Hospital Quality and Capacity Strain: Equilibrium Effects on Health Outcomes*

Gautam Gowrisankaran
Columbia University, NBER

Kelli Marquardt
Federal Reserve Bank of Chicago

Robert Town
University of Texas at Austin, NBER

October 30, 2025

Abstract

The striking observed variation in hospital performance suggests that alternative policy designs could meaningfully affect welfare. One important contributor to adverse outcomes is hospital capacity strain, which has been documented across many countries. This paper examines the empirical importance of capacity strain, patient information, and hospital learning in affecting outcomes, using data from Texas during the Covid-19 pandemic. Our main outcome is risk adjusted mortality, controlling for unobserved patient selection with a Bayesian estimator and distance as instruments. Quality varied substantially across hospitals: moving every patient out of a below-median quality hospital would save 62,000 lives if applied nationally. We find that intensive care unit (ICU) capacity strain increased hospital mortality—without capacity strain, deaths would fall 9.0%—but find no evidence of learning-by-doing. Choice estimation reveals that patients value hospitals with high baseline quality, but did not account for capacity strain in their choices. This caused enough quality degradation that quality *negatively* predicts equilibrium choice. Equilibrium counterfactual analyses show that if patients accounted for capacity strain and placed 10 times greater weight on hospital quality– more consistent with published estimates of the value of lives lost– Covid-19 deaths would drop 3.7% and mean travel time would rise only slightly. However, using more precise information but not accounting for capacity strain would cause Covid-19 deaths to *rise* 7.9%.

JEL Codes: I13, I18, D43, D61, L10.

Keywords: Mortality, capacity strain, intensive care units, selection, Bayesian estimation.

*We thank Cory Capps, Caitlin Carroll, Peter Hull, Keith Joiner, Ashley Langer, Pauline Mourot, Cristóbal Otero, 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 Theinova. The views here do not represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System. Results are preliminary and subject to change. All errors are our own.

1 Introduction

A large and growing body of research documents striking variation in hospital performance across many settings (e.g. Geweke et al., 2003, Doyle et al., 2019, Chandra and Staiger, 2020). Given that the U.S., for example, spends more than $1.5 trillion annually on hospital services—roughly 5.5% of GDP—this heterogeneity implies that alternative policy designs could meaningfully affect welfare. One potential contributing factor to this unevenness is that the hospital sector has undergone rapid structural change over the past two decades, with 195 closures, 612 mergers, and a 16% decline in bed capacity in the U.S. alone.[1] Given both the hospital sector's scale and dynamism, understanding the causes of variation in hospital performance is central to evaluating the welfare consequences of health policy.

Since hospital capacity is constrained in many countries (Kadri et al., 2021, Hoe, 2022, Avdic et al., 2024), capacity strain is a potentially important source of adverse patient outcomes. Capacity strain is a particularly problematic source of under-performance as it is difficult for patients to assess the level of strain when choosing hospitals. Patients' information about hospital quality may be based on hospital ratings and reputation[2]—both of which evolve over the long run—while capacity strain changes in the short run. Thus, capacity strain may also create adverse equilibrium responses as a patient's hospital choice can spill over and affect the outcomes of other patients when hospital capacity is strained. For instance, in periods of peak demand, patients may flock to high-quality hospitals, causing them to be overcrowded and thereby lowering their quality.

This paper examines the equilibrium effects of capacity strain, patient information, and hospital learning in affecting outcomes in the hospital sector. We focus on Covid-19 health outcomes from 2020–2022. This pandemic period provides a useful context for many reasons. First, the volume of hospitalizations during the pandemic allows us to understand the impact of capacity strain on hospital quality. Over our sample period, the United States experienced 104 million confirmed Covid-19 cases (Clarke, 2022) with 4.9 million inpatient hospitalizations[3] and over 1.1 million deaths (Silk et al., 2023). Second, hospital performance plays

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

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

[3] Hospitalization estimates are based on authors' calculations using estimates of the popu-

a crucial role in affecting Covid-19 mortality since approximately 65% of Covid-19 deaths occurred in inpatient settings (Ahmad et al., 2023). Third, medical protocols for Covid-19 changed rapidly during the pandemic (Gulick et al., 2024), implying that achieving high quality required learning about new protocols and translating them to the clinical setting. Thus, hospital learning-by-doing, which many researchers have found to be important for surgeries and other complex medical conditions (Gaynor et al., 2005, Avdic et al., 2019, James et al., 2022), may also be relevant for Covid-19 care. Finally, the *widespread* capacity strain during this period results in ambiguous equilibrium responses: if severely ill patients had obtained more accurate quality information, they may have chosen less crowded hospitals, but this may have caused those hospitals to become overcrowded and have worse outcomes.

While hospital performance is multidimensional and encompasses both clinical quality and patient experience, we focus on the expected rate of in-hospital mortality or discharge to hospice, controlling for observed and unobserved determinants of risk. For brevity, we refer to either outcome as a death. Our analysis accounts for the fact that quality may degrade from overcrowded hospitals or improve from experience, that patients may choose hospitals based on imperfect or inaccurate information, and that there may be multi-faceted equilibrium responses from better information or more capacity.

Although our focus is on Covid-19 deaths during the pandemic, our results also speak more broadly to productivity, heterogeneity, and information in the healthcare sector. Hospital capacity strain is a concern across many countries[4] and can have important distributional consequences (Singh and Venkataramani, 2024). Many other conditions—e.g., bacterial sepsis, pneumonia, and renal failure—are broadly similar to Covid-19 in that patients' health capital can depreciate over the matter of hours or days in the absence of medical attention, suggesting that the forces that we identify are applicable across settings.

Our study estimates hospital quality using Texas discharge data over the period Q2:2020 to Q3:2022.[5] Texas is a large and diverse state with a substantial urban and rural population,

lation 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`.

[4] `https://www.oecd.org/en/publications/2023/11/health-at-a-glance-2023_e04f8239/full-report/hospital-beds-and-occupancy_10add5df.html`.

[5] We drop hospitalization data from Q1:2020, the start of the pandemic, due to inaccurate reporting. Our analysis sample merges the discharge data with area demographic and hospital characteristic data.

offering a broad sample of hospitals with varying characteristics (e.g., rural, urban, suburban, academic, critical access, for-profit, and system-owned hospitals). Over the course of the pandemic, Covid-19 accounted for more than 92,000 deaths in Texas.[6] Using our analysis data, we estimate hospital quality for Covid-19 care, evaluate the extent of heterogeneity in hospital quality, and investigate the causes of performance differences. We then estimate a patient hospital choice model to quantify the impact of hospital quality in affecting hospital choice and evaluate the impact of counterfactual patient information and market structures.

Our estimation of hospital performance faces a number of challenges. 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. Additionally, a Covid-19 death is a relatively low-frequency event. The low signal-to-noise ratio can also render classical inference biased, reporting hospitals that randomly had few deaths as being high quality. Finally, hospital quality for Covid-19 may vary across time due to capacity utilization and learning-by-doing; i.e., overcrowded hospitals may perform relatively poorly, but hospitals may learn to treat Covid-19 better over time from experience.

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). Our Bayesian inference accounts for the fact that hospitals with low risk-adjusted mortality may disproportionately include ones that had favorable draws. We also control for time-varying aspects of hospital quality with a two-step process that first measures hospital quality with GGT at the quarterly level and then estimates the impact of this time-varying quality measure on capacity strain, learning-by-doing, and hospital fixed effects. The method accounts for the fact that our quarterly quality estimates are most often imprecise because there are relatively few patients.

We begin our analysis by estimating hospital quality with the GGT Bayesian inference method. We include patient demographic characteristics and indicators for over 20 comor-

[6] `https://www.dshs.texas.gov/covid-19-coronavirus-disease-2019/texas-covid-19-data`.

bidities as controls for observable patient severity of illness. We also account for county level vaccine take-up and Zip Code level demographics and income. Our baseline specification allows for unobserved (to us) and observed severity of illness to affect both hospital choice and mortality. We correct for the impact of unobserved severity on hospital choice by specifying that hospital choice is affected by instruments that, conditional on observables, do not affect mortality. We also estimate a more traditional model that assumes only observed severity of illness affects hospital choice. For each specification, we estimate models that output a separate quality level for each hospital in each quarter and ones that output the mean hospital quality over our 10 quarter sample period.

Focusing first on the hospital-level estimates over our sample period, we find large variation in hospital performance that translates into large population-level mortality impacts. The median hospital in our sample has an expected mortality rate of 0.129 when faced with the population distribution of (observed and unobserved) patient severity. The standard deviation in expected mortality across hospitals is 0.029. Moving all patients that were admitted from an above median mortality hospital to the median hospital (and assuming no change in quality) would reduce mortality by 8.5%. Applying this to the U.S. population would imply a reduction in over 62,000 deaths during the pandemic, or over $220 billion using Covid-19 value of life and income estimates (Silva et al., 2023).

Next, to evaluate the determinants of hospital quality, we start with the Bayesian estimation output, which provides posterior draws of hospital quality by quarter. We transform this measure to obtain draws of each hospital's mean mortality and regress these draws on hospital fixed effects, region-by-quarter fixed effects, and time-varying factors such as capacity utilization and previous Covid-19 patient volume. We find that the most significant time-varying predictor of hospital performance is intensive care unit (ICU) capacity strain: hospital mortality increases when ICU capacity utilization is over 100% of the reported number of ICU beds, which occurs for over three-fifths of the hospital-quarters in our sample.[7] Overall bed capacity utilization, on the other hand, is not a statistically significant predictor of hospital performance, suggesting that high-severity patients who need ICU care drive mortality reductions. Despite the large literature showing positive volume-outcome relation-

[7]ICU capacity utilization can exceed 100% as hospitals can temporarily expand ICU beds during surges.

ships (e.g Luft et al., 1987), we also find no evidence of learning-by-doing: hospitals that had more Covid-19 patient volume in the past did not exhibit better performance. Finally, we take the hospital fixed effects from this regression and ask whether time-invariant hospital characteristics are correlated with performance. System membership and pre-pandemic COPD mortality predict better baseline hospital quality, but other variables, such as size, ownership type, and critical access status, are not significant.

We then estimate multinomial logit patient hospital choice models to understand the extent to which patients accounted for hospital quality in their decisions of the hospital at which to seek treatment. We include measures of distance, hospital size, and ownership type as controls. We find that hospital mortality without capacity strain is a significant negative predictor of patient hospital choice. The logit coefficients on hospital mortality are approximately -5 across specifications with different controls, implying if a patient had the choice of two otherwise identical and equidistant hospitals where one had a 10 percentage point higher mortality rate than the other, then she would choose the better one 62% of the time.

In contrast, contemporaneous hospital mortality—which includes capacity strain—is mostly a *positive* predictor of patient hospital choice, with a logit coefficient of approximately 7 in a specification with controls. In other words, Covid-19 patients disproportionately sought care at hospitals that would have been high quality in the absence of high capacity utilization. But, these hospitals ended up being worse than other alternatives in equilibrium because their quality degraded due to overcrowding.

Finally, we perform counterfactuals that examine the impact of market structures and information on mortality and utility, accounting for spillovers on capacity strain. We first examine the infeasible counterfactual of Covid-19 treatments with every hospital being empty. We find that this would have increased quality enough to reduce mortality by 9%, even larger than our thought experiment of moving patients from below median-quality hospitals to the median one. We next examine what would happen if we added 5 ICU beds to each hospital. This would reduce mortality by 1.2%, increasing welfare by reducing deaths by $30 billion if applied nationally.

Turning to equilibrium simulations, we investigate the mortality impact of providing

patients with information regarding actual hospital quality, assuming first that their true preferences are given by the estimated patient hospital choice model that includes baseline hospital quality as a regressor. Despite providing patients more accurate quality information, equilibrium mortality goes *up* slightly, by 0.14%. The essentially null result occurs because this information causes patients seeking high quality care to select low occupancy hospitals—which might have been worse ex ante—and causes them to become more crowded.

The magnitude of our estimated coefficient on baseline quality suggests that patients may have only limited information even about baseline quality. If patients had more precise information, this would increase their weight on quality. We find that if patients valued actual quality 10 times as much as they valued baseline quality in our estimation—which would occur if they had both more accurate and more precise information—equilibrium mortality would decrease 3.7%, with travel time increasing by only 4.2%. However, if they valued baseline quality (not accounting for capacity strain) 10 times as much as they did in our estimation, equilibrium mortality would *increase* 7.9%, as patients crowded even more into hospitals that were already over ICU capacity.

Related Literature: This paper contributes to three main literatures. First, we contribute to a large literature that seeks to understand the role of market forces, technology, and consumer information on equilibrium outcomes in healthcare. A series of papers show that consumers respond to healthcare quality ratings (Chernew et al., 2008, Dafny and Dranove, 2008, Gaynor et al., 2016). However, this response is not uniformly observed. For example, Abaluck et al. (2021) finds that enrollees neither respond to Medicare Advantage health plan ratings nor do these ratings accurately reflect quality. A number of papers have focused on the role of learning-by-doing for surgeries, generally finding large positive effects of increased volume on outcomes (e.g. Birkmeyer et al., 2002, Gaynor et al., 2005). Other papers have found that capacity strain leads to increased hospital readmissions, obstetric trauma, or deaths (Kuntz et al., 2015, Bravata et al., 2021, Kadri et al., 2021, Hoe, 2022, Avdic et al., 2024). Finally, some studies have studied whether consumer information leads to positive equilibrium selection of high quality hospitals (Chandra et al., 2016). We contribute to this literature in three ways: (1) with our results showing that hospital quality depends upon capacity strain; (2) with our analysis of information, which shows that patients are

informed about baseline but not contemporaneous quality, leading to spillovers from capacity strain; and (3) with our equilibrium counterfactuals that analyze what policies might improve hospital outcomes in the face of capacity strain and limited patient information.

Second, a literature uses inpatient discharge data to estimate the distribution of hospital (or other medical facility) quality, its determinants, and its evolution. This literature dates at least to Luft and Hunt (1986). Recent research has proposed related solutions to overcoming the challenges to estimating hospital performance noted above, that unobserved patient severity and the low signal-to-noise ratio of hospital mortality can bias classical estimates of hospital performance (see McClellan et al., 1994, Gowrisankaran and Town, 1999, Geweke et al., 2003, Doyle et al., 2015, Hull, 2018, Olenski and Sacher, 2024). We contribute to this literature by showing how to use Bayesian estimation methods that address these two issues in a large panel data set with time-varying hospital quality, and by performing this estimation in a computationally simple way in Stata, using a companion paper and code (Marquardt et al., 2025).

Finally, we contribute to a literature that has examined the determinants of Covid-19 mortality with our frontier estimation of hospital quality and subsequent analyses of technology, information, and equilibrium. Allcott et al. (2024) and Alsan et al. (2021) consider the causes and presence of geographic variation in Covid-19 mortality. Asch et al. (2021) considers variation in Covid-19 hospital quality, Song et al. (2021) estimates the relationship between patient characteristics and in-hospital mortality, and Kunz and Propper (2023) evaluates the relationship between county-level outcomes for Covid-19 and pneumonia treatments.

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

We consider capacity strain using the Covid-19 pandemic as our setting. 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, 1.1 million deaths, and 1.2% of the US population was hospitalized. Our analysis focuses on Texas, which experienced 501,000 inpatient admissions and 90,000 deaths over this period.[8]

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 a pulmonary embolism or deep vein thrombosis 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.

[8] See `https://dshs.texas.gov/covid-19-coronavirus-disease-2019/texas-covid-19-data`.

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[9]. For patients with mild to moderate Covid-19 who are at high risk for further progression, the guidelines recommend the use of convalescent plasma (IDSA (2021)). The guidelines also recommend the use of antiviral agents such as remdesivir (FDA approved Octomber 2021), molnupiravir (FDA approved December 2021), dexamethason (for severe but not critical patients). Immunomodulatory treatments, including IL-6 inhibitors like tocilizumab are also recommended treatments for progressively severe patients.

Beyond pharmacological interventions, hospitalized Covid-19 patients often require thromboprophylaxis due to the elevated risk of blood clots. Low-molecular-weight heparin or direct oral anticoagulants are commonly used to prevent thrombotic complications, which are a major contributor to morbidity and mortality.

If a patient's condition deteriorates then intubation and the use of a mechanical ventilation may be appropriate. In our data, 13% of patients are intubated and there is a very high mortality rate for intubated patients (48%). There is also meaningful ambiguity on best practices in the use of ventilation for Covid-19 patients. As Pisano et al. (2020) states with regard to ventilation for Covid-19 patients, "[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 probability of death. This creates a challenging risk/benefit calculation for intubation and also increases costs for the hospital.

The quality of the evidence underlying the treatment guidelines is often weak, implying significant scope for differential interpretation of the evidence. Furthermore, the guidelines were updated several times over the pandemic as new treatments and evidence became available. For example, as noted above, remdesivir was approved in October of 2020. The FDA issued an emergency use authorization (EUA) for baricitinib in November 2020, for its use

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

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).

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. Treatment outcomes also likely varied over time, driven by the introduction of new therapeutics and by episodes of capacity strain. Such 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. 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).

2.2 Data

Our primary data source is the State of Texas Department of Health and Human Services patient discharge dataset.[10] These data contain virtually every inpatient hospital discharge in the state. 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. Over this period, 99,000 Covid-19 deaths occurred in Texas (9% of US total).

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

[10] 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

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 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 to construct measures of patient travel time from the patient's ZIP Code centroid to hospitals using OpenStreetMaps.

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

We construct detailed risk adjustment measures using Elixhauser Comorbidity Indices (approximately 25) (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)[11] 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, we gather information on the hospital's for-profit status, number of general acute care beds, number of intensive care unit beds (ICU), teaching status (major and minor), system membership, rural status, and critical access hospital status. From the Centers for Medicare and Medicaid Services (CMS), we import information on pre-Covid, hospital specific inpatient mortality rates for acute myocardial infarctions, coronary artery bypass graph surgery, COPD, heart failure, pneumonia, and stroke. We also collect the average per-capita Medicare spending associated with the hospital, the surgical complication rate and CMS's hospital star rating.[12] The star rating system assigns a hospital a rating from one to five based on their reported outcomes across five areas of quality and

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

[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.

maps those outcomes into a single star rating for each hospital.[13] Those areas are mortality, patient safety, readmission rates, patient experience, and the timeliness and effective care.

We restrict our final empirical sample to adult patients (those over 20 years of age) and to general/acute care hospitals (i.e., we remove specialty facilities such as psychiatric centers, heart and spine hospitals, and children's hospitals). We drop patients whose home county is either missing or outside of Texas and thus we cannot assign them a county-quarter level vaccine rate. We also drop patients whose home ZIP Code is either missing, defined as a P.O. Box, or otherwise outside of the OpenStreetMaps database and therefore unable to calculate travel distance from their ZIP Code to each hospital. This first patient drop amounts to about 9.25% of the adult Covid sample. We exclude a small minority (about 3.3%) of patients who traveled outside of their Texas public health region for care.[14]

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 estimated by our model.

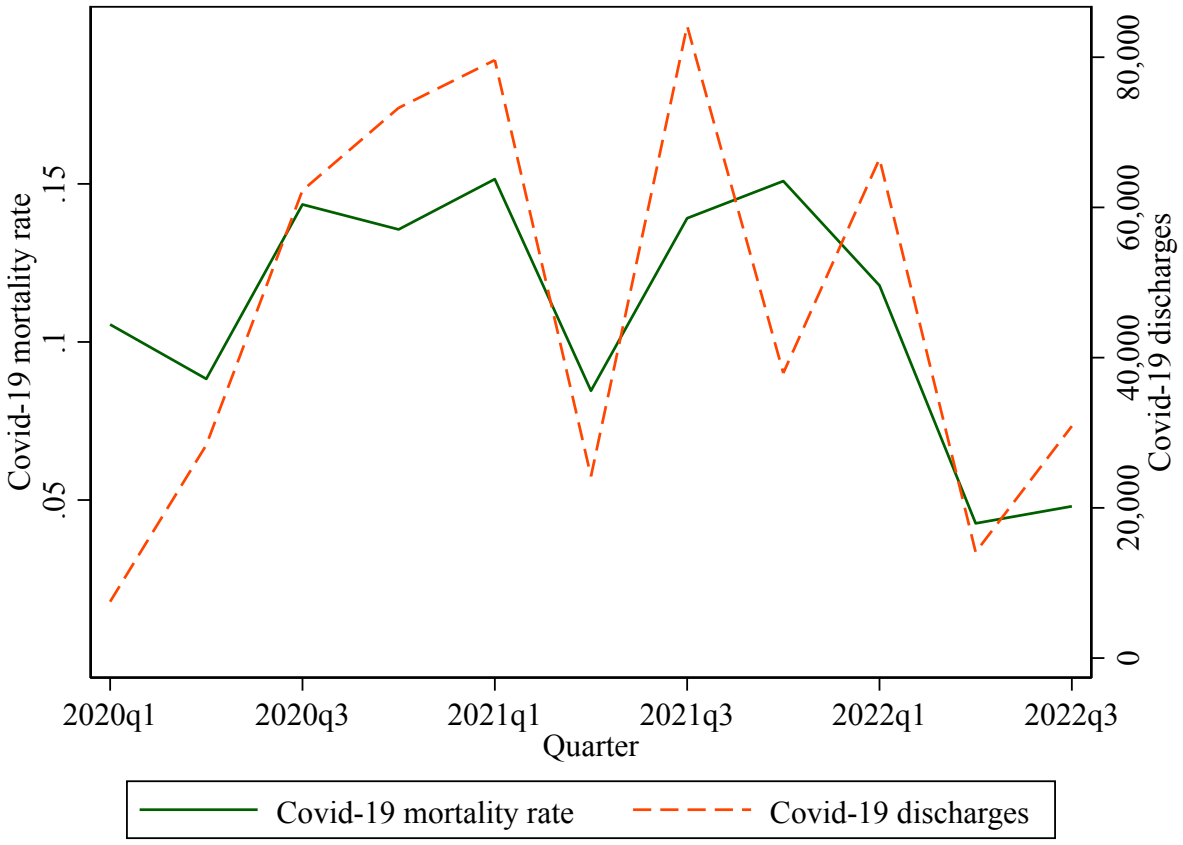
2.3 Trends and Summary Statistics

Figure 1 presents the time series of the total number of Covid-19 admissions 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 admissions and the mortality rate are 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.

[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 `https://www.dshs.texas.gov/center-health-statistics/texas-county-numbers-public-health-regions`.

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 the summary statistics of our patient sample broken down by year. Our analysis sample contains 407,081 inpatient observations with 54,844 deaths (13.4%). The mortality and the number of inpatient admissions peaked in 2021 with almost 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 6.5% are covered by Medicaid.

The average length of stay is 7.3 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 resident county vaccination rate was 44% in 2021 and 68% in 2022. A third of patients were admitted to the closest hospital and the average distance traveled from home to hospital was

25 minutes. This last fact suggests that patients make active hospital choice decisions.

Table 1: Summary Statistics on Patients

	Full Sample	2020	2021	2022
Died	0.134	0.138	0.148	0.100
Male	0.496	0.505	0.507	0.460
White (non-Hispanic)	0.447	0.396	0.456	0.505
Black (non-Hispanic)	0.132	0.134	0.130	0.133
Hispanic	0.339	0.382	0.333	0.284
Medicaid	0.065	0.059	0.063	0.080
Medicare	0.436	0.426	0.393	0.545
Commercial insurance	0.326	0.332	0.359	0.244
Other/No insurance	0.173	0.184	0.184	0.131
Age 65+	0.443	0.445	0.392	0.549
Age 80+	0.156	0.154	0.126	0.223
Number comorbidities	2.974	3.034	2.871	3.098
SNF transfer	0.008	0.010	0.005	0.011
Length of stay	7.343	7.056	8.105	6.164
Primarily in ICU	0.431	0.443	0.448	0.375
Minutes to closest	15.023	14.914	15.128	14.968
Minutes to chosen	24.591	24.790	24.527	24.415
Went to closest	0.344	0.346	0.344	0.338
Vaccination rate	0.344	0.003	0.440	0.676
Median income ($100,000)	0.379	0.366	0.384	0.387
N Patients	407,081	135,519	185,021	86,541

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 quarters quarters whereas the year 2022 includes patients with discharges during the first three quarters. All values are averages across patients.

Table 2 presents the hospital summary statistics. Our analysis sample includes data on 349 hospitals.[15] The average bed size is 170 beds but there is a wide distribution of bed size in our sample: the standard deviation is 215 beds. The average number of ICU beds is 16.7 with a standard deviation of 24.1 beds. Many hospitals do not have any ICU beds.

[15]We assigned time-invariant hospital characteristics to each hospital using the CMS Certification Number (CCN) based on a hospital-name match. For about 15% of hospitals, we did not find a name match, in which case we assigned hospital characteristics using a manual search. Some hospital characteristics remain missing, either because they are missing in the CMS and/or AHA databases, or because we could not identify them with a manual search. On-line Appendix Table A1 presents the hospital characteristics means (excluding those with missing values) for both the CCN Matched Sample and the CCN Unmatched Sample.

Forty percent of the hospitals are for-profit, over two-thirds are members of a system, 8% are teaching hospitals, and about one-sixth of hospitals are either rural or critical access hospitals.

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, the American Hospital Association (AHA) Annual Survey reports 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, we often observe hospitals utilizing more than 100% of their measured ICU capacity.

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 in their star rating—the standard deviation is one star. For spending, CMS attributes an average of $9,944 in per-capita Medicare spending to the average hospital and again there is a large variance across hospitals in this measure.

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, 44% of bed-days filled during the sample period, with a standard deviation of 26% of bed-days filled. The mean hospital has an average ICU utilization rate of 191%, 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% only for hospitals with 0 ICU beds. The mean hospital's *adjusted* ICU utilization rate is 153%, with a wide standard deviation of 165%.

Table 2: Summary Statistics on Hospitals

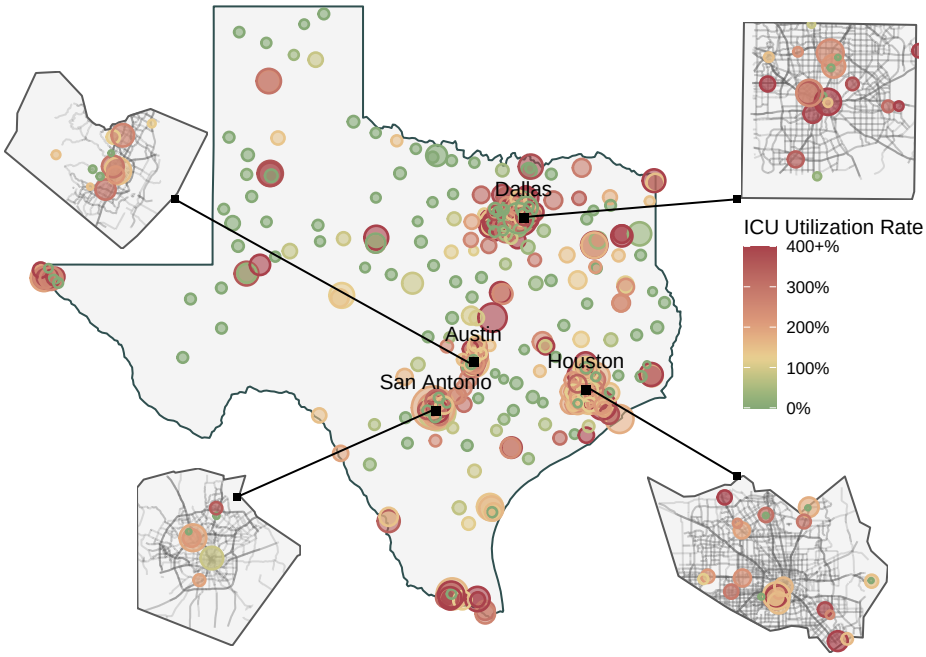
	Min	Mean	SD	Max	N
Total beds	4.00	169.76	215.24	1774.00	349
Total ICU beds	0.00	16.71	24.05	201.00	349
Teaching hospital	0.00	0.08	0.28	1.00	349
For-profit	0.00	0.40	0.49	1.00	349
System member	0.00	0.69	0.46	1.00	349
Rural	0.00	0.17	0.37	1.00	299
Critical access	0.00	0.15	0.36	1.00	349
AMI mortality rate	0.10	0.13	0.01	0.15	183
CABG mortality rate	0.02	0.03	0.01	0.07	107
COPD mortality rate	0.06	0.08	0.01	0.11	255
Heart failure mortality rate	0.07	0.11	0.01	0.16	243
Pneumonia mortality rate	0.10	0.15	0.02	0.22	273
Stroke mortality rate	0.11	0.14	0.02	0.17	178
Surgical complication rate	0.11	0.16	0.02	0.21	150
CMS star rating	1.00	3.33	1.02	5.00	257
Per-capita Medicare spending ($)	5,310	9,944	2,940	30,989	266
Average overall utilization rate	0.01	0.44	0.26	1.33	349
Average ICU utilization rate	0.00	1.91	4.00	58.73	349
Average ICU utilization rate (adj.)	0.00	1.53	1.65	10.60	349

Note: Table of hospital level summary statistics for the full sample. Some values are missing for certain hospitals, either due to insufficient patient volume for CMS reporting requirements or due to name-CCN match failure. See On-line Appendix Table A1 for hospital characteristics means (excluding those with missing values) for both the CCN Matched Sample and the CCN Unmatched Sample. 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%.

Figure 2 presents the geographic distribution of ICU utilization rate across hospitals for a peak quarter, Q1:2021.[16] 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%. Figure 3 presents the histogram of ICU bed-day utilization across hospital-quarters for all hospitals with a positive number of ICU beds. Many hospitals in this quarter operated at above 150% ICU utilization.

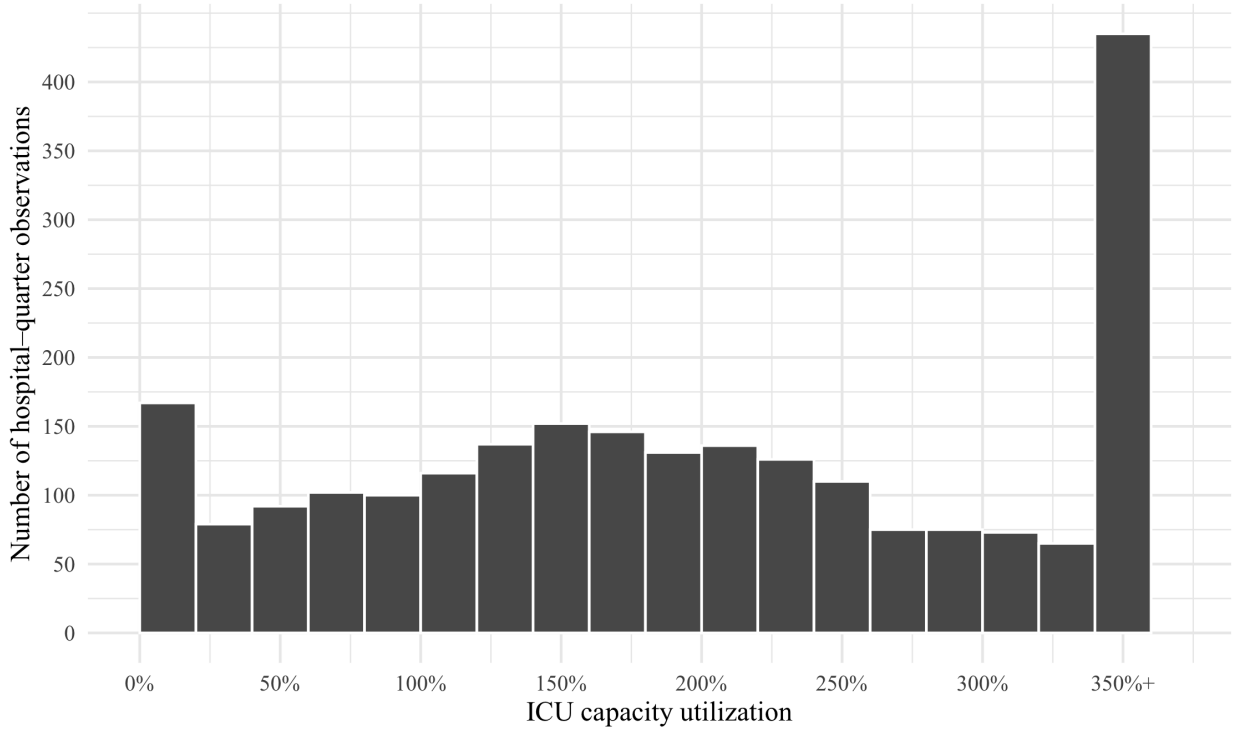
[16]See Appendix Figure A1 for the analogous map of ICU utilization rate in Q2:2020.

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.

Figure 3: ICU Capacity Utilization by Hospital/Quarter



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

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 determinants. 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 and 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.[17] 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 term, ε_{it}:

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

In (1), c_{it} is a vector of J_t indicators that specifies the hospital chosen by patient i, β_j is the (inverse) quality of hospital j, x_{it} are patient characteristics, and γ 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 skilled nursing facility and 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. We specify

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

the utility level for hospital j to be $\tilde{c}^*_{itj} = \tilde{Z}'_{itj}\alpha + \tilde{\eta}_{itj}$, 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 are coefficients, and $\tilde{\eta}_{itj} \sim N(0,1)$, *i.i.d.* across options.[18] 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 replacing $c^*_{itJ_t} = 0$. For $j < J_t$, we rewrite $c^*_{itj} = Z'_{itj}\alpha + \eta_{itj}$ and let Σ be the variance of the vector η_{it}. We then allow for selection based on unobservables by specifying that $\varepsilon_{it} = \eta'_{it}\delta + \zeta_{it}$, where $\zeta_{it} \sim N(0,1)$, and δ are $J-1$ 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 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, hospital, and region.[19] This allows for hospital quality to vary flexibly across time and for patients to select hospitals that are good at particular quarters based on the patient's unobserved health characteristics. We specify weak priors on hospital quality, using the defaults in **?**. 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 draws that form the posterior distribution of the parameters. Because the number of patients—and hence deaths—in a given quarter

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

[19] 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 $\delta = 0$.

is limited, the posterior hospital quality levels are generally imprecisely estimated. However, because we obtain the full posterior distribution of quality levels, we can use these hospital/quarter quality estimates as inputs in the hospital quality fixed effects and other quality determinants, which are more precisely estimated. Moreover, the estimator does not rely on a law of large numbers and asymptotic normality as in a frequentist estimator. It instead returns the simulated posterior even in cases where the data are limited. Thus, our use of this estimator in repeated samples with limited data is appropriate.

3.2 Decomposing Hospital Quality

This estimation starts with 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, $mort_{rtjd}$.[20] 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, $f(occupancy_{rtj})$ 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 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. Our baseline specification measures $f(occupancy_{rtj})$ with a spline in the occupancy rate for intensive care unit (ICU) beds for a hospital/quarter.[21]

[20] We use the mortality rate as our dependent variable instead of the quality fixed effects, q_{rtj}, for ease of interpretation.

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

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 was in the ICU.[22] 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 will not start out at the posterior distribution. Hence, we keep every 100th draw and drop 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 because residuals may be correlated across time for a hospital or across hospitals in a given region/quarter.

Our principal identifying assumption in (2) is that the time-varying component of volume is not correlated with the unobserved time-varying component of quality, e_{rtjd}. 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 (Gaynor et al., 2005, Gowrisankaran et al., 2006). As a general point, 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. In addition, in this case, we might find a positive coefficient on experience even in the absence of learning-by-doing. 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

[22]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.

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 plausibly 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 available hospitals in her region and quarter, but now with different regressors. We specify latent utility as:

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

where x^1_{ritj} indicate variables that relate to the interaction of patient and hospital characteristics, x^2_{rtj} are hospital characteristics, $Emort_{rtj}$ 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) and hospital. We include fixed hospital characteristics in some specifications. We also create a measure of each patient's mortality risk quantile, using the GGT posterior draws of γ and δ. Some specifications include this measure of observed patient severity interacted with distance as an additional term in x^1_{ritj}. Our measures in x^2_{rtj} are hospital characteristics, and include its size, whether the hospital in 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 $Emort_{rtj}$ across specifications. First, we consider mortality in the absence of capacity strain, i.e., when the hospital is empty. Consumers might be informed about this measure because it would be consistent with the long-run quality of the hospital. We calculate this measure from the GGT posterior draws by simulating the mortality rate at a hospital when faced with a simulated set of Texas patients, with randomly drawn observed characteristics, x_{rti}, and unobserved severity of illness, ε.

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, $\overline{mort}_{rj}$, to its additional mortality due to its current capacity, as estimated by $f(occupancy_{rtj})$. 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 by specifying the quality with the patient added to capacity utilization for hospitals not chosen, with the assumption that a given patient's length-of-stay and whether she is treated in the ICU would remain constant across hospitals.

For this 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 for these results. 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 results for the three main empirical exercises discussed in Section 3, in turn.

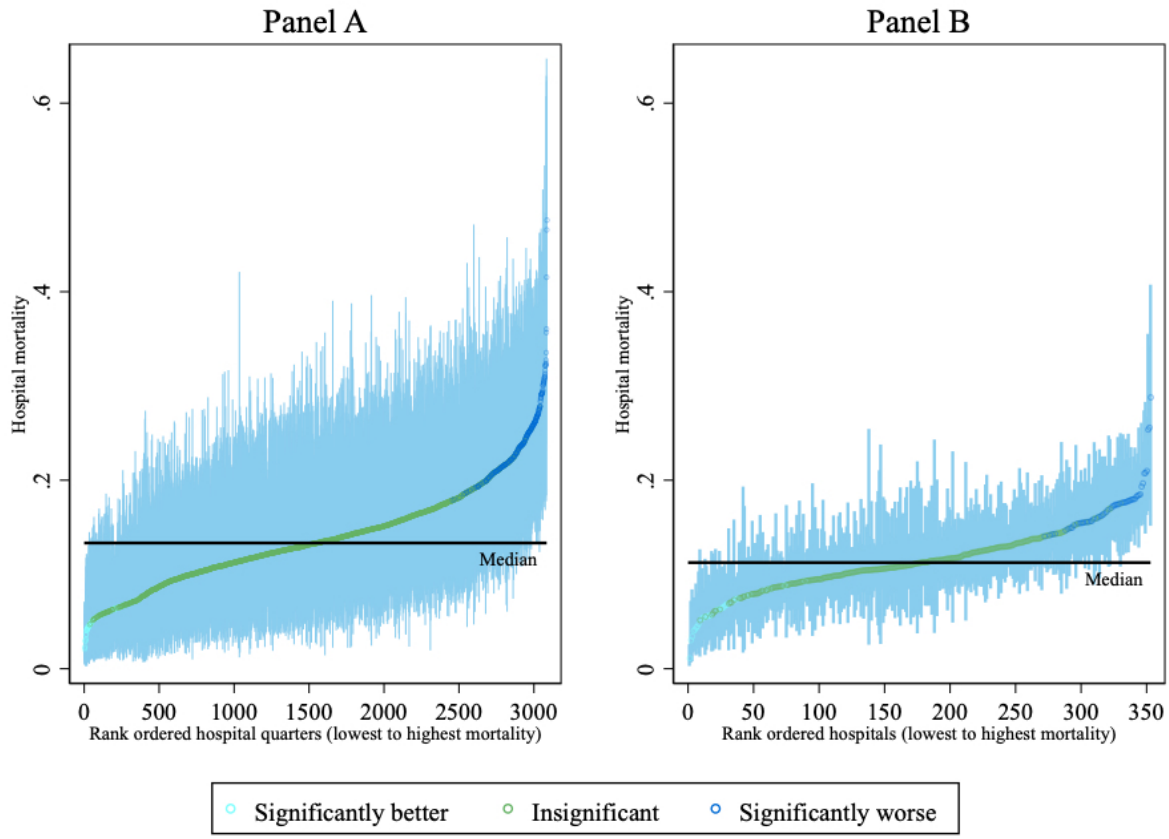
4.1 Hospital Quality Estimates

Our hospital quality estimates are the posterior distribution of hospital quality, q_{rtj}, as estimated by the GGT model. 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.[23] 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 4 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, 76 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 these posterior draws across quarters to obtain more

[23]See On-line Appendix Table A2 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.

Figure 4: 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.

precise inference on the determinants of hospital quality. Notably, our Bayesian estimator recovers posterior distributions that do not rely on a law of large numbers, implying that it provides appropriate inference on hospital quality with a limited number of patients that can then be combined across quarters.

For comparison purposes, we also estimate specifications that perform the estimation together for all patients and hospitals in a region.[24] Figure 4 Panel B provides analogous information 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

[24]These specifications include quarter fixed effects in addition to the risk adjusters reported in Table A2.

per hospital are much more precise, as we would expect given the larger sample of patients that contributes to each hospital fixed effect. About 33 percent of hospitals have posterior mean quality levels that are significantly different than the median hospital.

The heterogeneity in hospital quality from Figure 4 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 8.5% or have saved over 4,700 lives in our estimation sample. Extrapolating this number to the U.S. population, which had approximately 730,000 in-hospital deaths during the pandemic, would imply a reduction in over 62,000 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 over $220 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.

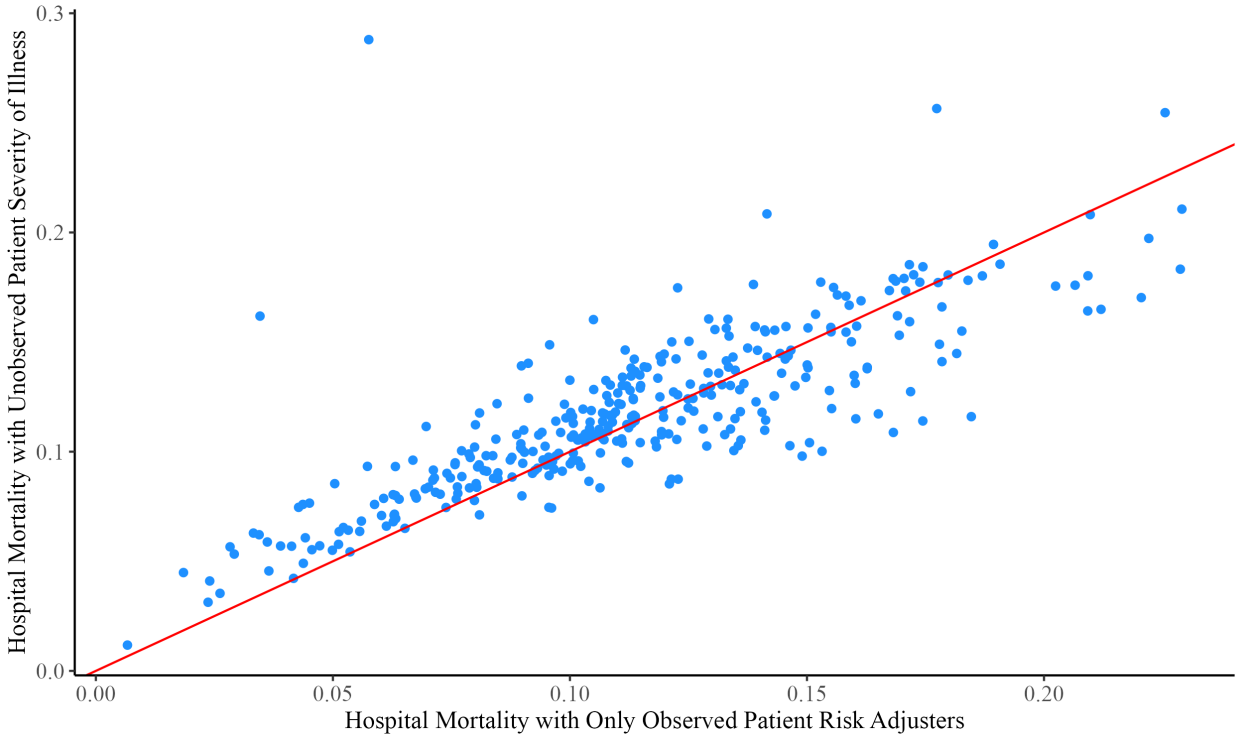
To understand the importance of selection based on unobservable patient risk, we estimate GGT controlling only for observed patient risk adjusters, by setting $\delta = 0$. For ease of exposition, we focus here on the time-invariant measure of quality, because it has many fewer hospital fixed effects.

Figure 5 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 the hospital mortality rate estimated at the quarterly level, $mort_{rtj}$. We regress this mortality rate on hospital fixed effects, region-by-quarter fixed effects, and

Figure 5: Effect on 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 seletion, i.e. with estimated values for δ. Both models include the same observed patient risk adjusters, x'_{it} from section 3.1.

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 in Column (1) while the overall utilization rate is negative, small, and statistically insignificant.

Table 3, column (3) investigates a linear spline for the ICU utilization rate with a knot at capacity utilization of 100%. We find that the increase in mortality from overcrowding

Table 3: Impact of Occupancy and Experience on Hospital Mortality

Dependent Variable:	Hospital expected mortality: $mort_{rtj}$					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Regressor:						
ICU utilization rate	0.007***	0.008***				
	(0.002)	(0.003)				
Overall utilization rate		-0.010		-0.010	-0.010	
		(0.014)		(0.014)	(0.014)	
ICU utilization rate spline: $[0, 1)$			0.006	0.008	0.008	
			(0.007)	(0.007)	(0.007)	
ICU utilization rate spline: $[1, \infty)$			0.007***	0.008***	0.008***	
			(0.002)	(0.003)	(0.003)	
log(Covid patients through t-1)					0.001	
					(0.001)	
ICU utilization rate spline: $[0, 0.5)$						-0.002
						(0.014)
ICU utilization rate spline: $[0.5, 1)$						0.010
						(0.013)
ICU utilization rate spline: $[1, 2)$						0.010*
						(0.005)
ICU utilization rate spline: $[2, \infty)$						0.007***
						(0.002)
Hospital/quarter/draw observations	2,776,500	2,776,500	2,776,500	2,776,500	2,776,500	2,776,500
R^2	0.44	0.44	0.44	0.44	0.44	0.44

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.

is close to linear but only statistically significant at ICU capacity levels over 100%. We use column (3) as our base specification. These basic results are corroborated by columns (4) and (6), which also show statistically significant spline coefficients for ICU capacity utilization greater than 100% after adding further controls. Column (6) shows a slightly negative effect (but not statistically significant) impact for capacity utilization less than 50% and then a nearly linear effect after that. On-line Appendix Figure A2 provides a visual representation of the spline coefficients in columns (3) and (6). In both cases, mortality increases by about 2 percentage points as ICU capacity utilization goes from 0% to 250%.

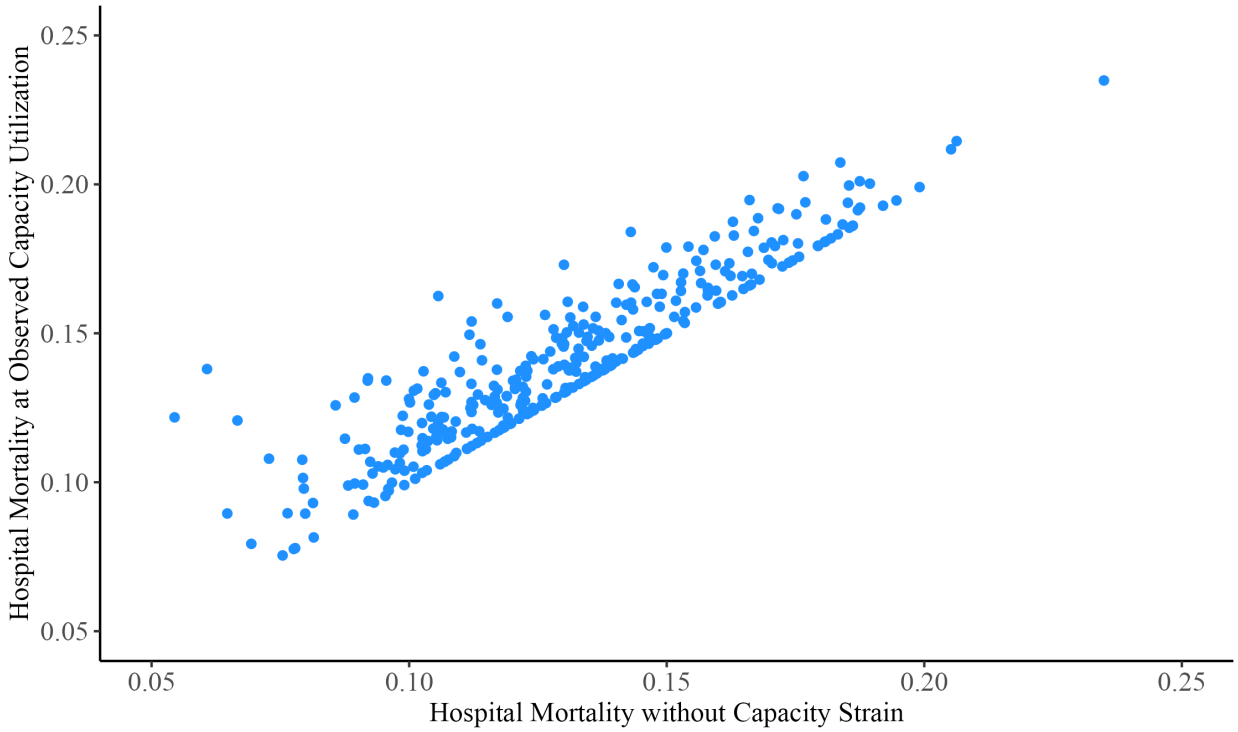
Table 3 column (5) investigates the presence of learning-by-doing by adding the log of the cumulative lagged Covid-19 patients at the hospital, up to the end of previous quarter. We find that previous experience is not statistically significant. Moreover, the point estimate is positive, implying that, if anything, more past Covid-19 patients increases 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 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 6 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 in the quality of Covid-19 treatments.

Table 4 provides a second look at the impact of capacity strain, on a hospital-by-hospital basis. We find that mortality at the mean sample capacity use is about 1.1 percentage points higher than it would be if the hospital faced no capacity strain. This amount of capacity strain is slightly higher than it would be if every hospital had an ICU that was at 150% of capacity.

Figure 6: 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.

Table 4: Mortality Variation from Capacity Strain

Measure	Mean	Std. Dev.	25th percentile	75th percentile
Mortality at mean sample capacity use	0.146	0.029	0.127	0.167
Mortality when empty	0.135	0.029	0.114	0.153
Mortality when use is at 100% of capacity	0.142	0.029	0.120	0.159
Mortality when use is at 150% of capacity	0.145	0.029	0.124	0.163

Note: Table 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.

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 5, column (1) considers the impact of hospital ownership. The only statistically significant regressor is an indicator for the hospital being a system member, which predicts a 0.8 percentage

point lower Covid-19 mortality. This coefficient remains significant and similar in magnitude across the specifications reported in this table.

Column (2) adds in the impact of hospital size on baseline mortality finding 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.[25] One measure stands out: 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 the complications of Covid-19 are respiratory, this spillover in treatment seems plausible.

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

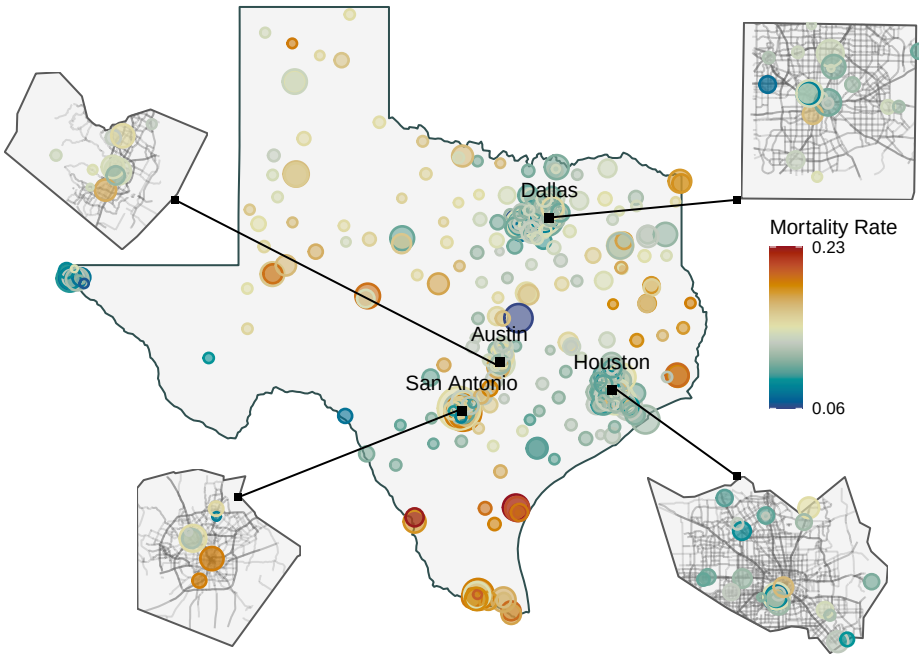
Table 5: Decomposition of Baseline Hospital Quality into Characteristics $\overline{mort}_{rj}$

Dependent Variable:	Hospital mortality fixed effects, $\overline{\text{mort}}_{rj}$		
Model:	(1)	(2)	(3)
Regressor;			
ICU bed share	-0.021 (0.025)	-0.025 (0.027)	-0.033 (0.027)
Major teaching	-0.009 (0.007)	-0.008 (0.008)	-0.004 (0.008)
For-profit	0.001 (0.004)	0.001 (0.004)	0.005 (0.004)
Public	0.016*** (0.006)	0.016*** (0.006)	0.013** (0.006)
System member	-0.002 (0.004)	-0.002 (0.004)	-0.002 (0.004)
Beds: <100		0.010 (0.008)	0.010 (0.008)
Beds: ≥100 and <500		0.007 (0.008)	0.012 (0.009)
Critical access			0.011 (0.007)
CMS star rating = 2			0.008 (0.014)
CMS star rating = 3			-0.002 (0.014)
CMS star rating = 4			0.002 (0.014)
CMS star rating = 5			-0.003 (0.014)
Star rating missing			-0.012 (0.014)
log(Medicare spend)			0.006 (0.007)
Spend missing			0.056 (0.065)
COPD 30-day mortality			0.775*** (0.207)
COPD mortality missing			0.066*** (0.018)
Observations	349	349	349
R^2	0.05813	0.06242	0.15119

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.

Finally, we show the geographic distribution of hospital mortality without capacity strain in Figure 7. South Texas has uniformly worse hospitals than average. The panhandle has approximately average quality hospitals. 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 worse hospitals. On-line Appendix Figure A3 shows the analogous map at the realized quality level accounting for capacity strain. It shows that the effects of capacity strain on degrading quality are prevalent throughout Texas.

Figure 7: 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.

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 understand the efficiency properties of this market (Chandra et al., 2016).

Table 6 presents our patient hospital choice model estimates, which are based on multinomial logit specifications from (3). Panel A shows results where patients' expectation of hospital mortality are given by quality when it is empty, i.e., $Emort_{rtj} = \overline{mort}_{rj}$. 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 -5.20 and a coefficient on minutes traveled of -0.10. Applying the mortality coefficients to the case of two hospitals that are identical except for mortality, a hospital with a mortality rate when empty that is 5.8 percentage points—or two standard deviations—lower than the other would be chosen 57.5% of the time, showing that baseline quality plays

Table 6: Impact of Quality Measures $Emort_{rtj}$ on Patient Hospital Choice

Panel A: Mortality evaluated when empty					
Regressor:	(1)	(2)	(3)	(4)	(5)
Hospital mortality	−4.989***	−4.892***	−4.319***	−4.168***	−5.100***
	(0.080)	(0.080)	(0.080)	(0.083)	(0.080)
Distance (minutes)	−0.100***	−0.129***	−0.134***	−0.133***	−0.124***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Distance^2/1000		0.282***	0.284***	0.284***	0.259***
		(0.001)	(0.001)	(0.001)	(0.001)
Closest		0.145***	0.181***	0.182***	0.177***
		(0.009)	(0.009)	(0.009)	(0.009)
Major teaching		0.246***	0.030***	0.039***	
		(0.005)	(0.005)	(0.005)	
System member		0.719***	0.235***	0.228***	
		(0.007)	(0.009)	(0.009)	
Critical access		−2.826***	−2.059***	−2.054***	
		(0.020)	(0.021)	(0.021)	
For-profit		−0.541***	−0.339***	−0.341***	
		(0.004)	(0.004)	(0.005)	
Government		0.244***	−0.056***	−0.072***	
		(0.008)	(0.010)	(0.010)	
Medium hospital			1.437***	1.432***	
			(0.007)	(0.007)	
Large hospital			1.884***	1.868***	
			(0.008)	(0.008)	
COPD				−1.805***	
				(0.198)	
COPD missing				−0.187***	
				(0.017)	
Severity × distance/100					2.767***
					(0.422)
Panel B: Mortality at realized occupancy (capacity adjusted)					
Regressor:	(1)	(2)	(3)	(4)	(5)
Hospital mortality	7.061***	3.884***	−1.253***	−0.828***	7.010***
	(0.084)	(0.087)	(0.092)	(0.095)	(0.084)
Distance (minutes)	−0.100***	−0.129***	−0.133***	−0.133***	−0.124***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Distance^2/1000		0.279***	0.283***	0.283***	0.256***
		(0.001)	(0.001)	(0.001)	(0.001)
Closest		0.172***	0.186***	0.186***	0.213***
		(0.009)	(0.009)	(0.009)	(0.009)
Major teaching		0.244***	0.040***	0.043***	
		(0.005)	(0.005)	(0.005)	
System member		0.667***	0.203***	0.196***	
		(0.007)	(0.009)	(0.009)	
Critical access		−2.687***	−2.033***	−2.027***	
		(0.020)	(0.021)	(0.021)	
For-profit		−0.551***	−0.365***	−0.376***	
		(0.004)	(0.004)	(0.005)	
Government		0.153***	−0.128***	−0.149***	
		(0.008)	(0.009)	(0.010)	
Medium hospital			1.450***	1.450***	
			(0.007)	(0.007)	
Large hospital			1.919***	1.908***	
			(0.008)	(0.008)	
COPD				−3.871***	
				(0.198)	
COPD missing				−0.334***	
				(0.017)	
Severity × distance/100					3.040***
					(0.424)

Note: "Medium hospitals" have 100–499 beds. Hospital mortality estimates from Table 3 column (3). The terms ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively. Number of observations: cols. (1), (2), (5): 19,179,978, cols. (3), (4): 14,240,399. Panel B mortality measure corrects for the fact that an ICU patient choosing a hospital affects its mortality.

a substantial role in affecting patient hospital choice. We also use these estimates to understand the value that patients place on mortality relative to travel time. A one-minute travel time decrease is equivalent to a 1.92 percentage point lower death rate, or $68,500 (Silva et al., 2023). This dollar number seems implausibly high if patients knew the true quality of a hospital, implying that they may be only partially informed about hospital quality, even if the measure they use is baseline quality.

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) show very similar coefficients on $Emort_{rtj}$ to Column (1). Column (5) adds an interaction between the patient's observed severity of illness quantile and distance.[26] It shows that more severely ill patients are less distance-sensitive than others, but the coefficient on $Emort_{rtj}$ is -5.32, which is very similar to the other specifications.

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, they prefer larger hospitals, and ones that do not have critical access status.

Next, panel B shows results where patients' quality expectations, $Emort_{rtj}$, are given by the actual realized quality of a hospital given the capacity strain in the current quarter, 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 7.15. If taken as causal, it shows that patients are choosing lower quality hospitals, all else equal. Columns (2), (3), and (5) show similar results to Column (1). However, Column (4), which includes the most hospital characteristics, reports a negative sign on $Emort_{rtj}$, albeit with a relatively small coefficient of -1.82.

Combining panels A and B, the implications of the choice model analysis are striking. Patients are choosing the "right" hospitals, if capacity strain were not important. In other words, they seem to have a reasonable sense of the identity of high quality hospitals for Covid-19 ignoring the equilibrium impact of other patients' choices. However, they do not account for this equilibrium impact: high quality hospitals are being chosen by many other

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

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.

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 what quality would be in the absence of capacity strain. These results suggest that alternative market structures or patient information can improve welfare. This section further considers the role of market alternatives and information on equilibrium outcomes.

In all our counterfactuals, we start with the estimated parameters from Table 6 panel A column (5). We use these parameters because 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).

Given the relatively low value placed on quality relative to distance in Section 4.3, it is likely that patients believe that their information about hospital quality is relatively inaccurate. For this reason, we also consider counterfactuals where we increase the weight that patients place on quality by a factor of 10, which would correspond to them have a higher precision for their quality signal.[27]

Table 7 reports results from six different counterfactual experiments, with each row displaying one experiment. First, we examine the baseline scenario. The baseline shows of our model predicts that 54,844 deaths occurred in sample (matching the data, by construction) and that patients traveled an average of 24.7 minutes from their home ZIP Code to the hospital. 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

[27]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 constrant term to match the sample mean mortality.

Table 7: 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	54,762.8	0.0	24.61	0.00
(2) No mortality increase from capacity strain	49,855.7	−9.0	24.61	17.52
(3) Five extra ICU beds per hospital	54,149.2	−1.1	24.61	2.19
(4) Patients informed about quality including strain	54,863.8	0.2	24.64	−0.36
(5) Informed about quality incl. strain with 10× weight	52,760.8	−3.7	25.66	7.15
(6) Baseline information but with 10× weight on quality	58,998.0	7.7	26.63	−15.12

Note: Deaths are predicted totals across our sample of 407,059 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 6, panels A and B, column (5).

number of deaths falls to 49,915, a 9.0% reduction. This would result in $17.6 billion in welfare gains from 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 635 in our sample. This would result in $2.3 billion in welfare increases from lower mortality just for these individuals.

Our three final counterfactuals considers the equilibrium impact of different patient information. For our fourth experiment, we continue to assume that the true preferences are given by Table 6 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. This counterfactual then asks what would occur in equilibrium if patients used the actual hospital quality in making their choices. To do this, we iterate on patient choices—which depend on quality— and update 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.

This experiment shows that giving patients better information by itself does not reduce in-sample deaths in equilibrium. The reason for this is that the information causes patients to select more hospitals that were empty relative to our baseline. However, these hospitals then fill up and end up being worse. In the end, because patients are now more likely to go to hospitals that are *ex ante* worse quality, this ends up being slightly worse for patients: deaths go up slightly—by 0.14%—relative to the baseline, as does travel time. The easiest way for deaths to decrease in our model is for patients to substitute from hospitals with utilization rates over one to less crowded ones, and this does not occur often enough.

Our fifth experiment is similar to the fourth experiment but multiplies the coefficient on quality by ten relative to the estimated parameter. In other words, we assume patients have both more accurate *and* more precise information on hospital quality. The combination of more accurate information and more weight on this information leads to significantly decreases in mortality. Specifically, in-sample mortality drops by 3.7% which is worth $7.3 billion in mortality declines. Thus, in this experiment, there is net substitution towards hospitals that have lower capacity utilization. Travel time does increase modestly, going from 24.7 to 25.7 minutes relative to the baseline.

Finally, we ask 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 its levels in the absence of capacity strain. Here, even more patients pile on to these *ex ante* good hospitals, not accounting for the fact that they are causing capacity strain. In this case, better but inaccurate information causes mortality to rise by 7.9% and travel time to increase by over 2 minutes relative to the baseline.

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. Indeed, a potential reason why patients' choices reacted only weakly to hospital quality when empty is because they might have recognized that this information was relatively imprecise.

6 Conclusion

A well-established literature documents substantial heterogeneity in hospital performance across conditions but also finds that consumers and firms respond to information about quality when available. Using Texas hospital discharge data, this paper considers the impacts of hospital performance in reducing deaths for the most important infectious disease outbreak in a generation, the Covid-19 pandemic.

We focus on understanding how capacity strain and learning-by-doing affected hospital quality, the extent to which patients responded to information about hospital quality in their hospital choices, and the resulting equilibrium implications for mortality under alternative information and market structures. We estimate time-varying measures of hospital quality using a Bayesian inference method that controls for the low signal-to-noise ratio of mortality for any hospital in a given quarter and for the fact that patients with an unobservably high severity of illness may choose different hospitals from others.

We find that there is substantial heterogeneity in quality across hospitals: moving every patient from a below-median quality hospital to a median quality hospital would reduce inpatient mortality by 8.5%, which would imply 62,000 lives saved nationally. Capacity strain in the intensive care unit played a central role in lowering quality: lowering ICU utilization rates from 200% to 100% of capacity would have reduced hospital mortality rates by 0.73 percentage points. In a counterfactual where all patients went to a hospital without any capacity strain, inpatient mortality would fall by 9.0%. In contrast, we do not find any evidence of learning-by-doing: lagged Covid-19 volume does not predict a hospital improving its quality (even controlling for capacity strain).

We estimate a patient choice model, finding that patients appear to choose hospitals partly on the basis of their quality in the absence of capacity strain. In equilibrium, this then leads hospitals that would be high quality without capacity strain to become low quality. When we estimate a patient choice model with actual quality including capacity strain, in most specifications, patients appear to value *high* mortality hospitals.

We investigate equilibrium counterfactuals that provide patients information about true hospital quality. If patients were informed about true quality and weight this information

ten times as much as they currently weight quality information, then in-sample Covid-19 deaths would have been reduced by 3.7%, implying $97 billion in welfare gains when scaled nationally. However, if patients value quality when hospitals are empty and weight this information ten times as much as they currently weight quality information, then in-sample Covid-19 deaths would have *increased* by 7.9% implying a cost of $259 billion when scaled nationally. Finally, we investigate increasing the number of ICU beds by 5 at every Texas hospital, and find that this would reduce mortality by 1.2 percentage points, implying $2.4 billion in welfare gains in-sample for Texas.

Overall, our results suggest that market forces can encourage the provision of high quality care. However, they were limited during the Covid-19 pandemic by the absence of quality information that accounted for capacity strain. While accurate quality information is helpful, providing individuals with inaccurate quality information—in our case, quality without capacity strain—may actually lead to worse outcomes. Our results also bear on the production function for hospitals, showing that investment in expensive resources such as intensive care unit beds may be warranted.

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., Cisewski, J. A., Xu, J., and Anderson, R. N. (2023). Covid-19 mortality update—United States, 2022. *Morbidity and Mortality Weekly Report*, 72(18):493.

Allcott, H., Boxell, L., Conway, J., Ferguson, B., Gentzkow, M., and Goldman, B. (2024). What explains temporal and geographic variation in the early us covid-19 pandemic? *Review of Economic Design*, pages 1–41.

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. (2019). Estimating returns to hospital volume: Evidence from advanced cancer surgery. *Journal of health economics*, 63:81–99.

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.

Bravata, D. M., Perkins, A. J., Myers, L. J., Arling, G., Zhang, Y., Zillich, A. J., Reese, L., Dysangco, A., Agarwal, R., Myers, J., et al. (2021). Association of intensive care unit patient load and demand with mortality rates in us department of veterans affairs hospitals during the covid-19 pandemic. *JAMA network open*, 4(1):e2034266–e2034266.

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.

Dafny, L. and Dranove, D. (2008). Do report cards tell consumers anything they don't already know? the case of medicare hmos. *The Rand journal of economics*, 39(3):790–821.

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.

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.

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.

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. 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.

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*.

James, S., Lagerqvist, B., Lundborg, P., and Vikström, J. (2022). Learning-by-doing and productivity growth among high-skilled workers: Evidence from the treatment of heart attacks. Technical report, Working Paper.

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 us 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.

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. *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., Duck, W. M., and et al. (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. (2024). Rationing by race. NBER working paper- 30380.

Song, Z., Zhang, X., Patterson, L. J., Barnes, C. L., and Haas, D. A. (2021). Racial and ethnic disparities in hospitalization outcomes among Medicare beneficiaries during the covid-19 pandemic. In *JAMA Health Forum*, volume 2, pages e214223–e214223. American Medical Association.

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	CCN Match Sample	CCN Unmatch Sample
Total beds	169.756	174.080	143.900
Total ICU beds	16.708	16.829	15.980
Overall utilization rate (avg)	0.438	0.433	0.467
Adj. ICU utilization rate (avg)	1.531	1.553	1.398
Teaching hospital	0.083	0.040	0.340
For-profit	0.398	0.395	0.420
System member	0.688	0.666	0.820
Rural	0.167	0.167	NaN
Critical access	0.155	0.171	0.060
AMI mortality rate	0.129	0.128	0.134
CABG mortality rate	0.033	0.033	0.035
COPD mortality rate	0.085	0.084	0.089
Heart failure mortality rate	0.110	0.110	0.111
Pneumonia mortality rate	0.152	0.153	0.150
Stroke mortality rate	0.135	0.135	0.139
Surgical complication rate	0.160	0.160	0.157
CMS 1 star rating	0.023	0.026	0.000
CMS 2 star rating	0.195	0.172	0.400
CMS 3 star rating	0.358	0.375	0.200
CMS 4 star rating	0.276	0.284	0.200
CMS 5 star rating	0.148	0.142	0.200
Per-capita Medicare spending ($)	9,944	10,039	9,110
N Hospitals	349	299	50

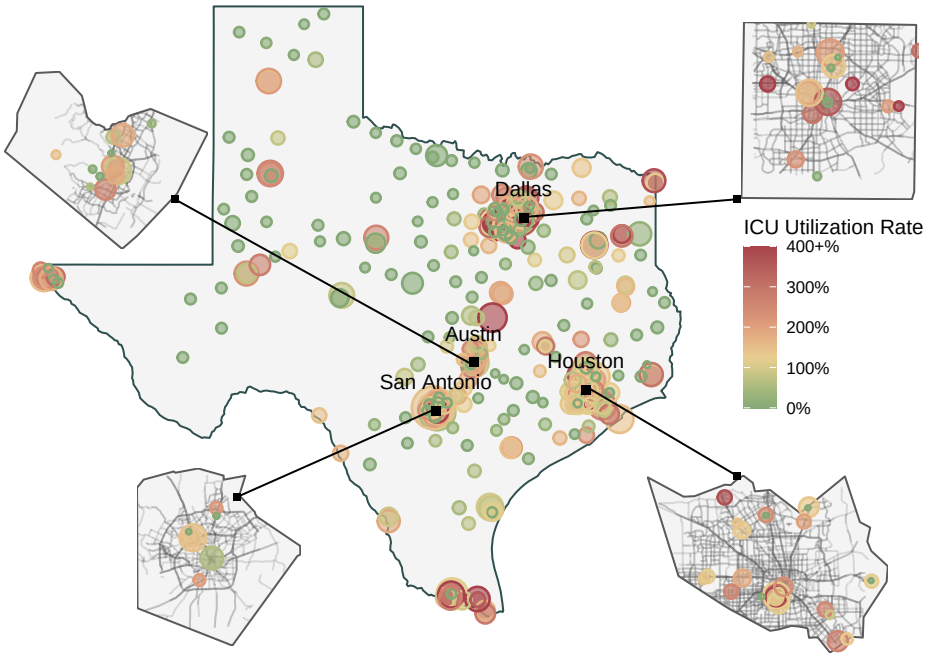
Note: Table of hospital characteristics means (excluding those with missing values) for the full sample of hospitals, for those with a CCN Match, and for those without a CCN match in which case hospital characteristics were determined via manual search. Hospital specific utilization rates are based on 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: Impact of Patient Characteristics on Mortality, γ

	Panel A			Panel B	
Risk adjuster	Coef.	SD	Risk adjuster (cont.)	Coef.	SD
Constant	-1.049	0.337	Pct Hispanic	-0.293	0.226
Male	0.128	0.139	Median Income	-0.299	0.219
White (Non-Hisp.)	-0.100	0.157	Dx: Blood Loss Anemia	-0.174	0.261
Black (Non-Hisp.)	-0.193	0.188	Dx: Coagulopathy	0.413	0.211
Hispanic	-0.089	0.235	Dx: COPD	-0.021	0.159
Age 25	-0.419	0.194	Dx: Depression	-0.090	0.193
Age 30	-0.338	0.251	Dx: Diabetes Type 1	-0.086	0.174
Age 35	-0.301	0.254	Dx: Diabetes Type 2	0.062	0.166
Age 40	-0.203	0.211	Dx: Fluid & Electrolyte Disorders	0.454	0.200
Age 45	-0.098	0.232	Dx: Hypertension	-0.203	0.211
Age 50	-0.038	0.246	Dx: Hypothyroidism	-0.042	0.163
Age 55	0.060	0.273	Dx: Iron Deficiency	0.134	0.229
Age 60	0.234	0.249	Dx: Kidney Disease	0.100	0.193
Age 65	0.361	0.325	Dx: Liver Disease	-0.001	0.189
Age 70	0.387	0.385	Dx: Lymphoma	0.017	0.270
Age 75	0.548	0.364	Dx: Metastatic Cancer	0.229	0.312
Age 80+	0.817	0.401	Dx: Neurological Disorders	0.124	0.201
Medicaid	-0.194	0.309	Dx: Obesity	0.063	0.236
Private (Managed Care)	-0.323	0.327	Dx: Paralysis	0.139	0.252
Medicare Advantage	-0.228	0.373	Dx: Peripheral Artery Disease	0.055	0.201
Blue Cross (Indemnity)	-0.218	0.250	Dx: Psychoses	-0.086	0.217
Medicare	-0.287	0.342	Dx: Pulmonary Circ. Disorders	0.151	0.242
Charity / Self-Pay	-0.289	0.286	Dx: Peptic Ulcer Disease	-0.084	0.324
Vac. Rate	-0.271	0.257	Dx: Solid Tumor	0.083	0.220
Transfer from SNF	0.357	0.456	Dx: Valvular Disease	-0.096	0.240
Pct 65+	-0.137	0.247	Dx: Weight Loss	0.266	0.239
Pct Non-Hisp. White	-0.321	0.219	Dx: Rheumatoid Arthritis	-0.034	0.224
Pct Non-Hisp. Black	-0.200	0.216	Dx: Chronic Condition (other)	0.173	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.

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.

Figure A2: Visualized Spline Coefficients from Table 3

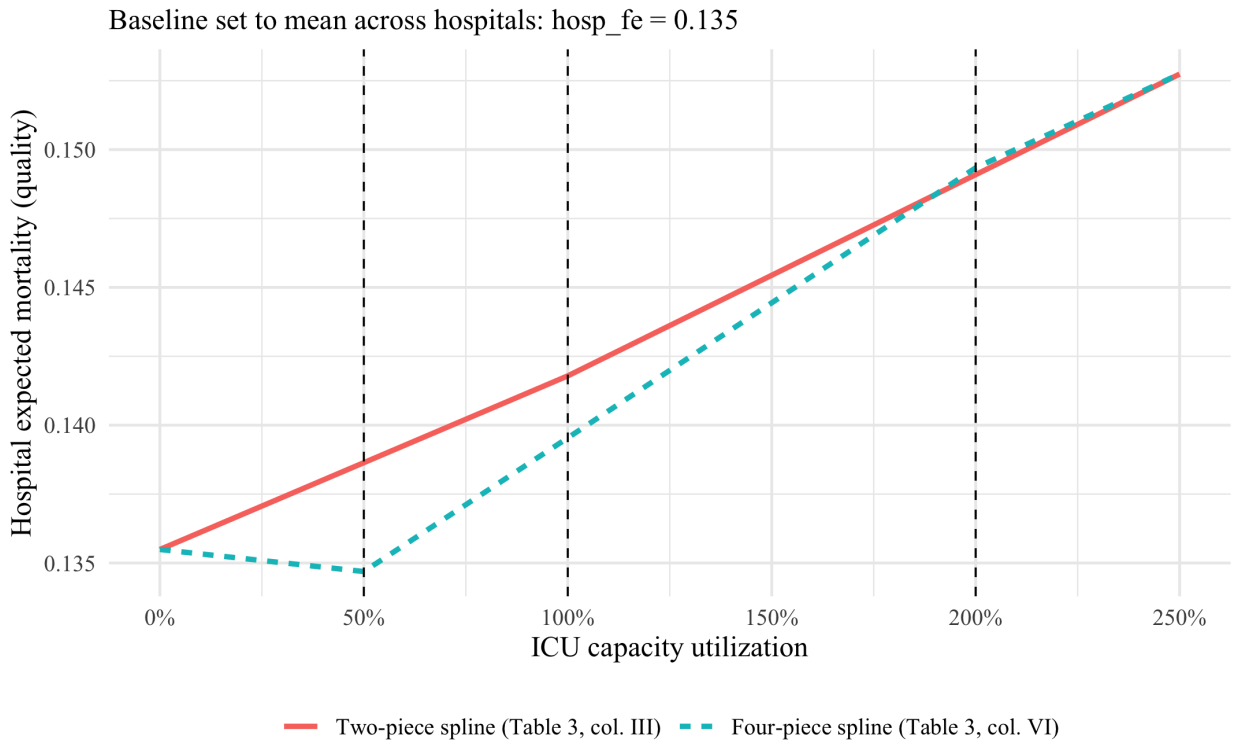
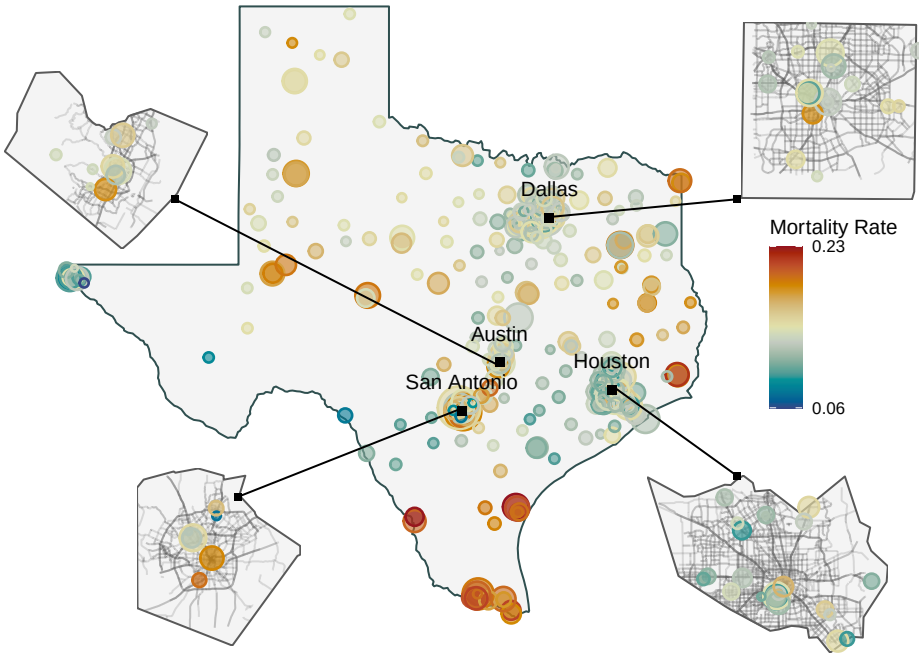


Figure A3: Hospital Mortality Across Hospitals at Observed Occupancy Levels



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.