Shocks to Hospital Occupancy and Mortality: Evidence from the 2009 H1N1 Pandemic*

Emilio Gutierrez† ITAM

Adrian Rubli‡ ITAM

May 2020

What is the effect of increased hospital occupancy on in-hospital mortality? Using data from public hospitals in Mexico, we exploit the shock in hospitalizations induced by the 2009 H1N1 pandemic, instrumenting hospital admissions due to acute respiratory infections (ARIs) with measures of ARIs at nearby healthcare facilities. We estimate that a one standard deviation increase in ARI admissions leads to a 38% increase in non-ARI in-hospital mortality. We show that these effects are non-linear in the size of the local outbreak, and are concentrated in hospitals with limited infrastructure. We discuss policy implications for the current pandemic.

JEL codes: I18, I14, I15, L23, H42

Key words: hospital occupancy; hospital mortality; pandemics; H1N1; health shocks; non-linearities

*This work was supported by the Asociación Mexicana de Cultura. All errors are our own.

†Instituto Tecnológico Autónomo de México, Department of Economics. Camino a Santa Teresa 930. Colonia Heroes de Padierna, Magdalena Contreras CDMX 10700, Mexico. Email: emilio.gutierrez@itam.mx

‡Corresponding author. Instituto Tecnológico Autónomo de México, Department of Business Administration. Camino a Santa Teresa 930. Colonia Heroes de Padierna, Magdalena Contreras CDMX 10700, Mexico. Email: adrian.rubli@itam.mx

1 Introduction

Hospital overcrowding may lead to lower quality of care due to smaller staff-to-patient ratios, fewer material resources per patient, and increased stress among personnel, all of which may be exacerbated in settings with limited hospital infrastructure. In the midst of the COVID-19 pandemic, estimates of the impact of overcrowding are a crucial piece of information for the managerial aspects of the mitigation strategy, such as designing the optimal allocation of patients to hospitals.

In this paper, we ask whether increased hospital occupancy has a causal impact on in-hospital mortality. We arguably overcome the empirical challenges faced by the existing literature by exploiting a dataset that is comprehensive of a large share of the Mexican hospital system, and leveraging arguably exogenous variation in hospital admissions driven by the 2009 H1N1 pandemic for identification.

Hospital management matters for hospital costs (Duggan, 2000; Schmitt, 2017; Capps et al., 2018), for performance (Dranove et al., 2003; Avdic et al., 2019), and ultimately, for patient outcomes. These issues may be of crucial importance during a pandemic, which induces important shocks in the flow of patients and their characteristics across the whole healthcare system. Hospital adjustments to these shocks may imply changes in the quantity and quality of inpatient care, resulting in adverse outcomes for patients.[1]

The causal estimation of overcrowding on patient outcomes is a difficult empirical task. The existing literature documents correlations between hospital overcrowding and mortality, mostly for developed countries (Madsen et al., 2014; Sprivulis et al., 2006), generally ignoring potential changes in patient selection, and treating variations in hospital admissions as exogenous (see Eriksson et al., 2017 for a review). The hospital management literature has done in-depth case studies of large hospitals (Kc and Terwiesch, 2012; Kim et al., 2015; Freeman et al., 2014). Their advantage lies in the level of detail in the information they exploit. However, they also tend to treat admissions as exogenous and, by focusing the analysis on single hospitals or hospital networks, may not be sufficiently informative in the context of a system-wide health shock.

[1] Among other adjustments, hospitals may use discretionary early discharge to manage patient overflow, which may lead to slight increases in mortality (Sharma et al., 2008; Berry Jaeker and Tucker, 2017; Eliason et al., 2018). Hospitals may also engage in capacity pooling by shuffling patients across different hospital areas in response to increased admissions, with an associated increase in length of stay and 30-day readmission (Song et al., 2019).

We exploit spatial and temporal variation in the severity of the 2009 H1N1 pandemic in Mexico as a source of arguably exogenous variation in hospital admissions to estimate the impact on non-ARI-related mortality. Focusing on the 2009 H1N1 pandemic in Mexico is important and informative for the current pandemic for several reasons. First, Mexico is a middle-income country with important limitations in healthcare infrastructure (OECD, 2005, 2016). Second, the H1N1 pandemic was a quasi-exogenous health shock, and may be informative of the expected effects of a similar ARI pandemic, such as COVID-19. Lastly, while we may be sacrificing precision due to the difference in timing and characteristics between H1N1 and COVID-19, the information that is currently available for the 2009 period is considerably more reliable than the real-time data which can be obtained as COVID-19 progresses.[2] Furthermore, the data for the H1N1 pandemic in Mexico has now been revised and vetted by researchers (Charu et al., 2011).

We follow an instrumental variables (IV) approach to estimate the impact of increased hospital admissions on mortality. For each public sector hospital in our data, we construct weekly measures of the number of ARI cases at nearby healthcare centers, using various criteria for defining the set of nearby centers. In all our regressions, coefficients are estimated off of within-hospital variation, by including hospital fixed effects, and account for over all common time trends with the inclusion of week fixed effects.

We begin by establishing a strong first stage, by regressing ARI hospitalizations on our instrument. Next, we present reduced-form results by estimating OLS regressions of non-ARI hospital admissions and mortality on neighboring facilities' ARI cases. We find no effects on non-ARI admissions, thus providing evidence that selection is unlikely to drive our results, and strong significant effects on mortality.

We then show IV estimates, instrumenting the hospitals' ARI admissions with the measures of nearby ARI cases. We find no effects on hospital admissions due to non-ARIs, which allows us to argue that selection and reverse causality are unlikely to be driving the rest of our results.[3] We then find that non-ARI hospital deaths increase by 38% when ARI admissions increase by one sd, or equivalently, a 12% increase in mortality due to a 50% increase in hospitalizations.

[2]See Manski and Molinari (2020) and Hortaçsu et al. (2020) for examples and discussions of the reliability of COVID-19 case load data for the US.

[3]The reverse causality issue may arise if increases in all-cause mortality lead to more medical emergency admissions (Jones, 2016).

We complement our main findings in two ways. First, we show that the mortality effect is non-linear in the size of the local ARI outbreak, with null effects for the first four quintiles and strong, significant effects for the top quintile. Moreover, we also show that the first stage is linear. These two results combined are consistent with the existence of tipping points in hospital management.[4] Second, we show that the mortality results are driven by smaller hospitals (fewer than 19 hospital beds) and by hospitals without an intensive care unit (ICU). Crucially for interpretation, all coefficients correspond to the marginal effect of a one sd increase in admissions, regardless of hospital infrastructure. Over all, these additional results suggest important non-linearities and emphasize the role of hospital infrastructure.

We contribute to the literature in at least three ways, all of which may be particularly relevant in light of the COVID-19 pandemic. First, we document the causal relationship between hospital occupancy and mortality in a developing country with limited healthcare infrastructure, which is likely to reach capacity as the COVID-19 pandemic develops.[5] Second, we exploit the variation in hospital admissions driven by the local ARI outbreaks associated with the 2009 H1N1 pandemic, which are uncorrelated with admissions due to other causes. Finally, we identify important non-linearities and heterogeneous impacts given hospitals' observable infrastructure, which may be crucial for the design of optimal patient allocation rules during the current pandemic.

Our findings have important policy implications for the management of pandemics. Quantifying the in-hospital mortality that is not directly related to the virus may provide valuable information on the costly negative externality of hospital admissions on mortality, which is vital for understanding the full death toll of a pandemic (Charu et al., 2011). Furthermore, given the non-linearities in the effect of admissions on mortality, there is an obvious potential mitigating policy: better allocation of patients across hospitals (Kuntz et al., 2015). This is particularly relevant in settings with limited hospital infrastructure, such as Mexico.

[4] We are not the first to show these non-linear effects. Kuntz et al. (2015) identifies safety tipping points in hospital utilization among 83 German hospitals, showing that high occupancy beyond the threshold increases in-hospital mortality. However, an important drawback is that this paper does not use an exogenous source of variation in hospital admissions.

[5] A somewhat similar paper, Guidetti et al. (2020) analyzes the effect of increased hospital occupancy due to air pollution shocks in Sao Paulo, Brazil, on hospital admissions for elective care, finding that hospitals absorb the excess demand by decreasing non-emergency procedures. The authors cannot look at mortality outcomes from unrelated hospitalizations precisely because in their setting the demand shock induces changes in hospital admissions.

2 The 2009 H1N1 Pandemic in Mexico

In March 2009, unusual increases in influenza-like illnesses were identified in Mexico. The initial outbreak in a small town in the southeastern state of Veracruz gave way to a generalized epidemic, with up to 32 thousand infected individuals by late April (Fraser et al., 2009; López-Cervantes et al., 2009). The cause of the rise in acute respiratory infections (ARIs) and complications in the form of atypical pneumonia was identified as a new strain of the influenza A virus called H1N1. Atypical pneumonia and H1N1 cases were also identified simultaneously in the United States, and the World Health Organization (WHO) declared a flu pandemic – the first of the 21st century – in June 2009 (Centers for Disease Control and Prevention, 2009).

The initial outbreak in Mexico evolved into a three-wave pandemic that affected the entire country with varying severity and at different moments of 2009 (Chowell et al., 2011). The first wave occurred during April and May, and mostly affected the Mexico City metropolitan area. The disease subsided, but then returned in a second, smaller wave in June and July, concentrated in the southeastern states. Lastly, there was a third wave, geographically widespread, during the last five months of the year.

Figure 1 shows nation-wide trends in ARIs for multiple years. Consistent with the literature (Charu et al., 2011), we focus on data for all ARIs since misdiagnosis and unconfirmed H1N1 cases may misrepresent the full extent of the pandemic. However, note that the overall pattern is maintained if we restrict to lab-confirmed H1N1 counts only. Figure 1 shows two large increases in ARIs during 2009 corresponding to the first and third waves of the pandemic. Importantly, these bouts are completely out of line with the usual epidemiological trends for ARIs in Mexico (Chowell et al., 2009).

At the onset, it was identified that individuals infected with H1N1 were typically characterized as young, with a median age in the low 40s (Chowell et al., 2009). Initial studies suggested that mortality rates among infected individuals that were suffering from critical illness were high, with around 40% of those patients dying (Domínguez-Cherit et al., 2009; Perez-Padilla et al., 2009).[6]

[6] Perez-Padilla et al. (2009) also estimates that around two thirds of critically ill hospitalized patients with H1N1 needed a ventilator to assist with breathing.

Over all, the H1N1 pandemic from April to December 2009 was associated with 12,000 excess deaths due to any cause and 4,200 pneumonia and influenza deaths (Charu et al., 2011).[7] This translates into a mortality burden up to 2.6 times greater than the typical influenza season, with important distributional changes due to the greater incidence on the younger population. Mortality in Mexico during the 2009 pandemic was among the highest in the world (Dawood et al., 2012).

Researchers identified early on that delayed access to inpatient treatment was a key risk factor for critically ill H1N1 patients (Echevarría-Zuno et al., 2009), leading to concerns about increased hospital occupancy and congestion. Many facilities implemented triage systems to better manage the inflow of patients (Rodriguez-Noriega et al., 2010). Various case studies point to the challenges faced by hospitals due to these increases in demand, and the potential mitigating policies like hospital reconversion to increase resources for the critically ill (Volkow et al., 2011; Serrano-Sierra et al., 2009). Note that this issue may be exacerbated by the low hospital supply and lack of human resources in Mexico (OECD, 2005).[8]

A series of policies were enacted by the government during the first wave of the pandemic to minimize the spread of the disease.[9] Social distancing measures – mostly in but not restricted to the Mexico City metropolitan area – included shutting down all schools, museums, libraries, night clubs, and generally any public spaces at the end of April for around 10-15 days (Córdova-Villalobos et al., 2009). Surveillance systems were implemented at transportation centers like bus terminals and airports (Alonso Reynoso, 2010). The government provided daily information briefings, emphasizing basic actions such as hand-washing and the use of face masks on public transportation (Agüero and Beleche, 2017; Condon and Sinha, 2010). The current response by the Mexican government to the COVID-19 pandemic is modeled around the strategy followed during the H1N1 pandemic of 2009.[10]

[7] Given the population in 2009, this is a mortality rate of 3.75 and 10.71 per 100,000, for ARIs and all causes, respectively.

[8] For example, according to the World Bank's Development Indicators, there were 1.7 hospital beds per 1,000 people in Mexico in 2010, compared to 3 in the US. Additionally, Mexico had 4.5 medical personnel (doctors, nurses, and midwives) per 1,000 people, compared to 15 for the US.

[9] The Mexican government had devised a preparedness plan in 2003, which was implemented in the form of the National Influenza Pandemic Plan, with coordination among different levels of government and the support of international organizations like the Panamerican Health Organization (Córdova-Villalobos et al., 2009).

[10] See, for example, http://187.191.75.115/gobmx/salud/documentos/manuales/24_Plan_Nal_Pandemia_Influenza.pdf, and https://businessinsider.mx/mexico-economia-pandemia-coronavirus-influenza-ah1n1/.

3 Data

We use three publicly-available datasets. First, we obtain data on all admissions at hospitals managed by the Ministry of Health (SSA). Note that the public healthcare system in Mexico is organized into separate and disjoint subsystems with their own set of providers, hospitals, and coverage plans. The three main public institutions are IMSS (for formal workers and their dependents), ISSSTE (for government workers and their dependents), and Seguro Popular (for informal workers and the unemployed). Seguro Popular provides service at SSA hospitals and clinics, and a small number of facilities run at the state level, which we include in our sample of SSA hospitals.[11]

There were a total of 660 SSA hospitals operating in 2009 according to our data. We drop 48 hospitals that had long strings (over 26 weeks) of zero hospitalizations, and construct a balanced panel of hospital-weeks using the date of admission (612 hospitals $\times$ 52 weeks). We construct counts of hospitalizations and deaths due to ARIs and all other causes (non-ARIs) from ICD-10 codes for the initial diagnosis recorded by physicians at the time of admission and the reason for discharge (death or improvement).

Our second dataset consists of epidemiological surveillance data from SSA recording all new cases of ARIs on a weekly basis (from ICD-10 codes). These data are gathered from each individual healthcare facility.[12] Note that all hospitals and clinics are geocoded.

We use this dataset to construct measures of ARI prevalence in the vicinity of each SSA hospital. In particular, we build these measures by assigning healthcare facilities to SSA hospitals in two ways. First, we consider the 10 healthcare facilities that are nearest to the SSA hospital, based on Euclidean distance. Second, we use data from all healthcare facilities that are located within a 5 km radius of the SSA hospital. For each definition of neighbors, we then count the total ARIs from these neighboring facilities for each week. In a robustness check in the online appendix, we further define neighbors as the five closest healthcare facilities, and as all facilities located within a 1 or 2 km radius.

[11] According to data compiled by the SSA, 50% of all public sector hospitalizations occurred at SSA hospitals. Note that the national health survey (ENSANUT, by its Spanish acronym) found that 21 and 17% of all hospitalizations occurred at privately-run hospitals in 2006 and 2012, respectively. This suggests that our data cover around 40-42% of all hospitalizations.

[12] Public healthcare facilities (both hospitals and clinics) are required to report by law. However, the data for the private sector may be less reliable since there is no enforcement for reporting.

For our last dataset, we recover data on hospital-level infrastructure for 2013. Unfortunately, this is the earliest year for which data at the hospital level are available. We will use the total number of hospital beds and the presence of an intensive care unit (ICU) for a heterogeneity analysis, noting that these characteristics are more likely to remain unchanged over time than staffing and equipment.

Table 1 presents summary statistics for various years. Panel A shows hospitalizations and deaths for the SSA hospitals by diagnosis (ARIs vs non-ARIs). Panel B shows prevalence of ARIs for different measures of hospital neighbors. Note that hospital outcomes and ARI prevalence were larger in 2009 than the previous two years. Lastly, Panel C shows infrastructure at SSA hospitals.

4 Empirical Strategy

We are interested in estimating the effect of increased hospital occupancy on hospital-level mortality outcomes, exploiting the H1N1 pandemic in 2009:

$$deaths_{jw} = \alpha\, hospARI_{jw} + \theta_j + \lambda_w + \varepsilon_{jw} \tag{1}$$

where $deaths_{jw}$ are non-ARI deaths in hospital j during week w, $hospARI_{jw}$ are hospital admissions due to ARIs, θ_j and λ_w are hospital and week fixed effects (FE), respectively, and ε_{jw} is the error term.

Focusing on ARI hospitalizations during the pandemic allows us to leverage shocks to hospital occupancy that were unexpected and uncorrelated with the underlying local demand prior to the shock. However, there is still an endogeneity issue that will potentially bias estimates of α since different hospitals in a given area may be attracting different types and quantities of patients (for example, hospitals with highly skilled physicians may be receiving more severely-ill patients).

To address this potential bias, we focus on our measures of ARIs at neighboring healthcare facilities instead of the hospitals' own ARI admissions. We first present reduced form estimates of the following form:

$$outcome_{jw} = \beta\, neighborARI_{jw} + \theta_j + \lambda_w + \nu_{jw} \tag{2}$$

where $outcome_{jw}$ is an outcome of interest, $neighborARI_{jw}$ are ARIs of neighboring healthcare facilities, and ν_{jw} is the error term. We cluster our standard errors at the hospital level to allow for serial correlation in the error term.

The inclusion of hospital FE implies that we identify β off of changes within each hospital, while the week FE account for over all seasonality throughout the year. Using ARIs from close neighbors allows us to isolate the effect of local ARI outbreaks from endogenous responses in patient load and composition at each SSA hospital.

As noted above, we are interested in the effect of increased hospital occupancy due to ARI outbreaks on mortality for non-ARIs, that is, the mortality spillovers from the pandemic. However, we first show evidence of a strong first stage (the effect of ARIs from neighbors on hospitals' own ARI admissions). We also show supporting evidence for the exclusion restriction, by estimating the impact of neighbors' ARIs on non-ARI admissions. If the pandemic is truly a quasi-random shock, neighboring ARIs should have no impact on SSA hospitals' non-ARI hospitalizations.

Given this approach, we then present instrumental variables (IV) estimates for equation 1, where the right-hand side variable ARI hospitalizations is instrumented with ARIs from neighbors. This estimate may be interpreted as causal if the exclusion restriction holds, namely that neighboring ARIs have no direct impact on mortality due to non-ARIs, except via their effect on ARI hospitalizations.

5 Main Results

For our main results, we present reduced-form estimates of equation 2, and IV estimates of equation 1 using generalized method of moments. Our sample consists of a balanced panel of 612 SSA hospitals that we observe over 52 weeks in 2009. As indicated, all regressions include hospital and week FE, and we calculate robust standard errors clustered at the hospital level. To account for size differences and patient burden across hospitals, we normalize all hospital admissions and deaths variables.[13]

[13]To be clear, we subtract the mean and divide by the standard deviation, so that the normalized variable has mean zero and a standard deviation of one. We use the mean and standard deviation across all years from 2007 to 2011 for this calculation. Results are very similar to just using 2009 data for the normalization procedure, as well as limiting to pre-pandemic data.

Table 2 presents our main results. Panel A assigns weekly ARIs from neighboring facilities by considering the sum over the 10 nearest healthcare centers. Panel B uses the sum of ARIs over all healthcare facilities within a 5 km radius. Columns 1-4 present our reduced-form estimates. Column 1 shows the effect of our (normalized) measure of neighboring ARIs on ARI hospitalizations at SSA hospitals. We estimate a strong first stage, indicating that a one standard deviation (sd) increase in ARIs at neighboring facilities leads to a 0.11 sd increase in the hospitals' own ARI admissions. Given the distribution of the (non-normalized) dependent variable, this implies an average increase of about 17%.

Column 2 shows an extension of the first stage, by estimating the effect of neighboring ARIs on in-hospital mortality due to ARIs. If neighboring ARIs lead to more ARI admissions, this should also lead to increases in ARI deaths (with potentially differing effects if there is patient selection and sorting). Our estimates show that this is indeed the case. A one sd increase in ARIs at nearby healthcare facilities leads to a 0.05 sd increase in ARI deaths, or around 15% on average.

Column 3 presents suggestive evidence that the exclusion restriction holds. For this, we consider hospital admissions due to non-ARIs. If the variation in neighboring ARIs is correlated with local demand conditions, then we would also see effects on non-ARI hospitalizations. Our estimates are an order of magnitude smaller, negative, and statistically indistinguishable from zero. This provides reassurance that the increases in ARIs as measured by counts at nearby facilities are exogenous to determinants of non-ARI deaths, except for their indirect effect via ARI hospitalizations that increase hospital occupancy.

Lastly, column 4 shows our main finding by focusing on non-ARI deaths as the dependent variable. We find a significant positive coefficient, indicating that a one sd increase in neighboring ARIs leads to a 0.02 sd increase in deaths due to non-ARIs. This estimate implies an average in-hospital mortality increase of around 4.5%.

We now turn to our IV estimates in columns 5 and 6 of Table 2, where we instrument the endogenous variable of ARI hospital admissions with our measure of neighboring ARI cases. We first replicate our reduced-form findings for hospital admissions due to non-ARIs, finding again null effects. Column 6 then shows a positive and significant IV estimate for the effect of increased ARI hospitalizations on in-hospital mortality due to non-ARIs. This coefficient implies that a one sd

increase in ARI hospital admissions leads to a 0.17 sd increase in deaths for non-ARIs, or on average an increase of about 38%.

We interpret this as the local average treatment effect, that is, the effect for compliers: SSA hospitals that increased their ARI admissions in response to an increase in nearby outbreaks of ARIs. We cannot speak to the effects for hospitals that would never (or always) increase their ARI patient load, regardless of the amount of ARI cases at neighboring facilities. Although this local effect may – to an extent – lack external validity, we believe that this is the parameter of interest for the expected effects during an ARI pandemic.

Our results are robust to alternative specifications. We explore alternative definitions of neighboring ARIs as described above. We also address two potentially confounding factors. First, one may worry that our results are driven by differential reporting patterns (and epidemiological trends) by states, particularly since each state is responsible for gathering and reporting local information. To address this, we include state-week fixed effects, allowing for a fully flexible differential time trend for each state. Second, one may worry about differential effects due to the composition of neighboring healthcare facilities, especially since the public healthcare system is made up of different subsystems. To address this, we interact indicators for quintiles of the share of neighboring facilities that are also managed by SSA with indicators for each week. See the online appendix for these IV estimates.

6 Non-Linearities and Heterogeneity

Having established in-hospital mortality spillovers from increases in ARI hospitalizations, we now ask whether these results may be non-linear in the size of the local ARI outbreak. Understanding if there is a tipping point for this effect is important for optimal policy design in terms of hospital management during a pandemic. In practice, we are interested in estimating the effect by quintile of our measure of neighboring ARIs. We present reduced-form evidence for this exercise, since an IV estimation would require a strong instrument for each quintile. For this, we estimate equation 2, substituting the explanatory variable with indicators for each quintile of the normalized measure of neighboring ARIs. We treat the third quintile as our excluded category.

We present our results graphically in Figure 2. Each coefficient series corresponds to a different definition of neighboring ARIs. The bars show 90 and 95% confidence intervals from robust standard errors clustered at the hospital level. The plot on the left shows the first stage, with ARI hospitalizations as the dependent variable. Our estimates are significant and increasing. This indicates that SSA hospitals that were surrounded by larger ARI outbreaks experienced larger increases in ARI admissions themselves. Furthermore, the effects are fairly linear.

The plot on the right in Figure 2 shows the reduced-form estimates for deaths due to non-ARIs. Our coefficients for the first, second, and fourth quintile are all small and statistically indistinguishable from zero. However, the estimate for the fifth quintile is positive and significant. This indicates that SSA hospitals with large increases in ARIs at nearby healthcare facilities were associated with increases in in-hospital mortality for non-ARIs. Given the first stage results, this is not driven by non-linearities in hospitalizations. Hence, this result suggests the existence of important non-linearities in the effect of increased hospital occupancy on mortality, consistent with the literature that identifies tipping points for the effect of overcrowding (Kuntz et al., 2015).[14]

We show further evidence by considering heterogeneity by hospital infrastructure. For these results, we present IV estimates of equation 1 by stratifying the sample. First, we consider quintiles of the total number of hospital beds, as a measure of the size of the hospital. We then stratify the sample by whether there is an ICU. Although hospitals without an ICU also tend to be smaller, there is sufficient variation suggesting that this is an informative result beyond the quintiles of hospital beds.[15]

Table 3 shows our estimates. As before, each panel corresponds to a different definition of the instrument. Columns 1-5 stratify the sample by quintiles of total hospital beds. Our IV estimate is large and significant for the first quintile, with smaller and insignificant estimates for the remaining four quintiles. This means that for the same increase in ARI hospitalizations (a one sd), smaller hospitals (with 18 or fewer beds) were the ones that experienced increases in non-ARI mortality. This suggests that infrastructure plays a key role in the damaging effects of hospital overcrowding.

[14]Note that the medical literature has found overcrowding at hospitals that are at 85% capacity (Madsen et al., 2014). Although this is not a one-to-one mapping, our non-linearities in the top quintile are in line with these findings.

[15]Nurses and doctors at hospitals without an ICU may also lack training for providing intensive care. For example, a case study of a hospital in the southern state of Chiapas recounts the challenges of treating critically ill patients without an ICU, and without medical personnel trained in ICU protocols (Volkow et al., 2011).

To the extent that smaller hospitals may also be understaffed and perhaps lacking in medical supplies, this might also reflect other infrastructure shortcomings.

Columns 6 and 7 show the results separately for hospitals without and with an ICU, respectively. We find large and significant effects for SSA hospitals without an ICU, and smaller and insignificant estimates for those with an ICU. Part of this effect may be driven by the correlation with total hospital beds. However, this also suggests that hospitals that were unprepared for critically ill patients (since lacking an ICU also likely means a lack in protocols and training for treating these types of patients), were the ones that experienced higher death counts for non-ARIs.

A final word of caution is necessary for the interpretation of these last results as we may be misclassifying hospitals in this exercise since we observe hospital infrastructure in 2013. While we can only conjecture about the potential bias in our estimates induced by this shortcoming in the data, we argue that the most likely implication of only observing infrastructure in 2013 is that our estimates suffer from attenuation bias. First, we are more likely to misclassify hospitals in contiguous quintiles (i.e., it is unlikely that a hospital would be in the first quintile of size in 2009 and in the fifth in 2013). Given that the effects are concentrated in the first quintile, the only bias is likely coming from hospitals that became somewhat larger over time (or were downsized). Second, it is unlikely that a hospital would lose its ICU, so that our estimate captures a weighted average of mostly hospitals without an ICU and some with an ICU.

7 Discussion and Conclusion

This paper asks whether the challenges that arise from sudden increases in hospital occupancy may lead to increases in in-hospital mortality. We analyze the 2009 H1N1 pandemic in Mexico, and use the sum of ARI cases at neighboring healthcare facilities as an instrument for increases in hospital admissions. Our main finding is that a one sd increase in hospital occupancy leads to a 38% increase in in-hospital mortality for conditions unrelated to ARIs, or equivalently, that a 50% increase in hospital admissions implies a 12% increase in non-ARI hospital deaths. Importantly, we show that our instrument is only predictive of increased hospital admissions for ARIs, and not of other conditions.

We complement this result by showing that the relationship with mortality is non-linear in the size of the local ARI outbreak, even though there is a linear relationship between size of the outbreak and ARI hospitalizations. This is consistent with the negative impacts arising from hospital overcrowding. We also show that our effects are concentrated in hospitals that are small and that do not have an ICU, suggesting that hospitals lacking appropriate infrastructure are the ones overwhelmed by the shock.

Over all, our findings are particularly relevant for understanding the mortality externalities from hospital overcrowding during a pandemic. Since the onset of the COVID-19 pandemic, there have been many anecdotal reports of hospital congestion,[16] and temporary delays in critical care for non-COVID-19 patients.[17] Direct deaths from infected individuals are not the only death toll from pandemics, and identifying excess mortality matters for better policy design (Charu et al., 2011). Specifically, our findings are suggestive of policies such as providing more information for assigning patients to hospitals (especially in a public and centralized healthcare system like Mexico's), and calculating metrics of hospital congestion to help facilities identify and avoid reaching overcrowding. These policies may contribute to improving outcomes during the ongoing COVID-19 pandemic.

[16] See https://www.nytimes.com/interactive/2020/03/17/upshot/hospital-bed-shortages-coronavirus.html?searchResultPosition=18 for initial concerns in mid-March 2020 for the US, as well as https://www.nytimes.com/2020/05/08/world/americas/mexico-coronavirus-count.html?searchResultPosition=3 and https://aristeguinoticias.com/0805/mexico/39-hospitales-covid-19-saturados-en-la-zona-metropolitana-de-la-cdmx/ for Mexico.

[17] See, for example, https://www.nytimes.com/2020/05/09/health/hospitals-coronavirus-reopening.html for accounts of hospitals in the US shutting down elective procedures, and https://www.nytimes.com/2020/04/15/opinion/coronavirus-cancer-surgery.html?searchResultPosition=3 and https://www.nytimes.com/2020/04/20/health/treatment-delays-coronavirus.html for anecdotes of life-saving surgeries for conditions like cancer.

References

Agüero, J. M. and T. Beleche (2017). Health shocks and their long-lasting impact on health behaviors: Evidence from the 2009 H1N1 pandemic in Mexico. *Journal of health economics 54*, 40–55.

Alonso Reynoso, C. (2010). La influenza A (H1N1) y las medidas adoptadas por las autoridades sanitarias. *Desacatos* (32), 35–52.

Avdic, D., P. Lundborg, and J. Vikström (2019). Estimating returns to hospital volume: Evidence from advanced cancer surgery. *Journal of health economics 63*, 81–99.

Berry Jaeker, J. A. and A. L. Tucker (2017). Past the point of speeding up: The negative effects of workload saturation on efficiency and patient severity. *Management Science 63*(4), 1042–1062.

Capps, C., D. Dranove, and C. Ody (2018). The effect of hospital acquisitions of physician practices on prices and spending. *Journal of health economics 59*, 139–152.

Centers for Disease Control and Prevention (2009). Outbreak of swine-origin influenza A (H1N1) virus infection - Mexico, March-April 2009. *MMWR. Morbidity and mortality weekly report 58*(17), 467.

Charu, V., G. Chowell, L. S. Palacio Mejia, S. Echevarría-Zuno, V. H. Borja-Aburto, L. Simonsen, M. A. Miller, and C. Viboud (2011). Mortality burden of the A/H1N1 pandemic in Mexico: a comparison of deaths and years of life lost to seasonal influenza. *Clinical infectious diseases 53*(10), 985–993.

Chowell, G., S. M. Bertozzi, M. A. Colchero, H. Lopez-Gatell, C. Alpuche-Aranda, M. Hernandez, and M. A. Miller (2009). Severe respiratory disease concurrent with the circulation of H1N1 influenza. *New England Journal of Medicine 361*(7), 674–679.

Chowell, G., S. Echevarría-Zuno, C. Viboud, L. Simonsen, J. Tamerius, M. A. Miller, and V. H. Borja-Aburto (2011). Characterizing the epidemiology of the 2009 influenza A/H1N1 pandemic in Mexico. *PLoS medicine 8*(5).

Condon, B. J. and T. Sinha (2010). Who is that masked person: the use of face masks on Mexico City public transportation during the influenza A (H1N1) outbreak. *Health Policy 95*(1), 50–56.

Córdova-Villalobos, J. A., E. Sarti, J. Arzoz-Padrés, G. Manuell-Lee, J. R. Méndez, and P. Kuri-Morales (2009). The influenza A (H1N1) epidemic in Mexico. Lessons learned. *Health research policy and systems 7*(1), 21.

Dawood, F. S., A. D. Iuliano, C. Reed, M. I. Meltzer, D. K. Shay, P.-Y. Cheng, D. Bandaranayake, R. F. Breiman, W. A. Brooks, P. Buchy, et al. (2012). Estimated global mortality associated with the first 12 months of 2009 pandemic influenza A H1N1 virus circulation: a modelling study. *The Lancet infectious diseases 12*(9), 687–695.

Domínguez-Cherit, G., S. E. Lapinsky, A. E. Macias, R. Pinto, L. Espinosa-Perez, A. de la Torre, M. Poblano-Morales, J. A. Baltazar-Torres, E. Bautista, A. Martinez, et al. (2009). Critically ill patients with 2009 influenza A (H1N1) in Mexico. *Jama 302*(17), 1880–1887.

Dranove, D., D. Kessler, M. McClellan, and M. Satterthwaite (2003). Is more information better? The effects of "report cards" on health care providers. *Journal of political Economy 111*(3), 555–588.

Duggan, M. G. (2000). Hospital ownership and public medical spending. *The Quarterly Journal of Economics 115*(4), 1343–1373.

Echevarría-Zuno, S., J. M. Mejía-Aranguré, A. J. Mar-Obeso, C. Grajales-Muñiz, E. Robles-Pérez, M. González-León, M. C. Ortega-Alvarez, C. Gonzalez-Bonilla, R. A. Rascón-Pacheco, and V. H. Borja-Aburto (2009). Infection and death from influenza A H1N1 virus in Mexico: a retrospective analysis. *The Lancet 374*(9707), 2072–2079.

Eliason, P. J., P. L. Grieco, R. C. McDevitt, and J. W. Roberts (2018). Strategic patient discharge: The case of long-term care hospitals. *American Economic Review 108*(11), 3232–65.

Eriksson, C. O., R. C. Stoner, K. B. Eden, C. D. Newgard, and J.-M. Guise (2017). The association between hospital capacity strain and inpatient outcomes in highly developed countries: a systematic review. *Journal of general internal medicine 32*(6), 686–696.

Fraser, C., C. A. Donnelly, S. Cauchemez, W. P. Hanage, M. D. Van Kerkhove, T. D. Hollingsworth, J. Griffin, R. F. Baggaley, H. E. Jenkins, E. J. Lyons, et al. (2009). Pandemic potential of a strain of influenza A (H1N1): early findings. *science 324*(5934), 1557–1561.

Freeman, M., N. Savva, and S. Scholtes (2014). Decomposing the effect of workload on patient outcomes: An empirical analysis of a maternity unit. Technical report, Working paper.

Guidetti, B., P. Pereda, and E. R. Severnini (2020). Health shocks under hospital capacity constraint: Evidence from air pollution in Sao Paulo, Brazil. Technical report, IZA Working Paper.

Hortaçsu, A., J. Liu, and T. Schwieg (2020). Estimating the fraction of unreported infections in epidemics with a known epicenter: an application to covid-19. Technical report, NBER Working Paper.

Jones, R. (2016). Hospital bed occupancy and deaths (all-cause mortality) in 2015. *British Journal of Healthcare Management 22*(5), 283–285.

Kc, D. S. and C. Terwiesch (2012). An econometric analysis of patient flows in the cardiac intensive care unit. *Manufacturing & Service Operations Management 14*(1), 50–65.

Kim, S.-H., C. W. Chan, M. Olivares, and G. Escobar (2015). ICU admission control: An empirical study of capacity allocation and its implication for patient outcomes. *Management Science 61*(1), 19–38.

Kuntz, L., R. Mennicken, and S. Scholtes (2015). Stress on the ward: Evidence of safety tipping points in hospitals. *Management Science 61*(4), 754–771.

López-Cervantes, M., A. Venado, A. Moreno, R. L. Pacheco-Domínguez, and G. Ortega-Pierres (2009). On the spread of the novel influenza A (H1N1) virus in Mexico. *The Journal of Infection in Developing Countries 3*(05), 327–330.

Madsen, F., S. Ladelund, and A. Linneberg (2014). High levels of bed occupancy associated with increased inpatient and thirty-day hospital mortality in Denmark. *Health Affairs 33*(7), 1236–1244.

Manski, C. F. and F. Molinari (2020). Estimating the COVID-19 infection rate: Anatomy of an inference problem. *Journal of Econometrics*.

OECD (2005). Reviews of Health Systems Mexico. Technical report, Organization for Economic Cooperation and Development, Paris.

OECD (2016). Reviews of Health Systems Mexico. Technical report, Organization for Economic Cooperation and Development, Paris.

Perez-Padilla, R., D. De La Rosa-Zamboni, S. Ponce de Leon, M. Hernandez, F. Quiñones-Falconi, E. Bautista, A. Ramirez-Venegas, J. Rojas-Serrano, C. E. Ormsby, A. Corrales, et al. (2009). Pneumonia and respiratory failure from swine-origin influenza A (H1N1) in Mexico. *New England journal of medicine 361*(7), 680–689.

Rodriguez-Noriega, E., E. Gonzalez-Diaz, R. Morfin-Otero, G. F. Gomez-Abundis, J. Briseño-Ramirez, H. R. Perez-Gomez, H. Lopez-Gatell, C. M. Alpuche-Aranda, E. Ramírez, I. López, et al. (2010). Hospital triage system for adult patients using an influenza-like illness scoring system during the 2009 pandemic Mexico. *PLoS One 5*(5).

Schmitt, M. (2017). Do hospital mergers reduce costs? *Journal of health economics 52*, 74–94.

Serrano-Sierra, A., V. Pérez-Robles, M. Nava-Frías, and M. Villa-Guillén (2009). Experiencia en el Hospital Infantil de México Federico Gómez ante la epidemia por el virus FluA/SW H1N1: Reporte preliminar. *Boletín médico del Hospital Infantil de México 66*(4), 301–305.

Sharma, R., M. Stano, and R. Gehring (2008). Short-term fluctuations in hospital demand: implications for admission, discharge, and discriminatory behavior. *The Rand journal of economics 39*(2), 586–606.

Song, H., A. L. Tucker, R. Graue, S. Moravick, and J. J. Yang (2019). Capacity pooling in hospitals: The hidden consequences of off-service placement. *Management Science*.

Sprivulis, P. C., J.-A. Da Silva, I. G. Jacobs, G. A. Jelinek, and A. R. Frazer (2006). The association between hospital overcrowding and mortality among patients admitted via Western Australian emergency departments. *Medical Journal of Australia 184*(5), 208–212.

Volkow, P., E. Bautista, M. de la Rosa, G. Manzano, M. V. Muñoz-Torrico, and R. Pérez-Padilla (2011). La respuesta de la unidades de terapia intensiva frente a la pandemia de influenza AH1N1 2009: la experiencia en Chiapas, México. *Salud Pública de México 53*(4), 345–353.

Table 1:
Summary Statistics

	2009 only	2007-2008	2010-2011	2013
A. Hospital-week outcomes				
Total hospitalizations	80.68	77.41	83.98	
	(97.61)	(86.78)	(101.23)	
Hospitalizations due to ARIs	3.60	3.08	3.39	
	(5.68)	(4.82)	(5.70)	
Hospitalizations due to non-ARIs	75.43	67.31	76.59	
	(93.96)	(83.02)	(97.07)	
Total deaths	1.69	1.67	1.74	
	(3.64)	(3.40)	(3.60)	
Deaths due to ARIs	0.23	0.17	0.21	
	(0.70)	(0.53)	(0.63)	
Deaths due to non-ARIs	1.43	1.35	1.45	
	(3.22)	(3.00)	(3.18)	
B. Clinic-week ARIs of assigned neighbors				
Main definitions:				
ARIs of 10 nearest neighbors	498.89	378.34	428.38	
	(626.11)	(444.46)	(518.49)	
Neighboring ARIs within 5 km	1143.56	857.16	975.15	
	(2000.61)	(1452.30)	(1643.84)	
Alternative definitions:				
ARIs of 5 nearest neighbors	299.64	231.84	260.69	
	(372.55)	(272.80)	(316.77)	
Neighboring ARIs within 1 km	171.69	134.43	150.65	
	(278.72)	(204.46)	(235.71)	
Neighboring ARIs within 2 km	350.75	269.28	302.50	
	(520.04)	(380.30)	(438.95)	
C. Hospital-level infrastructure				
Total beds				72.47
				(98.57)
Hospital has ICU				0.28
				(0.45)
Total beds in ICU				2.68
				(6.74)
Observations	31,824	63,648	63,648	612

Notes: This table presents summary statistics for 2009 and other years. Means are shown, with standard deviations in parentheses. Panel A shows hospitalizations and deaths at the hospital-week level. Panel B shows total weekly ARI cases for various definitions of neighboring healthcare facilities. Panel C shows infrastructure at the hospital level (these data are only available for 2013).

Table 2:
Effect of ARI Outbreaks on Hospitalizations and Mortality

	OLS				IV	
	Hosp. ARIs	Deaths ARIs	Hosp. Non-ARIs	Deaths Non-ARIs	Hosp. Non-ARIs	Deaths Non-ARIs
	(1)	(2)	(3)	(4)	(5)	(6)
A. 10 nearest neighbors						
Neighboring ARIs	0.112***	0.049***	-0.001	0.019***		
	(0.009)	(0.008)	(0.008)	(0.006)		
Hospitalizations due to ARIs					-0.001	0.172***
					(0.068)	(0.059)
Observations	31,824	31,824	31,824	31,824	31,824	31,824
R-squared	0.167	0.040	0.393	0.055		
F statistic					355.3	355.3
B. Neighbors within 5 km						
Neighboring ARIs	0.113***	0.051***	-0.005	0.020***		
	(0.009)	(0.008)	(0.008)	(0.006)		
Hospitalizations due to ARIs					-0.043	0.179***
					(0.068)	(0.059)
Observations	31,824	31,824	31,824	31,824	31,824	31,824
R-squared	0.167	0.040	0.393	0.055		
F statistic					350.8	350.8
Mean dependent variable	3.60	0.23	75.43	1.43	75.43	1.43
SD dependent variable	5.68	0.70	93.96	3.22	93.96	3.22

Notes: This table presents estimates of the effect of ARI outbreaks. Panel A assigns neighboring ARIs based on the 10 nearest healthcare facilities. Panel B uses all healthcare facilities within a 5 km radius. Columns 1-4 show OLS estimates of the outcome variable on the (normalized) measure of neighboring ARIs. Columns 5-6 present IV estimates, instrumenting the hospital's own (normalized) ARI admissions with the (normalized) measure of neighboring ARIs. Regressions include hospital and week FE. Standard errors are clustered at the hospital level. The mean and standard deviation of the dependent variable are shown.

Table 3:
Effect of ARI Outbreaks on Non-ARI Mortality by Hospital Infrastructure

	Quintiles of number of beds					Has ICU	
	q1	q2	q3	q4	q5	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. 10 nearest neighbors							
Hospitalizations due to ARIs	0.560**	0.049	0.180	0.171	0.182	0.242***	0.107
	(0.271)	(0.147)	(0.124)	(0.105)	(0.145)	(0.079)	(0.101)
Observations	6,500	6,240	6,552	6,188	6,344	22,984	8,840
F statistic	22.7	42.4	88.3	100.5	84.4	190.7	142.0
B. Neighbors within 5 km							
Hospitalizations due to ARIs	0.541*	0.124	0.183	0.164	0.134	0.250***	0.107
	(0.277)	(0.154)	(0.130)	(0.105)	(0.137)	(0.082)	(0.098)
Observations	6,500	6,240	6,552	6,188	6,344	22,984	8,840
F statistic	21.4	36.4	86.6	108.6	78.1	179.5	140.2
Mean dependent variable	0.34	0.11	0.51	1.35	4.87	0.50	3.83
SD dependent variable	0.98	0.39	0.90	1.56	5.64	2.00	4.33

Notes: This table presents IV estimates of the effect of ARI outbreaks, stratifying the main sample by measures of hospital infrastructure. Panel A assigns neighboring ARIs based on the 10 nearest healthcare facilities. Panel B uses all healthcare facilities within a 5 km radius. Columns 1-5 show results for each quintile of the distribution of hospitals by total number of beds. Columns 6-7 stratify hospitals by whether they have an ICU. Regressions include hospital and week FE. Standard errors are clustered at the hospital level. The mean and standard deviation of the dependent variable are shown.

Figure 1:
Epidemiological Trends of ARIs

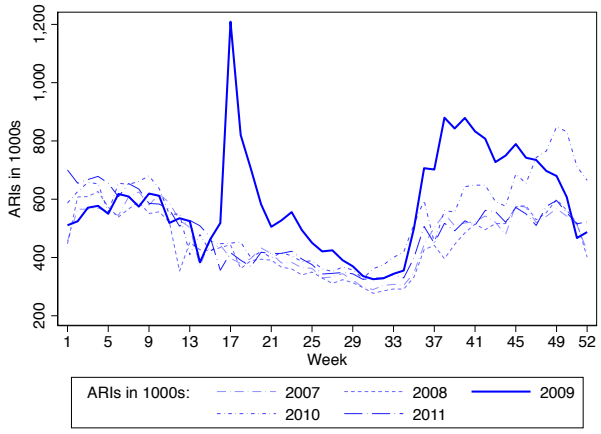

Notes: This graph shows the nation-wide epidemiological trends of ARIs for each year from 2007 to 2011.

Figure 2:
Effect of ARI Outbreaks on Non-ARI Mortality by Size of the Outbreak

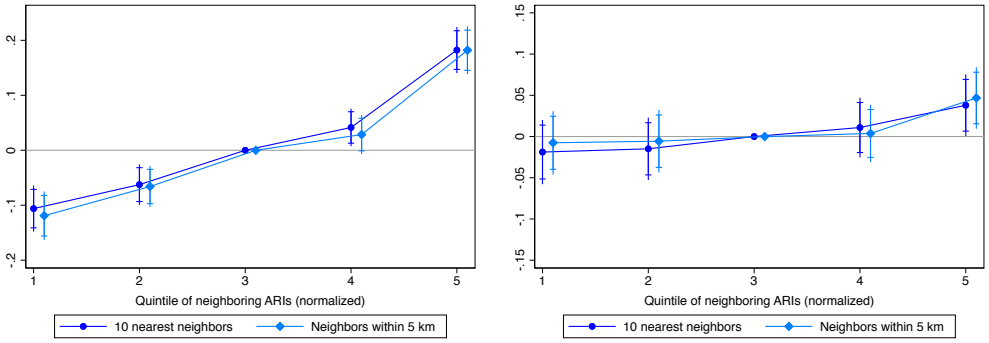

(a) Hospitalizations due to ARIs (b) Deaths due to Non-ARIs

Notes: These graphs show OLS estimates of hospitalizations due to ARIs (left) and deaths due to non-ARIs (right) on the (normalized) measure of neighboring ARIs by quintiles. Each coefficient series corresponds to a different definition of neighboring healthcare facilities. Regressions include hospital and week FE. Bars correspond to 90 and 95% confidence intervals, from standard errors clustered at the hospital level.

Appendix for Online Publication

Supplementary Tables and Figures

Table A1:
Descriptives of Assigned Neighbors

	Main definitions		Alternative definitions for robustness		
	10 nearest neighbors	Neighbors within 5 km	5 nearest neighbors	Neighbors within 2 km	Neighbors within 1 km
Total neighbors assigned	6.58	12.67	4.25	4.49	2.57
	(3.18)	(14.72)	(1.21)	(3.56)	(2.10)
Average distance to neighbors	1.86	2.25	1.53	1.00	0.54
	(0.94)	(1.01)	(0.92)	(0.43)	(0.24)
Share SSA neighbors	0.55	0.55	0.54	0.52	0.50
	(0.27)	(0.25)	(0.29)	(0.33)	(0.40)
Share IMSS neighbors	0.28	0.28	0.29	0.30	0.29
	(0.22)	(0.21)	(0.24)	(0.29)	(0.35)
Share ISSSTE neighbors	0.11	0.11	0.12	0.12	0.15
	(0.13)	(0.13)	(0.15)	(0.18)	(0.26)
Share private healthcare neighbors	0.03	0.02	0.02	0.03	0.03
	(0.07)	(0.06)	(0.09)	(0.10)	(0.12)
Share neighbors other public institutions	0.04	0.03	0.03	0.03	0.04
	(0.09)	(0.08)	(0.10)	(0.09)	(0.14)
Observations	612	612	612	573	455

Notes: This table presents summary statistics of the neighbors assigned to hospitals under each of the two main definitions used in the text, and the three alternative definitions presented as robustness checks. Means are shown, with standard deviations in parentheses. SSA are healthcare facilities run by the Ministry of Health (same as the hospitals in our sample). IMSS and ISSSTE are government healthcare systems for formal workers and government workers, respectively.

Table A2:
Correlation between ARI Hospitalizations and Outcomes

	Deaths ARIs (1)	Hosp. Non-ARIs (2)	Deaths Non-ARIs (3)
Hospitalizations due to ARIs	0.226***	0.042***	0.019***
	(0.014)	(0.008)	(0.007)
Observations	31,824	31,824	31,824
R-squared	0.090	0.394	0.055
Mean dependent variable	0.23	75.43	1.43
SD dependent variable	0.70	93.96	3.22

Notes: This table presents the correlation between (normalized) ARI hospitalizations and other (normalized) outcomes from an OLS regression with hospital and week fixed effects. Standard errors are clustered at the hospital level. The mean and standard deviation of the dependent variable are shown.

Table A3:
Robustness Checks

I. Alternative definitions of neighbors

	Hosp. Non-ARIs			Deaths Non-ARIs		
	5 nearest neighbors	Neighbors within 1 km	Neighbors within 2 km	5 nearest neighbors	Neighbors within 1 km	Neighbors within 2 km
	(1)	(2)	(3)	(4)	(5)	(6)
Hospitalizations due to ARIs	0.005	-0.065	-0.027	0.144**	0.249**	0.122*
	(0.077)	(0.113)	(0.079)	(0.061)	(0.100)	(0.067)
Observations	31,824	23,660	29,796	31,824	23,660	29,796
F statistic	309.2	118.7	261.5	309.2	118.7	261.5
Mean dependent variable	75.43	77.60	77.27	1.43	1.51	1.47
SD dependent variable	93.96	97.01	95.51	3.22	3.42	3.28

II. Additional controls

	Hosp. Non-ARIs		Deaths Non-ARIs	
	State ×week FE	SSA share ×week FE	State ×week FE	SSA share ×week FE
	(7)	(8)	(9)	(10)
A. 10 nearest neighbors				
Hospitalizations due to ARIs	-0.0195	-0.0158	0.207**	0.169***
	(0.114)	(0.070)	(0.100)	(0.059)
Observations	31,824	31,824	31,824	31,824
F statistic	114.1	339.4	114.1	339.4
B. Neighbors within 5 km				
Hospitalizations due to ARIs	-0.0818	-0.0205	0.224**	0.180***
	(0.121)	(0.069)	(0.106)	(0.061)
Observations	31,824	31,824	31,824	31,824
F statistic	98.0	328.3	98.0	328.3
Mean dependent variable	75.43	75.43	1.43	1.43
SD dependent variable	93.96	93.96	3.22	3.22

Notes: This table presents a series of robustness checks on the main results. The first part of the table shows IV estimates under different definitions of neighbors, for non-ARI hospitalizations (columns 1-3) and deaths (columns 4-6), instrumenting (normalized) hospitalizations due to ARIs with the (normalized) measure of neighboring ARIs. The second part of the table returns to the baseline definitions of neighbors (10 nearest neighbors in panel A, and all neighbors within 5 km in panel B). Columns 7 and 9 include additional controls in the form of indicators for each week interacted with indicators for each state. Columns 8 and 10 include indicators for each week interacted with indicators for quintiles of the share of neighboring healthcare facilities that belong to SSA. Regressions include hospital and week FE. Standard errors are clustered at the hospital level. The mean and standard deviation of the dependent variable are shown.

Table A4:
Effect of ARI Outbreaks on Non-ARI Hospitalizations by Hospital Infrastructure

	Quintiles of number of beds					Has ICU	
	q1	q2	q3	q4	q5	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. 10 nearest neighbors							
Hospitalizations due to ARIs	0.239	-0.0964	0.176	-0.112	-0.0237	0.0563	-0.0445
	(0.222)	(0.176)	(0.129)	(0.144)	(0.173)	(0.086)	(0.127)
Observations	6,500	6,240	6,552	6,188	6,344	22,984	8,840
F statistic	22.7	42.4	88.3	100.5	84.4	190.7	142.0
B. Neighbors within 5 km							
Hospitalizations due to ARIs	0.171	-0.148	0.174	-0.105	-0.115	0.0260	-0.0740
	(0.220)	(0.197)	(0.131)	(0.141)	(0.166)	(0.089)	(0.120)
Observations	6,500	6,240	6,552	6,188	6,344	22,984	8,840
F statistic	21.4	36.4	86.6	108.6	78.1	179.5	140.2
Mean dependent variable	30.03	21.69	51.27	88.78	186.73	42.09	162.11
SD dependent variable	40.04	16.12	31.00	54.60	142.74	52.10	119.68

Notes: This table presents IV estimates of the effect of ARI outbreaks on non-ARI hospitalizations, stratifying the main sample by measures of hospital infrastructure. Panel A assigns neighboring ARIs based on the 10 nearest healthcare facilities. Panel B uses all healthcare facilities within a 5 km radius. Columns 1-5 show results for each quintile of the distribution of hospitals by total number of beds. Columns 6-7 stratify hospitals by whether they have an ICU. Regressions include hospital and week FE. Standard errors are clustered at the hospital level. The mean and standard deviation of the dependent variable are shown.

Figure A1:
Spatial Distribution of Hospitalizations, Deaths, and Neighboring ARIs

(a) Total hospitalizations

(b) Total deaths

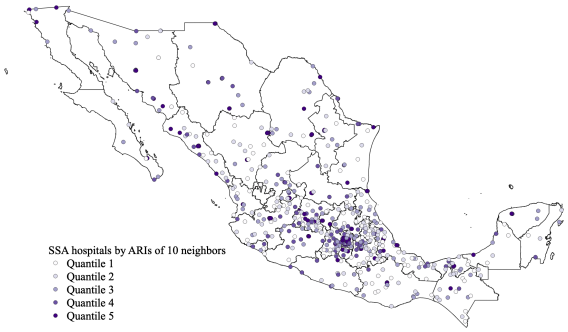

(c) ARIs from 10 nearest neighbors

(d) ARIs from neighbors within 5 km

Notes: These maps show the spatial distribution of SSA hospitals used in the estimating sample. The top left map classifies hospitals by quintiles of the total number of hospitalizations in 2009, the top right map by the total number of deaths, the bottom left map by ARIs from the 10 nearest neighboring healthcare facilities, and the bottom right map from all neighbors within 5 km.

Figure A2:
Distribution of Distance from Hospital to Assigned Neighbors

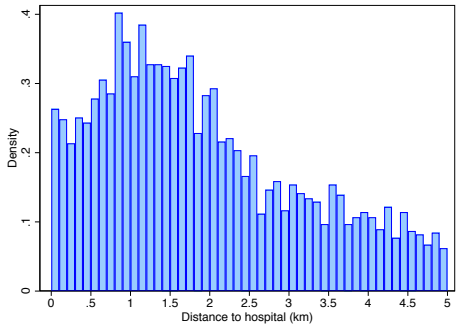

(a) 10 nearest neighbors

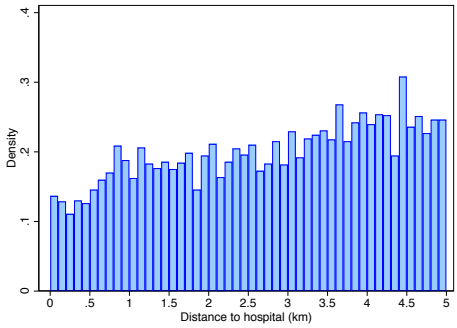

(b) Neighbors within 5 km

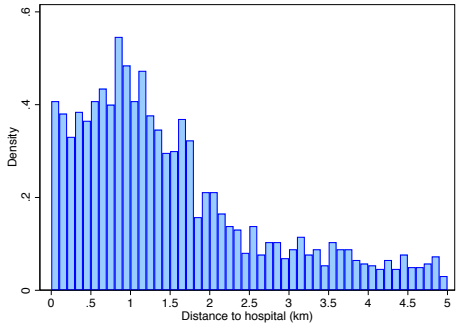

(c) 5 nearest neighbors

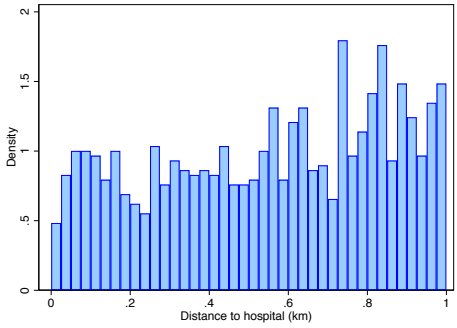

(d) Neighbors within 1 km

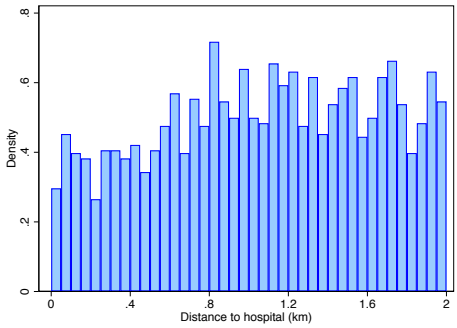

(e) Neighbors within 2 km

Notes: These graphs show histograms for the distribution of distance from the SSA hospitals to their assigned neighbors, under each of the five definitions. The first two definitions are the main ones used in the text, while the other three are used in the robustness checks.

Figure A3:
Correlation between ARIs from 10 Nearest Neighbors and All Neighbors Within 5 km

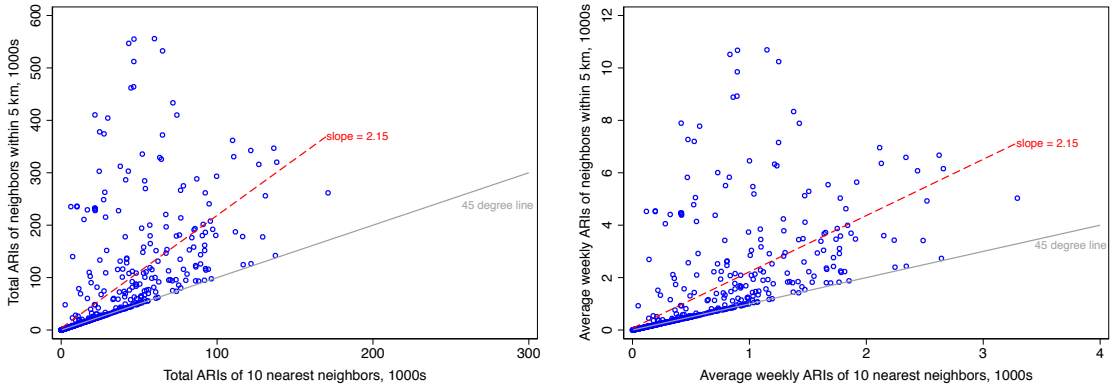

(a) Total ARIs in 2009 (b) Average weekly ARIs in 2009

Notes: These plots correlate the ARIs assigned from the 10 nearest neighbors to the hospital with those from all neighboring healthcare facilities within 5 km. The plot on the right uses the total ARIs reported in 2009, while the plot on the left calculates the average ARIs per week. The 45 degree line is shown, as well as the line of best fit from a simple OLS regression.

Figure A4:
Effect of ARI Outbreaks on Non-ARI Hospitalizations by Size of the Outbreak

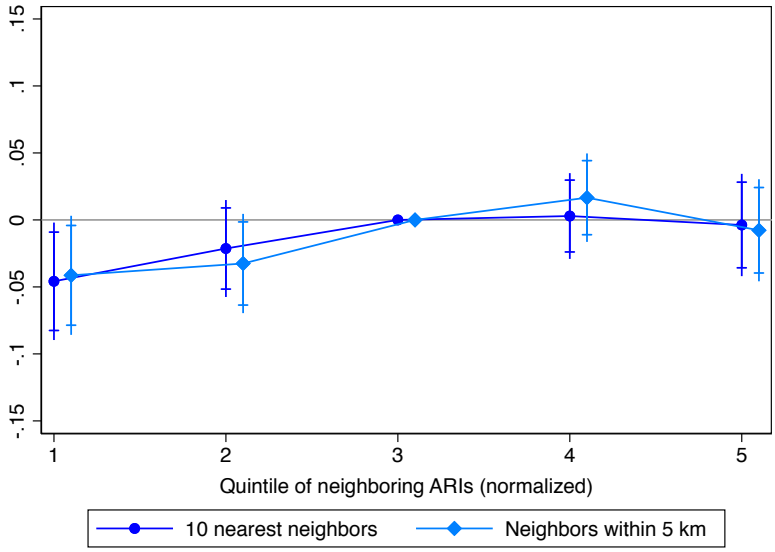

Notes: This plot shows OLS estimates of hospitalizations due to non-ARIs on the (normalized) measure of neighboring ARIs by quintiles. Each coefficient series corresponds to a different definition of neighboring healthcare facilities. Regressions include hospital and week FE. Bars correspond to 90 and 95% confidence intervals, from standard errors clustered at the hospital level.

Figure A5:
Total Hospital Beds and ICU Beds

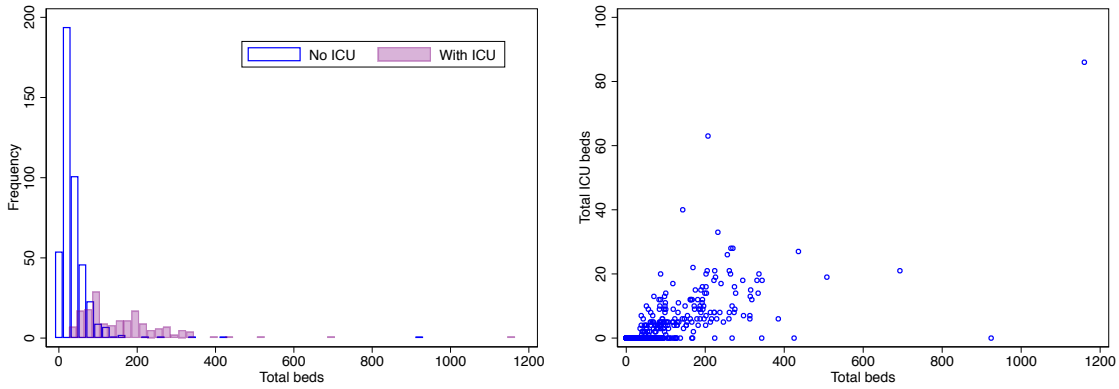

(a) Total hospital beds

(b) Total beds and total ICU beds

Notes: These plots describe hospital capacity as measured in 2013. The plot on the left shows histograms of total hospital beds, stratifying the sample between hospitals with and without an ICU. The plot on the right shows the correlation between total beds and total ICU beds for all hospitals.

Figure A6:
Correlation between Total Hospital Beds and Hospitalizations

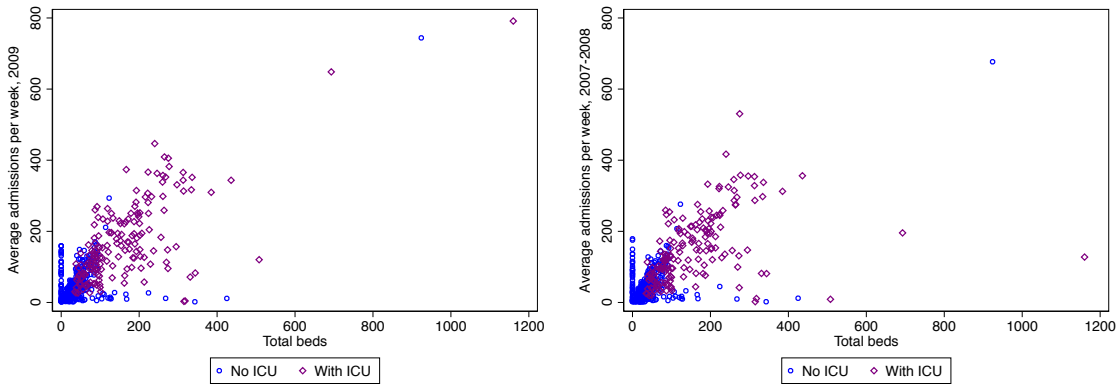

(a) Data for 2009

(b) Data for 2007 and 2008

Notes: These plots correlate the total number of hospital beds with the average weekly hospital admissions, distinguishing between hospital with and without an ICU. The plot on the left uses data from 2009, while the plot on the right uses pre-pandemic data.