Quality, Selection, and Congestion: Equilibrium Hospital Outcomes During Covid-19*

Gautam Gowrisankaran
Columbia University, CEPR, NBER

Kelli Marquardt
Federal Reserve Bank of Chicago

Robert Town
University of Texas at Austin, NBER

July 6, 2026

Abstract

We examine the importance of hospital quality, patient selection, and congestion in affecting equilibrium health outcomes, using Covid-19 pandemic data. Our main outcome is mortality, controlling for unobserved selection with distance instruments. Intensive care unit strain increased hospital mortality: without congestion, deaths would fall 9.6%. Patients value hospitals with high baseline quality, but their choices did not reflect capacity strain information. Equilibrium counterfactuals show that if patients were informed about capacity strain and believed quality signals were more precise, deaths would drop 3.8%. However, using more precise information but not accounting for congestion would cause Covid-19 deaths to *rise* 3.8%.

JEL Codes: L15, I11, D80
Keywords: Mortality, capacity strain, pandemic, intensive care units, learning-by-doing, information, Bayesian estimation

*We thank Cory Capps, Caitlin Carroll, Peter Hull, Keith Joiner, Ashley Langer, Pauline Mourot, Cristóbal Otero, Sebastián Otero, Pietro Tebaldi, and seminar/conference participants at numerous institutions for valuable comments and suggestions. We also acknowledge excellent research assistance from Anne Fournier, Taegan Mullane, and Nicole Theinová. The views here are solely those of the authors and do not represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System. All errors are our own.

1 Introduction

Outcomes at many institutions—e.g., hospitals, schools, and courts—reflect not only the *time-invariant quality* of the institution itself, but also the *composition* of individuals who select into it and the *congestion* imposed by the set of other users simultaneously demanding its resources. A central challenge in empirical economics is measuring and understanding the causal impact of these separate factors on outcomes. A common approach to addressing this challenge disentangles quality and selection, sometimes using an instrument to control for sorting. But congestion introduces a third, conceptually distinct channel that is empirically confounded with the other two. Disentangling all three channels requires separate identification of selection and congestion, alongside a statistical framework that is robust to the low signal-to-noise ratios that often characterize outcomes data.

This paper proposes a framework for identifying these three channels and applies it to an important setting: hospital performance during the Covid-19 pandemic. Our framework involves finding shocks to overall demand that are unrelated to patient sorting. We control for patient selection with distance-based instruments and use pandemic demand surges to identify congestion separately from selection. We implement a Bayesian estimator that leverages both sources of variation and provides appropriate inference even with few outcome observations per institution-period. This allows us to separately recover hospital baseline time-invariant quality from selection and congestion effects. We then use hospital choice decisions to understand patient information regarding baseline hospital quality and congestion. Finally, we examine the welfare consequences of counterfactual information environments and hospital capacity investments, accounting for equilibrium spillovers from congestion.

Although similar empirical challenges are prevalent in many settings, the hospital sector provides a particularly compelling application. The sector is large, exhibits substantial heterogeneity in performance, and has undergone major structural transformations that could affect performance. Specifically, hospitals account for over $1.5 trillion in annual spending in the United States—approximately 5.5% of GDP—while a large and growing empirical literature documents substantial dispersion in hospital performance.[1] This implies that even

[1] See, e.g., Geweke et al. (2003), Doyle et al. (2019), Chandra and Staiger (2020).

a modest reallocation of patients towards higher-performing providers could generate sizable benefits. Finally, congestion is potentially important: over the past two decades, there have been 195 hospital closures, 612 mergers, and a 16% reduction in bed capacity in the U.S. alone,[2] while in many other countries, hospitals operate near capacity even under baseline conditions (Hoe, 2022, Avdic et al., 2024).

Capacity strain—defined as when utilization approaches or exceeds available resources necessitating costly reallocation of inputs away from their planned uses—is the primary congestion mechanism we study. Such strain commonly arises during demand surges associated with natural disasters or epidemics, when patients present with life-threatening illness and timely access to quality care is critical. Capacity strain is similar to peer effects in many ways, in that a patient's hospital choice can spill over and affect the outcomes of other patients when a hospital is capacity strained. Unlike peer effects, capacity strain may be highly nonlinear, with quality degrading convexly as utilization approaches or exceeds recommended capacity. A value of our paper is in identifying the potentially nonlinear relationship between congestion and quality.

Capacity strain is a particularly problematic source of underperformance because it may be difficult for patients to assess when choosing hospitals. Information is critical here since patients' beliefs about hospital quality may be based on hospital ratings and reputation[3]—both of which evolve over the long run—while capacity strain changes in the short run.

We focus on Covid-19 health outcomes from 2020–2022. During this period, the U.S. experienced dramatic increases in hospital demand and high-stakes mortality outcomes, with 104 million confirmed Covid-19 cases (Clarke, 2022), 4.9 million inpatient hospitalizations,[4] and over 1.1 million deaths (Silk et al., 2023), about two-thirds of which occurred in inpatient settings (Ahmad, 2023).The pandemic provides a useful laboratory to study the impact of quality, selection, and congestion at scale because the substantial number of deaths and demand fluctuations generate substantial and plausibly exogenous variation in capacity strain

[2]See Brand et al. (2023), Leuchter et al. (2025), and `https://www.shepscenter.unc.edu/programs-projects/rural-health/rural-hospital-closures/`.

[3]See, e.g., Luft et al. (1990), Pope (2009), Chandra et al. (2016).

[4]Hospitalization estimates are based on authors' calculations using estimates of the population hospitalization rates and case rates. See `https://data.cdc.gov/Case-Surveillance/Weekly-United-States-COVID-19-Cases-and-Deaths-by-/pwn4-m3yp/data_preview` and `https://www.cdc.gov/covid/php/covid-net/index.html`.

across hospitals and over time.

Moreover, the *equilibrium* impact of many counterfactual policies is theoretically ambiguous, implying that empirical evidence is valuable. For instance, a potential policy would give patients more accurate capacity information. This might then cause patients to choose different hospitals that were ex ante better. At an individual level, this would clearly be beneficial. However, if more patients choose hospitals with slack as a result of the more accurate information, this could cause those hospitals to become congested, resulting in relatively worse outcomes at those hospitals.

Our analysis uses Texas discharge data over the period Q2:2020 to Q3:2022, with our main observed outcome being death in the hospital. We estimate time-varying hospital quality for Covid-19 care and investigate how congestion and other factors such as hospital learning-by-doing may affect this measure. We estimate a patient hospital choice model to understand patient information regarding hospital quality and capacity strain. Finally, we evaluate the equilibrium impact of counterfactual patient information and hospital characteristics.

As noted above, our estimation of hospital performance faces three interrelated challenges. The standard challenge is that illness severity may affect the choice of hospital. For example, patients who are more severely ill may be willing to travel further to reach better-quality hospitals. If illness severity is not perfectly observed, this may bias estimated quality, making high-performance hospitals look relatively bad because they treat more severely ill patients. This interacts with our second empirical challenge: we need to distinguish baseline quality from congestion. We recover the impact of capacity strain by examining whether the same hospital performs worse when capacity is strained than when there is slack. Finally, selection and congestion interact with the fact that most hospital-quarters have relatively few Covid-19 deaths, making classical statistical inference challenging: hospitals that randomly had few deaths tend to be reported as high quality.

We address these empirical challenges with a Bayesian estimation procedure that controls for unobserved patient selection, using distance to a hospital as an instrument for its choice, in a model with a discrete choice of hospital and binary outcome of mortality (Geweke et al., 2003, henceforth, GGT). We control for time-varying aspects of hospital quality with a two-step process: we measure hospital quality with GGT at the quarterly level and then estimate

the impact of capacity strain on this time-varying quality measure, controlling for hospital fixed effects. As with empirical Bayes shrinkage approaches, our hospital-quarter estimates move hospitals with extreme outcomes towards the overall mean, enabling appropriate statistical inference despite the fact that our quarterly quality estimates are most often imprecise because there are relatively few patients per hospital-quarter and many fewer deaths.

Our empirical analysis allows us to identify quality and disentangle it from selection and congestion. Specifically, the instruments allow for hospitals with high risk-adjusted mortality to have disproportionately treated patients with high severity of illness, identifying outcomes from patients who live near the hospital, rather than who are treated by the hospital. Distance is widely used as an instrument for healthcare provider choice (e.g., McClellan et al., 1994, Geweke et al., 2003, Cutler, 2007, Card et al., 2023, Einav et al., 2025), but to our knowledge, not in combination with congestion. The pandemic period saw massive surges in Covid-19 cases and hospitalizations, which created variation in congestion across hospitals. This provides exogenous variation in capacity strain at the hospital level to the extent that individuals did not have information regarding capacity strain at a particular hospital in a particular quarter (which we find to be true in our choice modeling). Even if patients did have information regarding capacity strain, our instruments will control for this source of endogeneity when we estimate hospital quality at the quarter level.

Focusing first on average hospital quality over time, we find large variation in hospital performance that translates into large population-level mortality impacts. Moving all patients admitted from an above-median mortality hospital to the median hospital (and assuming no change in quality) would reduce mortality by 14.9%. Applying this to the U.S. would imply a reduction of $109,000$ deaths during the pandemic, or $390 billion using Covid-19 value of life estimates (Silva et al., 2023).

To evaluate the impact of congestion on hospital quality, we regress the posterior hospital-quarter quality draws on hospital fixed effects, region-by-quarter fixed effects, and time-varying factors such as capacity utilization and previous Covid-19 patient volume.[5] The most significant time-varying predictor of hospital performance is intensive care unit (ICU)

[5]If one were concerned that patients were selecting hospitals based on their contemporaneous change in quality, we could also instrument for capacity utilization with volume predicted from a choice model in a baseline period (as in, e.g., Gowrisankaran et al., 2006).

capacity strain: hospital mortality increases with ICU capacity utilization, with a statistically significant impact when utilization is over 75% of the reported number of ICU beds in spline specifications. Despite the large literature showing positive volume-outcome relationships (e.g., Luft et al., 1987, Birkmeyer et al., 2002), we find no evidence of learning-by-doing: hospitals that had more previous Covid-19 patient volume did not have lower mortality.

We then estimate multinomial logit patient hospital choice models to understand the extent to which patients were informed about hospital quality. We find that baseline hospital mortality (i.e., not accounting for capacity strain) is a significant negative predictor of patient hospital choice, even after controlling for distance to the hospital, hospital quality for similar diseases to Covid-19, and hospital ownership structure. The magnitude is consistent with patients observing a relatively weak signal of true quality and rationally reacting to that signal. We also estimate the impact of contemporaneous hospital mortality—which includes the impact of capacity strain—with a leave-one-out estimator. Here, we find that contemporaneous mortality is a *positive* predictor of patient hospital choice in many specifications. These results are consistent with patients not having information about congestion, them seeking care at hospitals that would have been high quality in the absence of congestion, and hospitals ending up worse than alternatives in equilibrium because of overcrowding.

Finally, we perform equilibrium counterfactuals that examine the impact of information and market structure on mortality and utility, accounting for equilibrium spillovers through capacity strain. We investigate providing patients with information regarding actual hospital quality, assuming their true preferences are given by the estimated hospital choice model that includes baseline hospital quality as a regressor. As above, two countervailing forces determine the overall effect. First, better information induces patients to select higher-quality hospitals, which we find reduces mortality by 0.16%. Second, reallocating patients increases capacity strain at the newly selected hospitals and reduces it at those previously selected. The net externality raises mortality by 0.28%, resulting in equilibrium mortality increasing slightly despite improved information and individually optimal re-sorting.

Given our finding that patients may view their hospital quality signals as relatively weak, we investigate how equilibrium outcomes would vary if patients believed they had more precise information. We find that if patients were both informed about capacity strain

and believed the quality signals to be 10 times as precise as in our estimation, equilibrium mortality would decrease 3.8%, with mean travel time increasing by less than one minute. Decomposing this result, the direct choice effect reduces mortality by 8.5% but the externality effect from capacity strain increases mortality by 4.8%. These net gains from increasing the quality signal weight continue as perceived accuracy increases, asymptoting to a mortality decrease of 6.7%. However, if patients instead continued to be uninformed about capacity but still believed quality to be 10 times as precise as in our estimation, equilibrium mortality would *increase* 3.8%, as patients crowded even more into hospitals already suffering from ICU capacity strain.

Related Literature: This paper contributes to three broad literatures. First, it relates to a longstanding literature that estimates institutional performance. Many papers in this literature proceed by separating quality from endogenous user selection in settings including schools, judicial outcomes, and health plans (see, e.g. Hoxby, 2000, Dobbie et al., 2018, Geruso et al., 2023, respectively). Other papers separate institutional performance from congestion in courts, broadband, and ride-hailing (see Coviello et al., 2010, Malone et al., 2021, Castillo et al., 2025, respectively). Allende et al. (2026) separates quality from both selection and congestion, though they model congestion with fixed capacity constraints that cause assignment to different institutions, rather than as directly affecting outcomes. We build on this literature by separating baseline quality from both selection and congestion.

Second, our paper relates to a literature that examines how market forces, information, and experience shape healthcare outcomes. Research has shown that patients respond imperfectly—and in some cases not at all—to available quality signals (Chernew et al., 2008, Gaynor et al., 2016, Abaluck et al., 2021, Brown and Jeon, 2024, Vatter, 2025). Complementing these findings, other works study equilibrium patient sorting in response to quality information (Dranove et al., 2003, Chandra et al., 2016). Several papers have focused on learning-by-doing, finding large positive effects of increased *surgical* volume on outcomes (e.g. Birkmeyer et al., 2002, Gaynor et al., 2005). In contrast, studies of capacity strain find that congestion can worsen patient outcomes (Kuntz et al., 2015, Eriksson et al., 2017, Kadri et al., 2021, Hoe, 2022, Avdic et al., 2024), particularly for minority patients (Singh and Venkataramani, 2026). Finally, a related literature documents substantial variation in

Covid-19 mortality across hospitals and geographic areas (Allcott et al., 2024, Alsan et al., 2021, Asch et al., 2021, Kunz and Propper, 2023). Our paper bridges these strands of the literature by jointly analyzing the effects of congestion and patient information on hospital quality within an equilibrium framework.

Third, our paper relates to the literature that uses inpatient discharge data to estimate the distribution of hospital quality, its determinants, and its evolution over time. This literature dates back at least to Luft and Hunt (1986). A central challenge in this work is that hospital performance is difficult to estimate accurately because unobserved differences in patient severity and the inherently low signal-to-noise ratio of mortality outcomes can bias conventional estimates of facility quality. Recent research has developed methods to address these identification challenges, improving the estimation of provider performance in administrative data (see McClellan et al., 1994, Gowrisankaran and Town, 1999, Geweke et al., 2003, Doyle et al., 2015, Hull, 2018, Olenski and Sacher, 2024, Einav et al., 2025). We build on this work by providing a tractable Bayesian approach for large panels with time-varying quality, implemented in standard software (Marquardt et al., 2025), and by jointly identifying congestion alongside baseline quality and selection.

The remainder of this paper has the following structure. Section 2 provides background on Covid-19 and hospital care and describes our data. Section 3 proposes our empirical framework. Section 4 presents our estimation results. Section 5 discusses counterfactuals and Section 6 concludes.

2 Background and Data

2.1 Covid-19 and Inpatient Care

The first case of Covid-19 in the US was diagnosed on Jan. 20, 2020. Between 2020 and 2022, there were over 103 million reported cases and 1.1 million deaths, and 1.2% of the US population was hospitalized. Our analysis focuses on Texas, which experienced approximately 500,000 inpatient admissions and 92,000 deaths over this period.[6]

[6]Authors' calculations from `https://dshs.texas.gov/covid-19-coronavirus-disease-2019/texas-covid-19-data`.

Typically, Covid-19 is not life-threatening and does not require hospitalization. However, in some cases, the disease can progress from a mild illness to a severe disease necessitating hospitalization. The disease progression is driven by several biological and clinical mechanisms. One primary factor is a dysregulated immune response, where an excessive inflammatory reaction—often referred to as a "cytokine storm"—leads to widespread tissue damage and organ dysfunction (Mehta et al., 2020). Patients with pre-existing conditions such as diabetes, hypertension, and cardiovascular disease are at heightened risk due to underlying endothelial dysfunction and chronic inflammation, which exacerbate the severity of infection (Richardson et al., 2020). Additionally, viral load and a host of genetic factors play a role in determining disease severity, with some individuals exhibiting more aggressive viral replication and immune evasion that contribute to rapid respiratory deterioration (Gottlieb et al., 2022).

Respiratory failure is a predominant cause of hospitalization among Covid-19 patients. The Covid-19 virus primarily targets the ACE2 receptors in the respiratory epithelium and, in severe cases, leads to alveolar damage and acute respiratory distress syndrome (ARDS) (Wu et al., 2020). As inflammation and fluid accumulation impair gas exchange, patients experience hypoxia, necessitating supplemental oxygen or mechanical ventilation. In addition to respiratory complications, thrombotic events such as pulmonary emboli or deep vein thromboses are frequently observed in hospitalized Covid-19 patients. These factors combined—immunopathology, respiratory failure, and coagulation abnormalities—explain why certain patients rapidly deteriorate and require hospital-level care.

Once admitted, patients receive a combination of supportive care and targeted therapies based on disease severity. Oxygen therapy is a cornerstone of treatment, ranging from low-flow nasal cannula to high-flow oxygen and mechanical ventilation for those with severe hypoxemia.[7] The guidelines recommend the use of convalescent plasma for Covid-19 patients with a high severity of illness.[8] They also recommend the use of antiviral agents such as remdesivir (FDA-approved October 2020), molnupiravir (FDA-approved December 2021), and dexamethasone (for severe but not critical patients). Immunomodulatory treat-

[7]See `https://www.covid19treatmentguidelines.nih.gov`.

[8]See `https://www.who.int/news/item/07-12-2021-who-recommends-against-the-use-of-convalescent-plasma-to-treat-covid-19`.

ments, including IL-6 inhibitors like tocilizumab, are also recommended for severe cases. Hospitalized Covid-19 patients often require thromboprophylaxis—blood clot prevention—due to the elevated risk of clots, which are a major contributor to morbidity and mortality. Low-molecular-weight heparin or direct oral anticoagulants are commonly used treatments.

If a patient's condition deteriorates, intubation and the use of a mechanical ventilator may be appropriate. In our analysis sample, 14% of patients are intubated, and there is a very high mortality rate for intubated patients (51%). There is also meaningful ambiguity on best practices in the use of ventilation for Covid-19 patients. As Pisano et al. (2020) states, "[t]his is a rather challenging task. In fact, current guidelines say nothing or are vague, with little adherence to clinical reality, and not based on strong evidence about the indications for tracheal intubation in patients with Covid-19." Furthermore, caring for intubated patients is complex with significant risk of pneumonia and lung injury, risks associated with sedation and immobility, and ultimately a probability of death. This creates a challenging patient risk/benefit calculation that underlies the intubation decision. Additionally, intubating a patient increases hospital costs.

The quality of the evidence underlying the treatment guidelines is often poor, implying significant scope for differential interpretation of the appropriate medical interventions. Furthermore, medical protocols for Covid-19 changed rapidly during the pandemic (Gulick et al., 2024). For example, the FDA issued an emergency use authorization (EUA) for baricitinib in November 2020, for its use in combination with remdesivir. The agency then revised the emergency use authorization in July, 2021 to authorize baricitinib as a stand-alone treatment, at which point it became the recommended treatment (Rubin, 2022). These changes imply that achieving high quality required learning about new protocols and translating them to the clinical setting.

The complexity of clinical decision-making, ambiguity in the interpretation of evolving guidelines, risks associated with interventions—particularly when implemented suboptimally—and the multifactorial nature of Covid-19 management imply that hospitals likely exhibited heterogeneity in treatment effectiveness. Furthermore, cognitive biases and the dissemination of misinformation contributed to the adoption of treatments lacking clinical efficacy, and in some cases associated with harm. Examples include glucocorticoids,

lopinavir/ritonavir, hydroxychloroquine, inhaled corticosteroids, famotidine, ivermectin, and colchicine (IDSA, 2021).

Treatment outcomes also likely varied over time, driven by the introduction of new therapeutics and possibilities of learning-by-doing and episodes of capacity strain. Given that Covid-19 was a new disease where best treatments for a given patient were often ambiguous suggests a role for learning-by-doing, as a hospital's direct experience and observation of patient outcomes could inform its approach to care. Capacity strain limited the availability of critical physical resources (e.g., ICU beds, ventilators, and oxygen supply) and disrupted staffing patterns by increasing nurse–patient ratios, necessitating redeployment of personnel to unfamiliar roles, and contributing to fatigue and burnout. These factors, in combination, may have reduced the quality of clinical decision-making and bedside care.

2.2 Data

Our primary data source is the State of Texas Department of Health and Human Services patient discharge dataset.[9] These data contain virtually every inpatient hospital discharge in the state. Texas is a large and diverse state with substantial urban and rural populations, offering a broad sample of hospitals with varying characteristics (e.g., rural, urban, suburban, academic, critical access, for-profit, and system-owned hospitals). We focus on the period from Q2:2020 to Q3:2022. This period covers the pandemic from its beginnings to essentially its end in 2022. We drop data from Q1:2020, the start of the pandemic, due to inaccurate reporting.

The discharge data contain information on the identity of the hospital, discharge status, patient demographics (age, gender, race, and ethnicity), home ZIP code, admission source, insurance status (Medicare, Medicaid, commercial insurance, self-insured), up to 24 International Classification of Disease Version 10 (ICD10) disease codes, and up to 24 ICD10 procedure codes. The data also identify the patient length of stay in days along with indicators for whether the patient spent the most (or second most) number of days in the

[9]https://www.dshs.texas.gov/center-health-statistics/texas-health-care-information-collection/health-data-researcher-information/texas-hospital-emergency-department-research-data-file-ed-rdf/hospital-discharge-data-public-use-data-file

ICU. We identify Covid-19 patients as those whose primary or secondary diagnosis upon admission was coded as any of U071, U072, or B9729. We use the ZIP code information and OpenStreetMaps to construct measures of patient travel time from the patient's ZIP code centroid to hospitals.

Our primary outcome variable is inpatient mortality, which we define as a discharge status of either died or discharged to hospice (Asch et al., 2021). This measure of inpatient mortality excludes patients who are discharged to their home or a nursing home and then die as a consequence of the care they received in the hospital.

We construct approximately 25 detailed risk adjustment measures using Elixhauser Comorbidity Indices (Elixhauser et al., 1998) following the approach of Asch et al. (2021). These measures indicate the presence of known risk factors associated with Covid-19 mortality, including diabetes, heart failure, cancer, and chronic obstructive pulmonary disease (COPD). We merge in county-level Covid-19 vaccination rates from the Texas Department of State Health Statistics (DSHS)[10] and census tract-level median income and demographic composition from the American Community Survey, aggregated to the ZIP code level.

We also collect information on hospital characteristics from different sources. From the American Hospital Association (AHA), we gather information on the hospital's for-profit status, total number of beds, number of intensive care unit beds (ICU), teaching status (major and minor), system membership, rural status, critical access hospital status, exact latitude and longitude location coordinates, and the CMS Certification Number (CCN). We assign these characteristics to each hospital using a hospital-name match to the AHA database.[11]

From the Centers for Medicare and Medicaid Services (CMS), we use the CCN to import information on pre-pandemic, hospital-specific inpatient mortality rates for acute myocardial infarctions, coronary artery bypass graft surgery, COPD, heart failure, pneumonia, and stroke. We also collect the average per-capita Medicare spending associated with the hospital,

[10] https://github.com/shiruken/covid-texas-data/tree/main/AccessibleVaccineDashboardData

[11] For about 15% of hospitals, we did not find a name match, in which case we assigned hospital characteristics using a manual search. For bed counts specifically, we supplement AHA using total and ICU beds reported by CMS Cost Reports (HCRIS) and/or Texas State Directory of General and Special Hospitals. Generally, the former is used in cases where both the AHA bed data were unknown/missing/0 and the CCN maps to only one hospital location; the latter is used in cases where the CCN is not unique, and thus hospitals report beds in a consolidated fashion to CMS HCRIS. See Appendix B for details on bed size construction.

the surgical complication rate, and CMS' hospital star rating.[12] CMS assigns each hospital a rating from one to five based on its reported outcomes across five areas of quality and maps those outcomes into a single star rating for each hospital.[13] Those areas are mortality, patient safety, readmission rates, patient experience, and timely and effective care.

We restrict our sample of hospitals to general acute care hospitals; i.e., we remove specialty facilities such as psychiatric centers, heart and spine hospitals, and children's hospitals. We restrict our patient sample to adults—those over 20 years of age—both because Covid-19 resulted in few pediatric deaths and to avoid comparing children's hospitals to general acute care hospitals. We drop patients whose home county is either missing or outside of Texas because we cannot assign them a county-quarter level vaccine rate. We also drop patients whose home ZIP code is missing, defined as a P.O. Box, or otherwise outside of the OpenStreetMaps database, as we cannot calculate travel distance from their ZIP code to hospitals for these patients. These patient drops amount to about 9.8% of the adult general acute care Covid-19 sample. We exclude a small minority (about 3.3%) of patients who traveled outside of their Texas public health region for care.[14] Next, we drop hospitals (and their patients) with fewer than 20 Covid-19 patients over the 10-quarter sample period, which removes 32 hospitals but less than 0.1% of patients. To avoid double counting, we also drop admissions with a discharge status denoting a transfer to another hospital, as these will be counted at the hospital to which the patient was transferred.

Finally, we impute length of stay for patients for which this variable is missing. We do this for the full patient sample—and not just Covid-19 patients—as we need length of stay for the full sample to impute hospital capacity utilization rates. We perform this imputation using predictive mean matching (PMM) with 20 nearest neighbors. We match based on total charges and age-bin indicators. For our counterfactuals, we also perform a similar imputation exercise for just Covid-19 patients. Here we take advantage of our greater information on these patients and match based on a polynomial of the quantile of predicted severity, as

[12]The average per-capita Medicare spending measure includes all Medicare Part A and Part B payments made for services provided to a patient during an episode of care, which includes the 3 days prior to the hospital stay, the inpatient hospital stay, and the 30 days after discharge from the hospital.

[13]See https://data.cms.gov/provider-data/topics/hospitals/overall-hospital-quality-star-rating/.

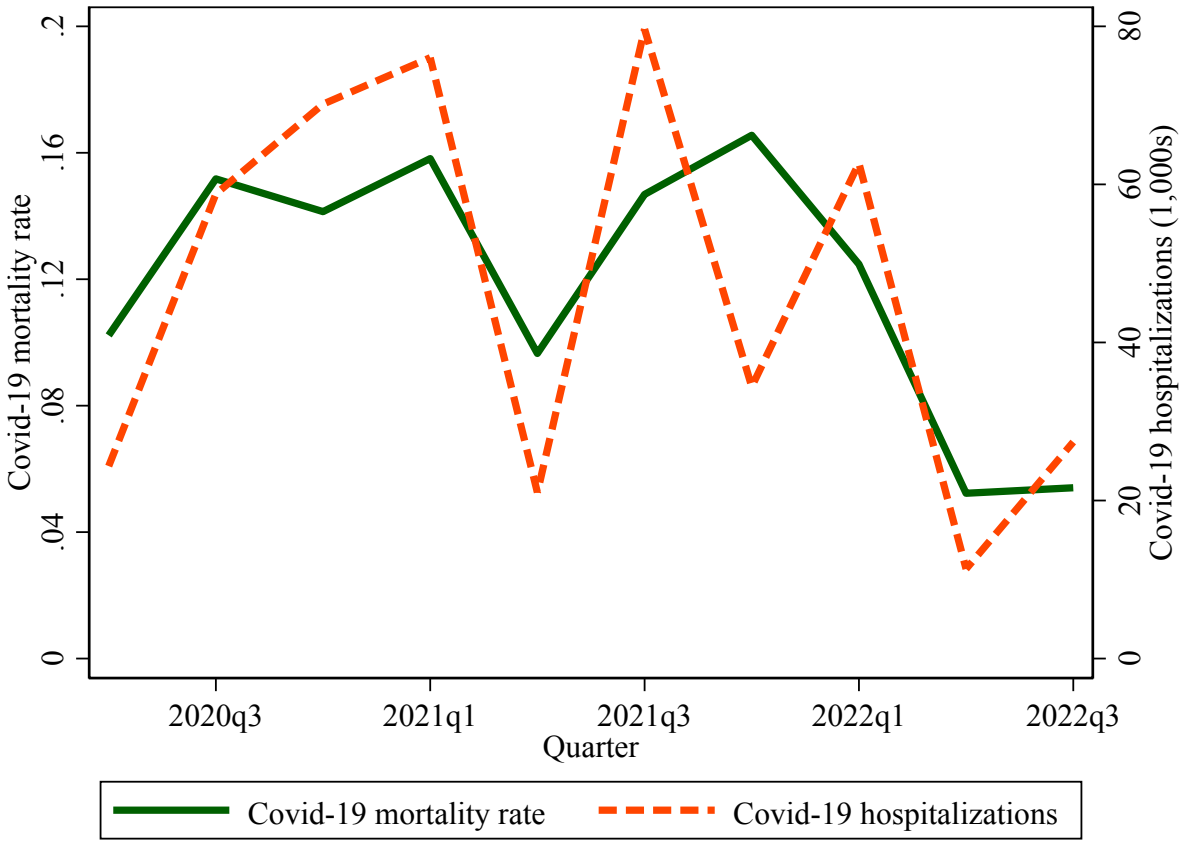
[14]Texas Health and Human Services divides each county in the state into one of 11 regions; see lhttps://www.dshs.texas.gov/center-health-statistics/texas-county-numbers-public-health-regions.

estimated by our model.

2.3 Trends and Summary Statistics

Figure 1 presents the time series of the total number of Covid-19 discharges in Texas and the Texas inpatient mortality rate over the pandemic. A couple of patterns are noteworthy. First, there were several surges in the pandemic. One occurred in Q1:2021, and the other occurred in Q3:2021. Second, the number of discharges and the mortality rate are positively correlated over time. This could be a result of the virulence of the strain circulating or the consequence of an overburdened care delivery system. Our analysis below will shed light on this issue.

Figure 1: Number of Covid-19 Discharges and Inpatient Mortality



Note: Figure reflects authors' calculations based on estimation sample of Texas inpatient discharge data.

Table 1 presents summary statistics of our patient sample broken down by year. Our analysis sample contains 397,631 inpatient observations with 54,651 deaths (13.7%). The mortality and the number of inpatient admissions were highest in 2021 with over 15% of patients dying. Approximately one-third of patients are Hispanic and 13% are Black (non-Hispanic). Forty-four percent of the sample has Medicare insurance, with a similar number over 65 years of age; a third have commercial/private insurance coverage, and 6.6% are covered by Medicaid.

We find that 0.8% of patients are transferred into a hospital from a skilled nursing facility, and 5% of patients are transferred from one hospital to another. The average length of stay is 7.4 days, and about 43% of patients spend the most or second most of their stay in the ICU. Patients have 2.98 comorbidities, on average. The mean county-level vaccination rate (weighted by patients) was 44% in 2021 and 68% in 2022. More than a third of patients were admitted to the closest hospital, and the average distance traveled from home to hospital was 24 minutes. This last fact suggests that patients make active hospital choice decisions.

Table 2 presents hospital summary statistics. Our analysis sample includes data on 333 hospitals.[15] The average bed size is 205 beds, but there is a wide distribution of bed size in our sample: the standard deviation is 267 beds. The average number of ICU beds is 22.7 with a standard deviation of 30.9 beds. Many hospitals (22%) have zero reported ICU beds. Thirty-nine percent of the hospitals are for-profit, while 71% are members of a system, 4% are major teaching hospitals, about one-sixth of hospitals are rural, and 14% have critical access status.

Hospitals can temporarily expand intensive care unit (ICU) capacity by converting general acute care beds into ICU beds, often through reallocation of staff, equipment, and monitoring capabilities during periods of surge demand (e.g., pandemics or mass-casualty events) (Abir et al., 2020). However, in standard hospital surveys and reports, we observe only the officially designated, static number of ICU beds, which does not capture such surge expansions and may itself be subject to reporting error or misclassification. For this reason,

[15] On-line Appendix Table A1 presents the hospital characteristics means (excluding those with missing values) for both the AHA Matched Sample and the AHA Unmatched Sample, where the latter is based on supplementary data sources and manual look-up of hospital characteristics.

Table 1: Summary Statistics on Patients

	Full Sample	2020	2021	2022
Died	0.137	0.142	0.152	0.101
Male	0.495	0.503	0.505	0.459
White (non-Hispanic)	0.447	0.396	0.456	0.504
Black (non-Hispanic)	0.132	0.134	0.131	0.133
Hispanic	0.338	0.382	0.332	0.284
Medicaid	0.066	0.059	0.063	0.080
Medicare	0.437	0.426	0.393	0.545
Commercial insurance	0.326	0.332	0.360	0.244
Other/No insurance	0.172	0.183	0.184	0.131
Age 65+	0.444	0.445	0.392	0.549
Age 80+	0.157	0.155	0.127	0.224
Number comorbidities	2.976	3.032	2.876	3.100
SNF transfer	0.008	0.011	0.005	0.011
Hosp transfer	0.047	0.055	0.045	0.040
Length of stay	7.358	7.094	8.129	6.136
Primarily in ICU	0.432	0.444	0.451	0.375
Minutes to closest	15.152	14.972	15.301	15.114
Minutes to chosen	24.427	24.623	24.341	24.308
Went to closest	0.362	0.364	0.363	0.357
Vaccination rate	0.346	0.003	0.440	0.676
Median income ($100,000)	0.379	0.366	0.384	0.387
N Patients	397,631	131,642	180,530	85,459

Note: Table of patient summary statistics for the full sample and for those admitted in a given year, 2020, 2021, and 2022. The year 2020 includes patients with discharges during the last three quarters, whereas the year 2022 includes patients with discharges during the first three quarters. All values are averages across patients.

we often observe hospitals utilizing more than 100% of their measured ICU capacity.[16]

There is meaningful variation across hospitals in the clinical outcome measures that we use as controls in some specifications. For example, the COPD mortality rate is 0.08 with a standard deviation of 0.01. It is also important to note that we have several missing values for the clinical measures because the hospital did not have enough volume of those procedures to meet the CMS reporting requirements. The CMS star rating system ranges from one to five, and the average hospital earned 3.33 stars. There is also significant variation across hospitals

[16]See Appendix B for additional details on how we construct hospital bed measures.

in their star rating—the standard deviation is one star. For spending, CMS attributes an average of $\$9,615$ in per-capita Medicare spending to the average hospital, and again there is a large variance across hospitals in this measure.

Table 2: Summary Statistics on Hospitals

	Mean	SD	Min	Max	N
Total beds	205.36	266.69	5.00	1703.00	333
Total ICU beds	22.67	30.91	0.00	211.00	333
Major teaching hospital	0.04	0.19	0.00	1.00	333
For-profit	0.39	0.49	0.00	1.00	333
Public	0.20	0.40	0.00	1.00	333
System member	0.71	0.46	0.00	1.00	333
Rural	0.15	0.35	0.00	1.00	333
Critical access	0.14	0.35	0.00	1.00	333
AMI mortality rate	0.13	0.01	0.10	0.15	199
COPD mortality rate	0.08	0.01	0.06	0.11	276
CMS star rating	3.33	1.03	1.00	5.00	274
Per-capita Medicare spending ($)	9,615	2,525	5,310	30,989	271
Average overall utilization rate	0.41	0.20	0.01	1.02	333
Average ICU utilization rate	1.32	1.55	0.00	22.21	333
Average ICU utilization rate (adj.)	1.25	1.04	0.00	4.24	333

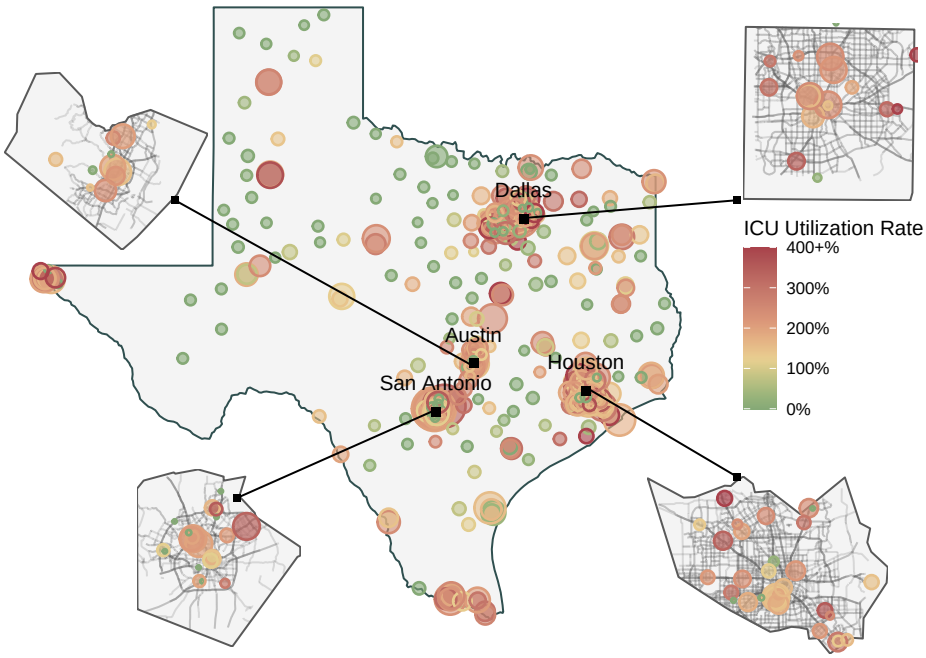
Note: Table of hospital level summary statistics for the full sample. Some values are missing for certain hospitals due to insufficient patient volume for CMS reporting requirements. See On-line Appendix Table A1 for hospital characteristics means (excluding those with missing values) for both the AHA Matched Sample and the AHA Unmatched Sample. See Appendix B for additional details on how we construct hospital bed measures. Hospital-specific utilization rates are based on the within-hospital average across quarters in the sample. Overall utilization rate is defined as the sum of patient length of stay divided by total bed-days in the quarter. ICU utilization rate is defined as the sum of ICU patient length of stay divided by ICU bed-days in the quarter. Hospitals without any ICU beds are assigned 1 ICU bed for purposes of utilization rate calculation. The adjusted ICU utilization rate caps the rate for hospitals with 0 ICU beds at 150%.

Finally, this table also presents the summary statistics for the overall and ICU utilization rates for hospitals, averaged across all 10 quarters. The mean hospital has, on average, 41% of bed-days filled during the sample period, with a standard deviation of 20% of bed-days filled. The mean hospital has an average ICU utilization rate of 132%, though it is measured imperfectly given the non-negligible share of hospitals with 0 ICU beds (to which we assign 1 bed for purposes of calculating the ICU utilization rate). In our main analysis, we consider an *adjusted* ICU utilization rate, in which we cap the utilization rate at 150% for hospitals

with 0 ICU beds. The mean hospital's *adjusted* ICU utilization rate is 125%, with a wide standard deviation of 104%.

Figure 2 presents the geographic distribution of ICU utilization rate across hospitals for a peak quarter, Q1:2021.[17] It shows that capacity strain was prevalent for hospitals in all regions of Texas at this peak period, with mostly smaller hospitals having a capacity utilization rate under 100%. On-line Appendix Figure A2 presents the histogram of ICU bed-day utilization across hospital-quarters for all hospitals with a positive number of ICU beds. Over 50% of the hospital-quarter observations have an ICU utilization greater than 150%.

Figure 2: Hospital ICU Utilization Rate, Q1:2021



Note: Map of hospital ICU utilization rate during Q1:2021. Each circle represents one hospital and has area proportional to total bed size.

The State of Texas took few measures to actively manage capacity strain. We investigate the empirical importance of these efforts by examining whether patient transfers lessened capacity strain. To start, we find that transfers are relatively uncommon. The mean transfer-in rate (i.e., the share of patient discharges in a hospital-quarter that were transferred in from

[17]See Appendix Figure A1 for the analogous map of ICU utilization rate in Q2:2020.

another hospital) was 4.0% and the mean transfer-out rate is 5.1%. The bed-size weighted averages were 6.7% and 2.0%, respectively, suggesting that larger hospitals were more likely to accept transfers and less likely to transfer out. Importantly, we do not find a systematic relationship between hospital capacity utilization strain and transfer rates. On-line Appendix Table A2 presents results from regressions of hospital-quarter level transfer rates on general and ICU utilization rates, with controls for hospital and region-by-quarter fixed effects. We do not find a systematic or strong statistical relationship between utilization rates and transfer rates. Our overall takeaway is that patient transfers were a relatively unimportant predictor of patient hospital choice in our data.

3 Empirical Framework

Our empirical approach proceeds in three parts. We first recover hospital quality at the quarter level. We next decompose hospital quality to evaluate its correlates. Finally, we estimate patient hospital choice models to understand the role of patient information in choice. We discuss each of these estimation exercises in turn.

3.1 Estimation of Hospital Quality

Our model and estimation of hospital quality build on the GGT approach as well as the related literature. We index patients by i, time—measured in quarters—by t, hospitals by j, and Texas Public Health Regions by r. At each time t, there are I_t patients and J_t hospitals in each region.[18] In our model, a patient contracts severe Covid-19 that requires hospitalization. She must choose from one of the hospitals in her region that are operating in the quarter. There is no outside option.

We model a health status, m^*_{it}, as a latent random variable, with an observed patient death occurring when $m^*_{it} > 0$, so $m_{it} = \mathbb{1}\{m^*_{it} > 0\}$. The latent health status depends on the patient's choice of hospital, the patient's individual characteristics, x_{it}, and an unobservable

[18] For ease of notation, we suppress the region subscript in this subsection since we perform the estimation separately by region.

term, ε_{it}:

$$m^*_{it} = c'_{it}\beta_t + x'_{it}\gamma_t + \varepsilon_{it}. \tag{1}$$

In (1), c_{it} is a vector of J_t indicators that specifies the hospital chosen by patient i, β_t is the vector of fixed effects for treatment at time t, x_{it} are patient characteristics, and γ_t are the coefficients that specify the impact of these characteristics on m^*. Patient characteristics, x_{it}, include gender, race/ethnicity, age bins, insurance type, indicators for transfer from a skilled nursing facility, 26 Elixhauser Comorbidity Indices (Asch et al., 2021), the county-level vaccine rate, and ZIP code-level demographics.

We allow for patient selection based on unobserved severity of illness with a flexible correlation structure between ε_{it} and the hospital choice indicators, c_{it}. Specifically, we model the patient as choosing one of the available hospitals in her region and quarter.[19] We specify the utility for hospital j to be $\tilde{c}^*_{itj} = \tilde{Z}'_{itj}\alpha_t + \tilde{\eta}_{itj}$, where $\tilde{Z}_{itj}$ includes characteristics that affect patient i's choice of hospital, specifically distance from the patient's ZIP code to the hospital in hours and its square, α_t are coefficients, and $\tilde{\eta}_{itj} \sim N(0,1)$, *i.i.d.* across options.[20] The patient chooses the hospital $j \in 1, \ldots, J_t$ that yields the highest utility level $\tilde{c}^*_{itj}$. Since there is no outside option, we normalize utility by subtracting the utility for the J_tth hospital from all others and letting the J_tth hospital have utility 0. For $j < J_t$, we rewrite $c^*_{itj} = Z'_{itj}\alpha_t + \eta_{itj}$ and let Σ be the variance of the $J_t - 1$ dimensional vector η_{it}. We then allow for selection based on unobservables by specifying that $\varepsilon_{it} = \eta'_{it}\delta_t + \zeta_{it}$, where $\zeta_{it} \sim N(0,1)$, and δ_t are $J_t - 1$ selection coefficients that we estimate, that capture the endogeneity between choice and quality.

Because our measures Z_{itj} are based on distance to a hospital and these do not enter

[19] A potential concern is that patients may face restricted choice sets because of narrow networks. We do not believe this assumption meaningfully affects our analysis for two reasons. First, the majority (53%) of patients in our sample have traditional Medicare, indemnity Blue Cross/Blue Shield, or have no coverage and hence do not have network restrictions. Second, Texas and federal Covid-19 policies substantially reduced the role of insurer networks. In March 2020, Texas regulators explicitly encouraged insurers to waive cost sharing and "penalties...for necessary out-of-network services and emergency-related care, alongside broader access expansions." (`https://gov.texas.gov/news/post/governor-abbott-tdi-ask-health-insurance-providers-to-waive-costs-associated-with-coronavirus?`) Accordingly, as an example, Blue Cross and Blue Shield of Texas waived deductibles, copays, and coinsurance for Covid-19 treatments, effectively eliminating patient liability. Medicare Advantage issued parallel federal guidance. For Medicaid, federal Section 1135 waivers ensured providers—including out-of-network or out-of-state hospitals—would be reimbursed and administrative barriers relaxed.

[20] We also experimented with an indicator for the closest hospital.

into the mortality equation (1), distance acts as an instrument for the choice of hospital, in a non-linear model. The identifying assumption is then that the unobserved illness shock—or equivalently, the unobserved patient-specific probability of death—is not correlated with distance to different hospitals.

As in GGT, we can write (the negative of) hospital quality as $q_{tj} = \beta_{tj}/\sqrt{1 + \delta_t' \Sigma \delta_t}$. This term q_{tj} specifies the quality fixed effect, normalized by the standard deviation of ε_{itj}. It can be used to analyze expected mortality when unobserved illness shocks are rescaled to have a standard normal distribution. For our purposes, the main output of the GGT algorithm is the posterior distribution of q_{tj}, which differs by quarter, hospital, and region.

We estimate a separate GGT specification by quarter and region.[21] This allows for hospital quality to vary flexibly across time and for patients to select hospitals that are good at particular quarters based on their unobserved health characteristics. We specify weak priors on hospital quality, using the defaults in Marquardt et al. (2025). While GGT allows for correlations in priors between hospitals of the same type, we treat all hospitals symmetrically.

GGT estimates parameters with a Markov-chain Monte Carlo (MCMC) estimation algorithm. The output of the estimator is a series of serially correlated draws from the parameters' posterior distribution. Since the number of patients—and hence deaths—in a given quarter is limited, the posterior hospital-quarter quality distributions will generally be quite noisy. We decompose the quality posterior distribution draws to recover mean hospital quality over time and other quality determinants. These measures are more precisely estimated. Importantly, our MCMC estimator does not rely on a law of large numbers and asymptotic normality as a frequentist estimator would. It instead returns draws from the posterior distribution even when the data are limited. For this reason, it is appropriate to use posterior draws as noisy measures of quality in our decomposition exercises, as long as we account for the fact that our MCMC estimation generates serially correlated draws.

[21] We also report estimates from GGT estimated with a fixed effect for each quarter, which makes the assumption that quality is fixed over time, and with GGT estimated without the correction for unobserved selection, i.e. with $\delta_t = 0$.

3.2 Decomposing Hospital Quality

This estimation starts with the posterior distribution draws of the hospital-quarter quality estimates. We perform this estimation for all hospitals and quarters in our sample, and hence we add the region fixed effect, r, to our notation. For each draw, d, we first simulate the posterior expected mortality rate for each hospital if it were faced with the population distribution of observed and unobserved severity of illness, which we denote $mort_{rtjd}$.[22] We specify that $mort_{rtjd}$ is a function of hospital fixed effects and time-varying characteristics:

$$mort_{rtjd} = \overline{mort}_{rj} + f(occupancy_{rtj}) + g(experience_{rtj}) + FE_{rt} + e_{rtjd}. \tag{2}$$

In (2), $\overline{mort}_{rj}$ indicates baseline hospital mortality without any capacity strain or learning-by-doing, $f(occupancy_{rtj})$ indicates the impact of capacity strain, $g(experience_{rtj})$ indicates the impact of learning-by-doing, FE_{rt} are fixed effects for each combination of region and quarter, and e_{rtjd} are unobserved time- and draw-varying quality shifters. We estimate (2) with linear regression.

Our baseline specification measures $f(occupancy_{rtj})$ with a spline in the occupancy rate for intensive care unit (ICU) beds for a hospital-quarter.[23] Our main measure of experience is the lag of cumulative Covid-19 volume up to the previous quarter. We include quarter and region fixed effects to capture differences in Covid-19 treatment diffusion or patient severity over time and region.

We calculate the occupancy rate as the number of ICU patient-days for discharges in the hospital-quarter divided by the number of ICU beds times 91 (the approximate number of days in a quarter). Our measure of ICU patient-days is imperfect: while we do observe patient length-of-stay, we do not know the exact number of days each patient stayed in the ICU during her stay. Instead, for each discharge, we know whether an ICU stay comprised the most days, or second-, third-, or fourth-most days over the stay. We assign the entire inpatient stay to the ICU for patients for whom the most or second-most number of days

[22] We use the expected mortality rate as our dependent variable instead of the quality fixed effects, q_{rtj}, for ease of interpretation. For each hospital, we derive this rate by drawing a random subsample of 500 patients, assuming that they have the population distribution of unobserved severity of illness, and calculating the posterior mean probability of death at the hospital across these patients and posterior parameter draws.

[23] We also consider occupancy based on overall hospital utilization, not ICU utilization.

was in the ICU.[24] In addition, patient stays may start and finish in different quarters, and hence our allocation of inpatient days to the quarter of discharge is imprecise.

Because the GGT algorithm returns MCMC draws, we have many draws d from the posterior distribution of hospital quality. Given the nature of the MCMC algorithm, these draws are highly serially correlated, and the initial draws will not reflect the posterior distribution. Hence, we keep every 100th draw and drop the first 10,000 burn-in draws. In our regressions based on (2), we two-way cluster at the hospital and region/quarter levels to obtain standard errors that are robust to the remaining serial correlation and are robust to correlation in the errors across time for a hospital or across hospitals in a given region/quarter.

The effects of capacity strain and experience on mortality are identified by within-hospital, over-time variation in patient volume. Our principal identifying assumption in (2) is that the time-varying component of hospital volume is not correlated with the unobserved time-varying component of quality, e_{rtjd}. Recall that this specification has time-by-region fixed effects that will soak up any correlation between regional demand shocks and unobserved patient severity. This assumption would be violated if, for instance, hospitals that improved their quality for exogenous reasons publicized their improvement in a way that drew Covid-19 patients to them or there were very localized infection severity surges. This source of endogeneity seems unlikely because it would have been difficult for hospitals to advertise in such a targeted way during the pandemic, when Covid-19 cases and occupancy were rapidly fluctuating. Also, we are unaware of any evidence of hyper-localized (e.g., ZIP code-level) severity surges. In addition, in this case, we might find a positive coefficient on experience even in the absence of learning-by-doing (Gaynor et al., 2005, Gowrisankaran et al., 2006). We would also expect to see overcapacity appear to *increase* hospital quality, since hospitals would be drawing more patients and improving at the same time. Thus, to the extent that we estimate no learning-by-doing or quality degradation from overcrowding, these findings would be unlikely to be explained by this type of endogeneity.

Finally, we also regress the hospital fixed effects, $\overline{mort}_{rj}$, on time-invariant hospital characteristics, notably ownership, system membership, and size. These estimates are less plau-

[24]This ICU measure does not include neonatal intensive care unit utilization. We also examined assigning a patient to the ICU only when the ICU stay comprised the most days. Our qualitative results are robust to this alternate definition.

sibly causal than the estimates on learning-by-doing and capacity strain, but nonetheless informative about which types of hospitals would have delivered high-quality Covid-19 care when not operating at high capacity utilization levels.

3.3 Patient Hospital Choice Model Estimation

Our final estimation exercise recovers patient hospital choice parameters accounting for the measures of quality that we estimated in Section 3.2. Broadly, the literature suggests that higher performing hospitals attract more patients all else equal (e.g., Luft et al., 1990). If this pattern holds for Covid-19 patients, then the deleterious impact of relatively poor-performing hospitals on population mortality rates can be mitigated by patients avoiding those hospitals. Conversely, if patients are more likely to choose those hospitals, then the impact of poor-performing hospitals is exacerbated. It is unclear how Covid-19 patients will select hospitals, as there was little public information available on Covid-19 hospital performance *per se*.

The model we estimate to recover the impact of hospital quality on choices is similar to the "first stage" hospital choice specifications we estimated to understand hospital mortality in Section 3.1. However, our specifications differ. Our goal here is causal: to understand what quality information patients are using in making hospital choice decisions. In the GGT hospital mortality estimation, our goal was to create variation in choice probabilities using distance without imposing assumptions on patients' information regarding hospital quality.

We again model each patient as choosing one of the hospitals in her region and quarter, but now with different regressors. We specify latent utility as:

$$u_{ritj} = \lambda_1 x^1_{ritj} + \lambda_2 x^2_{rtj} + \lambda_3 E[mort_{rtj}] x^3_{rit} + \nu_{ritj}, \tag{3}$$

where x^1_{ritj} indicates variables that relate to the interaction of patient and hospital characteristics, x^2_{rtj} are hospital characteristics, $E[mort_{rtj}]$ is patients' expectation of expected hospital mortality (the negative of quality), which we allow to interact with patient characteristics x^3_{rit}, and λ_1, λ_2, and λ_3 are parameters to estimate. We also include a type 1 *i.i.d.* extreme value unobservable utility shock ν_{ritj}. Thus, the probability of an individual

choosing a particular hospital takes a multinomial logit form, which allows us to estimate (3) with maximum likelihood.

Our principal regressors in x^1_{ritj} are functions of the travel time between the patient (based on her ZIP code centroid) and hospital. We also create a measure of each patient's mortality risk quantile, using the GGT posterior draws of γ_t and δ_t. Some specifications include this measure of observed patient severity interacted with distance as an additional term in x^1_{ritj}. Our characteristics measures in x^2_{rtj} include hospital size, whether it is a system member, and ownership type.

A priori, it is unclear which type of Covid-19 hospital quality information, if any, patients might respond to in their hospital choices. For this reason, we consider two different quality measures in $E[mort_{rtj}]$ across specifications. First, we consider mortality without capacity strain, i.e., when the hospital is empty, $E[mort_{rtj}|occ_{rtj} = 0] \equiv \overline{mort}_{rj}$, where occ_{rtj} denotes the ICU occupancy rate used to calculate mortality. Patients might be informed about this measure because it would be consistent with the long-run quality of the hospital.

Second, we construct the expected mortality at the realized occupancy level for each hospital in each quarter. We construct this measure by adding the hospital's baseline mortality to its additional estimated mortality due to its current capacity, $E[mort_{rtj}|occ_{rtj} = occ^{obs}_{rtj}] \equiv \overline{mort}_{rj} + f(occ^{obs}_{rtj})$, where occ^{obs}_{rjt} denotes observed ICU capacity utilization. For this measure of quality, for patients treated in the ICU, the choice of hospital may degrade quality slightly by increasing ICU overcrowding. We account for this endogeneity with a leave-one-out type estimator, by adding the patient's impact on $f(occ^{obs}_{rtj})$ to capacity utilization for hospitals not chosen, with the assumption that her length-of-stay and whether she is treated in the ICU would remain constant across hospitals.

In our patient hospital choice estimation, our first measure of hospital quality information—baseline quality without capacity strain—is constant over time for a hospital. For this reason, we do not include hospital fixed effects in the results based on equation (3). Thus, we will identify a positive impact on quality information to the extent that patients choose a given hospital more often if it has high quality information, after controlling for hospital attributes as well as patient/hospital attributes, such as distance. Given the importance of controlling for hospital attributes, we investigate the robustness of our findings to variation in the

hospital characteristics used as controls.

4 Empirical Results

This section presents the results for the three main empirical exercises discussed in Section 3. Specifically, we report the estimates of hospital performance in treating Covid-19; the determinants of hospital quality; and the impact of hospital quality on hospital choice.

4.1 Hospital Quality Estimates

Our hospital quality estimates are draws from the posterior distribution of hospital quality, q_{rtj}, estimated by the MCMC process developed in GGT. Our base specification estimates this measure separately by hospital and sample quarter. It controls for observed patient risk adjusters and selection based on unobserved severity of illness.[25] As discussed in Section 3.2, for interpretability, we transform each posterior draw of q_{rtj} into a draw of $mort_{rtj}$, the mortality rate that a hospital would face if it treated a representative population sample of patients from Texas over our sample period.

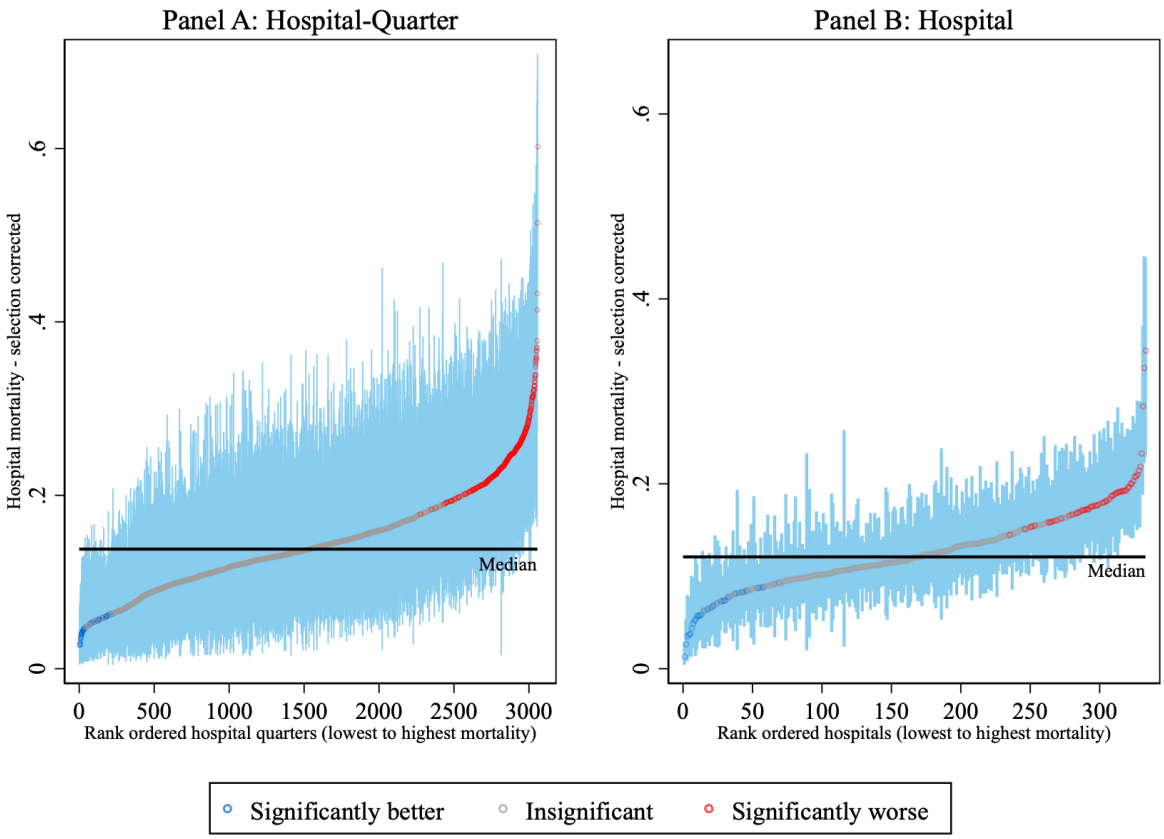
Figure 3 panel A displays the posterior mean of $mort_{rtj}$ and the 95 percent posterior confidence interval from our base GGT estimation, for each hospital-quarter in our sample. The large majority of hospital-quarters, 80 percent, have fixed effects for which the 95 percent posterior confidence interval includes the median estimated hospital mortality. This may reflect many hospitals having similar quality, but it may also be due to a lack of precision because of the limited number of Covid-19 patients by hospital and quarter. To control for the first explanation, Section 4.2 combines posterior draws across quarters to obtain more precise inference on the determinants of hospital quality, accounting for the serial correlation in the draws.

For comparison purposes, we also estimate specifications that perform the estimation together for all patients and hospitals in a region.[26] Figure 3 panel B provides analogous in-

[25] See On-line Appendix Table A3 for the means and standard deviations of the patient risk adjuster posterior distribution means γ across regions and quarters. The coefficients generally align with what we would expect. For instance, age is a strong predictor of Covid-19 mortality, as are fluid and electrolyte disorders. Conditional on hospital admission, men are more likely to die than women from Covid-19.

[26] These specifications include quarter fixed effects in addition to the risk adjusters reported in Table A3.

Figure 3: Hospital Mean Mortality $mort_{rtj}$



Note: Each point shows a posterior mean hospital mortality. Each vertical line displays a 95 percent confidence interval for the posterior with 5% significance indicated. Panel A shows this information by hospital-quarter, $mort_{rtj}$, and panel B shows this information by hospital over the whole sample.

formation to panel A but for this time-invariant measure. These estimates are valid expected mortality estimates if Covid-19 hospital quality were time invariant over our sample. The reported 95% posterior confidence interval shows that the estimates with one fixed effect per hospital are much more precise, as we would expect given the larger sample of patients that contributes to each hospital fixed effect. About 30 percent of hospitals have posterior mean quality levels that are significantly different than the median hospital.

The heterogeneity in hospital quality from Figure 3 panel B is significant. To quantify the variation in hospital quality across hospitals in our sample, we simulate moving all patients who were admitted to a hospital with above-median mortality rate to the median hospital. We find that this would have reduced mortality by 14.9%. Extrapolating this

number to the U.S. population, which had approximately 733,000 in-hospital deaths during the pandemic, would imply a reduction of 109,000 deaths during this time period. Silva et al. (2023) combine consumer surplus and lost income contributions to estimate a value of $3.57 million per Covid-19 life lost, implying that a back-of-the-envelope welfare gain would be about $390 billion. Of course, we interpret this result as descriptive and not causal because hospital quality may vary over time. Section 4.2 decomposes the variation in hospital quality over time to obtain more causal estimates.

To understand the importance of selection based on unobservable patient risk, we estimate GGT controlling only for observed patient risk adjusters, by setting $\delta = 0$. For ease of exposition, we focus here on the time-invariant measure of quality, because it has many fewer hospital fixed effects. On-line Appendix Figure A3 reports the relationship between the posterior mean of the hospital fixed effects with and without controls for selection. We find that there is a substantial correlation between the sets of estimates, with a correlation coefficient of 0.81. Nonetheless, the correlation is not perfect, implying that patient selection based on unobserved severity of illness plays a role in affecting observed hospital mortality.

4.2 Hospital Quality Determinant Estimates

We next turn to decomposing the impact of hospital quality. We focus first on understanding the role of capacity strain and learning-by-doing in affecting hospital quality. Our dependent variable is posterior draws of the hospital mortality rate estimated at the quarterly level, $mort_{rtjd}$. We regress this mortality rate on hospital fixed effects, region-by-quarter fixed effects, and within-hospital time-varying factors.

Table 3 provides our regression results on the impact of capacity strain and learning-by-doing. All specifications include hospital fixed effects—which control for baseline hospital quality differences—and region-by-quarter fixed effects—which control for heterogeneity on these dimensions. Column (1) includes a regressor for the ICU utilization rate, which we find significantly increases hospital mortality (decreases quality). Column (2) includes both the ICU utilization rate and the overall utilization rate (based on each hospital's total beds rather than ICU beds). The ICU utilization rate coefficient is similar to that in column (1), while the overall utilization rate is negative, small, and statistically insignificant.

Table 3: Impact of Occupancy and Experience on Hospital Mortality

Dependent Variable:	Hospital expected mortality draw: $mort_{rtjd}$					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Regressor:						
ICU utilization rate	0.011***	0.013***				
	(0.002)	(0.003)				
Overall utilization rate		-0.018		-0.017	-0.018	
		(0.012)		(0.012)	(0.012)	
ICU utilization rate spline: $[0, 0.75)$			0.008	0.011	0.012	
			(0.008)	(0.009)	(0.009)	
ICU utilization rate spline: $[0.75, \infty)$			0.012***	0.013***	0.013***	
			(0.002)	(0.003)	(0.003)	
log(Covid patients through t-1)					0.001	
					(0.001)	
ICU utilization rate spline: $[0, 0.75)$						0.007
						(0.009)
ICU utilization rate spline: $[0.75, 1.5)$						0.015**
						(0.007)
ICU utilization rate spline: $[1.5, 3)$						0.009***
						(0.003)
ICU utilization rate spline: $[3, \infty)$						0.016***
						(0.004)
Hospital/quarter/draw observations	2,752,200	2,752,200	2,752,200	2,752,200	2,752,200	2,752,200
R^2	0.4803	0.4805	0.4803	0.4805	0.4805	0.4804

Note: All regressions include hospital and region-by-quarter fixed effects. We calculate standard errors with two-way clustering at the hospital and region-by-quarter levels. The terms ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 3, column (3) investigates a linear spline for the ICU utilization rate with a knot at capacity utilization of 75%. We find that the increase in mortality from overcrowding is higher and statistically significant at ICU capacity levels over 75%. We use column (3) as our base specification. These basic results are corroborated by columns (4) and (5), which also show statistically significant spline coefficients for ICU capacity utilization greater than 75% after adding further controls.[27] Column (6), which adds more knots to the spline, also

[27]Appendix Table A4 further shows robustness of our baseline specification (column 3) to those that use alternative total and ICU bed source data.

shows statistically significant impacts only when utilization is greater than 75%. On-line Appendix Figure A4 provides a visual representation of the spline coefficients in columns (3) and (6). In both cases, mortality increases by over 2.1 percentage points as ICU capacity utilization goes from 0% to 200%.

Table 3 column (5) investigates the presence of learning-by-doing by adding the log of the cumulative lagged Covid-19 patients at the hospital, up to the end of the previous quarter. We find that previous experience is not statistically significant. Moreover, the point estimate is positive, implying that, if anything, treating more past Covid-19 patients is also associated with higher mortality, even after controlling for the fact that more previous patients may imply more capacity strain. In contrast, the capacity strain spline coefficients are similar to those in column (3).

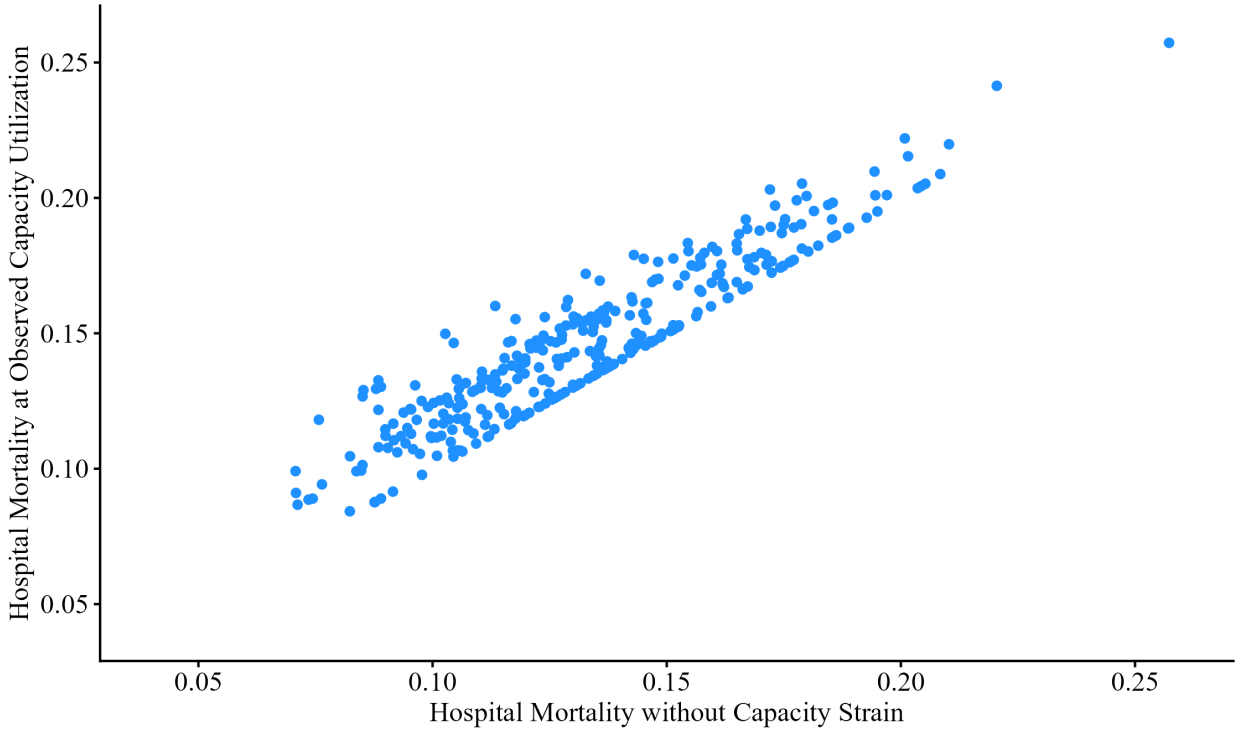
Thus, our overall takeaway from Table 3 is that ICU capacity strain was an important component leading to a degradation of Covid-19 hospital quality. The capacity strain that affected Covid-19 mortality was occurring in the ICU, rather than elsewhere in the hospital. In contrast, we find no evidence that learning-by-doing from greater patient volume led to lower mortality.

We report the magnitudes implied by the capacity strain with two exhibits, using the coefficient estimates from our base specification, Table 3 column (3). First, Figure 4 shows the hospital mortality means without any capacity strain—i.e., when empty—and at the realized level of capacity strain. By construction, all points lie above the 45-degree line. Many of the points are significantly higher than the 45-degree line, providing a visualization of the effect of capacity strain on the quality of Covid-19 treatments.

On-line Appendix Table A5 provides a second look at the impact of capacity strain, on a hospital-by-hospital basis. On average, mortality at the sample's mean level of capacity utilization is about 1.3 percentage points higher than it would be with no strain. This figure is also 0.4 percentage points higher than the mortality that would occur if each hospital were operating at 100% of its ICU capacity.

Our second empirical decomposition starts with the baseline hospital quality fixed effects, $\overline{mort}_{rj}$, as dependent variables, evaluating the predictors of these fixed effects. Table 4,

Figure 4: Mean Hospital Mortality When Empty and at Realized Occupancy



Note: Figure based on authors' estimation sample and regression estimates presented in Table 3 column (3). The reported mortality numbers are the mean of the estimated mortality level by hospital across all quarters.

column (1) considers the impact of hospital ownership and the ratio of ICU to total beds. For-profit and public hospitals are statistically worse than not-for-profit hospitals. These coefficients remain significant across the specifications reported in this table.

Column (2) adds in the impact of hospital size and critical access status on baseline mortality. Critical access status—which is only given to small, rural hospitals—is associated with a 1.1 percentage point higher Covid-19 mortality rate. Besides this, we find little evidence of a correlation between size and mortality.

Finally, column (3) adds in the impact of a variety of quality ratings and other measures.[28] Mortality for chronic obstructive pulmonary disease (COPD) is large and highly significant, with a higher (or missing) COPD mortality rate predicting higher Covid-19 mortality. This suggests a spillover from existing COPD quality to Covid-19 quality. Given that many of

[28]These measures have many missing values, and hence we add in indicators for whether each measure is missing.

Table 4: Decomposition of Baseline Hospital Quality into Characteristics $\overline{mort}_{rj}$

Regressors	Dependent Variable: Hospital mortality fixed effects, $\overline{\text{mort}}_{rj}$ (1)	(2)	(3)
Ratio ICU to total beds	-0.016 (0.027)	-0.001 (0.028)	-0.000 (0.028)
Major teaching	-0.015* (0.009)	-0.012 (0.010)	-0.010 (0.010)
For-profit	0.008** (0.003)	0.008** (0.003)	0.009*** (0.003)
Public	0.023*** (0.005)	0.020*** (0.005)	0.015** (0.006)
System member	-0.000 (0.004)	0.001 (0.004)	-0.003 (0.005)
Beds: ≥100 and <500		0.001 (0.004)	-0.001 (0.005)
Beds: ≥500		-0.001 (0.007)	-0.002 (0.008)
Critical access		0.011** (0.005)	0.019** (0.007)
CMS star rating = 2			0.006 (0.011)
CMS star rating = 3			-0.003 (0.011)
CMS star rating = 4			0.002 (0.011)
CMS star rating = 5			0.001 (0.012)
Star rating missing			-0.003 (0.012)
log(Medicare spend)			0.001 (0.008)
Spend missing			0.005 (0.075)
COPD 30-day mortality			0.660*** (0.175)
COPD mortality missing			0.047*** (0.016)
AMI 30-day mortality			0.403* (0.223)
AMI mortality missing			0.059** (0.029)
Observations	333	333	333
R^2	0.09740	0.11459	0.21871

Note: The terms ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. Each observation is weighted by the inverse of the standard error of the estimated $\overline{mort}_{rj}$ fixed effect.

the complications of Covid-19 are respiratory, this spillover seems plausible. Mortality for acute myocardial infarction predicts weakly higher Covid-19 mortality.

Finally, On-line Appendix Figure A5 shows the geographic distribution of hospital mortality without capacity strain. The four most populous regions—Austin, Dallas, Houston, and San Antonio—each have very large heterogeneity in their Covid-19 treatment quality, with some of the best and worst hospitals. On-line Appendix Figure A6 displays the analogous map at the realized quality levels accounting for capacity strain. It shows that the effects of capacity strain on degrading quality are prevalent throughout Texas.

4.3 Patient Hospital Choice Model Estimates

Having shown evidence that hospital quality varies over time due to capacity strain, we next seek to understand the extent to which patients use information about hospital quality in making their decisions of where to seek treatment. This will allow us to gain insight into the efficiency properties of this market (Chandra et al., 2016).

Table 5 presents our patient hospital choice model estimates, which are based on multinomial logit specifications from (3). Panel A shows results where patients' expectations of hospital mortality are given by quality when its ICU is empty, i.e. using $E[mort_{rtj}|occ_{rtj} = 0]$. Column (1) provides the simplest specification, where choices are a function of distance to the hospital and hospital mortality. We find that hospital mortality is significantly negative, with a coefficient of -3.331 and a coefficient on minutes traveled of -0.099.

The remaining specifications in panel A largely corroborate these results. Column (2) adds further measures of distance and columns (3) and (4) add more hospital characteristics. Columns (2)–(4) all show significantly negative coefficients on $E[mort_{rtj}]$. The non-quality estimated parameters are largely as expected across specifications. Patients prefer admission to hospitals with less travel time, with the closest hospital giving patients an additional increase in utility. Patients prefer hospitals that are system members, larger hospitals, and ones that do not have critical access status.

Column (5) adds an interaction between the patient's observed severity of illness quantile and distance.[29] It shows that more severely ill patients are less distance-sensitive than others, but the base coefficient on $E[mort_{rtj}]$ is -3.358, which is similar to the other specifications.

Despite the statistical significance, hospital quality when empty is a relatively weak predictor of patient choice compared to other hospital attributes. For example, using column (1), a one-minute travel time decrease is equivalent to a 3.0 percentage point lower mean death rate. Using estimates from Silva et al. (2023), this is equivalent to a social value of $\$106,600$, with some of this gain accruing to society and not to the patient, implying a lower value to the patient. Regardless, if patients fully believed our estimated mortality measures to be the true treatment effects, this would represent an implausibly low value of life lost. For instance, the $\$106,600$ figure is 426 times the \$250 estimated value in Gowrisankaran

[29] We calculate the patient's observed severity of illness using the posterior mean of $x'_{it}\gamma/\sqrt{1+\delta'\Sigma\delta}$.

Table 5: Impact of Quality Measures $E[mort_{rtj}]$ on Patient Hospital Choice

Panel A: Mortality evaluated when empty					
Regressor:	(1)	(2)	(3)	(4)	(5)
Hospital mortality, $E[mort_{rtj}\|occ_{rtj}=0]$	-3.331^{***}	-3.363^{***}	-3.652^{***}	-4.208^{***}	-3.358^{***}
	(0.086)	(0.087)	(0.086)	(0.090)	(0.087)
Distance (minutes)	-0.099^{***}	-0.122^{***}	-0.134^{***}	-0.134^{***}	-0.125^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Distance^2/1000		0.258^{***}	0.283^{***}	0.286^{***}	0.258^{***}
		(0.001)	(0.001)	(0.001)	(0.001)
Closest		0.116^{***}	0.122^{***}	0.113^{***}	0.116^{***}
		(0.009)	(0.009)	(0.009)	(0.009)
Major teaching			0.022^{***}	0.074^{***}	
			(0.007)	(0.007)	
System member			0.382^{***}	0.342^{***}	
			(0.009)	(0.010)	
Critical access			-1.915^{***}	-1.700^{***}	
			(0.022)	(0.023)	
For-profit			-0.352^{***}	-0.344^{***}	
			(0.004)	(0.005)	
Public			0.100^{***}	0.080^{***}	
			(0.011)	(0.011)	
Beds: ≥100 and <500			1.626^{***}	1.319^{***}	
			(0.008)	(0.009)	
Beds: ≥500			2.065^{***}	1.702^{***}	
			(0.008)	(0.010)	
COPD mortality				-3.350^{***}	
				(0.212)	
AMI mortality				7.732^{***}	
				(0.212)	
Severity × distance/100					0.483^{***}
					(0.043)
Panel B: Mortality at realized occupancy					
Regressor:	(1)	(2)	(3)	(4)	(5)
Hospital mortality, $E[mort_{rtj}\|occ_{rtj}=occ^{obs}_{rtj}]$	5.597^{***}	5.549^{***}	-1.570^{***}	-2.644^{***}	5.551^{***}
	(0.083)	(0.083)	(0.090)	(0.095)	(0.083)
Distance (minutes)	-0.099^{***}	-0.122^{***}	-0.134^{***}	-0.134^{***}	-0.125^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Distance^2/1000		0.257^{***}	0.283^{***}	0.286^{***}	0.256^{***}
		(0.001)	(0.001)	(0.001)	(0.001)
Closest		0.133^{***}	0.123^{***}	0.114^{***}	0.133^{***}
		(0.009)	(0.009)	(0.009)	(0.009)
Major teaching			0.048^{***}	0.085^{***}	
			(0.008)	(0.008)	
System member			0.387^{***}	0.354^{***}	
			(0.009)	(0.010)	
Critical access			-1.925^{***}	-1.708^{***}	
			(0.022)	(0.023)	
For-profit			-0.352^{***}	-0.341^{***}	
			(0.004)	(0.005)	
Public			0.091^{***}	0.083^{***}	
			(0.011)	(0.011)	
Beds: ≥100 and <500			1.642^{***}	1.329^{***}	
			(0.008)	(0.009)	
Beds: ≥500			2.070^{***}	1.702^{***}	
			(0.009)	(0.010)	
COPD mortality				-3.960^{***}	
				(0.212)	
AMI mortality				6.937^{***}	
				(0.211)	
Severity × distance/100					0.500^{***}
					(0.043)

Note: Hospital mortality estimates are from Table 3 column (3). The terms ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. Number of observations: $18,305,684$ in all columns. Panel B mortality corrects for ICU patients' choice affecting hospital mortality with a leave-one-out approach. Specifications with mortality for other diseases also include indicators for these variables being missing.

et al. (2015) for one minute of patient travel time. [30]

On the other hand, our results are consistent with patients being only partially informed about hospital quality. For instance, if patients formed posterior beliefs regarding hospital quality using a common prior across hospitals and signals that are our estimated measures of quality when empty, then they would optimally weight these signals less than known attributes, such as travel time. Indeed, a potential reason why patients may have reacted only weakly to hospital quality when empty is that they might have recognized that this information was imprecise, including by not accounting for capacity strain.

Next, panel B shows results where patients' quality expectations are given by the actual realized quality of a hospital given the capacity strain in the current quarter, $E[mort_{rtj}|occ_{rtj} = occ_{rtj}^{obs}]$, using the coefficient estimates from our base specification, Table 3 column (3). The results are starkly different. Column (1) shows a *positive* coefficient on mortality, with a coefficient of 5.597. If taken as causal, it shows that patients are choosing lower quality hospitals, all else equal. Columns (2) and (5) also show a positive coefficient on hospital mortality. However, columns (3) and (4), which include more hospital characteristics, report negative signs on $E[mort_{rtj}]$, albeit with relatively small coefficients of -1.570 and -2.644, respectively.

Combining panels A and B, the implications of the choice model analysis are striking. Patients are choosing the "right" hospitals (albeit with weak priors) if capacity strain were not important. In other words, patients' choices are consistent with them having a reasonable sense of which hospitals would have a high baseline quality for Covid-19. However, these choices do not appear to account for equilibrium quality degradation: high baseline quality hospitals are being chosen by many other Covid-19 patients, causing them to have capacity strain, which in turn causes them to no longer be high quality. Once equilibrium capacity strain is taken into account, patients are choosing the "wrong" hospitals. Section 5 further explores the equilibrium implications of the fact that patients are not incorporating capacity strain into their choice decisions.

[30] We obtain \$250 by converting the \$167 in 2005 dollars reported in this study to 2021 dollars.

5 Counterfactuals

Our findings show that inaccurate equilibrium expectations are playing a role in affecting Covid-19 quality. In particular, during the pandemic, patients were basing their expectations of hospital quality on quality in the absence of capacity strain, suggesting that alternative market structures or patient information could improve welfare. This section further considers the role of market alternatives and information on equilibrium outcomes.

Our concept of equilibrium is Bayesian Nash equilibrium. Specifically, we assume that patients make simultaneous choices of a hospital within their region. In some specifications, they perceive that quality will degrade from equilibrium capacity strain. Each patient observes her own latent utility shocks, ν_{ritj}, and uses the population distribution of other patients' shocks to predict the expected number of patients at each hospital and ultimately her probability of death conditional on a hospital choice.[31] We solve for an equilibrium by iterating on patient choices—which depend on quality—and updating the resulting quality until we reach a fixed point. We assume that a patient's length-of-stay and whether she is treated in the ICU would remain the same across her choices.

In all our counterfactuals, we start with the estimated parameters from Table 5 panel A column (5). We use these parameters because we view panel A as being more representative of consumer preferences than panel B, since they show that consumers value quality positively (albeit the wrong measure of quality).[32]

Table 6 reports results from six different counterfactual experiments, with each row displaying one experiment. First, we consider the baseline scenario. Our model's baseline prediction is that $54,651$ deaths occurred in sample and that patients traveled an average of 24.43 minutes from their home ZIP code to the hospital (both matching the data, by construction). The dollar welfare gain relative to the baseline is 0, by definition.

Second, we consider the (infeasible) situation of not having any capacity strain. The number of deaths falls to $49,414.7$, a 9.6% reduction. This would result in $18.69 billion

[31] As with our estimation results in Table 3 panel B, we let the patient incorporate the impact of her hospital choice on a hospital's capacity strain, with the assumption that her length-of-stay and whether she is treated in the ICU would remain constant across hospitals.

[32] While our estimation allowed for flexible time-varying impact of mortality through region-quarter fixed effects, for simplicity we assume here that treatment does not vary over time except from capacity strain. We calibrate the constant term to match the sample mean mortality.

Table 6: Counterfactuals on Market Structure and Information

Experiment:	Outcome: Total deaths	ΔMort. rate (%)	Mean travel time (min.)	ΔWelfare (bil. $)
(1) Baseline: patients use quality without strain	54,651.0	0.0	24.43	0.00
(2) No mortality increase from capacity strain	49,414.7	−9.6	24.43	18.69
(3) Five extra ICU beds per hospital	53,662.9	−1.8	24.43	3.53
(4) Patients informed about quality including strain	54,719.7	0.1	24.47	−0.25
(5) Informed about quality incl. strain with 10× weight	52,598.4	−3.8	25.15	7.33
(6) Baseline information but with 10× weight on quality	56,703.4	3.8	25.11	−7.33

Note: Deaths are predicted totals across our sample of 397,631 patients. Change in mortality rates are in sample, relative to baseline. Welfare changes are in sample, relative to baseline, and from mortality reductions only. Calculations use Table 5, panels A and B, column (5).

in welfare gains for patients in our sample. Third, we consider adding 5 ICU beds to every hospital in Texas, and thus reducing ICU capacity strain. This would also reduce mortality significantly, reducing the number of deaths by 988 (1.8%) in our sample. This would result in $3.53 billion in welfare increases from lower mortality just for these individuals.

For our fourth experiment, we continue to assume that the true preferences are given by Table 5 panel A column (5), but we now assume that patients were misinformed about hospital quality and that they actually valued quality given the realized capacity strain and not baseline quality when empty. Two countervailing forces determine the overall mortality impact. First, there is a direct effect: better information induces patients to select higher-quality hospitals. This effect reduces mortality by 86 lives, or 0.16%. Second, the reallocation of patients across hospitals generates congestion externalities. Outflows from initially chosen hospitals reduce capacity strain there, benefiting remaining patients, while

inflows to newly selected hospitals increase strain, adversely affecting incumbent patients. The net externality—whose sign depends on the slope of the mortality–capacity relationship and patient volumes—raises mortality by 155 lives or 0.28%. in this counterfactual. As a result, despite improved information and individually optimal re-sorting, equilibrium mortality *increases* slightly, by 69 lives.

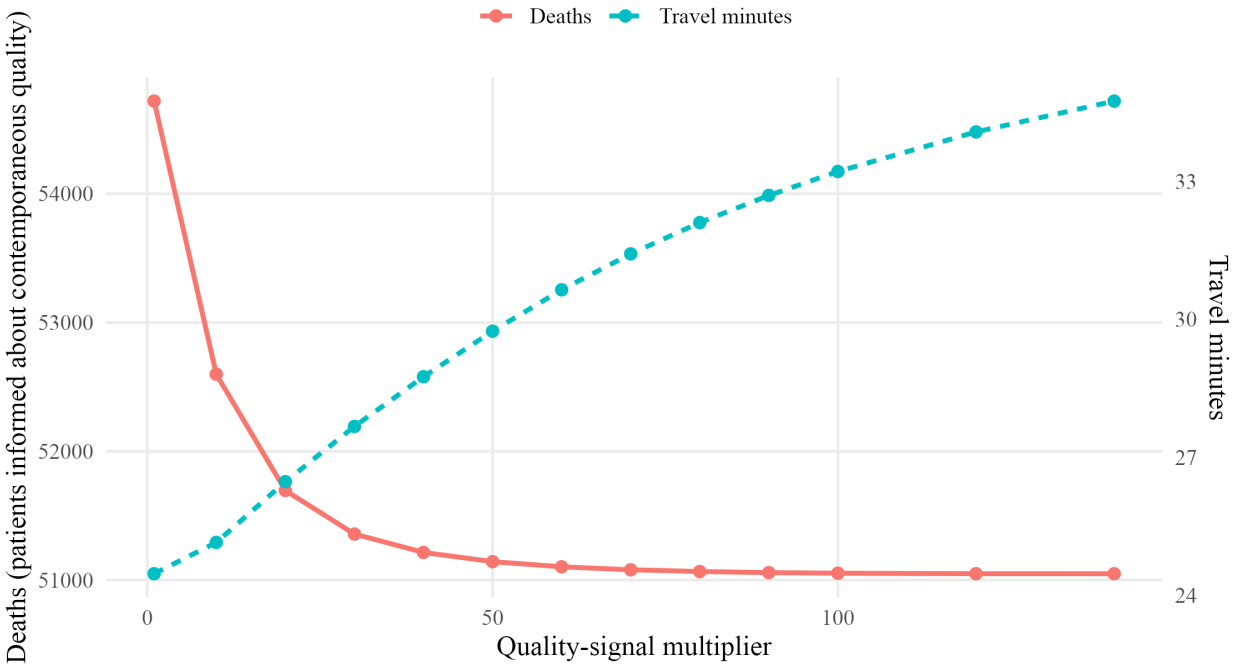
As discussed in Section 4.3, the relatively low value placed on quality relative to distance is consistent with patients believing that their hospital quality information is relatively imprecise or inaccurate. For this reason, we also consider increasing the weight that patients place on quality. Specifically, our fifth experiment is similar to the fourth experiment, except that it multiplies the coefficient on quality by ten relative to the estimates, corresponding to patients perceiving a higher precision on the quality signal, though lower than full information. The combination of more accurate information and more weight given to this information leads to quantitatively significant decreases in mortality. Specifically, in-sample mortality drops by 3.8%—which is worth $7.33 billion in mortality declines—due to less equilibrium capacity strain. Travel time increases modestly, going from 24.43 to 25.15. minutes relative to the baseline. As in Counterfactual 4, we decompose the impact of the more precise quality information on mortality into two parts: the direct impact on choice and the externality impact. The direct choice effect reduces mortality by 8.5%, while the externality effect increases mortality by 4.8%.

Our sixth experiment then asks what would occur if patients had valued quality ten times as much as in our baseline estimates but still erroneously believed that hospital quality levels were given by their levels in the absence of capacity strain. Here, even more patients seek care at these *ex ante* good hospitals, not accounting for the fact that they are hugely increasing capacity strain. This more precise but inaccurate information causes mortality to *rise* by 3.8% and travel time to increase by about 0.7 minutes relative to the baseline.

Finally, Figure 5 investigates the impact of further increases in the precision of the quality signal, assuming that patients make hospital choices using the hospital's actual quality, as in our fourth and fifth experiments from Table 5. We consider increasing precision up to almost the value that would indicate that patients believed the quality signals to be full information, as discussed in Section 4.3. The biggest drops in mortality occur when the signal precision

increases by 10 or 20 times relative to the baseline. Specifically, at 20 times precision, mortality drops by 5.5% relative to a signal of weight 1 (which is the fourth specification of Table 5). Further increases beyond this yield only small drops in equilibrium mortality but relatively large increases in travel time.

Figure 5: Impact of Increasing Information Precision when Signal Includes Strain



Our overall takeaway is that better information can be helpful if patients strongly believe and react to that information. But, if patients were to react more strongly to information about baseline characteristics—and, as we find, these characteristics do not reflect equilibrium quality—this may actually reduce welfare.

6 Conclusion

We study the impact of congestion on equilibrium institutional outcomes, focusing on hospitals during the Covid-19 pandemic. In many markets, including those for schools, transportation systems, and digital platforms, users who benefit more may sort toward institutions perceived to be high quality, but their presence at these institutions changes overall quality. Both

measuring quality in the presence of congestion and selection based on unobservable individual attributes and understanding the impact of information provision on equilibrium interactions are central to designing policies that improve welfare. Thus, our framework applies more generally to these settings.

We offer three takeaways from our results. First, congestion is a first-order determinant of hospital performance in environments with limited capacity. ICU capacity strain had large effects on mortality during the pandemic that were not offset by learning-by-doing. As a result, hospitals that were perceived to be high quality attracted additional patients precisely when congestion was eroding their performance. Lowering a hospital's ICU utilization rates from 200% to 100% of capacity would reduce its mortality rates by 1.2 percentage points. If hospital performance were unaffected by capacity strain, inpatient mortality would fall by 9.6%, worth $251.2 billion if applied nationally to Covid-19 patients during the pandemic.

Second, patients appear to be at least somewhat informed about long-run quality, but not about the short-run impact of congestion. Specifically, patients' hospital choices are positively affected by baseline hospital quality in the absence of congestion, including after controlling for a variety of other factors that might affect these decisions. However, their choices do not have any systematic relationship with equilibrium hospital quality measures that include congestion.

Finally, the equilibrium interactions between quality and congestion imply that patient re-sorting can generate substantial spillovers across hospitals, so that individually rational choices need not improve aggregate outcomes. In particular, if patients used information about true quality including capacity strain to make their hospital choices, deaths would actually have *risen* by 0.1% due to these equilibrium spillovers. However, if patients used this information and weighted it ten times as much as they currently weight quality information, then in-sample Covid-19 deaths would have been reduced by 3.8%, implying $99 billion in welfare gains when scaled nationally.

More broadly, our research demonstrates an approach that can jointly account for congestion and selection even when outcomes are observed infrequently at the institution-period level. Approaches similar to ours may prove useful in a wide range of settings where congestion and selection are important, particularly in cases where sparse outcome data

have traditionally limited empirical analysis. Our findings also underscore the importance of incorporating both congestion and information into analyses of hospital quality during the pandemic, suggesting that these mechanisms may also be central to understanding the effects of public policies in other settings.

References

Abaluck, J., Caceres Bravo, M., Hull, P., and Starc, A. (2021). Mortality Effects and Choice Across Private Health Insurance Plans. *The Quarterly Journal of Economics*, 136(3):1557–1610.

Abir, M., Nelson, C. D., Chan, E. W.-M., Al Ibrahim, H., Cutter, C., Patel, K., and Bogart, A. (2020). *Critical Care Surge Response Strategies for the 2020 COVID-19 Outbreak in the United States*. Rand.

Ahmad, F. B. (2023). Covid-19 Mortality Update—United States, 2022. *MMWR. Morbidity and Mortality Weekly Report*, 72.

Allcott, H., Boxell, L., Conway, J., Ferguson, B., Gentzkow, M., and Goldman, B. (2024). What Explains Temporal and Geographic Variation in the Early US Coronavirus Pandemic? *Review of Economic Design*, pages 1–41.

Allende, C., Gallego, F., and Neilson, C. (2026). Approximating the Equilibrium Effects of Informed School Choice. Technical report.

Alsan, M., Chandra, A., and Simon, K. (2021). The Great Unequalizer: Initial Health Effects of COVID-19 in the United States. *Journal of Economic Perspectives*, 35(3):25–46.

Asch, D. A., Sheils, N. E., Islam, M. N., Chen, Y., Werner, R. M., Buresh, J., and Doshi, J. A. (2021). Variation in US Hospital Mortality Rates for Patients Admitted With COVID-19 During the First 6 Months of the Pandemic. *JAMA Internal Medicine*, 181(4):471–478.

Avdic, D., Lundborg, P., and Vikström, J. (2024). Does Health Care Consolidation Harm Patients? Evidence from Maternity Ward Closures. *American Economic Journal: Economic Policy*, 16(1):160–189.

Birkmeyer, J. D., Siewers, A. E., Finlayson, E. V., Stukel, T. A., Lucas, F. L., Batista, I., Welch, H. G., and Wennberg, D. E. (2002). Hospital Volume and Surgical Mortality in the United States. *New England Journal of Medicine*, 346(15):1128–1137.

Brand, K., Garmon, C., and Rosenbaum, T. (2023). In the Shadow of Antitrust Enforcement: Price Effects of Hospital Mergers from 2009 to 2016. *The Journal of Law and Economics*, 66(4):639–669.

Brown, Z. Y. and Jeon, J. (2024). Endogenous Information and Simplifying Insurance Choice. *Econometrica*, 92(3):881–911.

Card, D., Fenizia, A., and Silver, D. (2023). The Health Impacts of Hospital Delivery Practices. *American Economic Journal: Economic Policy*, 15(2):42–81.

Castillo, J. C., Knoepfle, D., and Weyl, E. G. (2025). Matching and Pricing in Ride Hailing: Wild Goose Chases and How to Solve Them. *Management Science*, 71(5):4377–4395.

Chandra, A., Finkelstein, A., Sacarny, A., and Syverson, C. (2016). Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector. *American Economic Review*, 106(8):2110–2144.

Chandra, A. and Staiger, D. O. (2020). Identifying Sources of Inefficiency in Healthcare. *The Quarterly Journal of Economics*, 135(2):785–843.

Chernew, M., Gowrisankaran, G., and Scanlon, D. P. (2008). Learning and the Value of Information: Evidence from Health Plan Report Cards. *Journal of Econometrics*, 144(1):156–174.

Clarke, K. E. (2022). Seroprevalence of Infection-Induced SARS-CoV-2 Antibodies - United States, September 2021–February 2022. *MMWR. Morbidity and Mortality Weekly Report*, 71.

Coviello, D., Ichino, A., and Persico, N. (2010). Don't Spread Yourself Too Thin: The Impact of Task Juggling on Workers' Speed of Job Completion. Technical report, National Bureau of Economic Research.

Cutler, D. M. (2007). The Lifetime Costs and Benefits of Medical Technology. *Journal of Health Economics*, 26(6):1081–1100.

Dobbie, W., Goldin, J., and Yang, C. S. (2018). The Effects of Pretrial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges. *American Economic Review*, 108(2):201–240.

Doyle, J. J., Graves, J., and Gruber, J. (2019). Evaluating Measures of Hospital Quality: Evidence from Ambulance Referral Patterns. *Review of Economics and Statistics*, 101(5):841–852.

Doyle, J. J., Graves, J. A., Gruber, J., and Kleiner, S. A. (2015). Measuring Returns to Hospital Care: Evidence from Ambulance Referral Patterns. *Journal of Political Economy*, 123(1):170–214.

Dranove, D., Kessler, D., McClellan, M., and Satterthwaite, M. (2003). Is More Information Better? The Effects of "Report Cards" on Health Care Providers. *Journal of Political Economy*, 111(3):555–588.

Einav, L., Finkelstein, A., and Mahoney, N. (2025). Producing Health: Measuring Value Added of Nursing Homes. *Econometrica*, 93(4):1225–1264.

Elixhauser, A., Steiner, C., Harris, D. R., and Coffey, R. M. (1998). Comorbidity Measures For Use With Administrative Data. *Medical Care*, 36(1):8–27.

Eriksson, C. O., Stoner, R. C., Eden, K. B., Newgard, C. D., and Guise, J.-M. (2017). The Association Between Hospital Capacity Strain and Inpatient Outcomes in Highly Developed Countries: A Systematic Review. *Journal of General Internal Medicine*, 32(6):686–696.

Gaynor, M., Propper, C., and Seiler, S. (2016). Free to Choose? Reform, Choice, and Consideration Sets in the English National Health Service. *American Economic Review*, 106(11):3521–3557.

Gaynor, M., Seider, H., and Vogt, W. B. (2005). The Volume-Outcome Effect, Scale Economies, and Learning-by-Doing. *American Economic Review*, 95(2):243–247.

Geruso, M., Layton, T. J., and Wallace, J. (2023). What difference does a health plan make? evidence from random plan assignment in medicaid. *American Economic Journal: Applied Economics*, 15(3):341–379.

Geweke, J., Gowrisankaran, G., and Town, R. J. (2003). Bayesian Inference for Hospital Quality in a Selection Model. *Econometrica*, 71(4):1215–1238.

Gottlieb, R. L., Vaca, C. E., Paredes, R., Mera, J., Webb, B. J., Perez, G., Oguchi, G., Ryan, P., Nielsen, B. U., Brown, M., et al. (2022). Early Remdesivir to Prevent Progression to Severe Covid-19 in Outpatients. *New England Journal of Medicine*, 386(4):305–315.

Gowrisankaran, G., Ho, V., and Town, R. J. (2006). Causality, Learning and Forgetting in Surgery. *Federal Trade Commission, draft*.

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

Gowrisankaran, G. and Town, R. J. (1999). Estimating the Quality of Care in Hospitals Using Instrumental Variables. *Journal of Health Economics*, 18(6):747–767.

Gulick, R. M., Pau, A. K., Daar, E., Evans, L., Gandhi, R. T., Tebas, P., Ridzon, R., Masur, H., Lane, H. C., Panel, N. C.-. T. G., et al. (2024). National Institutes of Health COVID-19 Treatment Guidelines Panel: Perspectives and Lessons Learned. *Annals of Internal Medicine*, 177(11):1547–1557.

Hoe, T. P. (2022). Does Hospital Crowding Matter? Evidence from Trauma and Orthopedics in England. *American Economic Journal: Economic Policy*, 14(2):231–262.

Hoxby, C. M. (2000). The Effects of Class Size on Student Achievement: New Evidence from Population Variation. *The Quarterly Journal of Economics*, 115(4):1239–1285.

Hull, P. (2018). Estimating Hospital Quality with Quasi-Experimental Data. *Available at SSRN 3118358*.

IDSA (2021). IDSA Guidelines on the Treatment and Management of Patients with COVID-19. *Infectious Diseases Society of America: Arlington, VA, USA*.

Kadri, S. S., Sun, J., Lawandi, A., Strich, J. R., Busch, L. M., Keller, M., Babiker, A., Yek, C., Malik, S., Krack, J., et al. (2021). Association Between Caseload Surge and COVID-19 Survival in 558 U.S. Hospitals, March to August 2020. *Annals of Internal Medicine*, 174(9):1240–1251.

Kuntz, L., Mennicken, R., and Scholtes, S. (2015). Stress on the Ward: Evidence of Safety Tipping Points in Hospitals. *Management Science*, 61(4):754–771.

Kunz, J. and Propper, C. (2023). Is Hospital Quality Predictive of Pandemic Deaths? Evidence from US Counties. *Journal of Urban Economics*, 133(2).

Leuchter, R. K., Delarmente, B. A., Vangala, S., Tsugawa, Y., and Sarkisian, C. A. (2025). Health Care Staffing Shortages and Potential National Hospital Bed Shortage. *JAMA Network Open*, 8(2):e2460645–e2460645.

Luft, H. S., Garnick, D. W., Mark, D. H., Peltzman, D. J., Phibbs, C. S., Lichtenberg, E., and McPhee, S. J. (1990). Does Quality Influence Choice of Hospital? *JAMA*, 263(21):2899–2906.

Luft, H. S. and Hunt, S. S. (1986). Evaluating Individual Hospital Quality Through Outcome Statistics. *JAMA*, 255(20):2780–2784.

Luft, H. S., Hunt, S. S., and Maerki, S. C. (1987). The Volume-Outcome Relationship: Practice-Makes-Perfect or Selective-Referral Patterns? *Health Services Research*, 22(2):157.

Malone, J. B., Nevo, A., and Williams, J. W. (2021). The Tragedy of the Last Mile: Economic Solutions to Congestion in Broadband Networks.

Marquardt, K., Gowrisankaran, G., and Town, R. (2025). GGT: Stata Module to Implement Geweke, Gowrisankaran, and Town Model Quality Estimator.

McClellan, M., McNeil, B. J., and Newhouse, J. P. (1994). Does More Intensive Treatment of Acute Myocardial Infarction in the Elderly Reduce Mortality? Analysis Using Instrumental Variables. *JAMA*, 272(11):859–866.

Mehta, P., McAuley, D. F., Brown, M., Sanchez, E., Tattersall, R. S., and Manson, J. J. (2020). COVID-19: Consider Cytokine Storm Syndromes and Immunosuppression. *The Lancet*, 395(10229):1033–1034.

Olenski, A. and Sacher, S. (2024). Estimating Nursing Home Quality with Selection. *The Review of Economics and Statistics*, pages 1–31.

Pisano, A., Yavorovskiy, A., Verniero, L., and Landoni, G. (2020). Indications for Tracheal Intubation in Patients With Coronavirus Disease 2019 (COVID-19). *Journal of Cardiothoracic and Vascular Anesthesia*, 35(5):1276.

Pope, D. G. (2009). Reacting to Rankings: Evidence from "America's Best Hospitals". *Journal of Health Economics*, 28(6):1154–1165.

Richardson, S., Hirsch, J. S., Narasimhan, M., Crawford, J. M., McGinn, T., Davidson, K. W., Barnaby, D. P., Becker, L. B., Chelico, J. D., Cohen, S. L., et al. (2020). Presenting Characteristics, Comorbidities, and Outcomes Among 5700 Patients Hospitalized With COVID-19 in the New York City Area. *JAMA*, 323(20):2052–2059.

Rubin, R. (2022). Baricitinib Is First Approved COVID-19 Immunomodulatory Treatment. *JAMA*, 327(23):2281–2281.

Silk, B. J., Scobie, H. M., and Duck, William M., e. a. (2023). COVID-19 Surveillance After Expiration of the Public Health Emergency Declaration — United States, May 11, 2023. *MMWR. Morbidity and Mortality Weekly Report*, 72:523–528.

Silva, S., Goosby, E., and Reid, M. J. (2023). Assessing the Impact of One Million COVID-19 Deaths in America: Economic and Life Expectancy Losses. *Scientific Reports*, 13(1):3065.

Singh, M. and Venkataramani, A. (2026). Rationing by Race.

Vatter, B. (2025). Quality Disclosure and Regulation: Scoring Design in Medicare Advantage. *Econometrica*, 93(3):959–1001.

Wu, C., Chen, X., Cai, Y., Zhou, X., Xu, S., Huang, H., Zhang, L., Zhou, X., Du, C., Zhang, Y., et al. (2020). Risk Factors Associated With Acute Respiratory Distress Syndrome and

Death in Patients With Coronavirus Disease 2019 Pneumonia in Wuhan, China. *JAMA Internal Medicine*, 180(7):934–943.

Appendix

A Additional Tables and Figures

Table A1: Summary Statistics: Hospitals

	Full Sample	AHA Match Sample	AHA Unmatch Sample
Total beds	205.357	197.507	262.244
Total ICU beds	22.673	22.664	23.400
Overall utilization rate (avg)	0.410	0.419	0.350
Adj. ICU utilization rate (avg)	1.251	1.292	1.012
Major teaching hospital	0.036	0.042	0.000
For-profit	0.390	0.364	0.556
Public	0.204	0.238	0.000
System member	0.709	0.678	0.911
Rural	0.147	0.168	0.022
Critical access	0.144	0.168	0.000
CCN Duplicate	0.201	0.084	0.911
AMI mortality rate	0.129	0.128	0.133
COPD mortality rate	0.084	0.084	0.083
CMS 1 star rating	0.022	0.025	0.000
CMS 2 star rating	0.201	0.189	0.278
CMS 3 star rating	0.354	0.370	0.250
CMS 4 star rating	0.270	0.286	0.167
CMS 5 star rating	0.153	0.130	0.306
Per-capita Medicare spending ($)	9,615	9,668	9,320
N Hospitals	333	286	45

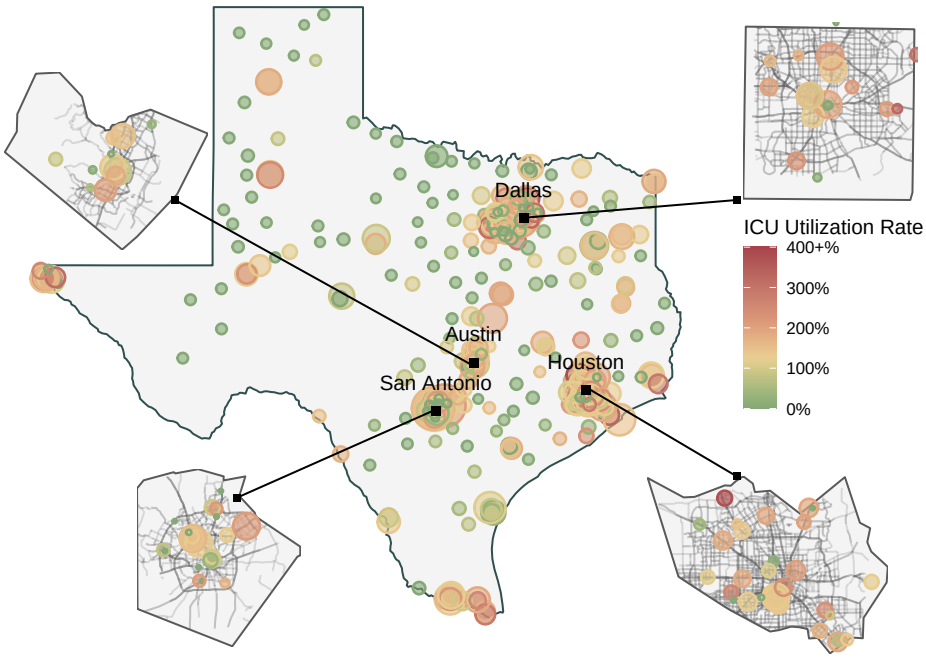
Note: Table of hospital characteristics means (excluding those with missing values) for the full sample of hospitals, for those with a AHA Match, and for those without a AHA match in which case hospital characteristics were determined via alternative sources. See Appendix B for additional details on how we construct hospital bed measures. Hospital-specific utilization rates are based on the within-hospital average across quarters in the sample. Overall utilization rate is defined as the sum of patient length of stay divided by total bed-days in the quarter. Adjusted ICU utilization rate is defined as the sum of ICU patient length of stay divided by ICU bed-days in the quarter, where hospitals without any ICU beds are assigned 1 ICU bed, and their utilization rate is capped at 150%.

Table A2: Transfers and Hospital Capacity Strain

	Transfer In Rate			Transfer Out Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
Overall utilization rate	0.0002	-0.024	-0.060	-0.004	0.014	-0.026
	(0.029)	(0.032)	(0.076)	(0.029)	(0.023)	(0.052)
ICU utilization rate	-0.002	0.007	0.017**	-0.005	-0.003	0.0002
	(0.004)	(0.004)	(0.008)	(0.006)	(0.004)	(0.007)
Weights	-	# patients	total beds	-	# patients	total beds
Observations	3,066	3,066	3,066	3,066	3,066	3,066
R^2	0.667	0.787	0.773	0.468	0.661	0.431

Note: All regressions include hospital and region-by-quarter fixed effects, and the unit of observation is at the hospital-quarter level. The dependent variable for the first three columns is the hospital-quarter *transfer in* rate (i.e., the share of Covid-19 patients in that hospital-quarter that were transferred in from another hospital). The dependent variable for the last three columns is the hospital-quarter *transfer out* rate (i.e., the share of Covid-19 patients in that hospital-quarter that were transferred out to another facility). Each model corresponds to a different regression weighting scheme- unweighted, weighted by the total number of patients admitted to the hospital in the quarter, or weighted by the total number of beds in the hospital. We calculate standard errors with two-way clustering at the hospital and region-by-quarter levels. The terms ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Figure A1: Hospital ICU Utilization Rate, Q2:2020



Note: Map of hospital ICU utilization rate during Q2:2020. Each circle represents one hospital and has area proportional to total bed size.

Table A3: Impact of Patient Characteristics on Mortality, γ

	Panel A			Panel B	
Risk adjuster	Coef.	SD	Risk adjuster (cont.)	Coef.	SD
Constant	-1.016	0.326	Pct Hispanic	-0.290	0.231
Male	0.138	0.138	Median Income	-0.301	0.234
White (Non-Hispanic)	-0.105	0.165	Dx: Blood loss anemia	-0.171	0.257
Black (Non-Hispanic)	-0.195	0.192	Dx: Coagulopathy	0.419	0.210
Hispanic	-0.084	0.228	Dx: COPD	-0.014	0.162
Age 25	-0.428	0.197	Dx: Depression	-0.088	0.204
Age 30	-0.345	0.244	Dx: Diabetes type 1	-0.079	0.175
Age 35	-0.315	0.253	Dx: Diabetes type 2	0.061	0.191
Age 40	-0.220	0.208	Dx: Fluid & electrolyte disorders	0.455	0.200
Age 45	-0.113	0.209	Dx: Hypertension	-0.206	0.217
Age 50	-0.045	0.241	Dx: Hypothyroidism	-0.032	0.170
Age 55	0.051	0.270	Dx: Iron deficiency	0.138	0.239
Age 60	0.225	0.242	Dx: Kidney disease	0.110	0.194
Age 65	0.351	0.320	Dx: Liver disease	0.000	0.180
Age 70	0.382	0.367	Dx: Lymphoma	0.019	0.268
Age 75	0.546	0.347	Dx: Metastatic cancer	0.223	0.322
Age 80+	0.812	0.385	Dx: Neurological disorders	0.118	0.203
Medicaid	-0.190	0.315	Dx: Obesity	0.069	0.244
Private (Managed Care)	-0.328	0.338	Dx: Paralysis	0.131	0.245
Medicare Advantage	-0.233	0.379	Dx: Peripheral artery disease	0.054	0.195
Blue Cross (indemnity)	-0.218	0.257	Dx: Psychoses	-0.085	0.215
Medicare	-0.291	0.342	Dx: Pulmonary circ. disorders	0.153	0.256
Charity / self-pay	-0.297	0.292	Dx: Peptic ulcer disease	-0.086	0.325
Vac. rate	-0.288	0.280	Dx: Solid tumor	0.082	0.220
Transfer from SNF	0.346	0.452	Dx: Valvular disease	-0.099	0.240
Pct 65+	-0.134	0.226	Dx: Weight loss	0.257	0.246
Pct Non-Hispanic White	-0.313	0.220	Dx: Rheumatoid arthritis	-0.030	0.231
Pct Non-Hispanic Black	-0.192	0.225	Dx: Chronic condition (other)	0.176	0.189

Note: Each entry provides mean of the posterior means of a risk adjuster, across sample regions and quarters. The standard deviations report the variation in these posterior means across sample regions and quarters.

Table A4: Impact of Occupancy on Hospital Mortality, Alternative Bed Source Data

Dependent Variable:	Hospital expected mortality draw: $mort_{rtjd}$			
Bed Data Source:	Baseline	AHA Only	HCRIS Only	TX Directory Only
Regressor:				
ICU utilization rate spline: $[0, 0.75)$	0.008	0.020**	0.017*	0.015*
	(0.008)	(0.009)	(0.009)	(0.008)
ICU utilization rate spline: $[0.75, \infty)$	0.012***	0.007***	0.009***	0.010***
	(0.002)	(0.002)	(0.002)	(0.002)
Hospital/quarter/ draw observations	2,752,200	2,383,200	2,738,700	2,671,200
R^2	0.480	0.485	0.480	0.479

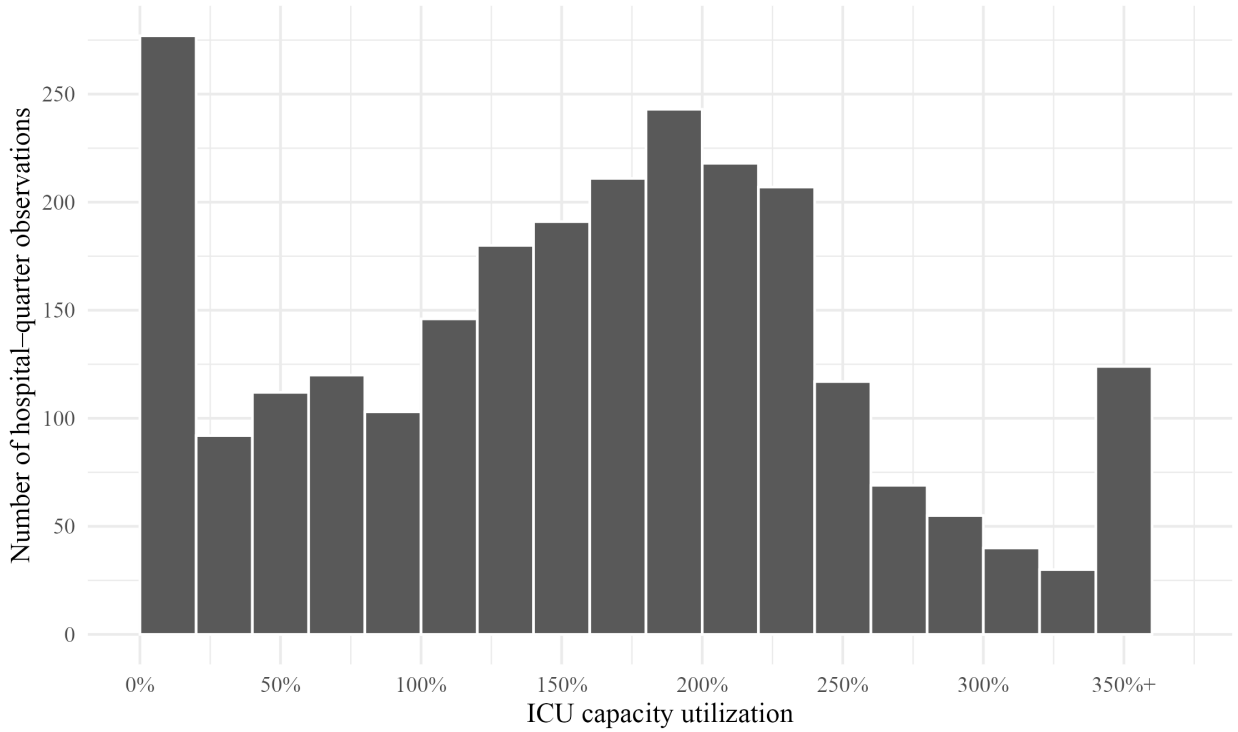
Note: All regressions include hospital and region-by-quarter fixed effects. We calculate standard errors with two-way clustering at the hospital and region-by-quarter levels. The terms ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. See Appendix B for additional details on how we construct hospital bed measures, which corresponds to the "Baseline" column. For each of the remaining columns, the ICU utilization rate is based on ICU beds as reported from the respective source, and thus will not include hospitals when either the hospital or their bed information is missing.

Table A5: Mortality Variation from Capacity Strain

Measure	Mean	Std. Dev.	25th percentile	75th percentile
Mortality at mean sample capacity use	0.154	0.030	0.131	0.175
Mortality when empty	0.140	0.031	0.117	0.164
Mortality when use is at 100% of capacity	0.150	0.031	0.126	0.173
Mortality when use is at 150% of capacity	0.155	0.031	0.132	0.179

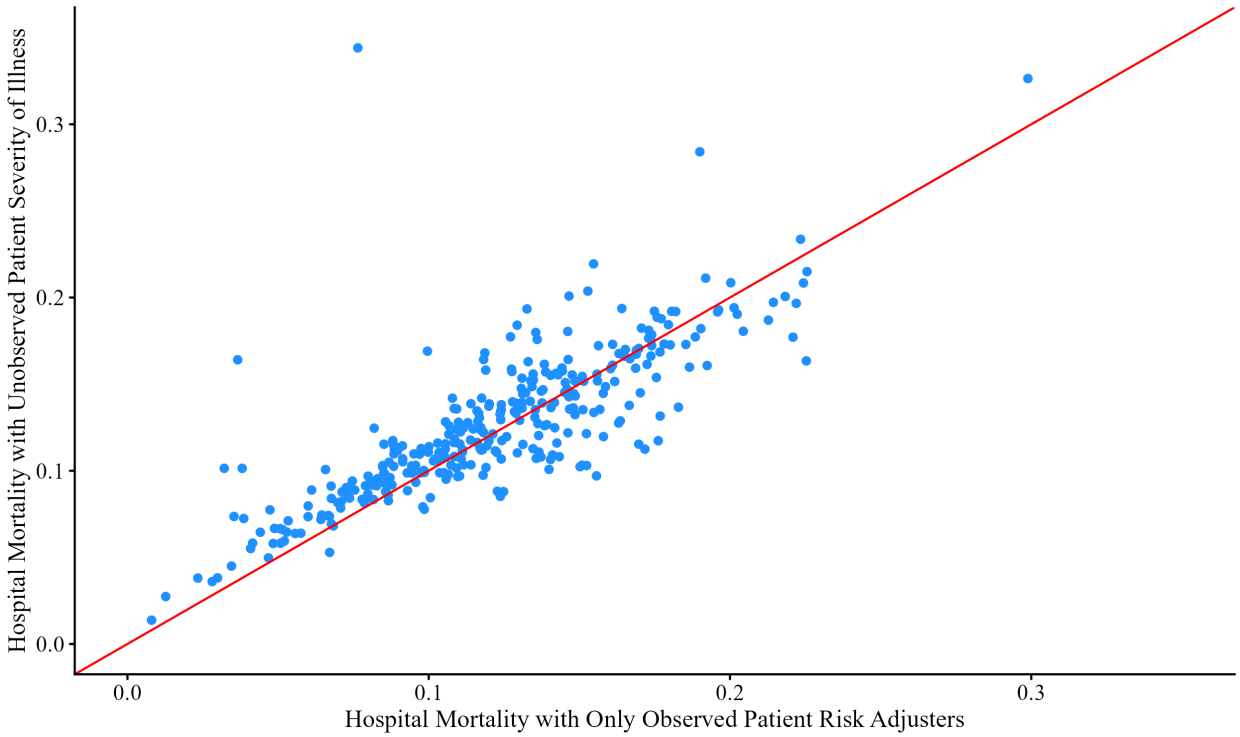
Note: Table based on authors' estimation sample and regression estimates presented in Table 3 column (3). We calculate the reported numbers using the unweighted sample of estimated hospital-quarter mortality levels. Mean sample capacity use is also unweighted.

Figure A2: ICU Capacity Utilization by Hospital-Quarter



Note: Histogram of hospital-quarter ICU utilization. Each observation is one hospital-quarter in our data. The figure does not display hospitals with 0 ICU beds.

Figure A3: Effect of Controlling for Unobserved Patient Severity on Mortality $mort_{rtj}$



Note: Each point denotes a hospital. The x-axis corresponds to the posterior mean of the hospital fixed effect (transformed to mortality rate for interpretation), obtained from the GGT estimation that does not allow for unobserved patient selection, i.e. with $\delta = 0$. The y-axis corresponds to the posterior mean of the hospital fixed effect (again transformed to mortality rate), obtained from the GGT estimation that adjusts for patient selection, i.e. with estimated values for δ. Both models include the same observed patient risk adjusters, x_{it} from Section 3.1.

Figure A4: Visualized Spline Coefficients from Table 3

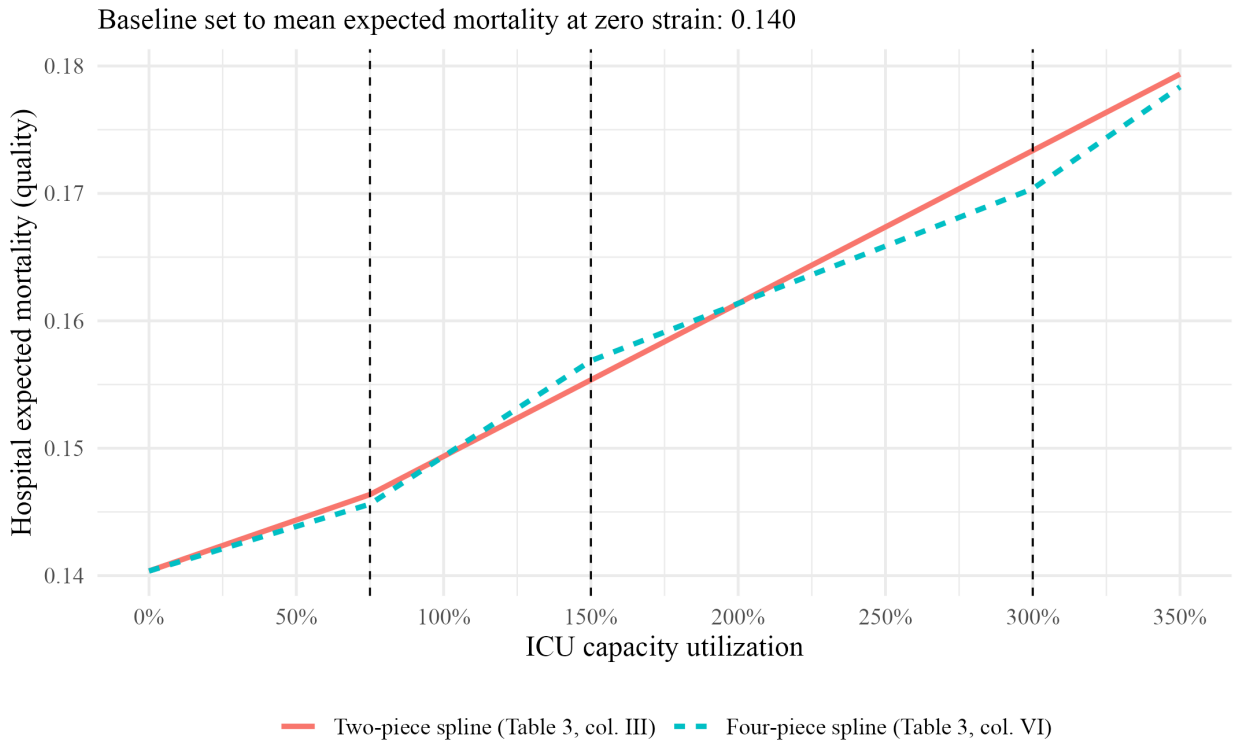
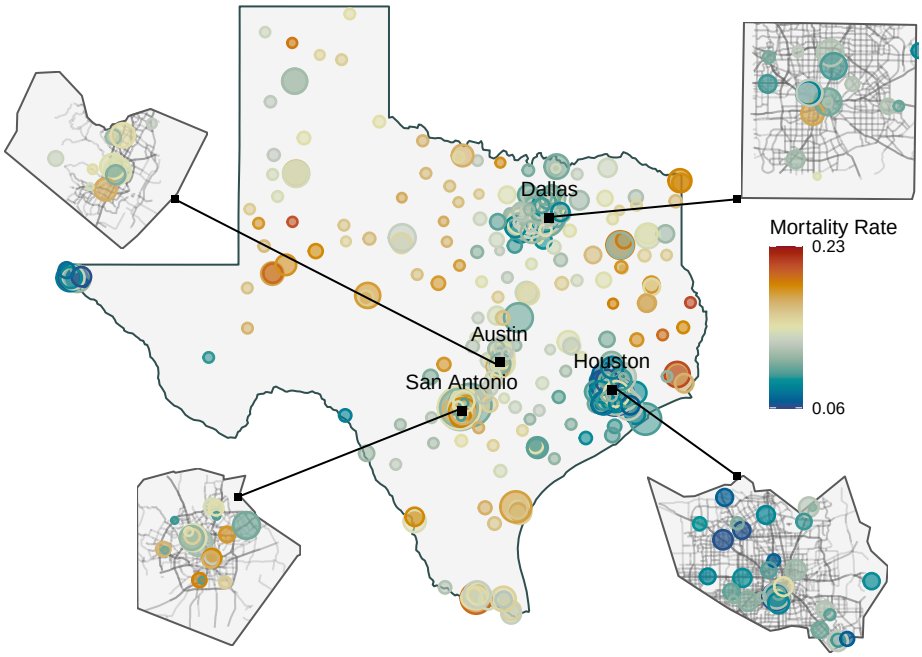
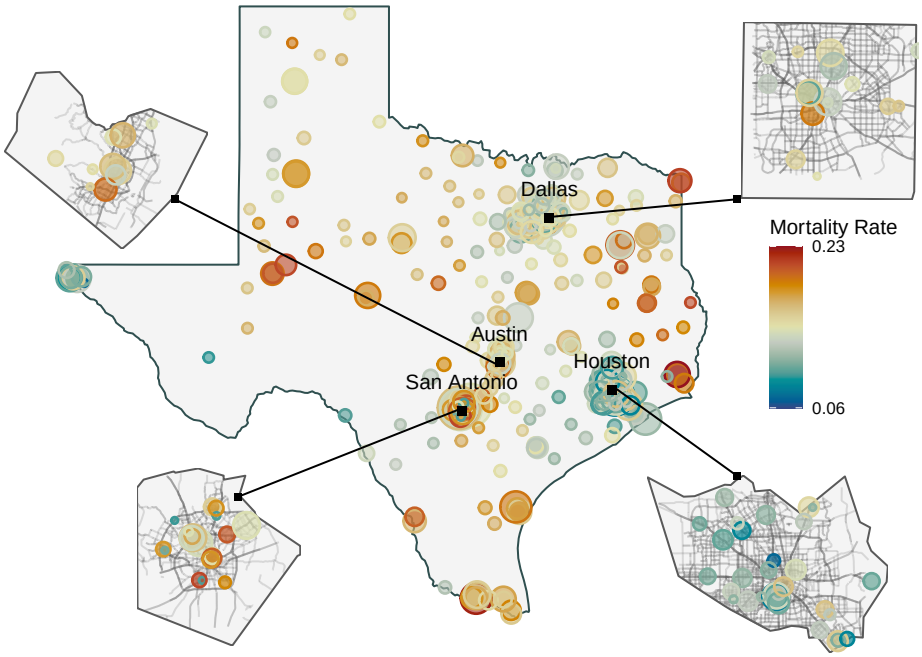


Figure A5: Hospital Mortality Across Hospitals Without Capacity Strain



Note: Map based on authors' estimation sample and regression estimates presented in Table 3 column (3). Each circle represents one hospital, and the circle area is proportional to total bed size.

Figure A6: Hospital Mortality Across Hospitals at Observed Occupancy Levels



Note: Map based on authors' estimation sample and regression estimates presented in Table 3 column (3). Each circle represents one hospital and has area proportional to total bed size.

B Construction of Hospital Bed Measures

We define overall utilization rate in a hospital-quarter as the sum of patient length-of-stay days divided by total bed-days (i.e., total beds times 91 days) and ICU utilization rate analogously as the sum of ICU length-of-stay days divided by ICU bed-days. The numerators of these measures come from the Texas patient discharge data, summing the total length of stay for all patients (or for ICU patients) in the hospital-quarter. The denominators require constructing a single hospital-level value for total beds and for ICU beds, which is the focus of this appendix. We first describe the underlying bed data sources, then the two-step algorithm we use to combine them.

We consider four sources of hospital-level total and ICU bed counts: (1) the American Hospital Association (AHA) annual survey for 2020, (2) the CMS Healthcare Cost Reports (HCRIS), (3) the January 2026 Texas Health and Human Services Directory of General and Special Hospitals, and (4) manual look-up for those missing bed size information in all other sources. We merge our hospitals to the AHA source using manual name match. In most cases, this gives us the CMS Certification Number (CCN), which we use, along with name, to merge with the other data sources.[33]

Our first algorithm assigns each hospital a single benchmark value for total beds and ICU beds by selecting across the four sources in a fixed order. For both total beds and ICU beds, we begin with the AHA value when the hospital has an AHA match, and the AHA value is non-missing and non-zero. When the AHA value is missing or zero, we use the HCRIS value provided the CCN is not flagged as a duplicate. If still not assigned a value, we use the Texas Directory value if provided, and use the hand-search value only if no other sources were available.

Our second algorithm checks the benchmark assignment against the hospital's quarterly utilization. The intuition is that if the benchmark bed count is too high or too low relative to the hospital's observed patient utilization days, the resulting capacity strain will appear

[33]For the HCRIS source, we use the merged synthetic calendar year data up to 2024 from `https://github.com/asacarny/hospital-cost-reports`. When assigning beds, we use the 2021 value when available and fall back to the nearest surrounding year otherwise. Because some CCNs report on a consolidated fashion, a single CCN can be mapped to multiple hospitals. Given this, it is possible that bed size reported in HCRIS represents the sum of beds (and ICU beds) across multiple hospitals. Where possible, we confirm bed size using HCRIS name match as well; otherwise we flag the CCN as a duplicate.

extreme relative to other hospitals in the same quarter, suggesting the benchmark bed size is inaccurate, and one of the alternative bed sources may yield a more plausible value. We implement this check separately for total beds and for ICU beds.

For each hospital-quarter jt and each candidate source $s \in \{$benchmark, AHA, HCRIS, Texas Directory$\}$, we compute the implied utilization rate, $\text{utilization_rate}_{jts} = \text{LOS}_{jt}/(\text{beds}_{js} \times 91)$, dropping quarters with zero length of stay as they are uninformative for bed size. We then construct a leave-one-out reference defined as the mean $\overline{u}_{(-j)t}$ and standard deviation $\tilde{\sigma}_{(-j)t}$ of the benchmark-source utilization rate across all *other* hospitals in quarter t. For each candidate source, we compute the normalized deviation $|\text{utilization_rate}_{jts} - \overline{u}_{(-j)t}|/\tilde{\sigma}_{(-j)t}$ and average it across the hospital's quarters to produce a single score per hospital-source pair. We retain the source with the lowest score, with the benchmark source kept in the event of a tie. The analogous calculation is performed for ICU beds, restricted to hospitals with positive ICU length of stay in at least one quarter; hospitals with zero ICU length of stay throughout the sample retain the benchmark assignment.

The final total-bed series draws on AHA for 145 hospitals (43.5%), the Texas Directory for 120 (36.0%), HCRIS for 67 (20.1%), and hand-search for 1; the final ICU-bed series draws on AHA for 189 hospitals (56.8%), the Texas Directory for 73 (21.9%), HCRIS for 56 (16.8%), and hand-search for 15 (4.5%). In Appendix Table A4, we show how the hospital mortality decomposition into baseline quality and utilization rates (Table 3) is broadly robust to alternative bed size measures.