

Data Release Information Sheet

Data Summary

Dataset name: Torch Insight COVID-19 Burden Index Projections

Summary: The Torch Insight COVID-19 Burden Index Projections presents current and projected data surrounding COVID-19 cases and growth. We present data at varying geographies, as well as hospitals with the ardent hope that clearer vision into what is really happening with COVID-19 will help policy makers, healthcare administrators, and citizens make wise decisions that ultimately save lives and improve our country.

For our publicly available COVID-19 tracking efforts, we have developed a “hotspotting” measure based on transmission rates and prevalence of active cases, with the goal of identifying regions where the disease is having a noticeable and growing impact. Using different thresholds of these two factors, we have developed five categories or levels of priority. We use effective transmission rates (R_t) and prevalence (population per 100,000) to create distinct hotspot categories. The hotspot categories are as follows:

Category	R_t		Prevalence	Description
Capacity Risk	≥ 2.2	AND	≥ 250	This level of prevalence, generally with high growth, is or imminently will be a significant strain on the region's health care system.
	(any R_t)	OR	≥ 500	
Hotspot	≥ 1.54	AND	≥ 175	This level of prevalence, combined with a fairly high growth rate, warrants particular concern and attention.
Warming	≥ 1.54	AND	≥ 40	With a fairly high growth or prevalence (but not both), these regions will likely become hotspots if not addressed soon.
	≥ 1.1	OR	≥ 250	
Watch Area	≥ 1.1	AND	≥ 40	These areas have growth and prevalence that warrant close monitoring, but perhaps no direct intervention yet.
No Current Concern	All other R_t and prevalence measures			Low growth and low prevalence of COVID-19 with low likelihood of becoming otherwise in the near term.

R_t is calculated on a rolling 7-day basis. Population data is sourced from the American Community Survey. Additional information on our hotspot methods can be found in the Torch Insight white paper “COVID-19 Hotspots”, published June 24, 2020.

Our projections model analyzes hospital and ICU burden and estimates the total hospital bed and ICU capacity in counties and hospitals given current and projected COVID-19 case levels. We extend projections up to 180 days from the most recent COVID-19 reported case numbers. The Burden Index Projections estimate if the hospital or county has exceeded its normal capacity. To adjust for current conditions in specific regions, the Burden Projections:

- Factor in county-specific infection rates: Our projections are based on growth curves of what is actually happening on the ground and update daily. Thus, when states implement social distancing, shelter-in-place, or other measures, our projections take these into account by updating accordingly as new case count data comes in. The models reflect *what is observed to be happening, not what is hoped or assumed will happen*.

- Adjusts for Hospital Service Area: Not all counties have hospitals, and sometimes patients cross county lines to go to different hospitals. We base our projections for counties and individual hospitals on hospital service areas (HSAs), which are local healthcare markets for hospital care, rather than simply on where the hospital is located. This means we have scores even for counties that have no hospitals in their borders, and that we account for hospital patients from outside the hospital's county.
- Adjusts for Hospital Baseline Capacity: Some hospitals were already near capacity before the COVID crisis. We base our projections of capacity on hospital- and ICU-bed utilization rates derived from the 2018 CMS Hospital Cost Reports.
- Eliminates Hospital Elective Procedures: Many hospitals can make additional beds available for COVID-related illness by postponing elective admissions. Based on a sample of all-payer claims and 100% of Medicare hospitalizations we estimate the percent of admissions in each hospital that are elective, allowing us to account for increased capacity over the baseline utilization.

Full Methodology:

Updated 6/24/2020

To calculate hospital and ICU burden, we first calculate the average normal hospital inpatient bed and ICU utilization based on 2018 hospital cost report data ([CMS form 2552-10](#)). This provides a baseline estimate of capacity in both settings and the percent of capacity utilized. Some hospitals are already over capacity on average, which we allow in the analysis. The cost reports provide data on total ICU patient days for all patients. When reporting on hospital bed capacity, we include ICU beds as well because our COVID hospitalization estimates technically include those treated in the ICU as well as non-ICU settings.

We adjust estimates of actual inpatient bed capacity by estimating the proportion of inpatient admissions for elective procedures and removing those from bed utilization estimates. Despite the advent of COVID-19, many patients for unrelated conditions will still need to be treated in the hospital, including emergency cases (such as a heart attack) or urgent cases (such as an important cardiac catheterization). From the 100% sample of Medicare inpatient claims data we are able to assess what proportion of admissions are elective, urgent and emergent. From an approximate 5% sample of all-payer claims (the [Healthcare Cost and Utilization Project](#)) we assessed the rate at which elective admissions occur for every diagnosis related group (DRG). We combined this data with hospital-level payer mix data from the hospital cost reports, which lets us estimate the rate and count of admissions for each hospital for each DRG, and then combined those rates to the Major Diagnostic Categories (MDC). We then adjusted estimated elective rates for Medicaid and Commercial payers based on the ratio for actual Medicare elective rates by admission category to average Medicare elective admission rates, using a log-transformation, at the MDC-level. Aggregate hospital elective rates were then calculated by combining the rates and counts of admissions from the MDC-level. Across 6,007 hospitals, the average rate of elective admissions is 31.5%, with a median of 19.5%, and an inter-quartile range of 8.9%-39.7%. For 851 hospitals, more than 90% of their admissions were elective, and 1,031 hospitals had fewer than 5% of their admissions as being elective. When adjusting for total admissions, we estimate that 22.0% of national hospital admissions are elective.

Next, we calculate county-level service areas for each hospital, based on the hospital service area file (HSAF) released by CMS for 2018. This file contains the count of patients discharged by each hospital from each zip code (based on patient residence). We aggregate this up to counties, excluding any counties with fewer than 5 discharges to limit rare events such as out-of-town emergency cases. Using the percent of a hospital's remaining discharges represented by each county, we apportion the total hospital bed and ICU capacity to the counties in its service area. We thus rely not on hospitals *in* a county, but on hospitals *serving* the county. This allows us to also conduct the county-level analysis, even when there are no ICU beds or even hospital beds in the county. We also are able to determine, within each county, a relative market share for each hospital serving the county at both levels of care.

The next step is to calculate the number of active COVID-19 cases in each county and estimate how many of these will require hospital beds or ICU care. We obtain county-level reported cases of COVID-19 from [Johns Hopkins University's](#) COVID-19 tracking tool (Lancet article about the tool [here](#)), which reports confirmed cases and deaths and is updated on a daily basis. Active cases are calculated as any case that began within the last 14 days, excluding deaths within the last 14 days. To this active case count we apply the age distribution for the county as reported in the [Census Bureau's](#) ACS 2014-2018 5-year average data. We use a susceptible-infected-recovered (SIR) epidemiological model to project cases into the future; the methodology for these projections, including estimates of doubling times at the county level, are described in more detail below.

We estimate the care utilization of the case using a couple of different sources. First, we collected data from state reports on the rates of COVID-19 infection and hospitalization within age bands. See the helpful [data source list](#) put together by the University of Minnesota team that put together the [COVID-19 Hospitalization Tracking Project](#) for a list of state reporting sites. We used ACS population data to calculate rates within standardized, consistent age bands across states. This data allowed us to calculate an effective hospitalization rate of confirmed cases within the state. Thirteen states had data on hospitalization by age band, so we calculated an effective hospitalization rate directly for those. For the other 37 states, we averaged the 13 states' rates by age band and applied these to the age distributions of the other 37 to obtain an effective hospitalization rate that varies according to the age distributions of those states. We apply these effective hospitalization rates directly to the COVID-19 case data and demographics to obtain percent of cases resulting in hospitalization. This approach is highly dependent on assumptions about what the data inputs represent, and our estimated hospitalizations may differ from what states report. See the section below on reasons for these discrepancies.

To obtain ICU cases, we use [CDC morbidity and mortality data](#) to determine the ratio of hospitalized cases to ICU cases and calculate percent of total cases needing an ICU bed. We present estimates based on both the upper and lower bounds presented in these studies.

We apply a reduction factor to the hospitalization and ICU cases to represent an average length of stay. The [CDC reported](#) a median ALOS for COVID-19 patients in 3 studies as 10-13 days, so we use the midpoint of that range (11.5 days). This takes into account that on average, a COVID-19 patient will not be in the hospital for the duration of the disease (assumed to be 14 days on average). We do apply the same ALOS to ICU stays as well as hospitalizations. We have not yet seen ALOS data on ICU stays specifically, so we are likely somewhat overestimating the ICU utilization of COVID-19 patients.

With the apportionment of hospital and ICU bed capacity to counties, we can then estimate the number of COVID-19 related hospital and ICU cases each hospital is likely to see based on the most recent

reports of active cases. This estimation includes an allowance for diversions, which will occur when one or more hospital portions serving a county is full but other hospitals still have capacity in the county. A proprietary algorithm distributes the diversion cases to hospitals evenly (but with respect to capacity limits) until all diversions have been assigned to hospitals or no more capacity remains. If the latter condition is realized, the remaining cases are divided among the hospitals serving the county according to their market shares within the county. In the algorithm, diversions are not allowed to take up capacity from different counties' hospital or ICU capacity, even if portions of the same hospital servicing a different county are not yet full.

Ventilator demand is done post hoc by applying a parameter of 75% of the ICU cases. The 75% figure comes from recent research on ventilator use by COVID-19 patients published in [The New England Journal of Medicine](#) by Bhatraju et al.

We express uncertainty in the number of COVID hospitalizations and ICU cases in two ways. First, we construct a 95% confidence interval around our estimates of COVID-19 cases and apply the hospitalization rate parameters to the upper and lower bounds of that confidence interval. Actual calculation of the confidence interval is detailed below. Second, we allow for the effective hospitalization rate to vary two standard deviations from the simple average of the state hospitalization rates. The effective hospitalization rate has a mean of 16.22% and a standard deviation of 3.71%. We apply the same uncertainty to the ICU admission rate (mean 3.15%, SD 0.78%). The lower bound of the hospitalization and ICU rates are applied to the lower bound of cases, and the upper bounds are similarly combined.

Projecting Cases by County

We employ the commonly used susceptible-infected-recovered (SIR) epidemiological model of infection rates to project county-level active cases up to 60 days from the date of the latest case reports. We largely follow the methodology for projecting COVID-19 infections used by University of Pennsylvania's [CHIME](#) modeling tool. This model relies on two parameters: the recovery time, which we assume to be 14 days based on the CDC guidelines on quarantine lengths, and the exponential growth rate in active cases. Rather than assume a uniform national growth rate, we estimate the growth rate by county wherever possible. This is accomplished by using the exponential growth rate formula:

$$growth\ rate_{t_2} = \log \left(\frac{cases_{t_2}}{cases_{t_1}} \right)$$

We apply the growth rate formula to all possible two-date combinations of reported cases data for the 7 most recent days (using the eighth day back when the growth rate on the seventh day back is being estimated), yielding a total of 28 possible number of observed growth rates (7!, or 7 + 6 + 5 + 4 + 3 + 2 + 1). The average of these observed growth rates constitutes our estimated current growth rate for the county. We also calculate a standard error and use the appropriate t-statistics for the number of observations to construct 95% confidence intervals around the growth rate parameter.

For some pairs of days in specific counties, we are unable to estimate a growth rate due to lack of COVID infections or a very small number of infections over one or two days, or to data inconsistencies in which case counts stay the same or decline for given two-day pair. In these cases, we exclude the affected observations from the average doubling time. When no valid doubling time observations exist for a

county, we use the population-weighted average of the county-level estimates we can calculate as the default doubling time.

With the recovery time and growth rate parameters, we are able to fit the SIR model to project cases for each county. We obtain both best-guess projections based on the average growth rate and upper and lower bound projections based on the upper and lower bounds of the growth rate 95% confidence interval. These are then used in the overall burden analysis as described above. As new studies are conducted on COVID-19, we may from time to time update the parameters used in the SIR model to obtain more accurate projections.

The SIR model also relies on county population estimates, which we take from the same ACS tables cited above for age distributions.

Discrepancies between State Hospitalization Reports and Our Estimates

We have attempted to use state-reported data to obtain hospitalization rates as close as possible to those observed by states. However, keep in mind the following:

- We are using confirmed cases in our estimates, whereas some states are reporting hospitalizations of both confirmed and suspected cases of COVID-19. In these cases, our estimates will be below what the state reports because we are estimating hospitalizations of confirmed COVID-19 cases.
- In many cases, states have not reported enough information for us to calculate the hospitalization rate among confirmed cases within age bands. In these cases, we assume that states with missing data will look like the average of states with such data, adjusting for the different age distributions in each state. States farther from this national average hospitalization rate will be significantly different from our estimates.
- Patients in many regions will receive care at hospitals outside the region. As noted above, we account for this dynamic in our estimates. However, state reports will often not make these distinctions, but will rather report the numbers of hospitalizations for hospitals in their state, regardless of where the patients came from. Take, for example, many of the New England states. Many patients residing in Rhode Island might go to hospitals in Massachusetts or Connecticut. Vermonters might be hospitalized in Massachusetts or New Hampshire. These will often be reported officially by the states where the hospitals are located, not by the patients' home states. Many other examples could certainly be found all across the country, and it is likely that all state reports are subject to some level of this phenomenon. Again, we actually attempt to estimate hospitalizations by citizens of a state, not by hospitals in the state.
- Some states have different definitions of what they are calling confirmed cases, and thus their reported caseloads will be different from what Johns Hopkins University reports. This will lead to slight differences between our estimates of hospitalizations and reported hospitalizations.
- Some states' hospitalization rates will change as they see more data, especially if testing of random samples of the population becomes more commonplace (as opposed to opportunistic testing of symptomatic individuals). We will try to keep abreast of such changes, but this is not part of the nightly update.

[Access current projections](#)

Acknowledgements

Suggested Citation:

Torch Insight. Torch Insight COVID-19 Burden Index Projections. Salt Lake City, UT, 2020

Variable Information

Data is available for download. The data format is as follows:

Variable	Variable Description
date	The date of the projected burden index score
cases	The estimated total hospital cases related to COVID-19 that will occupy a hospital bed.
casesLower	The lower limit of estimated total hospital cases related to COVID-19 that will occupy a hospital bed.
casesUpper	The upper limit of estimated total hospital cases related to COVID-19 that will occupy a hospital bed.
icu	The estimated total ICU cases related to COVID-19 that occupy ICU beds
icuLower	The lower limit of estimated total ICU cases related to COVID-19 that occupy ICU beds
icuUpper	The upper limit of estimated total ICU cases related to COVID-19 that occupy ICU beds
beds	The estimated number of available hospital beds, accounting for baseline hospital bed capacity
icuBeds	The estimated number of available ICU beds, accounting for baseline ICU bed capacity
ventilators	The estimated number of ventilators needed

Additional Information

All use of the Torch Insight Burden Index data is subject to our [Terms of Use](#). Commercial organizations wishing to acquire a license to access and download Torch Insight COVID-19 Burden Index Projections must contact Info.torchinsight@milliman.com prior to doing so.

Contact Information

For questions, concerns, or additional questions please contact:

Torch Insight

299 S. Main St. Suite 2300

Salt Lake City, UT 84102

USA

Email: Info.torchinsight@milliman.com

For additional information: TorchInsight.com

Author Contact

Robert Richards, PhD robert.richards@milliman.com