OM Forum: COVID-19 Scratch Models To Support Local Decisions

Edward H. Kaplan[*]

(Forthcoming: *Manufacturing & Service Operations Management*)

Abstract

This article is based on modeling studies conducted in response to requests from Yale University, the Yale New Haven Hospital and the State of Connecticut during the early weeks of the SARS-CoV-2 outbreak. Much of this work relied on scratch modeling, that is, models created from scratch in real time. Applications included recommending event crowd-size restrictions, hospital surge planning, timing decisions (when to stop and possibly restart university activities), and scenario analyses to assess the impacts of alternative interventions, among other problems. This paper documents the problems faced, models developed, and advice offered during real-time response to the COVID-19 crisis at the local level. Results include a simple formula for the maximum size of an event that ensures no infected persons are present with 99% probability; the determination that existing ICU capacity was insufficient for COVID-19 arrivals which led to creating a large dedicated COVID-19 negative pressure ICU; and a new epidemic model that showed the infeasibility of the university hosting normal spring and summer events, that lockdown-like stay-at-home and social distancing restrictions without additional public health action would only delay transmission and enable a rebound after restrictions are lifted, and that aggressive community screening to rapidly detect and isolate infected persons could end the outbreak.

[*]William N. and Marie A. Beach Professor of Operations Research, Professor of Public Health, Professor of Engineering, Yale School of Management, 165 Whitney Avenue, New Haven 06511, CT. email: edward.kaplan@yale.edu

1 Introduction

The novel coronavirus SARS-CoV-2 outbreak and accompanying COVID-19 pandemic have shocked the world, and nations around the globe are struggling to contain these outbreaks. Within the United States there is substantial variation in the extent and timing of the epidemic across states, cities and towns, and there is a need to help support decisions at the local level as institutions and local governments struggle to respond.

Yale University is located in New Haven, Connecticut, a small state with only 3.6 million residents. However, as of April 12, only four states have more COVID-19 cases per capita (New York, New Jersey, Louisiana and Massachusetts) and deaths per capita (New York, New Jersey, Louisiana and Michigan) than Connecticut (https://www.worldometers.info/coronavirus/country/us/).

At the end of February, the Yale President convened a group of public health and medical experts to advise university decision-making and provide public health guidance to the Yale Community. The Yale University COVID-19 Advisory Committee (henceforth YCC), chaired by Dr. Paul Genecin (the Director of Yale Health), includes the Deans of the Yale Schools of Public Health, Medicine and Nursing, the Chair of the Department of Epidemiology of Microbial Diseases, the Director of the Yale Institute for Global Health, other experts in infectious disease from the Yale School of Medicine, the Yale New Haven Hospital, and Yale Health, the Yale President's Chief of Staff, Yale's Vice Provost for Health, the Director of the City of New

Haven's Health Department, and the author of this paper. In total the YCC has thirteen members who convened daily via phone (including weekends) throughout the month of March into the first weeks of April.

The YCC has addressed many university issues, as summarized in Table 1. Some of these issues required modeling analysis to support the decisions

Yale COVID-19 Committee Issues Addressed To Date
• Infection control, social distancing, travel recommendations • Crowd capping; home monitoring/quarantine/isolation; student housing
• University operations and timing decisions • Moving classes online; keeping students off-campus; essential personnel/stay-at-home decisions; implications for University calendar
• Helping kickstart/run critical public health activity • COVID testing labs at Yale School of Medicine; contact tracing for state; application for healthcare workers to self-monitor for early signs of infection; stress test Yale New Haven Hospital ICU capacity; plannig for community COVID screening
• Advice on university communications

Table 1

under consideration. These analyses were conducted by the author in real-time using hastily-constructed "scratch models" tailored to the questions asked. While initially called upon to support the university via the YCC, this policy modeling effort expanded to provide information to the Yale New Haven Hospital, and more recently the State of Connecticut.

In this paper, my goal is to present the questions asked, models developed, and advice offered over a very intensive effort conducted during March to mid-April 2020 where modeling was influential. *This is not normal research*: due to the time pressure and nature of the decisions being made, it was not possible to spend the amount of time or provide the usual attention to statistical details the way one would in pure academic research. Rather this has been an adventure in policy modeling (Kaplan 2020a) conducted largely in solitaire using a laptop computer while working from home. Nonetheless, I feel it important to document this effort, in part to encourage other researchers to help support their own local communities during times of crisis.

The next section reviews simple models supporting decisions restricting the size of Yale University events. The following section focuses on a queueing analysis conducted to "stress test" hospital intensive care unit capacity. The paper than shifts to epidemic modeling, first in application to university timing decisions, and then in support of state efforts to monitor the outbreak while considering whether to continue restrictions and deploy aggressive community screening. The paper concludes with a brief summary and current issues on the horizon.

2 Restricting the Size of University Events

The first meeting of the YCC was March 3, 2020. This date coincided with the reported infection of a man from New Rochelle in Westchester County, the first recorded case of community transmission in New York State (Hill and Villeneuve 2020). It would only be another five days until the first COVID-19 case was reported in Connecticut (Lamont 2020a). The YCC was asked to consider restricting the size of university events in response to the possible spread of infection.

The quickly agreed-to goal of restricting attendance was that no new transmissions of infection should occur as a result of an event. It fell to the author to operationalize this goal and bring recommendations to the YCC. I did so as follows: there are two ways that no new infections would be transmitted at an event in an already sterile environment: either no infected persons attend the event, or some infected persons enter but all fail to transmit infection to others. If X is the random number of infected persons entering an event with n attendees in total (each of whom has probability p of being infected), then

$$\begin{aligned}\Pr\{\text{No Transmission}\} &= \Pr\{X = 0\} + \Pr\{X > 0 \ \cap \ \text{all } X \text{ fail to transmit}\} \\ &\geq \Pr\{X = 0\} = (1 - p)^n. \end{aligned} \tag{1}$$

Framing the issue this way changed the focus of discussion from infection control principles such as spacing between event attendees to the realization that *the best way to prevent transmission of infection is to ensure no infected persons enter the event.*

Producing a recommended crowd limit required two additional inputs: a comfortable lower bound for the probability that no transmission would ensue (the Chair of Yale's Epidemiology of Microbial Diseases department, Prof. Albert Ko, suggested 99%; this became known as the Ko Kriterion), and an estimate of p, the prevalence of infection among event attendees.

Of course, with no COVID-19 cases yet reported in Connecticut, nobody knew the value of p. I argued that one did not really need to know p; rather what was the largest value of p against which we wished to defend? Everyone believed that the underlying prevalence of infection was very small and likely no larger than 1 in 100,000 to 1 in 50,000. Given such small prevalence estimates, one could safely approximate $(1-p)^n \approx 1-np$ providing $np \leq 0.1$, and via the Ko Kriterion of 99%, the following very easily understood formulas took shape:

$$\Pr\{\text{No Transmission}\} \geq 1 - np \geq 0.99 \tag{2}$$

and thus

$$n \leq \frac{0.01}{p} \tag{3}$$

Thinking of *defending* against a prevalence of infection versus *estimating* the actual level, the YCC felt comfortable choosing a value of p that was five- to ten-times higher than *a priori* beliefs, thus p was set to 1 in 10,000, and the university administration accepted our recommendation to restrict events to at most $n = 100$ attendees. The university community was notified of this decision in an e-mail message on March 7, 2020 that was also posted to Yale's COVID-19 communications website (https://tinyurl.com/v7af7bd).

Fast forward to March 14, the date a Yale community member was first diagnosed with COVID-19. By then 126 cases had been diagnosed in Connecticut, 40 of whom had been admitted to the hospital, but as yet no COVID-19 deaths had been recorded. The YCC advised the university that given the rapid rise in cases over the previous two weeks, reducing event sizes further was required at a minimum, while many (including me) felt all group meetings should be abandoned. An argument to further restrict events to 20 or fewer participants went like this: early evidence suggested that SARS-CoV-2, the underlying infection responsible for COVID-19 disease, grew exponentially at a rate of 10% *per day* (Li et al, 2020). This rate implies a *quintupling* of the underlying incidence of infection over a 16 day period. Since we defended against a prevalence as large as 1 in 10,000 two short weeks ago, we should now defend against a prevalence as large as 5 in 10,000; setting p to 0.0005 in equation (3) led to a new maximum event size of 20.

Even with such a simple model, some misinterpreted the result as meaning that one event of 100 participants could be replaced by five events of size 20. While the likelihood that no infected person enters a group of 20 is greater than the same for a group of 100, one also must account for the factor-of-five increase in the number of events. Mathematically,

$$\Pr\{\text{No Transmission in 5 Groups of 20}\} \geq \left[(1-p)^{20}\right]^5 = (1-p)^{100} \quad (4)$$

so splitting the large event into smaller events does not help if one is only considering whether any infected persons participate.

There was some worry that the rule proposed was too conservative; suppose one knew the transmission probability q from an infected to an uninfected person given exposure. In an event with n people, X of whom are infected, a more complete model for the probability that no infections are transmitted is given by

$$\begin{aligned} \Pr\{\text{No Transmission}\} &= E_X[\{(1-q)^{n-X}\}^X] \\ &= \sum_{x=0}^{n} \binom{n}{x} p^x (1-p)^{n-x} \{(1-q)^{n-x}\}^x \end{aligned} \tag{5}$$

assuming that both already infected persons and transmission from them to others follow Bernoulli processes. While one estimate of $q \approx 1/200$ based on two infections among 445 home or hospital contacts of 10 patients who were infected outside of the United States was known (Burke et al, 2020), I did not feel comfortable relying on this number to increase n by using equation (5). In the end it did not matter, because on March 16, President Trump issued a national guideline asking Americans to avoid gatherings of more than 10 people (https://tinyurl.com/reg85wo), and the university immediately followed suit.

3 Stress Testing COVID-19 Intensive Care Capacity

On March 7, an interview with the director of intensive care in the Lombardy region of Italy appeared in the newspaper *Corriere Della Sera* (Ravizza 2020). The interview described the dire situation in Italy's thoroughly modern hospital system where extremely ill COVID-19 patients requiring negative pressure Intensive Care Unit (ICU) isolation and treatment with ventilators were being boarded in the hallways of the hospital due to the lack of ICU beds. This article set off alarm bells at the YCC meeting that day as attention shifted to the ability of the Yale New Haven Hospital (YNHH) to handle an explosion of COVID-19 cases. I volunteered to conduct a "stress test" of the hospital's ICU capacity with an eye towards determining the likely consequences of high COVID-19 ICU arrival rates. The Dean of Yale's Medical School immediately contacted YNHH's Chief Medical Officer, and a medical student was dispatched to take an inventory of all the ICU beds in the hospital by location and ability to maintain negative pressure.

I contacted my colleague Edieal Pinker, an expert in healthcare operations, to help design a very quick analysis. Rather than attempting to forecast downstream arrival rates, we instead focused on determining the maximum arrival rates the hospital could handle before significant COVID-19 ICU bed blocking would occur. Our approach was very simple: we used standard Erlang loss models as taught in elementary operations research courses to estimate the probability that all beds would be filled under a host of scenarios considering different number of available ICU beds. Hospital data suggested that 90% of ICU beds were presently occupied; combined with an assumed three day mean length of stay for such patients (Hunter, Johnson, Coustasse 2014), it was possible to back out the implied arrival rates for such patients given the number of ICU beds staffed assuming steady state conditions from the Erlang loss formula. Letting λ_N and P_m denote the arrival rate of non-COVID-19 ICU cases and the probability that all beds are filled when m beds are staffed as determined from the Erlang loss model (presuming a mean ICU length of stay equal to 3 days), λ_N was determined by solving

$$\frac{\lambda_N \times (1 - P_m) \times 3}{m} = 0.9. \tag{6}$$

If COVID-19 ICU patients were to arrive at rate λ_C with a mean COVID-19 length of stay given by ℓ days, and these patients had to compete for the same m (negative pressure) ICU beds, we approximated a new ICU-wide mean length of stay as

$$\overline{\ell} = \frac{\lambda_N \times 3 + \lambda_C \times \ell}{\lambda_N + \lambda_C} \tag{7}$$

We then employed the standard Erlang loss model to determine a new value for P_m that presumed daily arrivals at rate $\lambda_N + \lambda_C$ and mean length of stay $\overline{\ell}$. The number of COVID-19 patients blocked daily would then average $\lambda_C \times P_m$.

We considered COVID-19 lengths of stay ranging from three to ten days, and considered daily arrival rates of one or five COVID-19 ICU patients per day. We considered many different scenarios based on the ICU inventory and location data that were provided to us, and reported back to the hospital with our findings on March 10, three days after this issue was first raised.

Here are some sample findings, as presented in our communication to the hospital:

*If one is quite certain that a single COVID-19 ICU patient would arrive per day on average, existing MICU (*note*: medical ICU) capacity alone (25 negative pressure beds) is not sufficient to prevent blocking 30% – 40% of COVID-19 patients (presently 2 non-COVID-19 ICU patients are blocked and presumably boarded elsewhere; $\lambda_N = 10$ in this scenario).

*Using all 57 existing negative pressure ICUs cuts COVID-19 blocking in half to manageable levels if only one COVID-19 ICU patient arrives per day, but these 57 beds are not enough if five COVID-19 ICU patients arrive daily ($\lambda_N = 19$). Also, these 57 units are scattered throughout the hospital, which would create serious infection exposure problems.

*Using all 117 negative pressure beds (existing plus 3 floor Smilow surge (*note*: the Smilow Cancer Hospital is contained within YNHH)) would handle both COVID-19 and non-COVID-19 ICU patients, providing there are at most five COVID-19 patients per day and their LOS is at most ten days, though again that raises infection risk if COVID-19 patients are distributed throughout the hospital, and again requires relocating cancer patients ($\lambda_N = 19$).

After receiving this analysis, the hospital decided to re-purpose the top three floors of the Smilow Cancer Hospital to create a dedicated 84 bed negative pressure COVID-19 ICU. An additional 28 beds can be added by assigning COVID-19 patients two to a room on one of these floors. We showed that the 84 bed configuration, which now would not compete with non-COVID-19 ICU arrivals, would prevent significant bed blocking for arrival rates as high as 10 ICU cases per day while isolating all of these patients in one ward, significantly easing (but not erasing) infection control concerns.

4 University Timing Decisions

The transition from spring to summer is normally a celebratory time at universities, with graduation and alumni reunions followed by summer activities being the norm. Yale is no exception, and on March 19 I received a request to model the outbreak to determine whether graduation was feasible, with follow-up questions regarding a host of summer activities. Answering this required formulating and implementing a dynamic transmission model for SARS-CoV-2.

I decided to build upon a previous analysis of the early outbreak in Wuhan, China (Kaplan 2020b). In that model, an infectious person surrounded by susceptibles transmits new infections in accord with a time-varying Poisson process with intensity function $\lambda(a)$ denoting the transmission rate at infection "age" a (that is, a time units after infection). The expected total number of infections such a person transmits over all time, the *reproductive number* R_0, equals

$$R_0 = \int_0^\infty \lambda(a)da, \tag{8}$$

and as is well known, an epidemic can only emerge if $R_0 > 1$ (Anderson and May 1991, Britton and Tomba 2019). Analysis of the early outbreak data presented in Li et al (2020) led to the estimate $R_0 \approx 2.3$ which is consistent with values reported elsewhere (Ferguson et al 2020, Kissler et al 2020). The transmission intensity function $\lambda(a)$ was estimated from early outbreak data following "Euler-Lotka" logic as explained in Kaplan (2020b).

4.1 The Scratch Model

Turning this earlier work into a dynamic transmission model (the scratch model) required new analysis that was formulated and implemented over the course of a day or two as described below. Let:

$s(t) \equiv$ fraction of the population that is susceptible to infection at (calendar or clock) time t;

$\pi(a, 0) \equiv$ fraction of the population that has been infected for duration a at time 0 (this is the initial condition);

$\pi(a, t) \equiv$ fraction of the population that has been infected for duration a at time time t;

$\pi(0, t) =$ *incidence of infection* at time t (those newly infected have an age of infection equal to zero);

$\lambda(a) \equiv$ the transmission intensity as a function of age of infection introduced earlier.

With this notation, given the initial condition $\pi(a,0)$ and taking $s(0)=1$, the scratch model can be written as:

$$\pi(0,t) = s(t)\int_0^\infty \lambda(a)\pi(a,t)da,\ \ t>0 \tag{9}$$

$$\frac{ds(t)}{dt} = -\pi(0,t),\ \ t>0 \tag{10}$$

$$\pi(a,t) = \begin{cases} \pi(a-t,0) & 0<t\leq a \\ \\ \pi(0,t-a) & 0<a\leq t \end{cases} \tag{11}$$

Equation (9) sets SARS-CoV-2 incidence proportional to the fraction of the population that is susceptible and the infected population-weighted age-of-infection adjusted transmission intensity. Note how this generalizes the usual "βSI" term in the well known Susceptible-Infectious-Recovered (SIR) or Suceptible-Exposed-Infectious-Recovered (SEIR) models (Anderson and May 1991); in such models, a susceptible individual's transmission risk is directly proportional to the number of infectious persons, whereas in the scratch model a susceptible individual's transmission risk is given by the infected population-weighted age-of-infection transmission intensity $\int_0^\infty \lambda(a)\pi(a,t)da$.

Equation (10) says that susceptibles deplete with the incidence of infection. Equation (11) shows that once a person is infected, the infection just travels in an age-of-infection wave. If a person's age of infection is greater than the elapsed time in the model ($a > t$), that person must have been among those already infected at time 0. If a person's age of infection is less than the time in the model ($a < t$), that person was infected at time $t-a$. Note that there is no explicit probability distribution for the duration of infectiousness as in the usual SIR or SEIR models (Anderson and May 1991). That is because the distribution of infectiousness is implicitly accounted for in the definition of $\lambda(a)$.

The total fraction of the population infected over the life of the outbreak, that is, the *final size* ϕ, is given by

$$\phi = \int_0^\infty \pi(0,t)dt. \tag{12}$$

This can be directly computed via the following argument: imagine an uninfected person at the beginning of the outbreak ($t = 0$). The instantaneous infection hazard faced at time t equals $\int_0^\infty \lambda(a)\pi(a,t)da$ as stated earlier, and thus the integrated hazard faced over the life of the outbreak, Λ, is given by

$$\Lambda = \int_{t=0}^\infty \int_{a=0}^\infty \lambda(a)\pi(a,t)dadt. \tag{13}$$

The probability an initially uninfected person becomes infected over the duration of the outbreak, which is the same as the final size ϕ, is given by

$$\phi = 1 - e^{-\Lambda} \tag{14}$$

as follows from the presumed Poisson infection process. Substituting equation (11) into equation (13) and changing the order of integration yields

$$\begin{aligned} \Lambda &= \int_{a=0}^{\infty} \lambda(a) \left[\int_{t=0}^{a} \pi(a-t,0)dt + \int_{t=a}^{\infty} \pi(0,t-a)dt \right] da \\ &\approx \int_{a=0}^{\infty} \lambda(a)da \int_{u=0}^{\infty} \pi(0,u)du = R_0\phi \end{aligned} \tag{15}$$

where the first integral in the bracket involving those initially infected is negligible compared to the second term accounting for transmission over the entire course of the epidemic, which is easily recognized as equal to ϕ by changing the variable of integration from t to $u = t - a$, and R_0 is recognized from equation (8). Substituting equation (15) into equation (14) yields the elegant formula

$$\phi = 1 - e^{-R_0\phi} \tag{16}$$

which can be implemented numerically by always taking ϕ as the larger root of this equation (assuming $R_0 > 1$); this is a well-known result applicable so SIR, SEIR, and other mass action models where susceptibles experience common instantaneous transmission risk over the duration of an outbreak (Anderson and May 1991). Applying equation (16) using the previously noted value of $R_0 = 2.3$ yields a final size of 0.86, suggesting that in the absence of any intervention programs or behavioral changes, the worst case (i.e. unmitigated) fraction of the population that would be infected over the duration of the outbreak equals 86%, consistent with what was reported by Ferguson et al (2020).

I implemented the scratch model in Microsoft Excel by discretizing both time and age-of-infection in 1 day increments, noting that

$$\int_{a-1/2}^{a+1/2} \lambda(u)du \approx \lambda(a),\ 0 < a \leq 30 \tag{17}$$

for the transmission intensity function estimated in Kaplan (2020b), and assumed that $\lambda(a) = 0$ for $a > 30$ (again consistent with Kaplan (2020b)), which means one need not keep track of infection age beyond 30 days.

The actual computational formulas used in Excel are (following equations $(9) - (11)$ above with initial conditions $\pi(a, 0)$ for $a = 1, 2, ..., 30$ and $s(0) = 1$):

$$\pi(0, t) = s(t) \sum_{a=1}^{30} \lambda(a)\pi(a, t),\ t = 0, 1, 2, ... \tag{18}$$

$$s(t + 1) = s(t) - \pi(0, t),\ t = 0, 1, 2, ... \tag{19}$$

$$\pi(a + 1, t + 1) = \pi(a, t),\ \ a = 0, 1, 2, ..., 29;\ \ t = 0, 1, 2, ... \tag{20}$$

4.2 Accounting For Interventions

There are three ways to account for interventions in this model, depending upon if the interventions act on infected persons (e.g. isolation), susceptibles (e.g. lockdowns), or both (e.g. social distancing). Here is how they work.

4.2.1 Interventions Targeting Infected Persons: Isolation

Following Kaplan (2020b), imagine isolating infected persons at some random time during their infectious period. The basic idea is to define a new transmission function $\lambda_I(a) = \theta(a)\lambda(a)$ where $\theta(a)$ is the probability that transmission is *not* blocked by isolation at age-of-infection a. Then substitute $\lambda_I(a)$ for $\lambda(a)$ in the scratch model.

For examples, if there are no interventions at all and the epidemic just runs its course, then $\theta(a) = 1$ and the original scratch model results. Suppose infected persons are detected at time T into their course of infection and isolated indefinitely (Kaplan 2020b considers examples where T has the same distribution as the SARS-CoV-2 incubation time, twice the incubation time, and distributions reflecting contact tracing). Then

$$\theta(a) = \Pr\{T > a\} \tag{21}$$

since if $T \leq a$ there would be no transmission at age-of-infection a due to isolation before a. This formulation enables many different approaches to detecting and isolating infected persons.

4.2.2 Interventions Targeting Susceptibles: Lockdown and Quarantine

Suppose that at time t, a fraction $\rho(t)$ of susceptibles are "locked down" and removed from the population for duration τ after which restrictions are lifted. This effect can be incorporated by modifying equation (10) as

$$\frac{ds(t)}{dt} = -\pi(0,t) - \rho(t)s(t) + \rho(t-\tau)s(t-\tau). \tag{22}$$

As an example, consider a 4 week lockdown with 90% compliance starting at time ℓ (for ℓockdown). Setting $\rho(\ell) = 0.9$, $\rho(t) = 0$ for $t \neq \ell$, and $\tau = 28$ captures this policy, since at time $t = \ell$, 90% of suceptibles are removed from the population, and at time $t = \ell + \tau$, $\rho(\ell + \tau - \tau)s(\ell + \tau - \tau) = 0.9s(\ell)$ susceptibles are returned to the population (exactly the same number "locked down" at time ℓ) . Typically lockdowns (or stay-at-home orders) have been announced on a given starting date for a given duration, though the framework above is flexible, allowing lockdown implementation to be phased in or out. As another example, suppose that each detected infection leads via contact tracing to discovery of an average of 5 susceptibles who are quarantined for two weeks, but only 5% of infections are detected. Setting $\rho(t) = 0.05 \times \pi(0,t) \times 5/s(t)$ and $\tau = 14$ captures the policy, as $\rho(t)s(t) = 0.05 \times \pi(0,t) \times 5$ which is five susceptibles quarantined for each detected.

4.2.3 Social Distancing

Social distancing effects both circulating infecteds and uninfecteds. It reduces transmission opportunities. Letting δ be the fraction of transmission opportunities that are *not* distanced away, the incidence term is modified to

$$\pi_\delta(0,t) = \delta \times \pi(0,t). \tag{23}$$

This captures the effect of establishing distance and reducing transmission opportunities due to social distancing.

4.3 University Timing Scenarios

When the request for modeling to consider the prospects for commencement and summer activities was made (March 19), stay-at-home and social distancing restrictions had not yet been implemented though schools throughout the state had already closed. Starting with a February 19 initial condition of $\pi(a,0) = 1/300,000$ for $a = 1, 2, ..., 30$ (which yields a conservative starting prevalence of 1 in 10,000), I generated a graph showing the total prevalence over age-of-infection under several different scenarios (Figure 1). The isolation scenarios presume detection of the stated percentages of

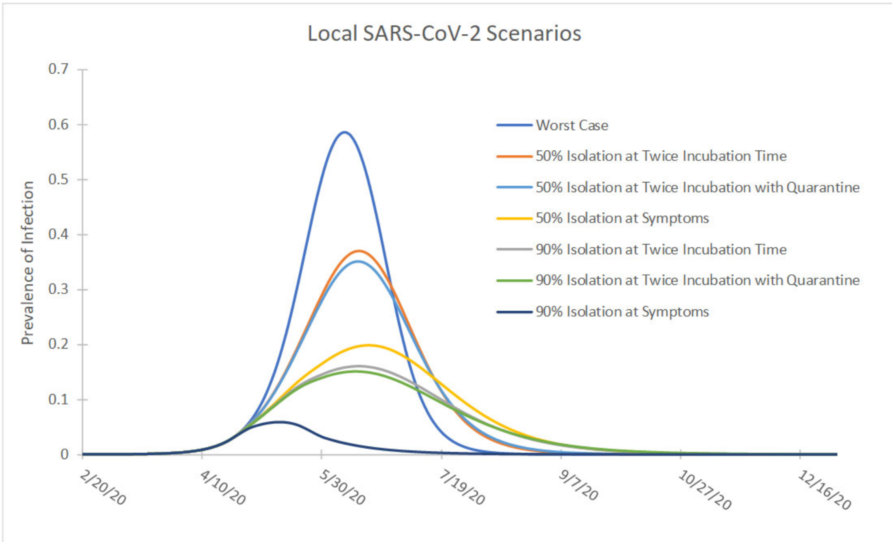

Figure 1

infected persons (50% or 90%) either at the incubation time (symptoms) or twice that time to account for delay in recognizing symptoms and seeking care. Quarantine presumed an average of five contact exposures located per detected infection and was implemented as described earlier. The total fractions of the population infected over these modeled outbreaks range from 85% in the worst case down to only 10% in the highly optimistic scenario where 90% of infected persons are isolated at the time of symptoms (a scenario that would require the ability for anyone with the slightest symptoms to receive a test, which is clearly not the case at present, Kaplan and Forman (2020)). These results were qualitatively very similar to results produced by Ferguson et al (2020) just days earlier using different methods (see Figures 1 and 2 in that paper).

The implication of these scenarios for the university was not the differing levels of infection over time, but rather the outbreak durations. All scenarios revealed the complete infeasibility of holding graduation ceremonies in May, while all but the most optimistic ruled out on-campus summer activities as well. Also, even though the fraction of the population infected would wane considerably by September, that fraction would still be much higher than the estimated level of infection at the time these scenarios were generated the weekend of March 21-22. No decision has been made regarding whether classes will resume on campus this fall as the situation has changed somewhat as will be discussed next, but the scenarios and attendant modeling, created from scratch in just a few days, had a major influence on university planning at the time they were submitted.

Anticipating that lockdown-like restrictions (e.g. stay-at-home orders) would soon be announced (as they were on March 23), I also produced lockdown scenarios, one of which presuming 95% compliance to a one month lockdown is shown in Figure 2. This analysis showed that the key qualitative feature of lockdowns absent other prevention activities is to buy time by shifting the outbreak forward. Such lockdowns only serve to reset the initial condition by preserving and then later releasing an almost completely susceptible population when infectious individuals remain undetected. This conclusion has also been reached by other modelers (Ferguson et al 2020, Kissler et al 2020) and reinforces the idea that lockdowns only serve to buy time unless consequential preventive actions are taken.

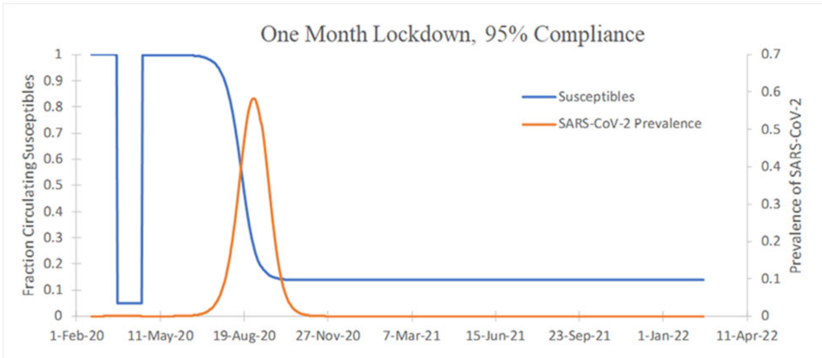

Figure 2

5 Tracking the Outbreak and Intervention Scenarios

On March 25 I received a request from the State of Connecticut's Chief Operating Officer asking for help modeling COVID-19 hospitalization bed needs in the state. The Connecticut Department of Public Health by now was producing daily updates reporting diagnosed cases, hospital admissions, deaths, and tests administered, while the Connecticut Hospital Association (CHA) was providing daily reports documenting COVID-19 bed occupancy. I was asked to help make sense of these trends, especially in response to recently tightened public health restrictions culminating with Governor Lamont's Stay Safe, Stay Home initiative that went into effect on March 23 (Lamont 2020b). Though less stringent than a complete lockdown, these measures have severely reduced person-to-person interactions with the goal of sharply reducing SARS-CoV-2 transmission.

The models described thus far had not made use of ongoing state-level data collection, but now with regular access to daily data, it was possible to track what was going on. I decided to focus on total hospitalizations for this purpose. The first purely statistical approach was to fit the growth in occupied beds reported by the CHA over two weeks starting from March 18 with an exponential curve and simply track daily COVID-19 bed occupancy relative to this trend as shown in Figure 3. Charts like Figure 3 have been provided to state officials on a daily basis. Had COVID-19 patients continued to occupy beds at the same rate over time, their occupancy needs would have exhausted Connecticut's limit of approximately 7,000 total hospital beds by April 10. Happily this did not occur, and the growth in hospitalized COVID-19 patients has slowed considerably in part to the restrictions enacted on March 23.

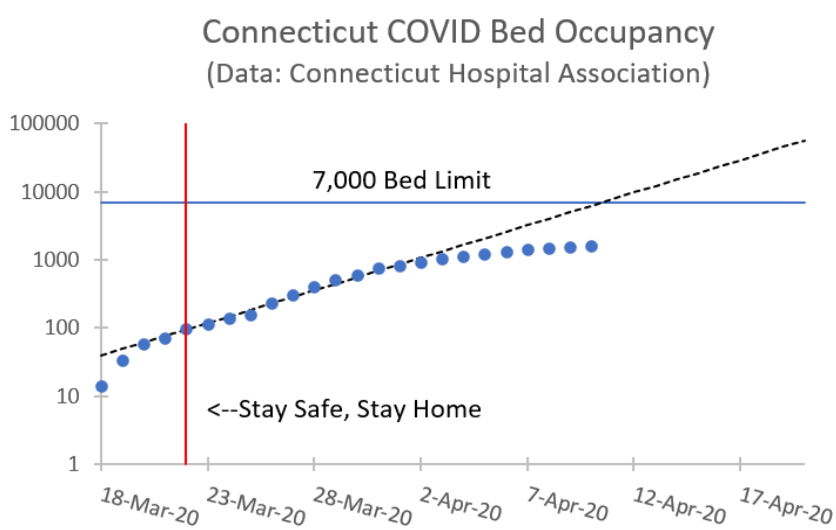

Figure 3: Observed (dots) and tracked (dashed line) bed occupancy.

The second modeling approach was to use the scratch model described previously to see how observed hospital data compared to various scenarios generated. To do this required turning infections into COVID-19 hospital bed occupancy. I presumed that early in the Connecticut outbreak, the incidence of infections would be growing exponentially as $k_0 e^{rt}$ for initial condition and growth rate k and r during the first few weeks of the outbreak, and thus the COVID-19 hospital admission rate, $h(t)$, should follow

$$h(t) = \int_0^t k e^{ru} f_L(t-u) du \tag{24}$$

where $f_L(t)$ is the probability density for the lag (L) from infection to hospitalization for those infections that result in hospitalization. Since not all infections lead to hospital admission (thankfully most do not), the parameter k accounts for the product of the initial incidence k_0 and the fraction of infections that lead to hospitalization. Connecticut's first COVID-19 hospitalization occurred on March 7, and I only had data through March 25, so letting March 6 represent clock time 0 and letting the lag L follow an exponential distribution with mean $1/\mu$ yields a mean hospitalization rate of

$$h(t) = \frac{k\mu}{r+\mu}(e^{rt} - e^{-\mu t}). \tag{25}$$

I presumed that observed hospitalizations would follow a Poisson distribution with mean as shown above, and estimated the parameters k and r conditional on mean delays ($1/\mu$) ranging from 7 through 11 days via maximum likelihood, which resulted in consistent estimates of the exponential growth rate $r \approx 0.14$ (the actual values ranged from 0.136 at an eleven day lag to 0.143 at a seven day lag; k was estimated to equal 5.7). I selected an average nine day delay from infection to hospitalization, roughly corresponding to the sum of the mean incubation time (5.2 days) plus another four days of delay from symptoms to hospital admisison, similar to values reported elsewhere (Li, Guan, Wu et al, 2020). Figure 4 shows the daily hospitalization data and estimated hospitalization curve associated with the presumed nine day lag.

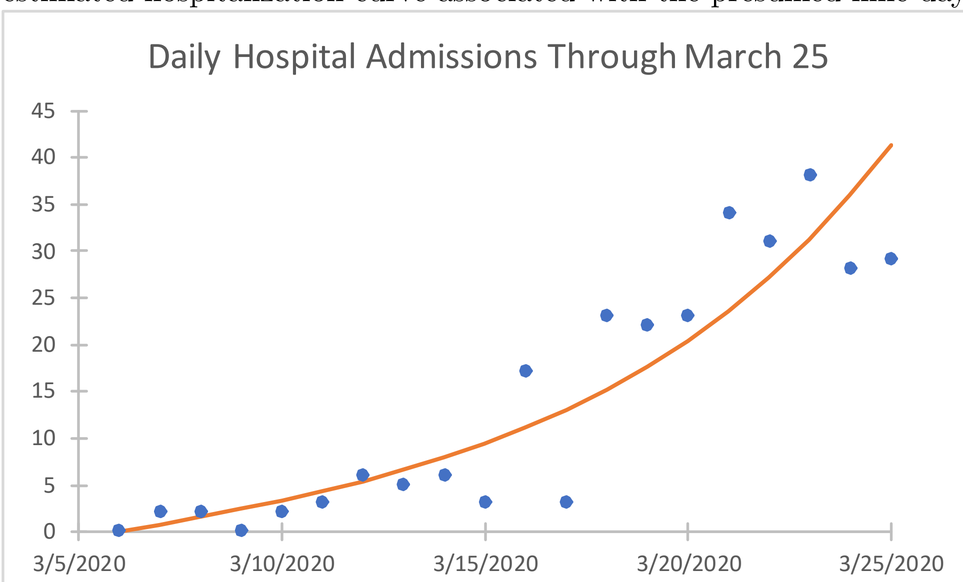

Figure 4

This newly estimated growth rate of 14% per day was higher than the 10% rate I had used in the scratch model until this point, so I modified the transmission function $\lambda(a)$ to account for this new rate following Kaplan (2020b).

Next, I noted that in the state health department data, the ratio of hospital admissions to diagnosed cases was approximately 15% even as cases continued to rise during the last week of March and the first week of April. The main CHA analyst informed me that COVID-19 patient length of stay over all admitted patients was averaging about a week. Since the fraction of infections diagnosed as COVID-19 cases depends upon testing limitations, I estimated what fraction of infections generated by the updated scratch model would have to be reported as cases to match the observed bed occupancies at the end of March as reported by the CHA, presuming that 15% of such cases ended up hospitalized nine days after infection. Doing so suggested that roughly 9.3% of infections were being recognized as cases, implying that there are more than ten times as many infections as diagnosed cases, a fraction similar to what has been reported elsewhere (Li et al 2020).

With this information it was possible to "splay" the SARS-CoV-2 incidence results from the scratch model into COVID-19 bed occupancy by first turning 9.3% of modeled infections into cases, then turning 15% of those cases into hospital admissions nine days later, and third maintaining a running sum of hospital admissions over the previous seven days to account for length of stay. Replacing these deterministic lags with geometric lags of the same mean durations did not appreciably change the results.

Figure 5 shows that an unchecked COVID-19 outbreak truly would wreak havoc on Connecticut hospitals, with almost twice as many patients needing hospital care as total beds at the modeled April 30 peak. Luckily, the actual situation is not nearly as bad, as the actual bed census data show. As in

Figure 3, actual COVID-19 hospital bed occupancy is growing more slowly than an unchecked outbreak. Instead, it is growing in a manner consistent with 35% social distancing (that is, $\delta = 0.65$); should this pattern continue, COVID-19 hospital occupancy will peak during the second week in May but at a level that is within the number of staffed beds, contrary to reports suggesting Connecticut would have the largest relative shortage of hospital beds in the nation due to the COVID-19 crisis (Hao 2020). However, COVID-19 bed occupancy is growing more quickly than a third reference scenario shown in Figure 5 which was meant to represent the Stay Safe, Stay Home initiative. In this scenario, a 70% compliant lockdown lasting until the end of April was implemented in the scratch model by setting $\rho(\text{March } 23) = 0.7$ and setting $\tau = 38$ days, while on top of this social distancing was implemented with 30% effectiveness (so $\delta = 0.7$). The social distancing choice was based on the most conservative (highest) value of δ inferred from studying Google Mobility reports for Connecticut (https://www.gstatic.com/covid19/mobility/2020-04-05_US_Connecticut_Mobility_Report_en.pdf), while the fraction of the population staying at home is merely suggestive. The point of this is twofold: first, as argued previously, a lockdown with no concerted effort to detect and isolate infected persons simply delays an outbreak until the end of the lockdown. Second, however, the observed COVID-19 hospitalizations are not following the pattern one would expect from a lockdown, but are consistent with social distancing, which calls into question just how much Connecticut residents really are staying at home.

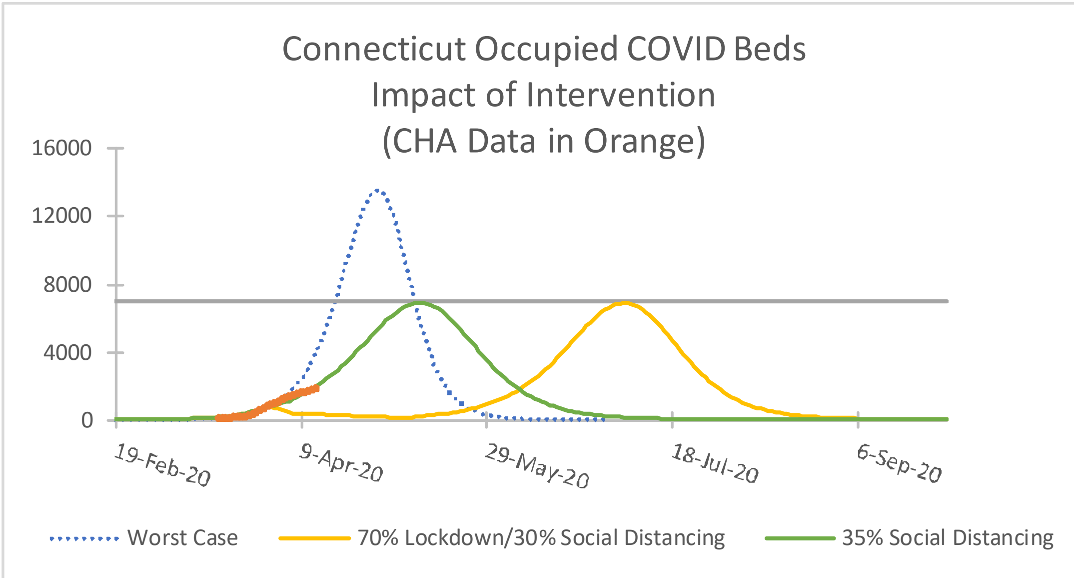

Figure 5

The Governor has extended Stay Safe, Stay Home at least until May 20, creating additional time to shore up needed hospital supplies and equipment. More importantly, as the supply of diagnostic and serologic tests is about to increase greatly, Connecticut public health officials are planning for an aggressive community screening effort to detect and isolate infected persons. The benefits of doing so are obvious from the optimistic scenario shown in Figure 1, and are consistent with the findings of Kaplan (2020b) showing that the rapid detection and isolation of infected persons could reduce the reproductive number R_0 well below unity. Unlike lockdown interventions or social distancing, both of which prevent transmission only while such measures are in force, aggressive detection and isolation of infected persons is not a band aid and could actually end this epidemic.

6 Summary

This article has reported the details of an intensive six week "scratch modeling" effort in support of university, hospital and state government public health decisions. With the national strategy for gaining control of the COVID-19 epidemic very much in doubt, the importance of local decision-making increases given the severity of this outbreak. Institutions, cities and states need help to do this. This article has documented how one operations researcher has tried to help university, hospital and state officials make public health decisions by providing insights from mathematical models. There are many, many researchers in our community who could contribute in similar ways.

The YCC at Yale continues to provide advice, with much of the attention now focusing on how to resume certain university operations, albeit in altered modes. Should the university re-open to students this fall, will courses continue to be offered online, or might the fall semester be delayed several months to be followed by a full academic year on a compressed schedule on campus? Critical research laboratories require scientists at the bench; how can their work processes be re-engineered to maximize infection control? Facilities physically designed to foster collaboration now find themselves in need of repurposing to isolate individual labs from each other, while very careful tracking of who is where at all times will be necessary in case local transmission occurs.

Trying to answer whether it is safe to re-start university activities requires having some ability to observe current levels of infection. Aggressive community screening has been proposed at the city and state level to detect and isolate infected persons, and should such operations take off, the resulting data would provide invaluable information in helping decide how to proceed. The need for immediate testing and isolation on campus is also under discussion.

A group of Yale researchers has launched a program sampling sewage from the New Haven treatment plant while also sampling effluent from the Yale New Haven Hospital's COVID-19 and children's wards as well as New Haven's Union Station, with the hope that testing such human output for SARS-CoV-2 will provide a signal of both the extent of infection locally as well as an early sign indicating the waning of the outbreak.

There is a need for more creative thinking when it comes to re-starting the economy as argued recently by Caulkins (2020). Operations researchers and management scientists can contribute to thinking about how to re-engineer critical economic activities, permitting them to re-open and operate in "safe mode" while respecting the need for infection control. While most of the discussion on the economy has occurred at the federal level, there is ample role for local decisions within individual companies and towns to safely bring back economic prosperity. Operations researchers and management scientists can provide important support in these crucial endeavors.

Acknowledgements

I wish to thank all of the members of the Yale University COVID-19 Advisory Committee chaired by Dr. Paul Genecin, Director of Yale Health. I also thank Edieal Pinker (Yale School of Management), Josh Geballe (Chief Operating Officer, State of Connecticut), Matt Cartter (State Epidemiologist, Connecticut Department of Public Health), Sydney Jones (Centers for Disease Control and Prevention, Epidemic Intelligence Service officer for Connecticut), John Brady and Paul Kidwell (Connecticut Hospital Association), Alex Hartzman (Dobson DaVanzo Health Economics Consulting), and David Greenslade, Matt Comerford, and Thomas Balcezak (Yale New Haven Hospital). I also thank the two reviewers for providing very helpful comments while sacrificing precious personal time to provide prompt reviews that greatly improved the paper. Any errors are my own.

References

[1] Anderson RM, May RM (1991) Infectious diseases of humans: Dynamics and control. Oxford University Press, Oxford

[2] Britton T, Tomba GS (2019) Estimation in emerging epidemics: Biases and remedies. J R Soc Interface 16:20180670. https://doi.org/10.1098/rsif.2018.0670 (accessed April 15, 2020)

[3] Caulkins JP (2020) With a Little Ingenuity, We Can Reopen Much of the Economy Right Now. Politico Magazine, https://tinyurl.com/vm2t87r (accessed April 12, 2020)

[4] Burke RM, Midgley CM, Dratch A et al (2020) Active Monitoring of Persons Exposed to Patients with Confirmed COVID-19 — United States, January–February 2020. Morbidity and Mortality Weekly Report (MMWR) 69(9);245–246, https://tinyurl.com/yxxtxwgg (accessed April 11, 2020)

[5] Ferguson N, Laydon D, Gemma N-G et al (2020). Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Imperial College COVID-19 Response Team, March 16, 2020, https://tinyurl.com/tcdy42y (accessed April 11, 2020)

[6] Hao K (2020) Here are the states that will have the worst hospital bed shortages. Technology Review, April 7, 2020, https://tinyurl.com/wpzmny5 (accessed April 13, 2020)

[7] Hill M, Villeneuve M (2020) Man hospitalized with New York's 2nd case of COVID-19 virus. Associated Press https://tinyurl.com/surombg (accessed April 11, 2020).

[8] Hunter A, Johnson L, Coustasse A (2014) Reduction of intensive care unit length of stay. The Health Care Manager 33(2):128-135 https://tinyurl.com/w5ta8vb (accessed April 12, 2020)

[9] Kaplan EH (2020a) Perspectives on Policy Modeling. Manufacturing and Service Operations Management, https://doi.org/10.1287/msom.2019.0849 (accessed April 12, 2020)

[10] Kaplan EH (2020b) Containing 2019-nCoV (Wuhan) coronavirus. Health Care Manag Sci https://doi.org/10.1007/s10729-020-09504-6 (accessed April 11, 2020)

[11] Kaplan EH and Forman HP (2020) End the coronavirus outbreak with aggressive screening to find and isolate infected people. USA Today, April 14, 2020 https://tinyurl.com/sfoug3a (accessed April 15, 2020)

[12] Kissler S, Tedijanto C, Lipsitch M, Grad YH (2020). Social distancing strategies for curbing the COVID-19 epidemic. Harvard School of Public Health, March 2020 http://nrs.harvard.edu/urn-3:HUL.InstRepos:42638988 (accessed April 11, 2020)

[13] Lamont N (2020a) Governor Lamont Announces First Positive Case of Novel Coronavirus Involving a Connecticut Resident, State of Connecticut Press Release, https://tinyurl.com/u2wctot (accessed April 11, 2020)

[14] Lamont N (2020b) Stay Safe, Stay Home: Latest Guidance. https://portal.ct.gov/Coronavirus (accessed April 11, 2020)

[15] Li R, Pei S, Chen B et al (2020) Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV2). Science https://tinyurl.com/tpx4css (accessed April 11, 2020)

[16] Li Q, Guan X, Wu P et al (2020) Early transmission dynamics in Wuhan, China, of novel coronavirus-infected pneumonia. N Engl J Med. https://doi.org/10.1056/NEJMoa2001316 (accessed April 15, 2020)

[17] Ravizza S (2020) Coronavirus: "Stiamo creando terapie intensive anche nei corrido" (Coronavirus: We are also creating intensive therapies in the corridor), Corriere Della Sera, March 7, 2020, https://tinyurl.com/qmhg4fq (accessed April 12, 2020)