

WORKING PAPER · NO. 2020-166

When Nurses Travel: Labor Supply Elasticity During COVID-19 Surges

Joshua D. Gottlieb and Avi Zenilman, R.N.

NOVEMBER 2020

When Nurses Travel: Labor Supply Elasticity During COVID-19 Surges*

Joshua D. Gottlieb
University of Chicago
and NBER

Avi Zenilman, R.N.
Yale University

November 19, 2020

Abstract

We study how short-term labor markets responded to an extraordinary demand shock during the COVID-19 pandemic. We study traveling nurse jobs—a market hospitals use to fill temporary staffing needs—to examine workers’ willingness to move to places with larger demand shocks. We find a dramatic increase in market size during the pandemic, especially for those specialties central to COVID-19 care. The number of jobs increased far more than compensation, suggesting that nurses’ willingness to travel is very responsive to compensation. To examine workers’ willingness to move across different locations, we examine jobs in different locations on the same day, and find even more responsive labor supply. We show that part of this supply responsiveness comes from workers’ willingness to travel longer distances for jobs when payment increases, suggesting that an integrated national market facilitates reallocating workers when demand surges. This implies that a simultaneous national demand spike might be harder for the market to accommodate rapidly.

*Gottlieb: jgottlieb@uchicago.edu. Zenilman: avi.zenilman@yale.edu. We are grateful to Health Carousel for providing data, to Daniel Sonnenstuhl, Scott Loring, and Wanran Zhao for excellent research assistance, and to the Becker-Friedman Institute for support. We thank Bill DeVille, Amie Ewald, Stephen Gottlieb, Wes Hamilton, Jeff Hicks, Brian LeCount, Neale Mahoney, Hugh Shiple, Janine Sulavik for helpful comments.

The flexibility to adjust staffing based on short-term shocks is an oft-cited reason for firms' increased use of outsourcing and temporary workers (Abraham and Taylor, 1996; Autor, 2003; Katz and Krueger, 2017). Alternatively, some analysts argue that the benefits of this short-term flexibility are overstated, and spot markets are popular because they enable employers to lower costs and avoid offering job security (Dube and Kaplan, 2010). Whether spot labor markets can really help accommodate transitory demand shocks depends on whether supply is also flexible—or can at least be reallocated to parts of the market where demand increases. We take advantage of the COVID-19 pandemic to study the flexible short-term labor market for nurses—a context where obtaining adequate staffing literally has life or death consequences.¹

The market for temporary nurses enables hospitals and other facilities to fill short-term nursing needs. In this market, hospitals post short-term job offers in a specified hospital unit, at a posted wage, and with specified work conditions. This market is worth \$10 billion in the United States (Landuis and Starkey, 2020), or approximately 7 percent of the overall labor market for nurses (Bureau of Labor Statistics, 2020). Travel nursing—where nurses are hired for multi-week stints (usually about 13 weeks), sometimes traveling across state lines or across the country—makes up the bulk of this market.

Travel nurses have the flexibility to accept or decline particular openings, so hospitals have to offer terms that can attract the staff they need. Nurses accept short-term postings at fluctuating prices and conditions they find acceptable. In this specific way, the markets for travel nurses and Uber drivers are similar. Unlike the ride-sharing market, the market for travel nurses is fragmented and intermediated by many different staffing firms, each finding nurses through its own network of recruiters. The hiring process also requires inter-

¹The literature on nurse staffing and quality finds mixed impacts on health outcomes (Aiken et al., 2002; Needleman et al., 2002, 2011; Cook et al., 2012; Spetz et al., 2013; Mark et al., 2013; Sloane et al., 2018), but the evidence from natural experiments comes from legal staffing ratio mandates at normal times, and the situation may be very different during a pandemic or other public health emergencies.

mediaries to evaluate the quality of the match, check nurses' skills, and manage regulatory hurdles such as state board licensing.

As COVID-19 surged in different parts of the country throughout the spring and summer of 2020, hospitals in affected regions needed additional nurses to manage the corresponding demand spike. We take advantage of these shocks to study the flexibility of this spot labor market. We look at different nursing specialties, some of which were central to coronavirus care, and others which were not. We find little-to-no increase in wages for nurses working in labor and delivery (L&D) units, which is unsurprising because the number of near-term pregnant women could not, and did not, exponentially increase between March and May 2020.

The patterns are different for intensive care unit (ICU) and emergency room (ER) nurses. Most acutely ill COVID-19 patients come to the hospital via the ER and many end up in the ICU. The number of job openings and the wage level for both specialties is positively associated with increased state-level COVID-19 case counts. Based only on the time series, one would estimate a labor supply elasticity to local wages of around 3; ICU jobs increased by 339 percent during the first wave of the pandemic, while compensation increased 50 percent. ER jobs increased by 89 percent while compensation increased by 27 percent.²

We then move to panel estimates, which capture nurses' ability to move across regions of the country. When we condition on state and calendar day fixed effects, we find an even larger supply elasticity of approximately 5 for ICU nurses. The difference with the time-series estimates is unsurprising; it suggests that nurses are quite elastic to compensation once they have decided to enter the market at a certain point in time.

But even this estimate would understate the local supply elasticity if part of the wage

²Using the numbers for ICU jobs, the implied labor supply elasticity is $\frac{\ln(339)-\ln(100)}{\ln(157)-\ln(105)} = 3$, and for ER jobs it is $\frac{\ln(189)-\ln(100)}{\ln(131)-\ln(103)} = 2.6$.

increase is compensation for increased risk of COVID-19 exposure. We design a simple empirical framework to adjust for this risk, by comparing high-demand specialties (such as ICU and ER) with L&D, whose demand does not increase in COVID-19 hotspots. We interpret increased compensation for L&D nurses as payment for the COVID-19 risk, which allows us to estimate supply based on the difference between compensation COVID-19 specialties and L&D nursing.

This simple framework suggests an even higher elasticity to local wages of at least 8 for ICU, and 6 for ER, and 7 for general hospital floor nurses. This highly elastic supply suggests that price signals are an effective way of reallocating labor to the parts of the country with increased staffing needs. To test this interpretation, we study the subset of jobs that are filled in our sample. For this subset, we can measure the distance between a nurse’s home and the job location. We find that workers accept postings farther from their homes when pay is higher. The large size of the United States, and nurses’ ability to work in different states, appears to be an important part of how this market adapted to the first waves of demand for COVID-19 nursing.

This raises serious questions about what happens in short-term staffing markets if there are simultaneous demand shocks in many different markets—such as when the pandemic surges simultaneously across the country. In this situation, reallocating workers to the highest-need places may be insufficient, and the market would need an increase in the total number of workers. This margin might be less responsive than the cross-region supply elasticity, but it still appears quite high.

While short-term work is thought to be widespread and growing (Autor, 2003; Collins et al., 2019; Katz and Krueger, 2019a,b), this market is difficult to study. Temporary workers and independent contractors are a blind spot in many labor market datasets (Abraham et al., 2017), and most of the evidence on staffing comes from employers’ perspective

(Abraham, 1988; Abraham and Taylor, 1996; Houseman, 2001; Dey et al., 2012). To study short-term workers’ behavior, the literature has turned to extremely short-term markets such as Uber drivers or online tasks (Hall et al., 2015; Angrist et al., 2017; Mas and Pallais, 2017; Chen and Sheldon, 2015; Chen et al., 2020; Farrell et al., 2018). But most spot markets for labor are generally for time spans longer than an Uber ride and may have deeper consequences for social welfare. This paper finds that labor supply is quite elastic in the temporary nursing market, and worker flexibility across states contributes to this responsiveness. These findings help understand how the market accommodates surges in healthcare demand due to COVID-19.

Section 1 discusses the institutional context of travel nursing and introduces our data. Section 2 introduces the conceptual and empirical framework we use to analyze them. Section 3 presents descriptive facts and basic time series patterns in our data, and uses them to compute an initial tentative supply elasticity. Section 4 presents our core empirical results. In section 6 we discuss the interpretation of these results and conclude.

1 Setting and Data

Registered nurses are by far the largest source of hospital labor—in 2019, American hospitals employed over 1.8 million registered nurses, compared to 120,000 physicians—and there has long been concern about a mismatch between the nursing workforce needed and the available labor supply (Buerhaus et al., 2000, 2009). When a hospital or individual hospital units faces an acute and temporary nursing staff shortage, they frequently access contingent labor through the \$10 billion market for temporary nurses. There are multiple ways to access this market, but the essential structure is similar: the hospital sets a price (a “bill rate” in industry terminology) and working conditions. An intermediary, such as a supplemental staffing agency, then tries to match the position with a registered nurse

who can fill the temporary gap. The nurses are generally considered employees of the intermediary, and this staffing firm provides benefits, liability insurance, and quality check. Hospitals thus outsource recruiting, licensing, and other HR tasks. Nurses that accept contracts for multiple weeks are generally called “travel nurses,” even though many come from the hospital’s local area. (In some cases, the hospitals hire them full-time after they finish their contract.)³

The price the hospital offers for travel nurse labor (the “bill rate”) has to cover the nurse’s wages and the intermediary’s fee, which includes benefits, housing stipends, transportation, and administrative costs.⁴ If the intermediary hires a secondary recruiter, it will have to split its share of the bill rate. There is usually some room for nurses to negotiate, and hospitals may offer other benefits such as a bonus or a housing allowance. Compensation and location are very salient in the recruiting process; they are posted prominently, along with the specialty and start date, in online job listings. Recruiters regularly solicit potential nurse recruits by sending them text messages describing potential jobs, often including the compensation in the initial solicitation.

Health Carousel, one of the ten largest healthcare staffing firms in the United States, provided data on job postings made available to its recruiters, and jobs it filled. We use two years of data, from September 1, 2018 through August 31, 2020, though many of our analyses focus on the COVID-19 pandemic, beginning February 1, 2020.

The postings are for registered nurses only, and exclude other professions such as licensed practical nurses or nurse practitioners. The information includes including postings from all fifty states and Washington, D.C. For each posting, the data report the specialty of

³Travel nurses are more likely to have a bachelor’s degree, and observational studies have found that, after adjusting for hospital quality, there are no deleterious outcomes associated with their use (Xue et al., 2012a,b).

⁴Some hospitals have exclusive relationships with certain staffing agencies that essentially operate their own self-contained marketplace, while others may post to independent vendor management systems for an auction.

nurse requested, number of nurses requested, job location and the compensation per nurse, called “bill rate” in the industry. This measure of compensation reflects the total amount that a hospital offers to pay for a nursing contract. It incorporates both the nurse’s earnings and the recruiter’s fee, and we scale it as an index relative to the nationwide average in early 2020.⁵ For ease of comparison, we also convert total counts of both job openings and jobs filled by Health Carousel’s recruiters (“completed jobs”) into indices relative to their early-2020 averages.

We measure compensation, job openings, and completed jobs were nationally and within subsamples by specialty and state. We define six categories of nursing specialties: Emergency Room (ER), Adult Intensive Care (ICU), standard hospital floors (Medical/Surgical or Telemetry), Labor and Delivery (L&D), Operating Room (OR), and Other.⁶ Within each subsample, we rescaled the job openings index to the subsample’s February 1 through March 14 mean. (Estimates for all 50 states can be found in the appendix.)

For completed jobs, the data reported the residential zip code of the nurse who filled the opening. We used this to compute the geodesic (straight-line) travel distance between the nurse’s residence and job location. We measured state and national COVID-19 incidence as the number of new cases reported daily (or, for some analyses, weekly) by the Johns Hopkins Coronavirus Resource Center.

2 Framework

Consider the market for nurses of specialty i in location j . At time t , the natural log of COVID-19 cases in location j is denoted by c_{jt} . We describe the natural log of nurse supply

⁵The index is normalized such that the mean from February 1 through March 14 is 100.

⁶The “other” category includes a broad range of specialties, including psychiatry, cardiac catheterization, pediatrics, administration, dialysis, skilled nursing facility, urgent care, and pediatric ICUs. Some of these may be relevant to COVID-19 care, while others are not.

as:

$$s^i(w_{jt}^i, c_{jt}) = \bar{\alpha} + \alpha w_{jt}^i + \beta c_{jt} + e_{jt}^i \quad (1)$$

where w_{jt}^i is the natural log of the compensation and e_{jt}^i is an orthogonal supply shock. Supply presumably responds positively to wages, and negatively to COVID-19 risk, so $\alpha > 0$ and $\beta < 0$.

We model the natural log of nurse demand as:

$$d^i(w_{jt}^i, c_{jt}) = \bar{\gamma} + \gamma w_{jt}^i + \delta^i c_{jt} + u_{jt}^i \quad (2)$$

where i indexes specialties and u_{jt}^i is an orthogonal demand shock. Demand presumably slopes down in the cost, so $\gamma < 0$. For some specialties, such as emergency medicine and critical care, $\delta^i > 0$. For other specialties, such as labor and delivery, it seems plausible that $\delta^i = 0$; the number of births, and hence demand for labor and delivery nurses, is determined far in advance of realized COVID-19 cases.

Setting the supply equal to demand yields the following equilibrium in this market:

$$w_{jt}^i = \frac{-\beta + \delta^i}{\alpha - \gamma} c_{jt} + \psi_0 + \frac{u_{jt}^i - e_{jt}^i}{\alpha - \gamma} \quad (3)$$

$$q_{jt}^i = \frac{\alpha \delta^i - \beta \gamma}{\alpha - \gamma} c_{jt} + \bar{\gamma} + \gamma \psi_0 + \frac{\alpha}{\alpha - \gamma} u_{jt}^i - \frac{\gamma}{\alpha - \gamma} e_{jt}^i \quad (4)$$

where $q_{jt}^i = s^i(w_{jt}^i, c_{jt}) = d^i(w_{jt}^i, c_{jt})$ is the number of jobs and $\psi_0 = \frac{\bar{\gamma} - \bar{\alpha}}{\alpha - \gamma}$ collects parameters.

The two relationships between wages and quantities, respectively, and the number of cases, are governed by four parameters—the supply and demand elasticities with respect to wages (α and γ), and with respect to COVID-19 conditions (β and δ^i)—and we only

have two equations. To understand the behavior of nursing supply thus requires additional data or additional assumptions.

We use both. First, consider a specialty such as L&D where demand is plausibly independent of COVID cases, i.e. $\delta^0 = 0$. This reduces the number of relevant parameters to 3. If we then add a second specialty, such as intensive care, where $\delta^1 > 0$, we have another set of two equations but only one additional parameter (δ^1). This gives us 4 equations—each of (3) and (4) for two specialties—to solve for four parameters ($\alpha, \beta, \gamma, \delta^1$).

Equations (3) and (4) suggest the following estimating equations:

$$w_{ijt} = \tau_0 + \tau_1 c_{ijt} + \pi_i \mathbb{1}_i c_{ijt} + \theta_j \mathbb{1}_j + \phi_t \mathbb{1}_t + \varepsilon_{ijt} \quad (5)$$

$$q_{ijt} = \kappa_0 + \kappa_1 c_{ijt} + \mu_i \mathbb{1}_i c_{ijt} + \rho_j \mathbb{1}_j + \sigma_t \mathbb{1}_t + \nu_{ijt} \quad (6)$$

where θ_j and ρ_j are state fixed effects, and ϕ_t and σ_t^i are time fixed effects. The τ_1 and κ_1 coefficients on c_{ijt} represent the estimates of (3) and (4) for the specialty unaffected by COVID-19. The π_i and μ_i coefficients for each other specialty represent the differences in the wage and quantity coefficients between the specialty in question and the unaffected specialty. For each specialty i , we solve analytically for the four parameters of interest ($\alpha, \beta, \gamma, \delta^i$) from the four relevant coefficients ($\tau_1, \pi_i, \kappa_1, \mu_i$) on c_{ijt} .

The fixed effects in regressions (5) and (6) absorb some of the variation in u_{jt}^i and e_{jt}^i from the equilibrium equations, increasing the plausibility of the identifying assumption that $\varepsilon_{ijt}, \nu_{ijt} \perp c_{ijt}$. We also consider specifications with different granularity of fixed effects, and others that estimate separate time and/or state fixed effects for each specialty.

Conceptually, these fixed effects ensure that we are identifying the model within state and time. Given our motivation—understanding short-term reallocation and flexibility in this market—we want precisely this sort of variation. We do not need to understand why baseline compensation is higher in New York than in Arizona (unions? living costs? regu-

lations?) in order to understand the within-state variation over time. Similarly, we do not need to know extensive margin elasticities over time (due to nurse training in the long run, or short-run entry into travel nursing) in order to understand allocation across space at a given point in time.

Our interpretation rests on a few assumptions. First, we need the parameters aside from δ^i to be constant across specialties. While this assumption is strong, its failure would cause errors in a predictable direction. To see this, imagine that we incorrectly imposed $\beta = 0$ —i.e. we assumed that nurses do not require compensation for the risk of treating COVID-19 patients. We could then compute the supply elasticity by taking the ratio of the coefficients on c_{jt} in equation (4) and in equation (3), which is:

$$\frac{\delta^i \alpha - \gamma \beta}{\delta^i - \beta} < \alpha. \quad (7)$$

If $\beta = 0$, we will incorrectly conclude that this coefficient is equal to α . It would appear that this ratio directly reveals the supply elasticity, when in fact it yields something less than the supply elasticity; in other words, supply is more elastic than it immediately appears to be. When we assume that the non- δ^i parameters are constant across specialties, this is effectively the error we are likely to make. That is, we assume ICUa and ER nurses require the same compensation for risk as L&D nurses, even though ICU and ER nurses are likely to deal with sicker, COVID-positive patients. If they demand extra compensation for that risk, their β would be more negative than we assume when we hold β constant across specialties. Our framework thus runs the risk of underestimating the compensating differential COVID-19 specialties require, and hence underestimating the supply elasticity. But it is unlikely to overestimate the supply elasticity.

At the same time, our second assumption is that the market is segmented between specialties. While this may not be perfectly true, the fact that compensation does diverge

between specialties implies at least some segmentation—as we would expect based on the genuine skill differences the different specialties require. To the extent nurses can substitute between specialties, this may help to explain the high supply elasticities we find.

Third, we assume demand for L&D postings is invariant to COVID-19 conditions. This does not mean that the equilibrium number of L&D nurses can't change—our model allows for movement along the demand curve. But shifts in the curve itself must be orthogonal to COVID-19 cases, conditional on our fixed effects. The lag inherent in pregnancy provides most of the justification for this assumption—our time period ends less than nine months after COVID-19 arrived in the United States, so the pandemic is unlikely to have affected fertility during this time period. L&D unit demand for travel nurses could still change if home births increased, or if hospitals reduced staffing levels due to budgetary shortfalls (Khullar et al., 2020). Despite these theoretical concerns, we find below that L&D job postings do not vary substantially with respect to COVID-19 cases, conditional on our fixed effects.

Finally, empirically we interpret the listed compensation from job postings as an equilibrium outcome. Even though posting a job does not guarantee that the position will be filled, the use of job postings has a strong tradition in labor economics (Davis et al., 2013; Kuhn and Shen, 2013; Rothwell, 2014; Hershbein and Kahn, 2018; Deming and Kahn, 2018; Forsythe et al., 2020), and we are able to demonstrate its appropriateness in our data. For the subset of data when we see the positions filled, the patterns look very similar to our broader results. We cannot replicate our most granular analyses exclusively with filled positions, but we are able to confirm the basic patterns from the postings analysis.

3 Time Series Patterns

Before delving into the regressions, we introduce descriptive patterns and present the basic time series relationships apparent in the data. Table 1 displays the descriptive patterns, nationally and in key subsamples.⁷ The baseline national indices for both job openings and compensation are normalized to 100. Focusing on the first row in Panel B, we see that the national job openings index increased to 165 in spring, fell to 47 in early summer, and increased again to 167 in late summer. The national compensation index was 133 in spring, 103 in early summer, and 114 in late summer.

The next three rows show results for three selected states: New York, Massachusetts, and Arizona. We choose these states because of their different experiences during the first 7 months of the pandemic: Massachusetts had an early spike; Arizona had a spike in late summer, and New York had a uniquely extreme spike in the spring. The baseline job openings index is normalized to 100 at baseline for each subsample, and the first column shows each state's share of the full sample. In New York, job openings increased to 863 in spring, fell to 49 in early summer, and recovered to 101 in late summer. In Massachusetts, job openings increased to 174 in spring, returned to 94 in early summer, and fell to 70 in late summer. In Arizona, the job openings index was stable until late summer, when it increased to 273. Appendix Table B.2 shows the patterns for all states.

Panel A of Figure 1 shows a smoothed time series of the national job openings index, overlaid with the time series of new COVID-19 cases. COVID-19 cases began increasing sharply after March 15, 2020 exceeding 30,000 daily by April 2. They gradually declined to 20,000 in early June, before increasing again to 67,000 by July 15. The national job openings index increased gradually during the pre-COVID part of our sample. It began at around around 70 early in our sample (late 2018), and was fluctuating between 80 and

⁷Appendix Table B.1 shows a corresponding table for the February-August 2020 subperiod.

130 from February 1 through March 15, 2020. Prior to COVID-19, there are notable seasonal patterns in December and January, but otherwise the secular increase dominates the pattern. This changed rapidly after the pandemic arrived. By April 6, the job openings index reached 300. It then declined rapidly, falling to 35 in late May, and stayed low until July, when outbreaks in the southern and southwestern United States peaked. The job openings index reached 170 by July 27, and rose to 185 by the end of August. It is clear that the openings closely track spikes in COVID-19 cases.

The graph shows a corresponding trend for New York, which experienced the largest, and most rapid, increase in COVID-19 care needs. The The COVID-19 daily incidence in New York (not shown) grew from 300 cases on March 15 to 10,000 by April 3. The job openings index, stable near its baseline value of 100 until early March, 2020, increased to over 2,500 during the first week of April. (Note that the graph shows the New York index divided by 10.) By April 15, it was 1,100, and ultimately returned to baseline on May 7. The job openings index remained low until the end of August, when it rose to 130.

Panel B shows the compensation index. The national compensation average was quite stable until early March, 2020. It then increased significantly, hitting 110 by March 26 and rising to a peak of 125 by April 10. It began to decline in May, and returned to baseline in June. It increased again in the second half of July and reached 115 by early August. New York's compensation index, which was 105 in February, started to trend upwards in March, was 140 by March 24, and then reached 165 within a week. It remained at around 160 through May, and then gradually declined to 100 by the end of July.

These time series patterns already provide preliminary indications of the supply elasticity. If the supply curve were constant (i.e. COVID-19 only affected labor demand), and the rate at which posted jobs are filled were constant, we could directly measure the supply elasticity as $\frac{\Delta \ln(\text{job openings index})}{\Delta \ln(\text{compensation index})}$ relative to the baseline. Using the national changes

from baseline to spring, we estimate a supply elasticity of $\frac{\ln(165) - \ln(100)}{\ln(133) - \ln(100)} = 1.8$.⁸ Using the estimates from New York in spring, the supply elasticity would be $\frac{\ln(863) - \ln(100)}{\ln(185) - \ln(106)} = 3.9$.

This interpretation is subject to two important caveats. First, posted jobs are not filled jobs. We will study job completions below to ensure that these estimates are not misleading for completed jobs.

Second, the supply curve may respond to COVID-19 cases. More cases increase nurses' exposure risk and likely make the work environment more difficult and less pleasant. Both of these forces would tend to depress supply precisely when demand is higher. On the other hand, if some nurses volunteer to work in COVID-19 hotspots due to altruism or a sense of professional obligation, that would increase supply. But this force seems unlikely to overwhelm the need for a compensating differential due to the risk and working conditions, an assumption that our subsequent results will support. As a result, this time series calculation is likely to underestimate the supply elasticity with respect to wages.

In order to address this, we delve into differences across nursing specialties. Figure 1 Panel D decomposes the job postings by specialty. We see the pandemic associated with dramatic increases in postings for jobs that deal with COVID-19 patients: intensive care (ICU), emergency room, and standard hospital floors (Medical/Surgical). In contrast, operating room and L&D postings decline during the initial pandemic wave.

Returning to Table 1, the final six rows summarize the descriptive patterns by nursing specialty. The first column shows the specialty's share of all job postings. The baseline openings index for each specialty was scaled to 100, while the baseline compensation index was scaled relative to the national mean. ICU, ED, L&D, and OR nurse postings had

⁸We divide the indices by 100 for this calculation, so subtracting the baseline value amounts to subtracting $\ln(1) = 0$.

compensation indices above 100 in February. Med/Surg had a mean of 95. In spring, the ICU job openings index more-than-tripled to 339, while the compensation index rose to 157. The ER job openings index increased to 189 and the compensation index rose to 131. The OR job openings index fell to 56 as many elective surgeries were canceled across the country.

Most helpful for our purposes is labor and delivery (L&D). Demand for L&D nurses depends on the number of births taking place, which—for our entire sample period—would have been determined before COVID-19 arrived in the United States. We see a 22% decline in L&D job openings to 78 and a slight increase in compensation from 110 at baseline to 115 in spring. These changes are inconsistent with a demand increase, as quantities decline while wages increase. These estimates require a negative supply shock, which we interpret as nurses demanding risk compensation.

4 Results

Table 2 shows the results of our explicit supply estimation. Columns 1 and 2 regress log job postings, and log compensation, respectively, against log cases at the state-by-day level. These columns do not distinguish among specialties; this would be the right approach if markets for different specialties are fully integrated. We estimate an elasticity of 0.17 of job postings with respect to COVID-19 cases, and a compensation elasticity of 0.03. These regressions would make sense if COVID-19 is purely a demand shock for nursing, and supply is fixed. In the language of our framework, these regressions assume $\beta = 0$ —nurses do not require any compensation for the risk of working in a COVID-19 hotspot. Under these assumptions, the supply elasticity is simply the ratio of the quantity and compensation coefficients, which is 5.6.

Since nurses might plausibly require a compensating differential for the risk of work-

ing in a COVID-19 hotspot, the remaining columns implement specifications (5) and (6) suggested by the model. Columns (2) through (10) estimate different versions of these specifications. For each type, we present both a wage and a quantity specification. Each specification shows the relationship between the dependent variable and the log number of new cases, along with interactions for different nursing specialties. The omitted interaction is labor and delivery, so the coefficient on log cases alone can be interpreted as the relationship between the L&D market outcomes and COVID-19 cases. We cluster standard errors by state. Across the specifications, we see insignificant positive estimates between COVID-19 cases and L&D job postings or wages.

The subsequent rows show the incremental relationships for other specialties. Looking at column (3), we see significant positive interactions for the quantity of job postings in ICU, emergency room, and medical/surgical jobs. If we add the ICU coefficient of 0.17 to the baseline coefficient of 0.03, the total elasticity for ICU jobs is 0.2. This means that a doubling of COVID-19 cases in a state is associated with a 14 percent increase in ICU job postings ($\ln(2) \times 0.2 = 0.14$). The magnitudes of these relationships are somewhat smaller for emergency and medical/surgical jobs, consistent with COVID-19’s particular increase in critical care needs.

Figure 2 visually plots the relationships between labor market outcomes and COVID-19 cases. Each panel shows a binned scatterplot for each specialty, after conditioning on day and state fixed effects. Panel A shows the relationship for the quantity of job postings, and the differences across specialties are visually apparent. L&D and operating room job postings are nearly flat with respect to local COVID-19 conditions. The slopes for emergency room and miscellaneous other jobs is clearly positive, and the slopes for ICU and medical/surgical jobs are even steeper, as Table 2 column (3) shows quantitatively. This is consistent with our interpretation that demand for some specialties does not respond to

local COVID-19 cases, while that for others does.

Panel B shows a corresponding graph for compensation. Here we see that all of the specialties' compensation offers slope upwards with respect to COVID-19 cases. But the slope of the relationship differs across specialties: ICU, emergency room, and medical/surgical jobs have the steepest relationships, while L&D and operating room are the flattest. These differences, combined with the quantity slope differences from Panel A, will allow us to solve for the supply and demand parameters below.

The remaining columns in Table 2 show the robustness of our estimates to different sets of controls. Columns (5)–(6) add state-by-specialty fixed effects. Columns (7)–(8) further add day-by-specialty fixed effects; these columns are fully interacted, in the sense that all independent variables are interacted with specialty. These are thus equivalent to completely separate regressions by specialty, so offer the data maximum flexibility. The results are remarkably stable under these different sets of controls. In all cases, the case-quantity and case-compensation relationships are significantly steeper for ICU and medical/surgical postings than for the omitted category, labor and delivery. The magnitudes for emergency room postings are slightly lower, and sometimes indistinguishable from those for L&D.

Below the estimates, we show the supply parameters that our regressions imply. These calculations use our framework from section 2 to solve for the key parameters in equations (2) and (1). We interpret our regressions as empirical implementations of equations (3) and (4), for pairs of specialties at a time. We treat labor and delivery as the non-COVID specialty (i.e., $\delta^{L\&D} = 0$) and sequentially compare it with ICU, emergency, and medical/surgical markets.

Combining the estimates from columns (3) and (4) for ICU nurses implies a supply elasticity of 6.7 with respect to wages. The corresponding elasticities are 3.9 for ER jobs, and 6.1 for medical/surgical. The supply elasticities are slightly lower in columns (5) and (6),

and increase again in columns (7) and (8). These are our most flexible specifications, with state-by-specialty and date-by-specialty fixed effects, and the estimated supply elasticities return to 6.2, 4.2, and 5.1 for ICU, ER, and medical/surgical respectively.

To ease interpretation, consider a worker looking for a short term job at a given point in time, and facing two comparable job offers, but with one paying twice as much as the other. A supply elasticity of 6.2 implies that 4.3 times as many workers choose the higher-paid job.⁹ Put another way, in a pool of such choices, 81 percent of workers would choose the job paying double, while 19 percent would choose the job paying half.¹⁰

The higher elasticities here compared with the time series results make perfect sense: We are zooming in to capture supply elasticities across regions at a given point in time. This framework effectively considers a market with a fixed national supply of workers and labor demand—even within a given specialty. Thus the only choice for a worker is which job to accept. In this context, a large positive supply elasticity is not surprising. In fact, an infinite supply elasticity may be just as natural a benchmark for this margin of response as any other. The whole point of the temporary job market is to enable workers to maximize short-term earnings, so it should not be any surprise that they overwhelmingly choose the better-paid opening. This market is designed to encourage this behavior; compensation is very salient in the recruiting process.

These elasticities are higher than those estimated in many other settings, but the response margins available in those other settings are more limited. For instance, Angrist et al. (2017) estimate an intertemporal labor supply elasticity among Uber drivers of 1.2, but this requires drivers to change when they work. Using short-term variation in compensation, both Chen and Sheldon (2015) and Chen et al. (2020) estimate mean supply elasticities, centered around 0.5. Moving beyond Uber drivers, Mas and Pallais (2017) find

⁹ $\ln(2) \times 6.2 = 4.3.$

¹⁰ $\frac{4.3}{1+4.3} = 0.81.$

similar estimates for temporary remote work, and show that workers have a meaningful, decreasing valuation of non-work time. These elasticities do not measure decisions about which work to accept, conditional on taking a job—for instance, in which part of a city an Uber driver chooses to look for rides. But in our context, that simple location choice takes on first-order importance.

In contrast with this literature, our estimates do not require workers to substitute between working at different times, or away from home production. In our specifications, workers’ outside option is taking a similar job, for the same temporary time period, but in a different state. Given the similarity of the choices—at least once we have adjusted for the compensation required to offset the different COVID-19 risk—it makes sense that the different choices are close substitutes.

To test this interpretation, we now turn to the subset of data for which we observe the worker who fills the job.

5 Evidence from Job Completion

We use data on completed jobs for two purposes. First, we show that the key patterns from job postings data hold up for completed jobs. Second, we test our interpretation of the results so far as reflecting workers’ mobility across different states.

The first piece of evidence that completed jobs follow the same patterns as job postings comes from Table 2 columns (9) and (10). These columns report regressions analogous to those in columns (3) through (8), but with more aggregated data to account for the reduced sample. These regressions aggregate to the weekly instead of daily level, and combine specialties to further increase precision. Since demand for operating room jobs follows similar pattern to that for labor and delivery, we combine those two specialties. Since medical/surgical, ER, and ICU jobs all play key roles in COVID-19 care, we also

combine them.

The results from this analysis are similar to those for job postings. For the non-COVID specialties, quantities are flat with respect to COVID-19 cases. For the COVID-related specialties, the elasticity of jobs with respect to COVID-19 cases is 0.13. The compensation elasticity for non-COVID specialties is 0.04, while that for COVID-related specialties is around 0.08. When we solve for the supply and demand parameters, we find a supply elasticity of 3.9 with respect to compensation, and -0.17 with respect to disease prevalence.

Figure 3 Panel A shows the time series patterns for filled jobs. Using smoothed weekly data, the figure shows completed jobs more than doubling from early 2020 to the peak in April. Job completions fall during the early summer, reaching below the baseline value of 100 in June, before climbing again in late summer. Compensation followed a similar pattern, increasing during the pandemic's initial phase, falling during the summer, and beginning to increase again in August.

For completed jobs, we see an additional characteristic that can help us understand the market outcomes: the worker's home location and the location of the job that worker accepted. We compute the distance between these two points, and plot the average as another time series in Panel A. We see the average distance increasing from 500 to 700 miles at the beginning of the pandemic, before falling back to 600 throughout the summer. This provides initial evidence that the market accommodated the surge in demand through workers' willingness to travel across states and change their work locations.

The remaining panels dig into this phenomenon further. Panel B shows the time-series relationship of distance traveled with respect to the bill rate. Since this relationship does not rely on COVID-19 cases, here we use an expanded sample beginning in July 2018. The y -axis shows the average log distance traveled to a job posting, and the x -axis shows the average log bill rate for these jobs. We see a sharp positive relationship, with an elasticity

of distance with respect to compensation of around 1.3. The months when bill rates are higher, nurses travel farther.

Panels C and D show these relationships using panel variation. Panel C shows a binned scatterplot of log distance travel against the log compensation, after partialing out week and state fixed effects. Panel D changes the distance measure to an indicator for workers traveling from outside of their home state. We see large positive relationships in both panels. The elasticity of distance with respect to compensation is approximately 1.2, very similar to the time series value of 1.3. The semi-elasticity of out-of-state travel with respect to compensation is 0.4, so each 10 percent increase in compensation is associated with a 4 percentage point increase in the share of workers accepting a job from out-of-state. Since 60 percent of workers come from out-of-state on average, the elasticity of cross-state travel with respect to compensation is two-thirds ($0.4/0.6$).

To tie this back to demand for COVID-19 care, Panels E and F change the x -axis to log COVID-19 cases. Panel E shows an elasticity of 0.12 between distance traveled to completed jobs and COVID-19 cases in a state. Panel F shows a semi-elasticity of 0.03 between traveling across states and COVID-19 cases, which corresponds to an elasticity of 0.05.

6 Discussion

Putting these results together, the national labor market for short-term nurse staffing appears to have very elastic supply. In response to price signals, workers choose jobs in different parts of the country—but not with perfect flexibility. Given the nature of this choice, why is supply not perfectly elastic with respect to wages?

Even for short-term traveling work, workers appear to value proximity to home. The elasticities of distance traveled, and of leaving one's home state, with respect to com-

compensation are quite positive—but far from infinite. There may also be match-specific reasons that labor supply is not infinitely responsive to price signals. Even within the specialty categories we study, nurses can subspecialize in types of healthcare care that are not perfect substitutes. When staffing companies consider nurse placements in a particular position, they evaluate other aspects of the nurse’s skills, and whether they are a good fit—even something as banal as experience with the hospital’s brand of electronic health record. The labor supply elasticity may also be depressed by regulatory hurdles, such as state-specific licensure requirements, which can delay nurses’ mobility.¹¹ New York, California, and other states temporarily liberalized their licensure during the acute phase of the COVID-19 pandemic, perhaps contributing to the large supply elasticity we find.

Workers’ ability to travel does make labor supply quite elastic in this context, even if not infinitely so. While our more compelling identification comes from the panel context, and looking across specialties, the time series patterns remain important: compensation only increased by 25 percent in the early phase of the pandemic, when job openings tripled.

The cross-state variation showed even more flexibility, suggesting that a national staffing market does offer a great deal of flexibility to accommodate demand shocks. When demand increases in specific geographic areas, nurses’ ability to travel can help mitigate a local shortage. Their willingness to do so is quite striking in light of the secular decline in geographic mobility across the United States (Molloy et al., 2014; Bartik, 2018; Yagan, 2019). In contrast, nurses’ temporary mobility appears quite elastic to higher compensation, suggesting that this expansion can be valuable in increasing the labor supply available in a particular market.

But if numerous different regions experience simultaneous COVID-19 surges, the spike in staffing needs would not be confined to one area. Nurses’ ability to relocate cannot

¹¹The Nurse Licensure Compact eliminates these requirements among most states, though DePasquale and Stange (2016) find no effect on labor supply.

increase the total number of workers, so may be less powerful for addressing a simultaneous national shortage. In this case, aggregate national supply would have to expand through longer work hours, hiring nurses who normally work in less acute settings, or increased labor force participation.

References

- Abraham, Katharine G**, “Flexible Staffing Arrangements and Employers’ Short-Term Adjustment Strategies,” Working Paper 2617, National Bureau of Economic Research June 1988.
- **and Susan K Taylor**, “Firms’ use of outside contractors: Theory and evidence,” *Journal of labor economics*, 1996, 14 (3), 394–424.
- Abraham, Katharine, John Haltiwanger, Kristin Sandusky, and James Spletzer**, “Measuring the gig economy: Current knowledge and open issues,” *Measuring and Accounting for Innovation in the 21st Century*, 2017.
- Aiken, Linda H, Sean P Clarke, Douglas M Sloane, Julie Sochalski, and Jeffrey H Silber**, “Hospital nurse staffing and patient mortality, nurse burnout, and job dissatisfaction,” *Jama*, 2002, 288 (16), 1987–1993.
- Angrist, Joshua D, Sydnee Caldwell, and Jonathan V Hall**, “Uber vs. Taxi: A Driver’s Eye View,” Working Paper 23891, National Bureau of Economic Research September 2017.
- Autor, David H**, “Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing,” *Journal of labor economics*, 2003, 21 (1), 1–42.
- Bartik, Alexander W**, “Moving costs and worker adjustment to changes in labor demand: Evidence from longitudinal census data,” *Manuscript, University of Illinois at Urbana-Champaign*, 2018.
- Buerhaus, Peter I, David I Auerbach, and Douglas O Staiger**, “The Recent Surge In Nurse Employment: Causes And Implications: Recession effects that have eased the shortage of hospital nurses must be viewed as temporary, lest they distract policymakers from continuing to address longer-term indicators.,” *Health Affairs*, 2009, 28 (Suppl3), w657–w668.
- , **Douglas O Staiger, and David I Auerbach**, “Implications of an aging registered nurse workforce,” *Jama*, 2000, 283 (22), 2948–2954.
- Bureau of Labor Statistics**, “Registered nurses made up 30 percent of hospital employment in May 2019,” 2020. Available online at <https://www.bls.gov/opub/ted/2020/registered-nurses-made-up-30-percent-of-hospital-employment-in-may-2019.htm> (accessed November 17, 2020).
- Chen, Kuan-Ming, Claire Ding, John A List, and Magne Mogstad**, “Reservation Wages and Workers’ Valuation of Job Flexibility: Evidence from a Natural Field

Experiment,” Working Paper 27807, National Bureau of Economic Research September 2020.

Chen, M Keith and Michael Sheldon, “Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform,” 2015. Mimeo, UCLA Anderson. Available online at https://www.anderson.ucla.edu/faculty_pages/keith.chen/papers/SurgeAndFlexibleWork_WorkingPaper.pdf.

Collins, Brett, Andrew Garin, Emilie Jackson, Dmitri Koustas, and Mark Payne, “Is gig work replacing traditional employment? Evidence from two decades of tax returns,” *Unpublished paper, IRS SOI Joint Statistical Research Program*, 2019.

Cook, Andrew, Martin Gaynor, Melvin Stephens Jr, and Lowell Taylor, “The effect of a hospital nurse staffing mandate on patient health outcomes: Evidence from California’s minimum staffing regulation,” *Journal of Health Economics*, 2012, 31 (2), 340–348.

Davis, Steven J, R Jason Faberman, and John C Haltiwanger, “The establishment-level behavior of vacancies and hiring,” *The Quarterly Journal of Economics*, 2013, 128 (2), 581–622.

Deming, David and Lisa B Kahn, “Skill requirements across firms and labor markets: Evidence from job postings for professionals,” *Journal of Labor Economics*, 2018, 36 (S1), S337–S369.

DePasquale, Christina and Kevin Stange, “Labor Supply Effects of Occupational Regulation: Evidence from the Nurse Licensure Compact,” Working Paper 22344, National Bureau of Economic Research June 2016.

Dey, Matthew, Susan N Houseman, and Anne E Polivka, “Manufacturers’ outsourcing to staffing services,” *ILR Review*, 2012, 65 (3), 533–559.

Dube, Arindrajit and Ethan Kaplan, “Does outsourcing reduce wages in the low-wage service occupations? Evidence from janitors and guards,” *ILR Review*, 2010, 63 (2), 287–306.

Farrell, Diana, Fiona Greig, and Amar Hamoudi, “The online platform economy in 2018: Drivers, workers, sellers, and lessors,” *JPMorgan Chase Institute*, 2018.

Forsythe, Eliza, Lisa B Kahn, Fabian Lange, and David Wiczer, “Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims,” *Journal of public economics*, 2020, 189, 104238.

Hall, Jonathan, Cory Kendrick, and Chris Nosko, “The effects of Uber’s surge pricing: A case study,” *The University of Chicago Booth School of Business*, 2015.

- Hershbein, Brad and Lisa B Kahn**, “Do recessions accelerate routine-biased technological change? Evidence from vacancy postings,” *American Economic Review*, 2018, 108 (7), 1737–72.
- Houseman, Susan N**, “Why employers use flexible staffing arrangements: Evidence from an establishment survey,” *ILR Review*, 2001, 55 (1), 149–170.
- Katz, Lawrence F and Alan B Krueger**, “The role of unemployment in the rise in alternative work arrangements,” *American Economic Review*, 2017, 107 (5), 388–92.
- **and** – , “The rise and nature of alternative work arrangements in the United States, 1995–2015,” *ILR Review*, 2019, 72 (2), 382–416.
- **and** – , “Understanding trends in alternative work arrangements in the United States,” *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2019, 5 (5), 132–146.
- Khullar, Dhruv, Amelia M Bond, and William L Schpero**, “COVID-19 and the Financial Health of US Hospitals,” *Jama*, 2020, 323 (21), 2127–2128.
- Kuhn, Peter and Kailing Shen**, “Gender discrimination in job ads: Evidence from china,” *The Quarterly Journal of Economics*, 2013, 128 (1), 287–336.
- Landuis, Timothy and Curtis Starkey**, “Largest Healthcare Staffing Firms in the United States: 2020 Update,” 2020. Staffing Industry Analysts.
- Mark, Barbara A, David W Harless, Joanne Spetz, Kristin L Reiter, and George H Pink**, “California’s minimum nurse staffing legislation: results from a natural experiment,” *Health services research*, 2013, 48 (2pt1), 435–454.
- Mas, Alexandre and Amanda Pallais**, “Valuing alternative work arrangements,” *American Economic Review*, 2017, 107 (12), 3722–59.
- Molloy, Raven, Christopher L Smith, and Abigail K Wozniak**, “Declining Migration within the U.S.: The Role of the Labor Market,” Working Paper 20065, National Bureau of Economic Research April 2014.
- Needleman, Jack, Peter Buerhaus, Soeren Mattke, Maureen Stewart, and Katya Zelevinsky**, “Nurse-staffing levels and the quality of care in hospitals,” *New England Journal of Medicine*, 2002, 346 (22), 1715–1722.
- , – , **V Shane Pankratz, Cynthia L Leibson, Susanna R Stevens, and Marcelline Harris**, “Nurse staffing and inpatient hospital mortality,” *New England Journal of Medicine*, 2011, 364 (11), 1037–1045.
- Rothwell, Jonathan**, “Still searching: Job vacancies and STEM skills,” *Metropolitan Policy Program at Brookings Institution*, 2014.

- Sloane, Douglas M, Herbert L Smith, Matthew D McHugh, and Linda H Aiken**, “Effect of changes in hospital nursing resources on improvements in patient safety and quality of care: a panel study,” *Medical care*, 2018, *56* (12), 1001.
- Spetz, Joanne, David W Harless, Carolina-Nicole Herrera, and Barbara A Mark**, “Using minimum nurse staffing regulations to measure the relationship between nursing and hospital quality of care,” *Medical Care Research and Review*, 2013, *70* (4), 380–399.
- Xue, Ying, Joyce Smith, Deborah A Freund, and Linda H Aiken**, “Supplemental nurses are just as educated, slightly less experienced, and more diverse compared to permanent nurses,” *Health Affairs*, 2012, *31* (11), 2510–2517.
- , **Linda H Aiken, Deborah A Freund, and Katia Noyes**, “Quality outcomes of hospital supplemental nurse staffing,” *The Journal of nursing administration*, 2012, *42* (12), 580.
- Yagan, Danny**, “Employment hysteresis from the great recession,” *Journal of Political Economy*, 2019, *127* (5), 2505–2558.

Table 1: Data Descriptives**(a)** Full Sample Basic Summary Statistics

Measure	N	Mean	SD	P10	Median	P90
Job posting index	19439	94.5	260.1	7.8	31.2	218.2
Comp. index (posted jobs)	19439	105.8	23.8	87.8	99.6	129
Filled jobs index	2550	139.2	281.5	41.1	82.1	246.4
Comp. index (filled jobs)	2550	117.3	34.3	85.9	103.3	178.5
COVID-19 cases (thousands)	10694	.83	3.78	0	.14	1.3
Travel distance (miles)	2550	548.3	625.5	59.2	290.8	1469.9

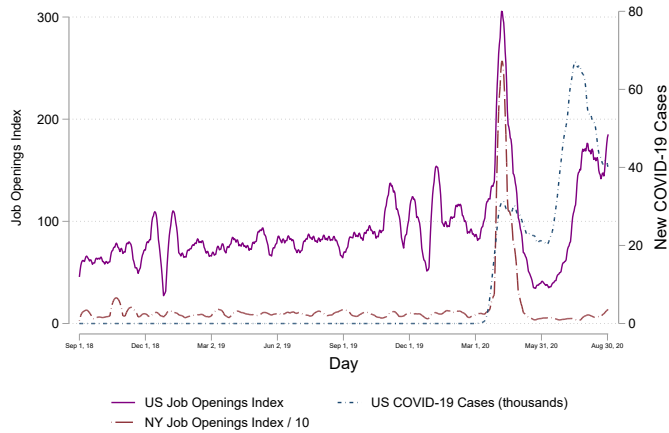
(b) Summary Values During COVID-19

Observation level		February 1 - March 14		March 15 - May 16		May 17 - July 18		July 19 - August 31	
		Job postings	Comp index	Job postings	Comp index	Job postings	Comp index	Job postings	Comp index
National	100	100	100	165	133	47	103	167	114
ICU	29	100	105	339	157	104	110	331	127
ER	10	100	103	189	131	22	102	121	110
LD	3	100	110	78	115	27	107	72	111
Med/Surg	39	100	95	130	119	39	97	168	111
OR	5	100	108	56	108	49	104	101	106
Other	15	100	100	162	125	37	99	105	100
NY	10	100	106	863	185	49	125	77	107
MA	3	100	106	174	131	94	106	70	103
AZ	2	100	93	145	113	75	104	273	127
Completed	4	100	104	324	162	94	120	266	126

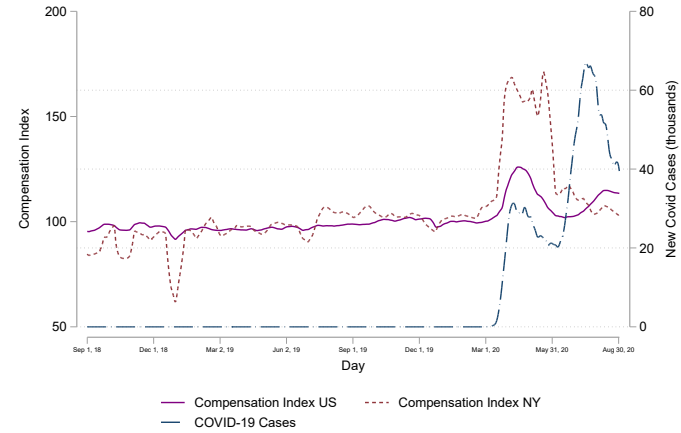
Data are from Health Carousel and are described in detail in the text. The unit of observation in Panel A is the state-by-day. The compensation index is normalized to national daily average from February 1–March 14, 2020 for all subsets, weighted by number of job postings. The job posting index is normalized average daily postings for February 1–March 14, 2020 for each subsample.

Figure 1: Time Series Patterns

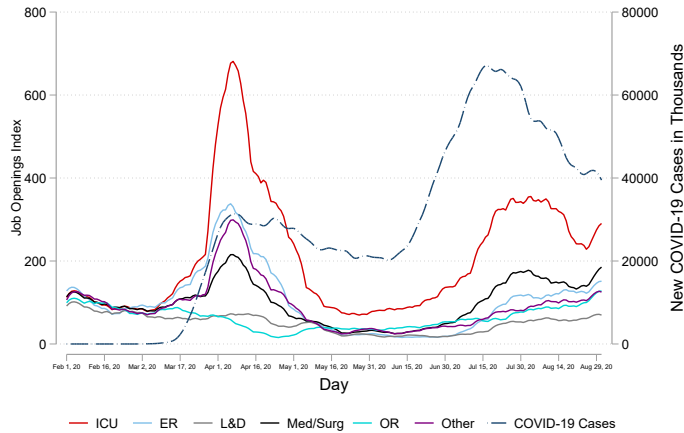
(A) Job Openings and COVID-19 Cases



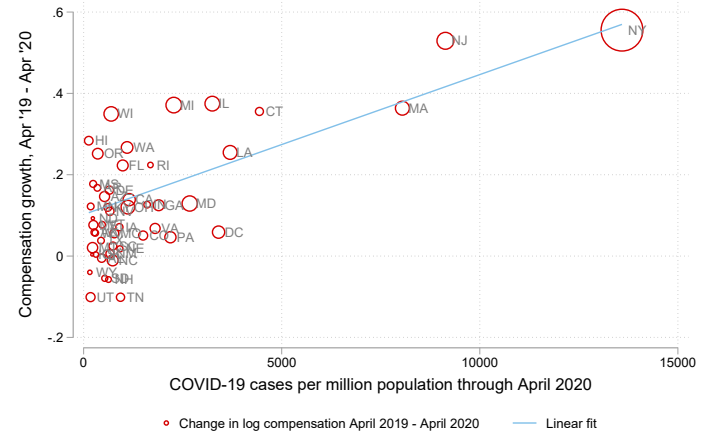
(B) Compensation and COVID-19 Cases



(C) Job Openings by Specialty



(D) Compensation Growth by State



29

Panel A shows job postings in the United States and in New York state from September 1, 2018 through August 31, 2020. Data are smoothed using an Epanechnikov kernel. The panel also shows (smoothed) national new COVID-19 cases. Panel B shows compensation trends, also nationally and for New York state, along with national COVID-19 cases. Panel C considers six specialty categories: intensive care (ICU), emergency room (ER), labor and delivery (L&D), standard hospital floors (Med/Surg), operating room (OR), and other. The panel shows smoothed time series of job postings by category from February 1 through August 31, 2020. The indices in Panels A, B, and C are normalized to a mean of 100 in February 2020. Panel D shows the change in log compensation from April 2019 to April 2020 by state, plotted against cumulative COVID-19 cases per million population by April 30, 2020.

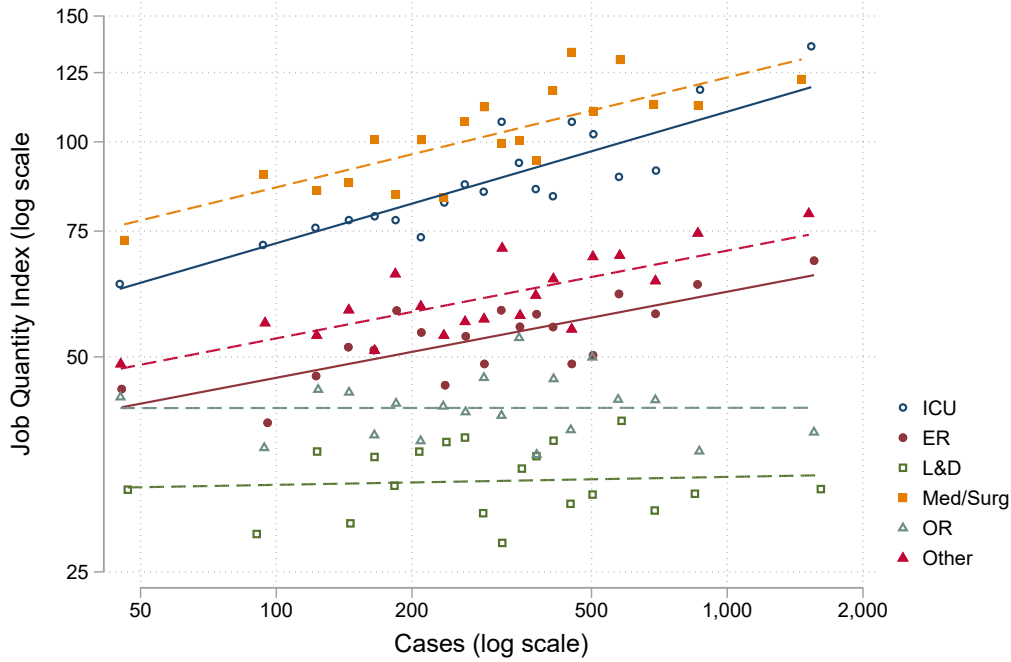
Table 2: Regression Estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Jobs	Comp.	Jobs	Comp.	Jobs	Comp.	Jobs	Comp.	Jobs	Comp.
ln(cases)	0.17*** (0.021)	0.030*** (0.0037)	0.036 (0.036)	0.011 (0.0079)	0.059 (0.033)	0.0094 (0.0075)	0.019 (0.033)	0.0055 (0.0077)	0.015 (0.035)	0.039* (0.018)
ICU \times ln(cases)			0.17*** (0.028)	0.026** (0.0087)	0.14*** (0.028)	0.028*** (0.0068)	0.16** (0.057)	0.026* (0.012)		
ER \times ln(cases)			0.059* (0.024)	0.015 (0.0094)	0.035 (0.027)	0.016* (0.0077)	0.11* (0.052)	0.025 (0.015)		
Med/Surg \times ln(cases)			0.11*** (0.031)	0.018* (0.0073)	0.074* (0.029)	0.020** (0.0058)	0.14** (0.049)	0.027* (0.011)		
OR \times ln(cases)			0.0019 (0.031)	-0.0040 (0.0047)	0.011 (0.036)	0.00015 (0.0057)	-0.029 (0.044)	-0.0025 (0.011)		
Other \times ln(cases)			0.056 (0.029)	0.0096 (0.0087)	0.033 (0.029)	0.0085 (0.0077)	0.11** (0.038)	0.012 (0.0092)		
COVID-19 \times ln(cases)									0.11* (0.044)	0.037** (0.012)
N	3860	3860	9413	9413	9413	9413	9413	9413	419	968
R^2	0.48	0.27	0.34	0.28	0.37	0.32	0.43	0.39	0.38	0.65
α (overall or ICU)		5.6		6.7		5.1		6.2		3.1
α (ER)				3.9		2.2		4.2		
α (Med/Surg)				6.1		3.8		5.1		
Date Fixed Effects	1	1	1	1	1	1	1	1		
Week Fixed Effects									1	1
State-Specialty FE					1	1	1	1		
Date-Specialty FE							1	1		

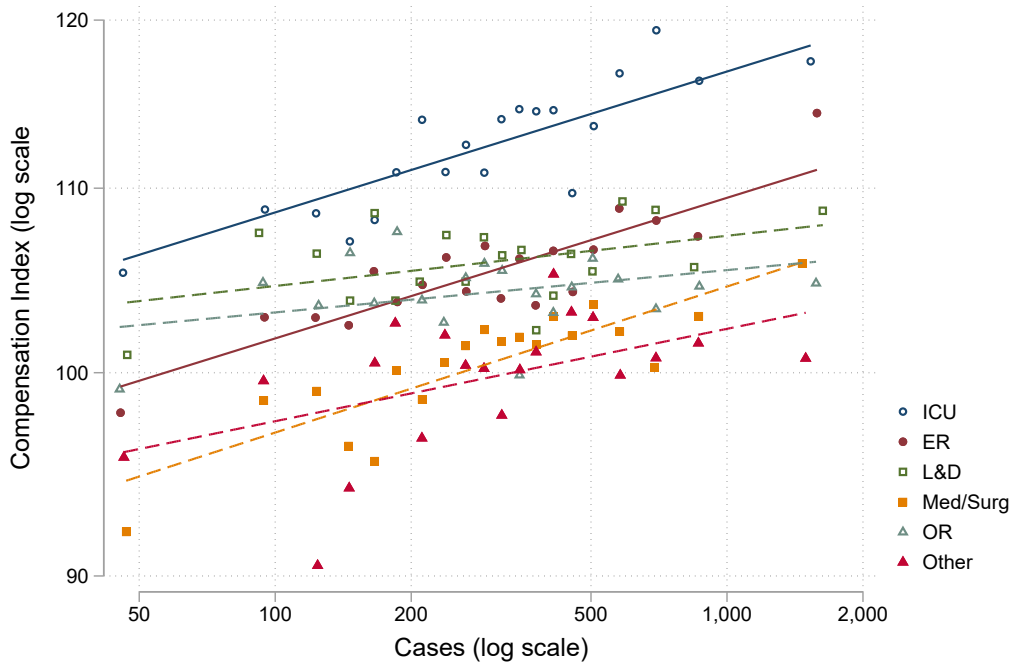
This table reports estimates of equations (6) and (5) on Health Carousel data on travel nursing job postings from February–August 2020. The dependent variable in odd-numbered columns is the log number of nursing jobs (nationally by day in col. 1; by specialty in cols. 3, 5, and 7; and filled by day in col. 9). The dependent variable in the even-numbered columns is the average log compensation for the jobs included in the prior column. Cols. (1)–(2) don’t distinguish among specialties, and the supply calculations assume that local supply is unaffected by local COVID-19 conditions ($\beta = 0$). In cols. (3)–(8), the omitted nursing specialty is labor and delivery. Cols. (9)–(10) combine ICU, ER, and Med/Surg together into "COVID-19 specialties", and combine OR with L&D into the omitted category. Regressions are estimated at the state-week level in cols. (9)–(10). All specifications include state fixed effects. Standard errors, in parenthesis, are clustered by state. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure 2: Graphical Supply Estimates by Specialty

(A) Job Postings vs. COVID-19 Cases | Fixed Effects



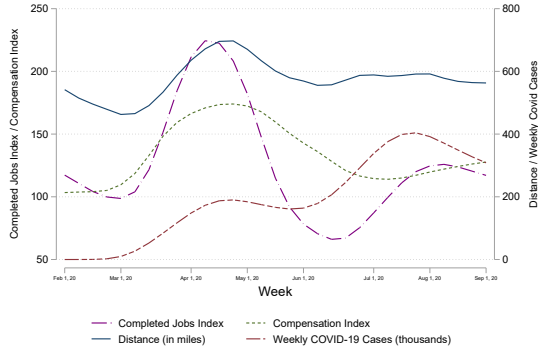
(B) Compensation vs. COVID-19 Cases | Fixed Effects



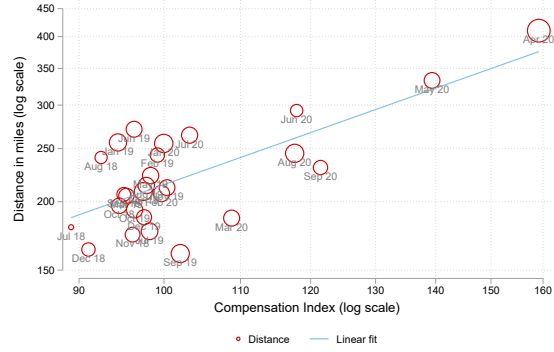
Panel A shows six binned scatterplots of job postings against COVID-19 cases by state/day, after conditioning on state and day fixed effects. We show separate scatterplots and corresponding log-linear fits for six specialty categories: intensive care (ICU), emergency room (ER), labor and delivery (L&D), standard hospital floors (Med/Surg), operating room (OR), and other. Panel B is analogous, but shows the mean compensation for the corresponding jobs rather than the quantity.

Figure 3: Workers' Travel Flexibility and Compensation

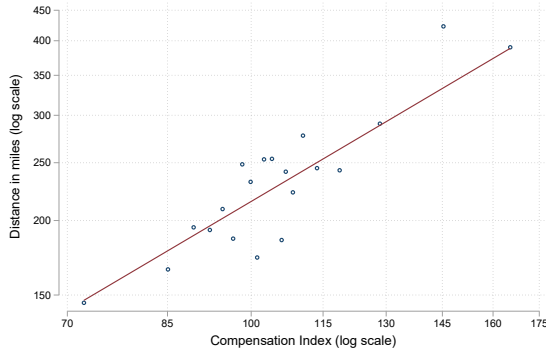
(A) Completed Jobs Over Time



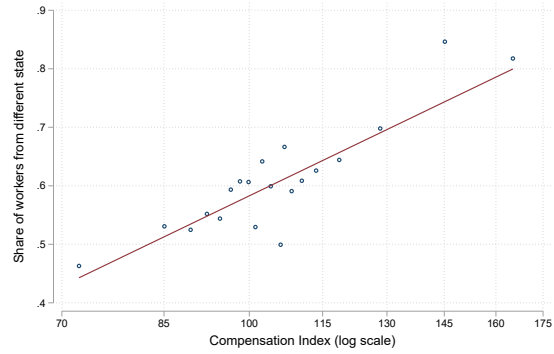
(B) Distance & Pay Over Time



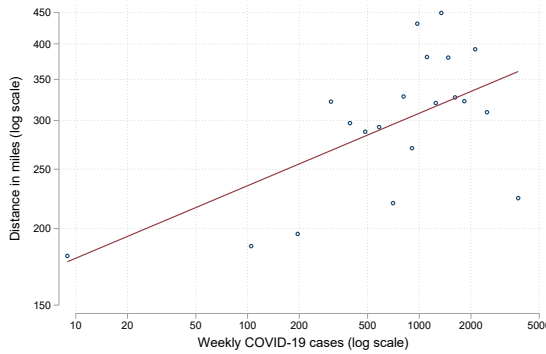
(C) Distance vs. Compensation



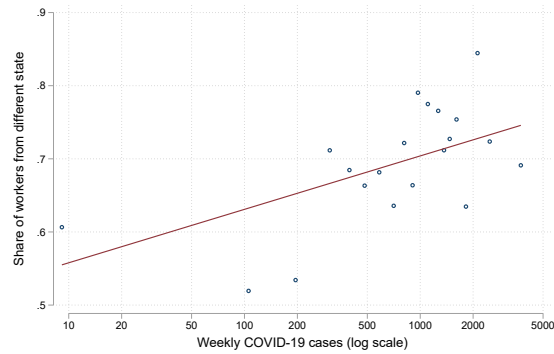
(D) Cross-State vs. Compensation



(E) Distance vs. COVID-19



(F) Cross-State vs. COVID-19



This figure uses data on jobs filled by Health Carousel. Panel A shows time series of completed jobs and compensation (both normalized to 100 in February 2020), average distance traveled from a nurse's home to the job location, and COVID-19 cases. Panel B shows the relationship between average distance traveled and the contemporaneous mean compensation index by month. Panels C, D, E, and F show binned scatterplots of the relationship between travel distance (Panels C and E) or share of jobs filled by out-of-state workers (Panels D and F) against the job's own compensation (Panels C and D) or the job location's weekly COVID-19 new case count (Panels E and F). The binned scatterplots are all conditional on state and specialty fixed effects.

A Model Solution

To back out the model's parameters from our estimates, we sequentially use the coefficients π_i and μ_i for each COVID-19-related specialty $i \in \{ICU, ER, \text{Med/Surg}\}$, along with the τ_1 and κ_1 coefficients for L&D. (That is, τ_1 and κ_1 remain the same as we move across specialties i .) Setting the empirical coefficients in (5) and (6) equal to the appropriate model-implied coefficients from (3) and (4), and then solving for the latter, yields:

$$\alpha = \frac{\mu_i}{\pi_i} \tag{8}$$

$$\beta = \kappa_1 - \frac{\mu_i \tau_1}{\pi_i} \tag{9}$$

$$\gamma = \frac{\kappa_1}{\tau_1} \tag{10}$$

$$\delta = \mu_i - \frac{\kappa_1 \pi_i}{\tau_1}. \tag{11}$$

Table 2 reports the estimates for the key supply parameter $\hat{\alpha}$ at the bottom of each pair of regressions.

B Additional Results

Table B.1: Summary Statistics During COVID-19 Pandemic

Measure	N	Mean	SD	P10	Median	P90
Job posting index	4977	126.7	391	7.8	39	288.3
Comp. index (posted jobs)	4977	118.8	32.2	91.9	109.3	160.7
Filled jobs index	669	219.9	481.6	41.1	82.1	492.8
Comp. index (filled jobs)	669	138.9	38.3	92.6	134.9	187.9
COVID-19 cases (thousands)	10694	.83	3.78	0	.14	1.3
Travel distance (miles)	669	588.5	604.8	68.1	344.3	1395.3

Data are from Health Carousel and are described in detail in the text. The unit of observation is the state-by-day. The compensation index is normalized relative to the national daily average from February 1–March 14, 2020, weighted by number of job postings. The job posting index is normalized average daily postings, relative to each subsample’s average from February 1–March 14, 2020.

Table B.2: State Summary

Observation level		February 1 - March 14		March 15 - May 16		May 17 - July 18		July 19 - August 31	
		Job postings	Comp index	Job postings	Comp index	Job postings	Comp index	Job postings	Comp index
AL	1.26	100	88	62	90	66	84	189	108
AK	.5	100	104	88	95	49	106	76	101
AZ	2.47	100	93	145	113	75	104	273	127
AR	1.33	100	100	86	101	40	96	341	111
CA	13.4	100	116	131	137	32	118	165	132
CO	.61	100	97	96	100	38	89	77	90
CT	.69	100	115	173	133	57	123	60	116
DE	.22	100	100	341	126	33	123	110	104
FL	5.7	100	91	140	117	17	90	382	115
GA	2.67	100	95	116	109	40	93	211	129
HI	.25	100	105	165	126	17	101	169	147
ID	.19	100	102	37	88	17	84	146	102
IL	2.82	100	93	203	117	65	96	83	100
IN	1.67	100	94	74	99	41	98	55	100
IA	1.06	100	102	104	101	66	99	112	114
KS	.33	100	100	52	95	19	87	103	104
KY	1.65	100	96	128	105	37	96	168	110
LA	.83	100	92	220	104	47	92	297	106
ME	.52	100	93	137	97	69	92	98	93
MD	3.65	100	93	208	105	118	107	155	111
MA	2.89	100	106	174	131	94	106	70	103
MI	1.62	100	110	243	137	32	90	55	89
MN	.41	100	96	106	95	48	124	78	101
MS	.31	100	65	242	100	77	83	1141	105
MO	2.84	100	98	112	105	30	94	199	105

State Summary (Continued)

Observation level		February 1 - March 14		March 15 - May 16		May 17 - July 18		July 19 - August 31	
		Job postings	Comp index	Job postings	Comp index	Job postings	Comp index	Job postings	Comp index
MT	.73	100	103	78	101	65	96	114	97
NE	.52	100	100	72	102	38	85	97	102
NV	1.75	100	100	168	114	14	96	311	130
NH	.62	100	100	108	104	55	99	124	102
NJ	2.4	100	120	513	182	34	117	37	115
NM	1.42	100	90	107	96	44	103	59	112
NY	10.25	100	106	863	185	49	125	77	107
NC	3.13	100	92	119	94	36	90	187	104
ND	.34	100	105	73	106	20	108	46	106
OH	1.96	100	92	135	112	62	98	187	101
OK	1.05	100	97	54	95	14	93	128	120
OR	1.46	100	106	150	111	37	100	117	106
PA	2.62	100	97	145	109	75	96	160	103
RI	.39	100	110	679	115	242	114	195	114
SC	2.34	100	92	106	93	31	93	353	116
SD	.48	100	98	121	100	103	97	125	98
TN	1.99	100	77	111	96	48	98	208	100
TX	4.74	100	96	53	97	63	101	573	116
UT	.11	100	101	247	89	250	172	285	81
VT	.86	100	99	98	101	31	97	143	107
VA	3.98	100	95	113	97	56	97	122	104
WA	3.25	100	105	152	128	23	106	74	110
WV	1.08	100	88	90	88	74	91	146	94
WI	1.56	100	92	178	120	22	99	61	100
WY	.21	100	88	64	88	48	83	71	98

36

Data are from Health Carousel and are described in detail in the text. The unit of observation is the state-by-day. The compensation index is normalized relative to the national daily average from February 1–March 14, 2020, weighted by number of job postings. The job posting index is average daily postings, normalized relative to each state’s average for February 1–March 14, 2020.