study designs: Robust and efficient estimation. *Working Paper* .

Boucly, Q., D. Sraer, and D. Thesmar. 2011. Growth LBOs. *Journal of Financial Economics* 102:432–453.

Brown, D. T., C. E. Fee, and S. E. Thomas. 2009. Financial leverage and bargaining power with suppliers: Evidence from leveraged buyouts. *Journal of Corporate Finance* 15:196–211.

Brown, Z. Y. 2019. An empirical model of price transparency and markups in health care. *Working Paper* .

Browning, M., F. Bourguignon, P. A. Chiappori, and V. Lechene. 1994. Income and outcomes: A structural model of intrahousehold allocation. *Journal of Political Economy* 102:1067–1096.

Bruch, J., D. Zeltzer, and Z. Song. 2021. Characteristics of private equity-owned hospitals in 2018. *Annals of Internal Medicine* 174:277–279.

Bruch, J. D., S. Gondi, and Z. Song. 2020. Changes in hospital income, use, and quality associated with private equity acquisition. *JAMA Internal Medicine* 180:1428–1435.

Burbidge, J. B., L. Magee, and A. L. Robb. 1988. Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association* 83:123–127.

Byrne, D. P. 2015. Testing models of differentiated products markets: Consolidation in the cable TV industry. *International Economic Review* 56:805–850.

Byrski, D., F. Gaessler, and M. J. Higgins. 2021. Market size and research: Evidence from the pharmaceutical industry. *Working paper* .

Cain, M. D., S. B. McKeon, and S. D. Solomon. 2017. Do takeover laws matter? Evidence from five decades of hostile takeovers. *Journal of Financial Economics* 124:464–485.

Capps, C., D. Dranove, and M. Satterthwaite. 2003. Competition and market power in option demand markets. *RAND Journal of Economics* 34:737–763.

Capps, C. S., D. W. Carlton, and G. David. 2020. Antitrust treatment of nonprofits: Should hospitals receive special care? *Economic Inquiry* 58:1183–1199.

Cerullo, M., K. K. Yang, J. Roberts, R. C. McDevitt, and A. C. Offodile. 2021. Private equity acquisition and responsiveness to service-line profitability at short-term acute care hospitals. *Health Affairs* 40:1697–1705.

Chandra, A., and D. O. Staiger. 2007. Productivity spillovers in health care: Evidence from the treatment of heart attacks. *Journal of Political Economy* 115:103–140.

Chernew, M., Z. Cooper, E. L. Hallock, and F. Scott Morton. 2021. Physician agency, consumerism, and the consumption of lower-limb MRI scans. *Journal of Health Economics* 76:102427.

Cicala, S., E. M. Lieber, and V. Marone. 2019. Regulating markups in US health insurance. *American Economic Journal: Applied Economics* 11:71–104.

Cohn, J. B., L. F. Mills, and E. M. Towery. 2014. The evolution of capital structure and operating performance after leveraged buyouts: Evidence from U.S. corporate tax returns. *Journal of Financial Economics* 111:469–494.

Cohn, J. B., N. Nestoriak, and M. Wardlaw. 2020. Private equity buyouts and workplace safety. *Review of Financial Studies* Forthcoming.

Collard-Wexler, A., G. Gowrisankaran, and R. S. Lee. 2019. 'Nash-in-Nash' bargaining: A microfoundation for applied work. *Journal of Political Economy* 127:163–195.

Cooper, Z., S. V. Craig, M. Gaynor, and J. Van Reenen. 2019. The price ain't right? Hospital prices and health spending on the privately insured. *Quarterly Journal of Economics* 134:51–107.

Craig, S. V., M. Grennan, and A. Swanson. 2021. Mergers and marginal costs: New evidence on hospital buyer power. *RAND Journal of Economics* 52:151–178.

Crawford, G. S., and A. Yurukoglu. 2012. The welfare effects of bundling in multichannel television markets. *American Economic Review* 102:643–685.

Cuesta, J. I., C. Noton, and B. Vatter. 2019. Vertical Integration between Hospitals and Insurers. *Working Paper* .

Cunningham, C., F. Ederer, and S. Ma. 2021. Killer Acquisitions. *Journal of Political Economy* 129:649–702.

Cutler, D. M., R. S. Huckman, and J. T. Kolstad. 2010. Input constraints and the efficiency of entry: Lessons from cardiac surgery. *American Economic Journal: Economic Policy* 2:51–76.

Cutler, D. M., and L. Sheiner. 1999. The geography of Medicare. *American Economic Review* 89:228–233.

Dafny, L., K. Ho, and R. S. Lee. 2019. The price effects of cross-market mergers: Theory and evidence from the hospital industry. *RAND Journal of Economics* 50:286–325.

Davis, S. J., J. Haltiwanger, K. Handley, R. Jarmin, J. Lerner, and J. Miranda. 2014. Private equity, jobs, and productivity. *American Economic Review* 104:3956–3990.

Davis, S. J., J. C. Haltiwanger, K. Handley, B. Lipsius, J. Lerner, and J. Miranda. 2019. The economic effects of private equity buyouts. *Working Paper* .

Dennis, P. J., K. Gerardi, and C. Schenone. 2021. Common ownership does not have anti-competitive effects in the airline industry. *Forthcoming, Journal of Finance* .

DiMasi, J. A., and L. B. Faden. 2011. Competitiveness in follow-on drug R&D: a race or imitation? *Nature Reviews Drug Discovery* 10:23–27.

DiMasi, J. A., H. G. Grabowski, and R. W. Hansen. 2016. Innovation in the pharmaceutical industry: new estimates of R&D costs. *Journal of health economics* 47:20–33.

Dou, W., L. A. Taylor, W. Wang, and W. Wang. 2020. Dissecting bankruptcy frictions.

Journal of Financial Economics Forthcoming.

Duggan, M. G. 2000. Hospital ownership and public medical spending. *Quarterly Journal of Economics* 115:1343–1373.

Eaton, C., S. Howell, and C. Yannelis. 2020. When investor incentives and consumer interests diverge: Private equity in higher education 33:4024–4060.

Eldar, O., J. Grennan, and K. Waldock. 2020. Common ownership and startup growth. *Working paper* .

Eliason, P. J., B. Heebsh, R. C. McDevitt, and J. W. Roberts. 2020. How acquisitions affect firm behavior and performance: Evidence from the dialysis industry. *Quarterly Journal of Economics* 135:221–267.

Ericson, K. M., and A. Starc. 2015. Measuring consumer valuation of limited provider networks. *American Economic Review Papers and Proceedings* 105:115–119.

Ewens, M., and M. Marx. 2018. Founder replacement and startup performance. *Review of Financial Studies* 31:1532–1565.

Ewens, M., R. Nanda, and M. Rhodes-Kropf. 2018. Cost of experimentation and the evolution of venture capital. *Journal of Financial Economics* 128:422–442.

Finkelstein, A., M. Gentzkow, and H. Williams. 2016. Sources of geographic variation in health care: Evidence from patient migration. *Quarterly Journal of Economics* 131:1681–1726.

Fracassi, C., A. Previtero, and A. Sheen. 2021. Barbarians at the store? Private equity, products, and consumers. *Journal of Finance* Forthcoming.

Fulghieri, P., and M. Sevilir. 2009. Size and focus of a venture capitalist's portfolio. *Review of Financial Studies* 22:4643–4680.

Gandhi, A., Y. Song, and P. Upadrashta. 2020. Private equity, consumers, and competition: Evidence from the nursing home industry. *Working Paper* .

Gao, J., M. Sevilir, and Y. S. Kim. 2021. Private equity in the hospital industry. *Working Paper* .

Gaynor, M., K. Ho, and R. J. Town. 2015. The industrial organization of health-care markets. *Journal of Economic Literature* 53:235–284.

Gaynor, M., and R. J. Town. 2011. Competition in health care markets. In *Handbook of Health Economics*, vol. 2, pp. 499–637. Elsevier.

Gaynor, M., and W. B. Vogt. 2003. Competition among hospitals. *RAND Journal of Economics* 34:764–785.

Ghili, S. 2016. Network formation and bargaining in vertical markets: The case of narrow networks in health insurance. *Working Paper* .

Gilje, E. P., T. A. Gormley, and D. Levit. 2020. Who's paying attention? Measuring common ownership and its impact on managerial incentives. *Journal of Financial Economics* 137:152–178.

Gompers, P. A., W. Gornall, S. N. Kaplan, and I. A. Strebulaev. 2020. How do venture capitalists make decisions? *Journal of Financial Economics* 135:169–190.

Gompers, P. A., and J. Lerner. 2004. *The venture capital cycle.* MIT press.

Gondi, S., and Z. Song. 2019. Potential implications of private equity investments in health care delivery. *Journal of the American Medical Association* 321:1047–1048.

González-Uribe, J. 2020. Exchanges of innovation resources inside venture capital portfolios. *Journal of Financial Economics* 135:144–168.

Gowrisankaran, G., A. Nevo, and R. Town. 2015. Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review* 105:172–203.

Gramlich, J., and S. Grundl. 2017. Testing for competitive effects of common ownership. *Working Paper* .

Grennan, M. 2013. Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review* 103:145–177.

Grennan, M., and A. Swanson. 2021. Bargaining in healthcare markets: Applications of Nash-in-Nash and extensions. In *Bargaining: Current Research and Future Directions.* Palgrave MacMillan.

Grossman, G. M., and C. Shapiro. 1987. Dynamic R&D competition. *The Economic Journal* 97:372–387.

Gulliver, J., and W. Jiang. 2020. The impact of private equity buyouts on productivity and jobs. *Working Paper* .

Gupta, A., S. Howell, C. Yannelis, and A. Gupta. 2021. Does private equity investment in healthcare benefit patients? Evidence from nursing homes. *Working Paper* .

Gutiérrez, G., and T. Philippon. 2017. Investmentless growth: An empirical examination. *Brookings Papers on Economic Activity* pp. 89–169.

Hackmann, M. B. 2019. Incentivizing better quality of care: The role of Medicaid and competition in the nursing home industry. *American Economic Review* 109:1684–1716.

Hall, B. H., J. Mairesse, and P. Mohnen. 2010. Measuring the Returns to R&D. In *Handbook of the Economics of Innovation*, vol. 2, pp. 1033–1082. Elsevier.

Hansmann, H. B. 1980. The role of nonprofit enterprise. *The Yale Law Journal* 89:835.

Harrington, S. E. 2012. Cost of capital for pharmaceutical, biotechnology, and medical device firms. *The Oxford handbook of the economics of the biopharmaceuticals industry* pp. 75–99.

He, J. J., and J. Huang. 2017. Product market competition in a world of cross-ownership: Evidence from institutional blockholdings. *Review of Financial Studies* 30:2674–2718.

Hege, U., and C. Hennessy. 2010. Acquisition values and optimal financial (in)flexibility. *Review of Financial Studies* 23:2865–2899.

Hennessy, C. A., and D. Livdan. 2009. Debt, bargaining, and credibility in firm-supplier relationships. *Journal of Financial Economics* 93:382–399.

Ho, K., and R. S. Lee. 2017. Insurer competition in health care markets. *Econometrica* 85:379–417.

Ho, K., and R. S. Lee. 2019. Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review* 109:473–522.

Ho, V. 2020. Revisiting states' experience with Certificate of Need. *Journal of the American Medical Association* 324:2033–2035.

Hombert, J., and A. Matray. 2017. The real effects of lending relationships on innovative firms and inventor mobility. *Review of Financial Studies* 30:2413–2445.

Hombert, J., and A. Matray. 2018. Can innovation help U.S. manufacturing firms escape import competition from China? *The Journal of Finance* 73:2003–2039.

Horn, H., and A. Wolinsky. 1988. Bilateral monopolies and incentives for merger. *RAND Journal of Economics* 19:408–419.

Huang, S. S., and J. R. Bowblis. 2019. Private equity ownership and nursing home quality: An instrumental variables approach. *International Journal of Health Economics and Management* 19:273–299.

Inderst, R., H. M. Mueller, and F. Münnich. 2007. Financing a portfolio of projects. *Review of Financial Studies* 20:1289–1325.

Ivanov, I., L. Pettit, and T. M. Whited. 2020. Taxes depress corporate borrowing: Evidence from private firms. *Working Paper* .

Ivy, S. P., L. L. Siu, E. Garrett-Mayer, and L. Rubinstein. 2010. Approaches to phase 1 clinical trial design focused on safety, efficiency, and selected patient populations: A report from the clinical trial design task force of the national cancer institute investigational drug steering committee. *Clinical Cancer Research* 16:1726–1736.

Jeffers, J. 2021. The impact of restricting labor mobility on corporate investment and entrepreneurship. *Working Paper* .

Jiang, W. 2017. Have instrumental variables brought us closer to the truth? *Review of Corporate Finance Studies* 6:127–140.

Johnston Ross, E., M. Puri, and S. Ma. 2021. Private Equity and Financial Stability: Evidence from Failed Bank Resolution in the Crisis. *Working Paper* .

Kale, J. R., E. Reis, and A. Venkateswaran. 2009. Rank-order tournaments and incentive alignment: The effect on firm performance. *Journal of Finance* 64:1479–1512.

Kaplan, S. 1989. The effects of management buyouts on operating performance and value. *Journal of Financial Economics* 24:217–254.

Kaplan, S. N., and J. Lerner. 2010. It ain't broke: The past, present, and future of venture capital. *Journal of Applied Corporate Finance* 22:36–47.

Kaplan, S. N., and P. Strömberg. 2009. Leveraged buyouts and private equity. *Journal of Economic Perspectives* 23:121–146.

Karpoff, J. M., and M. D. Wittry. 2018. Institutional and legal context in natural experiments: The case of state antitakeover laws. *Journal of Finance* 73:657–714.

Kennedy, P., D. P. O'Brien, M. Song, and K. Waeher. 2017. The competitive effects of common ownership: Economic foundations and empirical evidence. *Working paper* .

Kini, O., S. Lee, and M. Shen. 2022. Common institutional ownership and product market threats. *Working paper* .

Koch, A., M. A. Panayides, and S. Thomas. 2020. Common Ownership and Competition in Product Markets. *Journal of Financial Economics* Forthcoming.

Kornai, J. 2009. The soft budget constraint syndrome in the hospital sector. *Society and Economy* 31:5–31.

Kornai, J., E. Maskin, and G. Roland. 2003. Understanding the soft budget constraint.

Journal of Economic Literature 41:1095–1136.

Kostovetsky, L., and A. Manconi. 2020. Common institutional ownership and diffusion of innovation. *Working paper* .

Krieger, J. L. 2021. Trials and Terminations: Learning from Competitors' R&D Failures. *Management Science* 67:5525–5548.

Krieger, J. L., D. Li, and D. Papanikolaou. 2022. Missing novelty in drug development. *Review of Financial Studies* 35:636–679.

Lakdawalla, D., and T. Philipson. 2006. The nonprofit sector and industry performance. *Journal of Public Economics* 90:1681–1698.

Lerner, J. 1995. Venture capitalists and the oversight of private firms. *Journal of Finance* 50:301–318.

Lerner, J., M. Sorensen, and P. Strömberg. 2011. Private equity and long-run investment: The case of innovation. *Journal of Finance* 66:445–477.

Lewellen, J., and K. Lewellen. 2021. Institutional investors and corporate governance: The incentive to be engaged. *Journal of Finance* 77:213–264.

Lewellen, K., and M. Lowry. 2021. Does common ownership really increase firm coordination? *Journal of Financial Economics* 141:322–344.

Lewis, M. S., and K. E. Pflum. 2015. Diagnosing hospital system bargaining power in managed care networks. *American Economic Journal: Economic Policy* 7:243–274.

Lichtenberg, F. R., and D. Siegel. 1990. The effects of leveraged buyouts on productivity and related aspects of firm behavior. *Journal of Financial Economics* 27:165–194.

Lichtenberg, F. R., D. Siegel, D. Jorgenson, and E. Mansfield. 1987. Productivity and changes in ownership of manufacturing plants. *Brookings Papers on Economic Activity* 1987:643–683.

Liebman, E. 2016. Bargaining in markets with exclusion: An analysis of health insurance networks. *Working Paper* .

Lindsey, L. 2008. Blurring firm boundaries: The role of venture capital in strategic alliances. *Journal of Finance* 63:1137–1168.

López, A. L., and X. Vives. 2019. Overlapping Ownership, R&D Spillovers, and Antitrust Policy. *Journal of Political Economy* 127:2394–2437.

Lopez, E., T. Neuman, G. Jacobson, and L. Levitt. 2020. How much more than Medicare do private insurers pay? A review of the literature. *Kaiser Family Foundation: Medicare* .

Loury, G. C. 1979. Market structure and innovation. *The Quarterly Journal of Economics* 93:395–410.

Matray, A. 2021. The local innovation spillovers of listed firms. *Journal of Financial Economics* 141:395–412.

Matsa, D. A. 2010. Capital structure as a strategic variable: Evidence from collective bargaining. *Journal of Finance* 65:1197–1232.

Michal, M. H., M. S. L. Pekarske, and M. K. Mcmanus. 2006. Corporate practice of medicine doctrine 50 state survey summary. *Center to Advance Palliative Care and National Hospice and Palliative Care Organization* pp. 1–19.

Myers, S. C. 2001. Capital structure. *Journal of Economic Perspectives* 15:81–102.

Offodile, A. C., M. Cerullo, M. Bindal, J. A. Rauh-Hain, and V. Ho. 2021. Private equity investments in health care: An overview of hospital and health system leveraged buyouts, 2003-17. *Health Affairs* 40:719–726.

Pakes, A., J. Porter, K. Ho, and J. Ishii. 2015. Moment inequalities and their application. *Econometrica* 83:315–334.

Petersen, M. A., and R. G. Rajan. 1994. The Benefits of Lending Relationships: Evidence from Small Business Data. *Journal of Finance* 49:3–37.

Pradhan, R., R. Weech-Maldonado, J. S. Harman, M. Al-Amin, and K. Hyer. 2014. Private equity ownership of nursing homes: Implications for quality. *Journal of Health Care Finance* 42:1–14.

Prager, E., and N. Tilipman. 2020. Disagreement payoffs and negotiated prices: Evidence from out-of-network hospital payments. *Working Paper* .

Rajan, R. G. 1992. Insiders and Outsiders: The Choice between Informed and Arm's-Length Debt. *Journal of Finance* 47:1367–1400.

Raskovich, A. 2003. Pivotal buyers and bargaining position. *Journal of Industrial Economics* 51:405–426.

Raval, D., T. Rosenbaum, and N. E. Wilson. 2020. Using disaster induced closures to evaluate discrete choice models of hospital demand. *Working Paper* .

Regnier, S., and D. B. Ridley. 2015. Forecasting market share in the US pharmaceutical market. *nature reviews drug discovery* 14:594–595.

Rice, T., and P. E. Strahan. 2010. Does credit competition affect small-firm finance? *Journal of Finance* 65:861–889.

Roberts, T. G., B. H. Goulart, L. Squitieri, S. C. Stallings, E. F. Halpern, B. A. Chabner, G. S. Gazelle, S. N. Finkelstein, and J. W. Clark. 2004. Trends in the risks and benefits to patients with cancer participating in phase 1 clinical trials. *Journal of the American Medical Association* 292:2130–2140.

Robles-Garcia, C. 2020. Competition and incentives in mortgage markets: The role of brokers. *Working Paper* .

Saretto, A., and H. E. Tookes. 2013. Corporate leverage, debt maturity, and credit supply: The role of credit default swaps. *Review of Financial Studies* 26:1190–1247.

Schmitt, M. 2017. Do hospital mergers reduce costs? *Journal of Health Economics* 52:74–94.

Seyhan, A. A. 2019. Lost in translation: The valley of death across preclinical and clinical divide– ididentification of problems and overcoming obstacles. *Tranlational Medicine Communications* 4:1–19.

Shepard, M. 2016. Hospital network competition and adverse selection: Evidence from the Massachusetts health insurance exchange. *Working Paper* .

Shleifer, A., and L. H. Summers. 1988. *Breach of trust in hostile takeovers*, vol. I. Chicago: University of Chicago Press.

SØrensen, M. 2007. How smart is smart money? A two-sided matching model of venture capital. *Journal of Finance* 62:2725–2762.

Spang, H. R., G. J. Bazzoli, and R. J. Arnould. 2001. Hospital mergers and savings for consumers: Exploring new evidence. *Health Affairs* 20:150–158.

Stevenson, D. G., and D. C. Grabowski. 2008. Private equity investment and nursing home care: Is it a big deal? *Health Affairs* 27:1399–1408.

Stole, L. A., and J. Zwiebel. 1996. Intra-firm bargaining under non-binding contracts. *Review of Economic Studies* 63:375–410.

Strömberg, P. 2008. The new demography of private equity. *The Global Impact of Private Equity* 1:3–26.

Thomas, D. W., J. Burns, J. Audette, A. Carroll, C. Dow-Hygelund, and M. Hay. 2016. Clinical development success rates 2006–2015. *BIO Industry Analysis* https://www.bio.org.

Town, R., and G. Vistnes. 2001. Hospital competition in HMO networks. *Journal of Health Economics* 20:733–753.

Train, K. E. 2009. *Discrete choice methods with simulation.* Cambridge university press.

Weber, S., and M. Péclat. 2020. GEOROUTE: Stata module to calculate travel distance and travel time between two addresses or two geographical points. *Working Paper* .

Wessner, C. W. 2009. *Venture funding and the NIH SBIR program.* National Academies Press.

Wollmann, T. G. 2020. How to get away with merger: Stealth consolidation and its real effects on US healthcare. *Working Paper* .

Wong, C. H., K. W. Siah, and A. W. Lo. 2019. Estimation of clinical trial success rates and related parameters. *Biostatistics* 20:273–286.

Wooldridge, J. M. 2010. *Econometric analysis of cross section and panel data.* MIT press.