

of realistic problem-solving capability, he says. At Stanford, where many of his colleagues are choosing the open-book route, exams turn into a learning experience “with far less artificial hoop-jumping and guessing at instructors’ idiosyncrasies,” he says.

Many universities have adjusted their policies on course withdrawal and grades. “A lot of students come in well-prepared,” Stelzer says. “But there is a population for whom that is not true, and that tends to be the same people who don’t have access to the internet and have a harder situation at home—sharing computers, occupying cramped spaces, or taking on extra family responsibilities.” To accommodate such inequities, and to be sympathetic, he says, many instructors have loosened deadlines for labs and quizzes. And most universities have extended the deadlines for dropping classes to just before—or even after—final exams.

Many universities have ditched

grades this term in favor of pass/fail. A few, such as Georgia Tech, have retained grades, despite student complaints. And many institutions are giving students the choice of either a grade or a pass/fail. At UBC, the science dean issued a rare decree requiring faculty to calculate grades with two different weightings for the final exam—30% and 5%; students will receive the better grade. Or they can opt for pass/fail. In late March, the American Physical Society sent a letter to department heads urging their graduate admissions committees to treat this term’s grades “holistically.”

Overall, the wholesale transition to remote teaching created a mad scramble and a lot of improvisation. But many faculty say they’ve learned things they’ll take forward for future online teaching and for when in-person classes resume. Jonathan Wurtele of the University of California, Berkeley, notes that his campus occasionally closes due to smoke from nearby fires. “We will put the

knowledge of remote teaching to use in the future,” he says. Similarly, Karen Daniels of North Carolina State University says she’d be comfortable teaching remotely for a day or so if she leaves town to attend a conference. But, she says, “even if we have found replacements for all the parts of a normal face-to-face class, it’s not the same. We are not delivering what we need to.”

Online office hours, for which students choose a time slot for a video conference, could continue to work well especially for commuter students, according to several instructors. Many professors found that students were good at helping each other in the text chat boxes in video-conferencing software, and they hope to incorporate that type of help in their in-person courses. Andrew Loveridge of UT Austin notes that with the transition to remote teaching, “we are forced to think about every part of our courses. Nothing will survive on its own inertia.”

Toni Feder

COVID-19 pandemic modeling is fraught with uncertainties

Policymakers face a plethora of predictions on how the disease will proceed and when it might resurge.

A self-described optimist, Pinar Keskinocak doesn’t like to be the bearer of bad news. But the model she co-developed at Georgia Tech of the COVID-19 pandemic in that state paints a “really bleak” picture of what lies ahead when physical distancing slowly erodes after shelter-in-place and stay-at-home orders end.

The model, which forecasts the outbreak in Georgia at the census tract level—county subdivisions that average 4000 inhabitants—shows that even if lockdowns had been extended through mid-May instead of being lifted 1 May, the rate of new infections would come roaring back once people returned to their daily routines.

Georgia was one of the first states to end shelter-in-place orders and permit some businesses to reopen. Although continued adherence to social distancing guidelines will tamp down the state’s peak numbers of new infections, even strict compliance—including the voluntary quarantining of all persons in households where only one member is infected—

won’t prevent the outbreak from surging to levels far higher than those yet experienced. The real peak of new cases in Georgia, says the model, is predicted to come in June or July (see graphs on pages 26 and 27) and potentially overwhelm health-care facilities in some parts of the state.

The Georgia Tech model’s findings, which were shared with state government officials—Keskinocak won’t say exactly whom—before governor Brian Kemp’s decision to lift stay-at-home restrictions, presented a different portrait of the pandemic from the widely reported modeling results coming from the University of Washington’s Institute for Health Metrics and Evaluation. Until 26 April, after which it was significantly revamped, the IHME model had predicted that the peak daily death toll from COVID-19 in Georgia had already passed, even before the mid-April zenith it had forecast for the nationwide death rate. On 3 May, a new, hybrid version of the IHME model was projecting that daily deaths in Georgia would peak on 30 May, well after the forecasted 1 May peak in daily US deaths.

Another model, developed at Los Alamos National Laboratory (LANL), reported with 96% confidence as of 3 May that the daily rate of new cases in Georgia has peaked. Although LANL’s model doesn’t explicitly include the effects of interventions such as sheltering in place and social distancing, it assumes that some social distancing measures will continue through the forecast period. LANL modeler Dave Osthus says the model won’t be adjusted to account for the ending of lockdowns because the extent to which people will actually change their behavior is unknown.

Many other models forecast new infections and deaths at the international, national, and state levels. The Centers for Disease Control and Prevention (CDC) regularly compiles on its website the forecasts of nine COVID-19 models, including LANL’s. Some show the rate of new deaths slowing nationally; others show daily fatality numbers remaining flat. Most of the included models assume the continuance of the social distancing policies that were in place on the date of model calibration. A few make no such assumptions.

The unknowns about the disease and its transmission produce large error bars

in predictions. The LANL model, for example, estimates that by 10 June Georgia will have 40 000 total confirmed cases of COVID-19. But that number lies between a best-case scenario of 33 000 cases and a worst case of 64 000. The lab's earlier predictions for Georgia's COVID-19 case and death totals have been closer to the best-case scenario. Osthus says the model's skill varies widely from state to state, depending in large part on the quality of data used as input. It has performed well for New Mexico, for example, but poorly for Connecticut. The model has a built-in assumption that the growth rates of new cases and deaths will trend toward zero, but it is constructed so that it can adapt to new and unexpected input, he says.


Statistics and mechanics

At the most basic level, models are either statistical or mechanistic attempts to simulate the disease transmission cycle. Until early May, the IHME model was purely statistical; no assumptions were made about how the disease spreads. Instead, it assumed that the epidemic will follow a trend that is based on experiences in China, South Korea, Italy, and other places.

On 4 May, after many states had begun reopening businesses, IHME updated its model to include a disease-transmission component, which took into consideration the relaxation of states' stay-at-home orders. The hybrid model produced more somber conclusions, including a near doubling of the US death toll to 135 000 by 4 August, up from the 72 000 deaths the old model had predicted. As of 13 May, US fatalities from the new coronavirus stood at nearly 83 000, according to the authoritative COVID-19 tracker maintained by Johns Hopkins University.

A statistical model can be appealing, says Harvard University research fellow Alison Hill, "because it fits the data well. It describes what really happened, and that maybe it's reasonable to expect that the trend will keep going in a particular direction." Such models have been used to compare the anticipated number of COVID-19 cases to the regional capacity of hospitals.

Purely statistical models, however, have plenty of limitations. "There's no natural law that says epidemics have to follow that curve," says Hill. Statistical models can't predict the outcome of relaxing restrictions and a resurgence of the disease, nor can they be adjusted for varying



PROGRESSION OF COVID-19 INFECTIONS in Georgia, as simulated by researchers at Georgia Tech. Shown are new daily infections for three scenarios. The simulations begin on 18 February with increasing voluntary quarantining by infected individuals and households, followed by school closures on 16 March, sheltering in place beginning 3 April, and then different levels of voluntary quarantine. The actual rate of new infections (red line) is multiplied by 8 to reflect underreporting. Other simulations (not shown) varied the shelter-in-place duration. (Adapted from P. Keskinocak et al., medRxiv preprint, doi:10.1101/2020.04.29.20084764.)

degrees of adherence to social distancing measures. There's no clear way to adapt a statistical model to accommodate multiple disease peaks or to consider what a combination of widespread testing and contact tracing could do to alter the pandemic's course, she says.


Mechanistic modeling, on the other hand, seeks to emulate the disease transmission process. MIT's model, developed by a team led by business professor Dimitris Bertsimas, is one example of a commonly used mechanistic approach known as SEIR, named for the successive phases—susceptible, exposed, infected, and resistant—through which it tracks the progression of the disease. The IHME's sharp increase in mortality predictions reflected the addition of a new SEIR component to its model's statistical underpinnings.

Mechanistic models are akin to physical models: Both translate laws and processes into mathematical form. Philip Stamp is a condensed-matter theorist at the University of British Columbia. He was called in by Brett Finlay, a UBC microbiologist who was first to map the genome of the 2003 severe acute respiratory syndrome (SARS) coronavirus, to formulate equations for a COVID-19 forecasting model being developed for the provincial government. "This is the sort of thing that theoretical physicists do, whether they are

working in statistical mechanics or theoretical astrophysics," Stamp says. "Whereas for most people in medical science, it's very far away from their expertise."

Stamp says that many disease models initially failed to account for the large numbers of asymptomatic carriers of the new coronavirus. "A lot of the assumptions that went into those models were obviously wrong and could be seen to be wrong way back at the beginning of February," he asserts. From the way the disease progressed, initially in China and then elsewhere, "it was absolutely obvious that most carriers were asymptomatic. There was no other way to explain the way this evolved."

The Georgia Tech model is mechanistic and agent-based. Each of the model's 1 million agents is assigned an age, census tract, household, and peer group based on the state's demographic profile. At the start of a simulation, the infection is introduced randomly to agents according to the distribution of Georgia's confirmed cases. For infected agents, the disease progresses through its stages with predefined probabilities, and the infection spreads to healthy individuals at rates that depend on each agent's social contacts at home, in peer groups, and in the broader community. The model allows those interaction rates to be adjusted according to the degree of social distancing the agent follows.



CUMULATIVE COVID-19 DEATHS in Georgia, as simulated by researchers at Georgia Tech. The simulations begin on 18 February with increasing voluntary quarantining by infected individuals and households, followed by school closures on 16 March, sheltering in place beginning 3 April, and then different levels of voluntary quarantine. The red line indicates confirmed numbers of deaths through 20 April. (Adapted from P. Keskinocak et al., medRxiv preprint, doi:10.1101/2020.04.29.20084764.)

Agent-based models are particularly useful for predicting the geographical course of the disease. Julie Swann, an engineering professor at North Carolina State University, is adapting for that state an agent-based model that she helped build at Georgia Tech for forecasting the 2009 H1N1 influenza pandemic in Georgia. The effects of social distancing, voluntary quarantining, and school closures are being built in. Data on household sizes, ages, and commuting flows are taken from US Census information.

All models make assumptions based on many unknowns about the novel coronavirus. Among the most important is the number of underreported infections. "Every respectable model that I've seen is taking into account the fact that there are plenty of people out there who are infected and haven't been reported," says Hill. The Georgia Tech model assumes that the real number of infections is six to eight times that of confirmed cases. Keskinocak says some reports suggest that actual cases could be 10 times the confirmed number.

Another important unknown is whether recovered individuals will be immune to reinfection, and for how long that immunity lasts. In the Georgia Tech model, permanent immunity is assumed because there isn't definitive evidence to the contrary, Keskinocak says.

Most models, including Georgia Tech's, address what happens several

weeks ahead. Projecting the track of COVID-19 in the months and years beyond the initial wave of the pandemic was the objective of researchers led by Harvard's Marc Lipsitch and Yonatan Grad. Their model draws from observed experience with coronaviruses that cause the common cold and tries to estimate some of the epidemiological parameters for COVID-19 that will determine its long-term dynamics. Depending on factors such as the duration of immunity in recovered patients, the degree to which warmer weather causes COVID-19 to recede, and the extent to which immunity to cold viruses might extend to COVID-19, the model presents multiple scenarios for potential outbreaks through 2025. Absent effective vaccines or therapeutics, aggressive contact tracing, quarantining, and intermittent social distancing may need to be maintained into 2022, the authors conclude, at a substantial social and economic cost.

Decision-making tools

"All forecasts are wrong, because you are predicting the future," says Swann. "But if you know how to interpret them, you can come to understand what aspects to look at, and you can offer some real benefits for decision making." If used correctly, a model will indicate what to expect in terms of infections and deaths weeks after schools are reopened, she says. Policy-

makers will have to consider the modeling results in the context of other costs—unemployment, poverty, suicides, and so on—resulting from continuing stay-at-home orders. Such tradeoffs are made all the time by governors, legislators, mayors, and city council members, Swann notes.

"Models are absolutely crucial in the same way that economic models help in planning the economy," says Stamp. "But for the virus, the set of equations that tell you how this thing will evolve have not been there. [Virus modelers] have been using some simple equations that aren't working, and they're simply not an adequate description of what's out there."

With so many models to choose from, which should policymakers rely on to guide them? The CDC compiles an ensemble of the nine models that it tracks. It generally plots a middle course between the most dire and most hopeful predictions.

"It can be a good thing to use more than one model, because you don't want to be fully dependent on a set of assumptions without really challenging those assumptions," says Swann.

"This is just way too hard of a problem to think you are going to get it right, but it's also incredibly consequential," says Osthus. Given the model's varying success from state to state, LANL began including performance records in each state forecast. "Decision makers are looking at this model and trying to base meaningful decisions on it, and we think that they are entitled to this information," he says. "The user can look at our model performance and decide if they want to trust what it's saying."

One factor to look for in evaluating models is the training and experience of the people who make them, says Hill. "There are always people trying to jump in and give their two cents' worth on what they think will happen."

There's little expectation that a single, authoritative model will emerge. Hill says researchers will never agree on what assumptions should be given more weight relative to others. "People who make models and communicate them need to be more straightforward in explaining in an accessible way what assumptions they've made," she says. "We'll never get away from uncertainty and think that people will agree on what's certain and what's not."

David Kramer **PT**