STABILITAS

COVID-19: A Global Pandemic, Distributed Locally.

WHEN WE'RE IN THE RED

Even as a company constantly identifying critical events for some of the most recognized companies in the world, the first reports of coronavirus didn't raise significant alarms; it was far away and the reported numbers weren't particularly concerning. At the office, we brushed it off as a bad flu. Our CTO, whose sister is an MD, didn't seem particularly worried either. Mortality rates in China (read: questionable data were relatively low. Same in Japan. As early reporting from South Korea started coming in, alarm bells began ringing. We were starting to see implications for global supply chains.

The first COVID-19 case in the US was reported on January 21st, then several reports of local transmission. On February 28th, the first reported death in the US due to COVID-19. Most of that centered here, in Washington State.

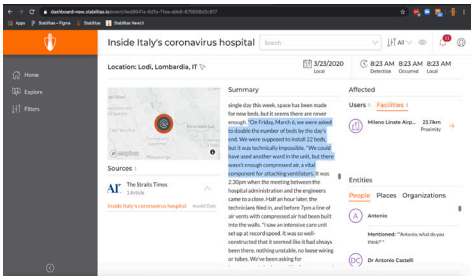

I saw reports from Italy; Doctors conducting triage in hallways, sending sick patients away given the lack of resources. Suddenly it got to the "why" behind mortality rates: hospitals were overwhelmed and lacked the resources to adequately prepare.

As COVID-19 progresses, patients require ventilators and ICU care. Recovery times are relatively long measured in weeks, not days. Demand for beds and equipment quickly outpaces supply.

As a data guy who started a data company, it got me thinking about what we would do in the US, and if we could develop a real-time gauge of hospital occupancy in our communities. I did a bit of analysis, found the number of hospital beds by zipcode on the Department of Homeland's Security's website. I then looked at reporting from the New York Times and a dataset that captured the ratio of ICU beds to normal beds posted as a project by the Harvard Global Health Initiative, as well as typical occupancy. What I found made me worried: if previous patterns held true, no major US city would not have the capacity to deal with the oncoming demand.

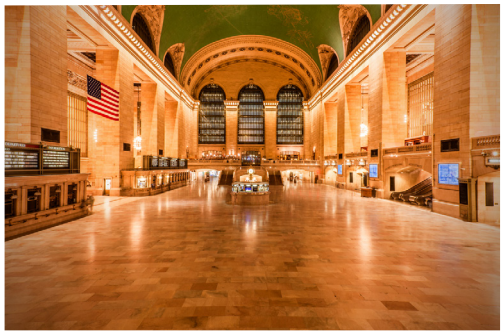

Conclusions

Here are my assumptions:

- The NYT piece reviewed ~300 hospitals, or about 5% of the overall number of hospitals in the US.
- With simple averages I found that 11.4% of overall beds are ICU beds.
- In the same NYT piece, I found that 63% of ICU beds are typically occupied.

Looking at US Department of Homeland Security's HFLID data, which covers all hospitals in the US, I applied the averages from the NYT piece.

What I found was that NYC has about 11,000 hospital beds. 11% of that 11,000 leaves about 1,250 ICU beds. Accounting for current occupancy reduces roughly 790 beds, leaving a total of 460 or so available ICU beds.

To take a swag at the number of confirmed COVID-19 cases that leads to a redline, take the percent of cases that end up in the ICU. From what I've found, that's 5% of total COVID-19 cases. For 9,400 positive cases in NYC, this implies that 470 end up in the ICU – pushing the system beyond its capacity of 460 available ICU beds.

As of Sunday morning, I see 6,211 confirmed cases in NYC. Even as hospitals divert resources to focus on COVID-19 and run above capacity, there are still issues associated with ventilators and other critical resources.

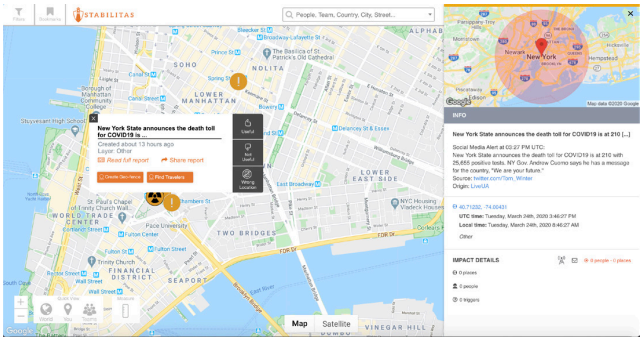

I hope I'm wrong about the ability to handle cases. But we should understand what the data is telling us and pay attention to what's happening in our own communities. If we don't pay attention and start preparing now, by the time this virus is in our neighborhoods, it's already too late.

As I mentioned previously, my focus is on the data, not on the medical science. I readily welcome feedback and input from others. If anyone out there can help refine this analysis, I welcome the chance to talk with you, your can reach me at ceo@stabilitas.io