Northeastern University Hospital Surge Capacity Planning Model: Bed, Ventilator, and PPE 1-30 Day Demand

Download tool and updates: https://www.hsye.org/covid-19-capacity-mgmt

1. Model Overview

The COVID pandemic is placing enormous surge demand and strain on health system capacity, staff, personal protective equipment (PPE), and other supplies, individually and regionally. Many hospitals and policy makers need real-time information about these evolving demands to make critical operational decisions. We developed the Hospital Surge Capacity Planning Model to help health systems estimate and visualize 1-to-30 day ahead hospital-specific demand for medical and ICU beds, ventilators, PPE, medications, and available staff on a rolling basis.

The tool is freely available to any health system worldwide and can be downloaded from our COVID models website at http://www.hsye.org/covid-19. We rapidly developed this model by adapting and integrating 10 years of prior research supported by the Agency for Healthcare Research and Quality (AHRQ), National Science Foundation (NSF), and National Institutes of Health (NIH) on three separate sets of new generalizable methods for modeling epidemics (originally focused on the U.S. opioid epidemic, and later vaping), detecting changes in hospital outbreaks, and predicting hospital daily occupancy on unit-by-unit under normal, changing, and surge conditions.

The resulting model provides 1-to-30 day ahead projections on a rolling basis of any individual hospital's (1) bed demand and occupancy (census) for medical beds, ICU beds, and ventilators, and other critical equipment; (2) 1-to-30 day ahead PPE consumption "burn" rates and stock-out dates as a dynamic function of predicted occupancy; and (3), shortly, staff needs and availability, which also can change dynamically as a function of bed occupancy, patient type mix, and caregiver exposure. The overall objective is to provide early signaling of capacity, supplies, and staffing concerns at hospital and system levels.

This model can complement more macro-level epidemic models informing public health policies, most using conventional susceptible-recovered concepts. While such large-scale models unquestionably are useful, more granular methods and data are needed for individual systems to make key operational day-today decisions about PPE and supply needs, opening ICU and medical space, modifying admission thresholds, allocation of dwindling supplies and medications, invoking makeshift PPE and equipment policies, and so on. Our approach blends theoretic and data-driven modeling methods to produce detailed actionable decision support, integrating current patient census by type, local new COVID case predictions, projected regular admissions, and probabilistic admission units (medical bed, ICU, ventilated), lengths of stay, and staff and PPE needs based on projected occupancy by type over time.

2. Examples of Use

The model can be used in a number of ways to help hospitals prepare for and manage capacity concerns from COVID like epidemics, including providing general information, operational decision making, expediting significant concerns, and in extreme situations to start making decisions on space to convert, admission criteria, use of makeshift PPE, patients who might need to be moved to different locations to free up space for patients in critical conditions. Examples of questions the model can help answer include:

a. General information

- How many patients will a hospital have in the ICU each day, and how many are going to be ventilated?
- How much spare medical/surgical unit capacity will be available each day should a hospital need it?
- When in the future is demand likely to exceed capacity? What is the expected timing of these events locally within this hospital?
- How fast will the hospital consume PPE and other supplies?
- Will there be enough capacity and supplies at an epidemic's peak, and if not, how large will be the shortage to help plan accordingly?
- How many staff, by type, are likely to be unavailable due to exposures, and how will this affect all the above?

b. Operational decisions

- When should a hospital convert routine space to ICU or isolated beds?
- When will a hospital need to decant, defer, or transfer non-COVID patients to other facilities and/or how much adaptive capacity will be needed?
- What percent of patients to decant, divert elsewhere, etc?
- What would be the impact of changes in admission criteria or location while waiting for test results?
- When to transition space back to its original use?

c. Best/worst case scenarios (sensitivity analysis)

- Given inherent variability (e.g., random lengths-of-stay) what is the probabilistic range of results a hospital might expect over the next week and month?
- Given lack of historical data and various input uncertainties (COVID intubation durations, ICU lengths-of-stay, in-hospital mortality, new case arrival rates), what do any of the above look like under different assumptions?
- What assumptions are a hospital's results most critically sensitive to (e.g., to inform data analysis and external benchmarking)?

d. Crisis management

- When to invoke makeshift PPE and equipment policies?
- When to enlist retired caregivers, primary care providers, and others in staffing routine healthcare acute care setting delivery needs?
- Early warning about needs to expedite supply and medication replenishments, and/or share resources between facilities?
- How to optimally use combined capacities and resources across a multi-facility health system?

e. Future planning

- How to prepare for future possible epidemic waves, such as a potential recurrence when the spread of COVID slows and we re-start normal activities?
- How to revise surge capacity policies, as a learning health system, to enable more pro-active surge management in the future?

3. Tool Description

The model is implemented in an Excel front-end for ease of use and sharing (**Figure 1**), with all mathematics computed in the back-ground, which take into account day-to-day census and bed occupancy flows, resources on hand, lengths-of-stay, patient escalations and step downs between medical and ICU beds, transfers and mortality. The file consists of five worksheet tabs – one for bed demand, PPE consumption, staff availability, input instructions, and calculation FAQs. Results are displayed graphically as run charts over time and tabularly, formatted for 8.5x11 printing to facilitate bed huddles, surge management meetings, and the like.

For basic functionality, the user enters various hospital and population specific information including historical lengths of stay for medical, ICU, and ventilated patients, admission rates and percentages by patient type, percentages that escalate and step down, and so on. Similar information is entered or estimated for COVID patients. Several op-

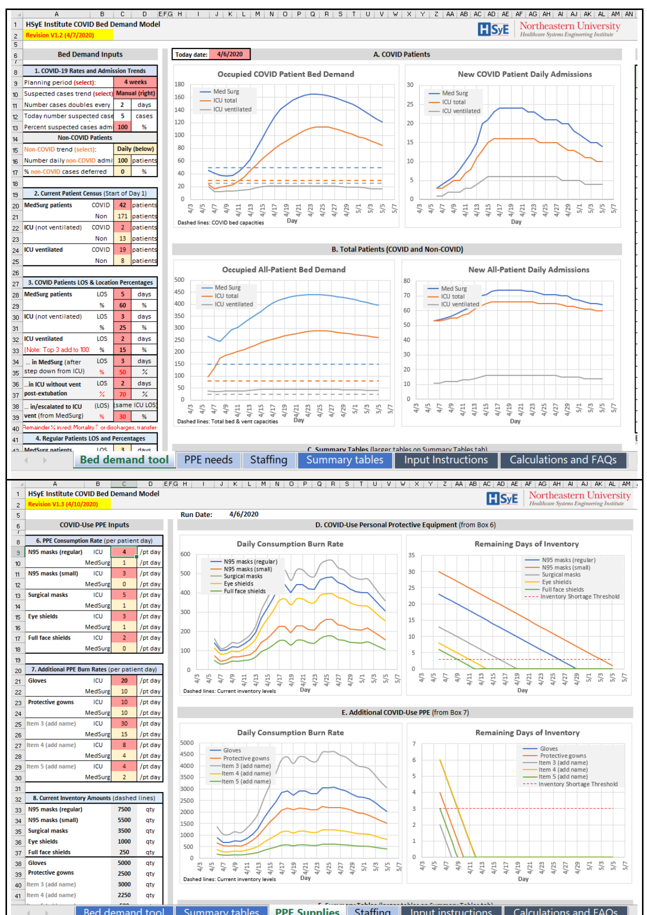

Figure 1. Hospital Surge Capacity Tool. (a) Projected daily new cases and total demand for COVID and non-COVID medical/surgical beds and ventilators, **(b)** Projected daily PPE consumption rates and days of remaining inventory

tions exist for new COVID and non-COVID patient projections under drop down menus, i.e. doubling rates, local estimates, external epidemic model projections, curve fitting (see FAQ for detailed explanation and use). The model then computes 1-30 day ahead COVID and non-COVID (1) new admissions by day and location (med/sur, ICU, intubated) and (2) total bed occupancy by day and location, and if desired corresponding PPE consumption rates and days remaining inventory corresponding to occupancy and status projections.

Optional functionality includes additional PPE and staff worksheet tabs. A hospital can enter key PPE, consumable supplies, or medications to track, along with historical, estimated, or benchmark consumption rates by patient type (COVID, non-COVID) and bed type (medical, ICU, ventilated). A similar prototype worksheet exists (completed shortly) for staff availability by type based on exposure rates and isolation durations and as a function of bed occupancy. A random simulation (Monte Carlo) option also exists to account for random admission volumes, admission units, lengths of stay, in-hospital transfers, and mortality. Since these all vary in actual practice, this can help a hospital visualize the range of possible futures that actually might occur to develop a sense for the likely range of scenarios they will experience.

We also recommend conducting sensitivity analysis on any inputs a hospital is unsure about in order to develop a similar sense of the range of possible futures. The model has been thoroughly debugged and is as accurate as the input assumptions. By example, health systems are reporting 85-95% accuracy in bed demand one-to-five days into the future using estimated historical lengths-of-stay, admission units, and other inputs.

4. Further Information and Resources

Further details on model logic, inputs, and calculations can be found in:

- The “Input Instructions” and “Calculations FAQ” tabs of the Hospital Surge Capacity Model spreadsheet.
- A seven-minute video on the tool website demonstrating its use and explaining the various inputs and outputs (surge capacity tool video)
- The below working paper (References section)
- There also is an evaluation form to provide feedback on use, suggestions, and accuracy. We will report back on the tool website case studies and accuracy data (only with a hospitals permission). Confidential feedback also is invaluable as we continue to expand the tool.

5. Development of Tool

The Hospital Surge Capacity Planning Model was developed by Principal Investigator James Benneyan and his team at the Healthcare Systems Engineering Institute at Northeastern University, in collaboration with Dr. Michael S. Rosenblatt at the Beth Israel Lahey Health-Lahey Hospital and Medical Center. Additional input was provided by several healthcare organizations

affiliated with our collaborative research Center For Healthcare Engineering Research (CHER). This project was indirectly supported by the Agency for Healthcare Research and Quality (P30HS02445301), National Institute of Drug Abuse (R21DA046776-01), and National Science Foundation (CMMI-1742521).

6. Feedback

We are continually updating the tool, with revisions and supporting materials posted to the same link (https://www.hsye.org/covid-19-capacity-mgmt). To provide feedback, improvement suggestions, or experiences and successes using the model, please contact us at hsye@coe.neu.edu.

7. References

Benneyan JC, Rosenblatt M, Bargal B, Yap S, Kaya Y (2020), *A hospital surge capacity bed, equipment, and staff demand planning model*, working paper.

8. Tool Citation

Benneyan JC, Bargal B, Yap S, Kaya Y. “Hospital Surge Capacity Bed, Equipment, and Staff Demand Planning Model”, Healthcare Systems Engineering Institute, Northeastern University; 2020. Available at: https://www.hsye.org/covid-19-capacity-mgmt

9. About the Developer

James Benneyan, PhD, is executive director of the Healthcare Systems Engineering Institute and a professor of industrial engineering at Northeastern University. Benneyan is a nationally recognized expert in healthcare systems engineering, Co-PI or Co-I on four AHRQ Patient Safety Learning Labs, and past president of the industrial engineering Society for Health Systems. His research focuses on development and application of systems engineering methods to improve healthcare broadly, including patient safety, care access, capacity modeling, rural disparities, epidemic modeling, and burnout. He and his team are funded by AHRQ, NIH, NSF, and members of the university/industry collaborative research Center for Healthcare Engineering Research (CHER).

HSyE | Northeastern University *Healthcare Systems Engineering Institute* (https://www.hsye.org/)

CHER (https://www.hsye.org/cher)

Appendix – Large View of Tool

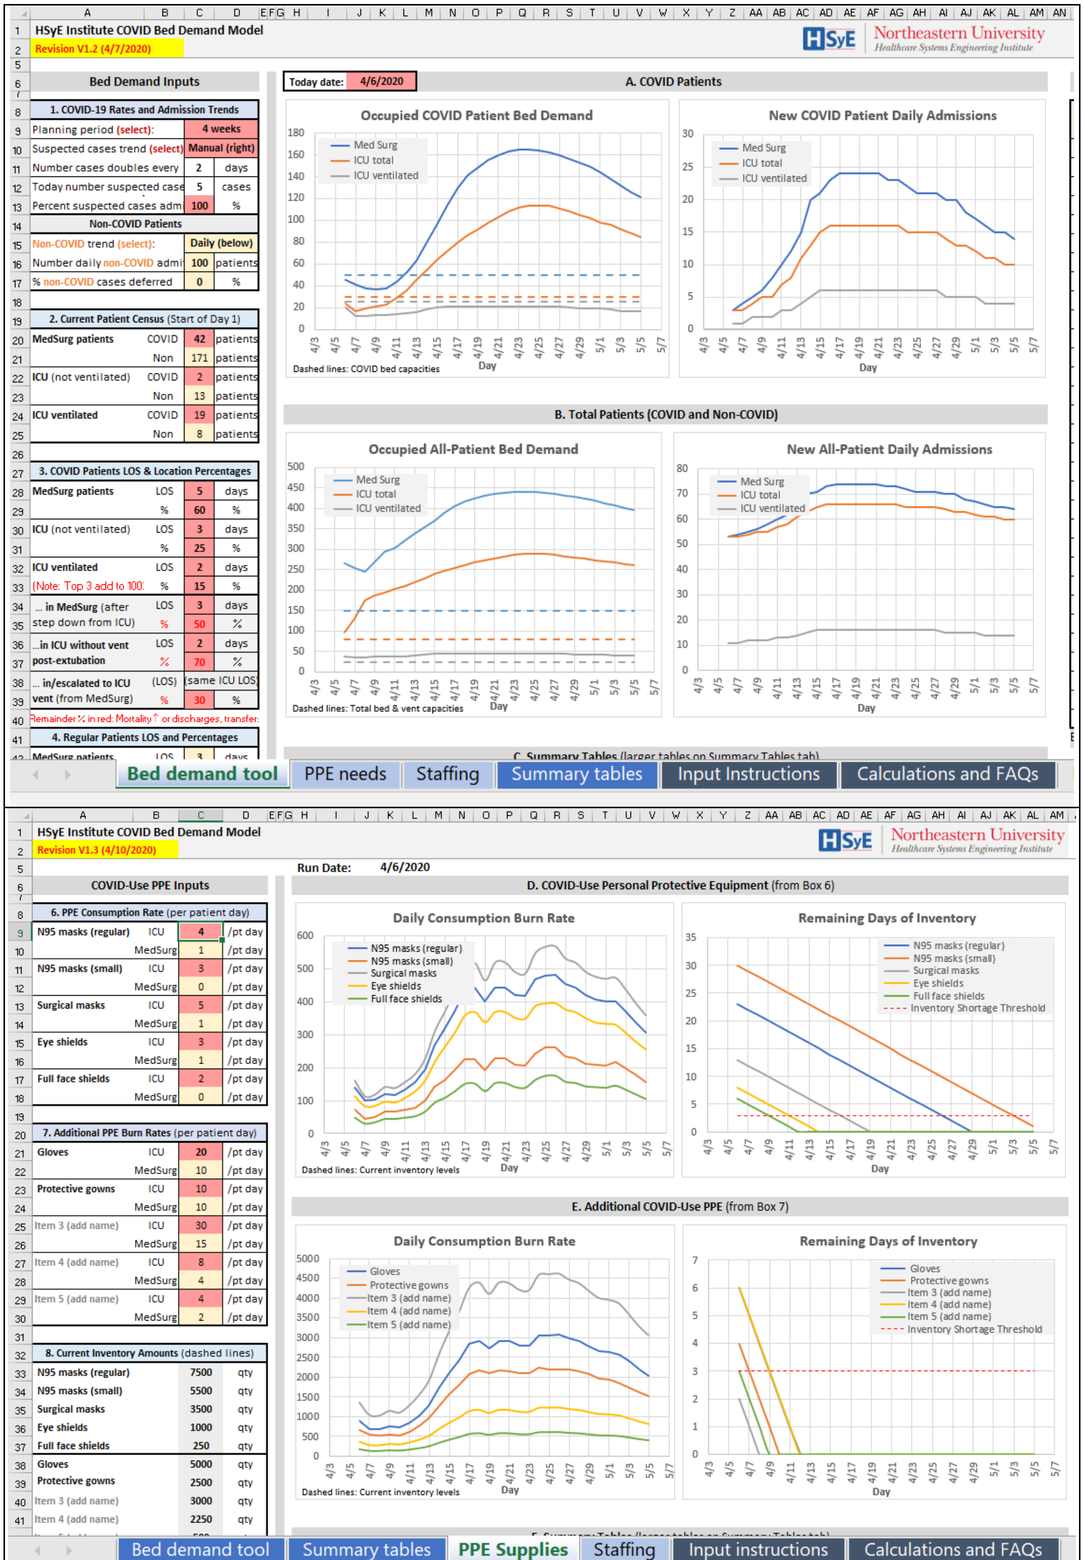

Figure 1. Hospital Surge Capacity Tool. (a) Projected daily new cases and total demand for COVID and non-COVID medical/surgical beds and ventilators, **(b)** Projected daily PPE consumption rates and days of remaining inventory

Appendix: Tool Calculations and FAQ Summary

Envisioned Uses

This tool computes the expected surge bed and ventilator demand for COVID and non-COVID patients for a given hospital, health system, or region. Several models exist for projecting the epidemic nationally or regionaly on a general scale (most using standard susceptible-infected-recovered epidemic spread models) and can be useful for general planning. This tool instead is intended to help a hospital make planning, readiness, and operational day-to-day decisions within the health system by considering day-to-day census and bed occupancy flows - e.g. resources on hand, length-of-stay, typical patient sequala (escalate from a MedSurg to ICU bed, step down, mortality, etc).

Questions: When in future is demand likely to exceed capacity? How much staff and PPE will be needed day-by-day as an epidemic increases exponentially and runs its course? What is the expected timing of these events locally? When to open up new MedSurg or ICU space for COVID cohorted patients? What percent of patients to decant, divert elsewhere, etc? When to transition space back to its original use?

Two Versions of Tool

Two versions of the hospital surge tool exist, one that is entirely macro-free and one that requires macros to be enabled (runnng visual basic). Both tools can be found in the same place on our website, along with other tutorial information, and are free to use, disseminate, cite, and share with others. While we strived to implemnent everything in a macro-free file, we implemented both tools because some organizations or users are uneasy or have a policy against enabling Excel macros. Nonetheless, the macro-enabled tool does allow us to provide a bit more fucntionality, and likely will be expanded over time. Currently this includes (1) full simulation replications and probability intervals on all results and (2) a self-calibration algorithm to improve prediction accuracy. Envisioned "extras" in the future include more advanced curve fitting of new admissions, data management features, real-time decision support and/or optimization of space modifications, staffing, and PPE make-shift decisions, and linking acute care surges to downstream capacity needs. The standard macro-free tool however provides most functionality that most users would typically use and is envisioned as the workhorse (rather than, say, a "light" version)

More Information and Downloads

More information on both hospital surge tools can be found on our website, including an overview paper and tutorial videos, t the following links:

- Tool website: https://www.hsye.org/covid-19
- Standard tool: Macro-free link
- Advanced tool: Macro-enabled link
- White paper: https://c7cc8594-7438-4dc2-bab5-1b4b1a258fd0.filesusr.com/ugd/8e76e9_1382769fa12c4cae95372dd47c9a577c.pdf
- News story: http://news.northeastern.edu/2020/04/15/hospitals-are-bracing-for-a-surge-of-covid-19-cases-this-model-can-help-them-get-ready-for-staff-and-equipment-shortages/
- Tutorial video 1: https://video.wixstatic.com/video/8e76e9_e5b6d7e5baf74e9b9af13becfb56f9a0/720p/mp4/file.mp4
- Tutorial video 2: https://video.wixstatic.com/video/8e76e9_e5b6d7e5baf74e9b9af13becfb56f9a0/720p/mp4/file.mp4

How to Cite:

- White paper: Benneyan JC, Rosenblatt M, Bargal B, Yap S, Kaya Y (2020), *A hospital surge capacity bed, equipment, and staff demand planning model*, working paper.
- Tool: Benneyan JC, Bargal B, Yap S, Kaya Y. "Hospital Surge Capacity Bed, Equipment, and Staff Demand Planning Model", Healthcare Systems Engineering Institute, Northeastern University; 2020. Available at: https://www.hsye.org/covid-19-capacity-mgmt

A. Bed and Ventilator Capacity

Basic logic and calculations: The general logic of how results in the tables and figures are computed is as follows, based on the user-inputs explained on the "Instructions" tab. Of the estimated new COVID and non-COVID cases each day (see below), a percentage of each are admitted to MedSurg or ICU beds (without or with ventilators) following given percentages for each. These patients join those already in each unit and remain in their admitted locations for the indicated number of days (length of stay, LOS), which may differ by type on unit and COVID vs non-COVID patient, after whcih each patient then follows indicated flow paths (see below).

Percent admitted: A user option allows something less than 100% of COVID suspected patients to be admitted, such as if symptoms are mild, the hospital is under siege, they are routed to another facility given lack or capacity, or other adaptive decision making. Similarly a percentage greater than 0% of regular (non-COVID) patients could be deferred elsewhere for similar reasons. Under routine conditions, these two percentages likely would be set to 100% COVID patients admitted and 0% non-COVID patients deferred respectively

Current census: A user would input the current census for COVID and non-COVID patients by location. For simplicity and to avoid needing more detailed input data, we assume that each of the current patients are uniformly distributed over their respective LOSs in terms of when they leave the hospital; for example if the current census of MedSurg non-COVID patients is 10 with a LOS of 5, two current patients are assumed to be discharged on each of the following 5 days. The census figures would be updated each day that the model is run

New cases: This is perhaps one of the most important inputs that affects projections and accuracy, after which most of the patient flow logic more-or-less simply plays out as above. There currently are four options for indicating the number of new COVID cases that will arrive each day of the planning period (1, 2, or 4 weeks) which the user selects from a drop down menu: "doubling rate", "manual entry", "exponential growth fit", and "growth curve fit". The first two options are based on user input assumptions as to expected upcoming growth (or shrinkage) or experienced recent growth in new cases, whereas the latter two options automatically fit growth curves to a hospital's historical recent data on a rolling basis. All methods work equally well when an epidemic is in growth mode, at a plateau, or in a decreasing mode. Options can be used in any combination for COVID and regular (non-Covid) new cases, e.g. exponential auto-fitting for COVID cases and manual entry for your plans to taper scheduled non-covid admissions (as one possible example).

We also recommend running the model under a few assumptions and scenarios for data and method in order to visualize and develop a sense for how these impact the range of possible hospital capacity of the next few weeks.

1. *Doubling Rate* - This is a standard epidemiology concept that describes the number of days until the current number of new cases per day doubles. Early in epidemics, population outbreak curves of daily cases tend to start fairly flat, then rise exponentially as the epidemic rockets skyward (when most surge conditions, admissions, and shortages occur), then flatten out as the epidemic is in full and constant activity (peak of the curve), then start to slope downwards. The doubling rate should follow accordingly - think of this rate as the slope, secant, or first derivative of the new case curve - starting long (eg new cases double every 14

days) then being short (2 days doubling), and then being long again (eg 1000 to indicate a flat slop when the epidemic is in full swing at the top of the curve).

- *What do I do when the early escalating upward slope is over?* As the epidemic starts to diminish, a negative slope of the new case curve can be represented by using negative numbers (eg a doubling (now halving) of -5 days means that the number of new days is reduced by 50% every 5 days). Note that while the epidemic is increasing up to the inflection point, this method slightly over-estimates new cases, which may be fine. Also see the similar option below called "exponential fit".

1. *Manual Entry* - Alternatively a user can enter a column of the estimated number of new cases each of the next 1-30 days to represent a specific 'curve'. These numbers could be based on (a) local expertise and insights, historical recent data trends, or regional predictions from epidemic models (e.g. taking regional or state results from any of the epidemic models (see links) and multiplying those estimates by the percentage of regional patients that your facility would likely see, such as based on your percentage of state beds or routing decisions regarding managing regional facilities in concert. Manual estimates of new cases can be used for either COVID patients, non-Covid patients, or both. To use this option, a user selects "Manual Entry" from the drop down menu in cells C10 (Covid patients) and C15 (non-Covid) and then enters your projected number of new cases in columns AO-AS. Note that you can save the data in these columns and just update them each day you run the model (deleting days that have already occurred and adding more days further out).
2. *Exponential Growth Curve Fit* - This option automatically fits an exponential curve to your recent data on new cases. As above, you can use this which you enter in columns AU-AW for your recent data on COVID new patients and regular/non-COVID new patients. Enter the date of the first day that you have data for in cell AU11 at the top of these columns and the rest of the dates will self-calculate. Note that you can simply keep adding more data for new days at the bottom of these columns as time goes on.

- *User Options* - The curve fitting will use all or the most recent *N* days of data on a rolling basis in the curve fitting depending on what you enter in AY11 and AZ11. The cells immediately below these automatically display what the date range is for the curve fitting, based on the "run date that you have entered in cell L6 (all dates through this tool and displayed in graphs and tables key off of this value each day that you run the tool.) If you pick a number of days (e.g. 10) and only have 8 days of past data, it will use all 8, then the next day when you had another row will use all 9, and then the after that use 10 days always going forward (unless you change the number in cells AY11 and AZ11. If you always want to use all past data, you can just set these cells to large numbers (e.g. 1000).
- *How many past days to use?* It depends - more data is better for a better curve fit IF the epidemic dynamics has not changed, but less is better if old data represent different conditions, such as pre-quarantining. Experientially, we have found 14 days seems a good choice. You might reduce this number if you feel the dynamics of the epidemic have changed such that old data no longer represent the same "causal system" (e.g. new policy on mandatory masks, school closures, reduced willingness to follow precautions, and so on - but not just changes in new case numbers based on "normal" epidemic dynamics of spread). Or you might increase this number if the "system" has been consistent.
- *Can I fit on different numbers of days for COVID patients versus regular admissions?* Yes. In fact, you might use a longer time range for non-Covid admissions, such as 10 days for COVID and 20 for non-COVID (just as an example)
- *What about curve fitting for new cases data that are plateauing and decreasing, or is not very "curved"?* This follows the same as above for "doubling rate" - during peaks and

plateaus, the curve will automatically fit an "exponential" curve that is essentially flat and during decreases will fit an exponential curve that slops downward. The amount of curve or straightness fit to the data also will take care of itself to best match your historical data (within what is possible)

4. *Logistic Growth Curve Fit* - The details and user inputs for this option all work the same as for option 3 (exponential fit), now fitting a non-linear epidemic growth curve instead using novel methods we have developed. See paper pre-print for mathematical details. Briefly, it is based on a closed form first order approximation to "SI" susceptible-infected epidemiology models much like the "S" curves frequently shown in the media and literature (without using complicated systems of differential equations etc). In the short term, this tends to do very well, especially as it is updated each day the model is run (dropping old data and adding new data - essentially we fit an S curve to 10 day rolling data rather than assume a linear doubling rate or exponential shape. In numeric experiments we have found accuracy approximately doubled (estimated versus actual 1-14 day ahead occupied beds)

<u>Patient paths</u>: To keep things somewhat simple, the following (6) patient flows are assumed. A percentage of COVID and non-COVID patients in MedSurg can be escalated to the ICU (following different percentages) either with or without a ventilator. A percentage (typically 100%) of patients in the ICU with a ventilator stay in the ICU for some number of additional days post-extubation, and a percentage of ICU patients with no ventilator can step down to a MdSurg bed. The remaining patients spend their full LOS in whatever bed and ventilator status they were admitted to. Since the above likely captures a large majority of patients, other nuances and possibilities have been omitted, which is likely fine in terms of negligible impact on model accuracy for the intended purposes. Each patient also is assumed to follow one of the above flow paths after their current number of days in their current location, again likely fine but may result in a slight over-estimate of bed demand.

<u>Mortality</u>: While the current version does not have a specific mortality input field, in-hospital death can be handled as follows, in input Boxes 3 and 4. The percentage of patients who do not step down from the ICU are assumed either direct discharge, transferred, or died. Similarly for the percentage of ventilated patients who do not remain in the ICU post-extubatation. (Med surg mortality to be added in V1.3).

<u>Health System Capacity</u>: The current number of available beds in each type of unit and of ventilators are input in Box 5. Note that these may change over time as a health system opens up new capacity, shifts functional use of some units, and stands up makeshift capacity based on emerging conditions. Thus these should be updated as a health systems makes any such changes. These values however do not affect the analysis but rather just the horizontal dashed lines on the figures for comparison purposes.

<u>Day 1</u>: *Dr. Smith is bracing for a surge of COVID patients to hit her hospital in the upcoming week and wants to figure out how to plan her beds accordingly. Her team does not have enough historical data to estimate the trend of incoming patients, and would like to use state estimates to do so. Her hospitals COVID policy states that any and all patients who are suspected of having the virus are admitted.*

(inputs) Planning period = 1 week, Suspected case trend = doubling forecast,

B. PPE Logic

Logic on the Equipment and PPE worksheet is based on resulting bed demand. Aside from user entries in Boxes 6-8, all results are calculated from the resulting bed demand on worksheet 1.

- *Types of PPE or Other Disposable Equipment* - In boxes 6-7, enter up to 10 types of disposable equipment or PPE that are important for you to track (column A) and their corresponding consumption rates (column C) per ICU COVID patient day and per Medical/Surgical COVID patient day.

- *Consumption Rates* - These can be estimated from your historical data, benchmarks, or estimates; for example, you can query your recent data on total number of COVID ICU patient days and total consumption of a given type of PPE over the same number of days - dividing the latter by the former produces consumption rate per patient day. While we envision these relating to COVID patients, they equally could be for any patients. (Durable equipment , such as dialysis machines, can be entered on the Bed Demand worksheet instead of ventilators.)
- *Current Inventory* - In Box 8 (optional) for each type of PPE you entered above, enter the current amount of items on hand on the run date of the tool. This input is optional and is used to compute the remaining days of inventory (ignoring incoming orders) currently on hand shown in Figures D and E, computed from the projected bed demand on the first worksheet and burn rates.
- *Stock-Out Warning Level* (optional) - This is optional and creates the horizontal dashed warning lines in Figures D and E, but is not used in any calculations.

C. Staffing Logic

Logic on the Staffing worksheet is based on resulting bed demand. Aside from user entries in Boxes 9-12, all results are calculated from the resulting bed demand on worksheet 1.

- *Types of Staff*- In Boxes 9-10, enter up to 10 types of staff (column A) and their corresponding exposure rates per patient day (column C) and patient-to-staff ratios (column D), both for ICU patients and medical/surgical patients.
- *Exposure Rates* - Similar to PPE consumption rates (see above), these can be estimated from your historical data, benchmarks, or estimates; for example, you can query your recent data on total number of COVID ICU patient days and total exposure of a given type of staff over the same number of days - dividing the latter by the former produces exposure rate per patient day. While we envision these relating to COVID patients, they equally could be for any patients. These inputs are used to compute the number of unavailable staff, byt type, either due to being isolated or wating for tests, shown in Figures G and H, based on the number of occupied beds computed in worksheet 1.
- *Patient-to-Staff Ratios* - In column D of Boxes 9 and 10 enter the required or desired staffing ratios that you ideally would like to be able to maintain. These are used to comput the number of needed staff shown in Figures G and H based on the number of occupied beds computed in worksheet 1.
- *Testing and Isolation Durations (Box 11)* - In Box 11 enter the test waiting period, the percent of tests that are positive, and the following isolation perdiod. For simplicity, we assume these are the same for all staff and types of beds (ICU vs med/surg). A good rule of thumb would be to simply use the average across all staff types and beds, or to use those that correspond to the most worrisome (presumably ICU beds and ICU staff).
- *Staff Pool Sizes and Status (Box 12)* - In Box 12 enter your total available pool of staff by type (which may vary over time as you adapt to the epidemic, status of your hospital, and local conditions) abd their current status - the number of each type of staff currently being tested and who have tested positive and are in isolation or otherwise unavailable to work. These results, along with the isolation and testing durations in Box 11, are used to compute the number of available staff shown in Figures G and H.

D. Simulation Functionality

The macro-free version of the tool has basic simulation capability (described here) and the macro-enabled version of the tool has full simulation capability (see below). Cell C20 has a drop down menu where you can select "Deterministic" or "Random". For reports and upper management, they may prefer to see the "nice smooth line" results that Deterministic mode produces, but to get a sense for what may REALLY happen the Random mode can also be useful.

- *Why Bother?* Most results you see in the media assume every thing happens deterministically, such as 10% of 100 patients escalating to ICU meaning exactly 10 (as opposed to each patient having a 10% random chance, basically a binomial random variable). In actual practice, lots of events will naturally vary - the exact number of new cases each day, the percent of these that need an ICU bed, each patients length of stay - rather than assuming they all are fixed numbers, all LOSs are equal to the average LOS, and so on. The Deterministic mode basically displays the most likely (expected value) outcome, whereas in reality any given "next week" could occur in lots of different ways
- *How to Use (Macro-Free Tool)?* Once you have input all the run conditions, just hit your F9 key several times. (On most keyboards this is the "recaluation" key - on some it may be different, or may be function-F9, or something else). Each time you hit the F9 key, a new set of random numbers is generated (number of new patients each day, each patient's LOS, and the percent that are admitted, escalate, or step down and all the table and figure results are recomputed. Hit F9 for example 25 times and you will see 25 possible futures - any one of which could actually occur! Since we'll only live through the next 4 weeks once (hopefully), this gives you a sense for the range of possibilities for occupied beds, staff needs, PPE burn rates, and so on. (You can be on any worksheet - beds, staff, equipment - when you hit F9 to see how results are affected.) Let your mind remember what it sees 10-25 times and in essence this gives you a sense for the probability interval within which the next few weeks will unfold.
- *Lengths of Stay* - In cell C21, enter the percent by which patient lengths of stay vary from their average roughly 95% of the time. For example if on average an ICU patients LOS is 14 days, with most (95% or so) patients varying between 7 and 21, you would enter 50% here (7 and 21 being 50% smaller and larger than 14). This is used to estimate the parameters of a negative binomial probability distribution (approximated by a somewhat equivalent gamma probability distribution), which is used to randomly generate each patient's length of stay number of days. For simplicity we assume these coefficients of variation are roughly the same for all patient types and locations (e.g. a med surg non-COVID patient might have an average LOS of 4 days with 95% between 2 and 6 - still 50% variance from the mean. Empirically this principle often occurs in practice, but if not you can either use the average about of percentage variation or the percent variation for the resources of greatest concern (presumably ICU space, associated staff, and equipment, and PPE).
- *Full Simulation Functionality (Macro Enabled Tool)* - See a below section that describes the alternate version of this tool, which uses enabled macros, that gives full simulation functionality to run any number of replications and plot actual probability intervals. We implemented this both ways because, while we like the full functionality, some users are uneasy or have a company policy against enabling Excel macros (and because it takes a little longer to run a full simulation). To use the simulation, be sure to download the correct version (Version V2.0 or higher). All the above inputs and functionality are identical to the non-macro version, and you can do everything described above (including use the "F9 simulation" feature above. You should also see two buttons in the upper right hand region of the Bed demand worksheet (at the top of columns AO-AU) called "Run Simulation" and "Reset Simulation". To run multiple replications and produce probability intervals on all the figures, click the Run Simulation button and enter the number of replications in the window that pops up and then hit

"Run". On an average computer it will take 10-60 seconds to run 5-100 replications, after which all the figures will be updated with 95% probability intervals shown around each deterministic (average) plotted line. To remove these results, click on the "Reset" button.

- *Integrated Use of Deterministic, Random, and Simulation Modes* - The way I usually use the overall tool is to first look at results in Deterministic mode to get an overall feel for results, and then hit F9 several times to get a feel for the likely variation in these results. Then I might test my input assumptions a bit, or run other assumptions on new cases, lengths of stay, and so on, and then if I want to be more exact (and I have a 60-90 seconds to wait) click the "Run Simulation" button to plot probability intervals. As analysts these make perfect sense to us, but some managers like the nice clean deterministic graphs - just remember that real results can differ a bit, potentially making actions such as opening up new beds too late.

E. Accuracy Tracker

We added an "Accuracy data" worksheet to help users track and develop a sense for input and model accuracy over time. The worksheet is a bit complex but self explanatory. Each day you run the model, copy the Table 1 bed demand data on the bed demand worksheet into the indicated cells on the accuracy data worksheet - what the model is predicting for the next 1-4 weeks of bed demand. Also copy the inputs you used in Box 2 for current census on the day you ran the model. Over time, this worksheet will be able to compare what the model predicted, for example, for today's census when you ran the model two weeks ago, giving us 14-day ahead accuracy summed over as many days as you have been using the tool. Results are automatically calculated and displayed on the "Accuracy Results" worksheet. By using this feature, you can learn over time about (1) how to change your input assumptions (and which method of estimating new cases to use) to improve accuracy and (2) how far into the future results in your particular hospital are accurate enough to act on. We have done this analysis for several hospitals, and results can vary by hospital (and improve over time) - usually due to poor initial estimated inputs but also just with some inherent differences between hospital's served populations, extraneous factors like changes in state policies, willingness to follow precautions, and so on. In many cases we have found the tool is, after some calibration and input learning, 90% accurate across the first 14-17 days and then decaying to 80% accuracy at around 15-30 days. (Aside -- Keep in mind that this is "deterministic accuracy" in the sense of error of actual census from predicted census, the the percent of "actual" data that fall within their probability intervals likely is much higher and closer to 95-100%)

G. Macro-Enabled Tool

More information to follow - for now please see above. While we strived to implement everything in a macro-free tool, we also developed this tool to be able to offer some more advanced functionality that will be expanded over time. These include:

- Full simulation capacity (completed)
- Accuracy tracking and self-calibration algorithms (to be released July 2020)
- Others to be developed - Data archiving, input suggestions, action decision support and optimization, and others