

What models can and cannot tell us about COVID-19

Alexander F. Siegenfeld^{a,b,1}, Nassim N. Taleb^c, and Yaneer Bar-Yam^b

The coronavirus disease 2019 (COVID-19) pandemic, caused by the novel coronavirus severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2), has already claimed more than 470,000 deaths worldwide at the time of this writing (1) and is likely to claim many more. Models can help us determine how to stop the spread of the virus.

But it is important to distinguish between what models can and cannot predict. All models' assumptions fail to describe the details of most real-world systems. However, these systems may possess large-scale behaviors that do not depend on all these details (2). A simple model that correctly captures these large-scale behaviors but gets some details wrong is useful; a

complicated model that gets some details correct but mischaracterizes the large-scale behaviors is misleading at best. The accuracy and sophistication of a model's details matter only if the model's general assumptions correctly describe the real-world behaviors of interest.

Carefully delineating models' strengths and shortcomings will not only clarify how they can help but also temper expectations among policymakers and members of the public looking to understand the full impact of the virus in the weeks and months ahead. More important even than prediction is the ability of models to guide actions that can change this impact, including actions that can potentially drive the virus to extinction.



Chaotic systems such as pandemics are fundamentally unpredictable. A constructive role for science is to identify interventions—including social distancing, mask wearing, policies of isolation, and travel restrictions—that will help the number of active infections to decline exponentially. Image credit: Shutterstock/Travelerpix.

^aDepartment of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139; ^bNew England Complex Systems Institute, Cambridge, MA 02139; and ^cTandon School of Engineering, New York University, Brooklyn, NY 11201

The authors declare no competing interest.

Published under the [PNAS license](#).

Any opinions, findings, conclusions, or recommendations expressed in this work are those of the authors and have not been endorsed by the National Academy of Sciences.

¹To whom correspondence may be addressed. Email: asiegenf@mit.edu.




Fig. 1. A collection of geographic regions can exist in one of two phases with respect to COVID-19. In the stable phase the disease is extinguished over time, whereas in the unstable phase it grows exponentially only to be bounded by the population size as a whole. Which regime a collection of regions is in depends on 1) the size of local outbreaks, which in turn depends on (among other factors) how quickly regions take action if they experience another uncontained outbreak, and 2), the rate of travel-driven transmission between regions (4–6).

Model Capability

Understanding what models cannot predict is sometimes more important than understanding what they can. For example, in a chaotic system such as the weather, only very short-term predictions are accurate; small changes in the present can result in very large changes in the future. Likewise in the case of the pandemic's trajectory: Because the number of infections depends exponentially on the growth rate of the epidemic, small inaccuracies in the prediction of the growth rate will lead to large changes in the number of deaths after enough time. Furthermore, the growth or decay rate of the epidemic depends on the precise implementation details of interventions, and a very small change in the strength of interventions could be the difference between two hugely different outcomes: exponential growth versus exponential decay. Gaining an approximate understanding of the trajectory of the epidemic is important. But given the considerable uncertainty arising from underlying disease and social dynamics—not to mention the uncertainty over exactly how interventions will be implemented—detailed refinements to models often create a misleading sense of certainty and precision.

More generally, trying to pin down details in models is futile if any accuracy gained is swamped by uncertainty in the measurements or by inaccuracies in the core model assumptions. What's the point of refining a model by 10% if there is a 50% uncertainty stemming from other aspects or assumptions of the

model? What's the point of a sophisticated adjustment to a model if there is a relevant large-scale behavior of the modeled system that the model fails to capture altogether?

Models that attempt to capture a system's small-scale detailed behavior (e.g., ref 3) will inevitably include some details and leave out others. Depending on which details are included, such models may mischaracterize the system's large-scale behavior. When they do work, it is often because their specific assumptions are a special case of a simpler, more general model. Thus, sometimes it is not the complicated models but the deceptively simple ones that are most effective for understanding a system's large-scale behavior.

Phase Transition

For COVID-19, one large-scale behavior is an exponential increase in infections in the absence of intervention (unless the number of people infected is approaching saturation); the exact growth rate depends on the location and the precise details of disease transmission. Interventions may change this growth rate. And robust interventions, such as lockdowns, may result in exponential decay rather than exponential growth.

Another large-scale behavior is the fact that transmission is predominantly local, with travel creating the possibility of long-range spread. The number of infections does not change uniformly all over the world at once but rather more or less independently in each region. The probability that the disease is transmitted from one region to another depends on the number of infections in the first region and the travel rate from the first to the second among contagious individuals. There are many small-scale details to the disease transmission process, but the large-scale dynamics seem to be captured by the rate of increase or decrease within a region and the rates of transmission between regions (both of which may change over time as a result of interventions, saturation effects, or other variations in external conditions).

Depending on the dynamics of these parameters, a collection of regions can exist in one of two phases: a stable phase, in which the disease dynamics tend towards a stable fixed point of elimination (i.e., no infected regions within the collection), and an unstable phase, in which the number of infected regions grows until a saturation point is reached (see Fig. 1).^{*} For a collection of regions to be in the stable phase, it is not necessary for regions to be under constant lockdown after they have been cleared of the virus but rather only for each region to be ready to lockdown in the event that it experiences another uncontained outbreak (6). If

^{*}A lockdown within a single geographic region can itself be analyzed using Fig. 1, if each household is considered as a "region." In this case, the mean size of an outbreak would be the average number of individuals within a household expected to get COVID-19 if one individual in the household is infected, with the disease transmission between "regions" corresponding to the probability that an infected individual in one household has of infecting an individual in a different household. The primary purpose of a lockdown is to control this probability.

the virus is introduced or reintroduced into a collection of regions in the unstable regime, the number of infected regions will exponentially grow. But if the virus is introduced or reintroduced into a collection of regions in the stable regime, the collection of regions will return to its uninfected state.

Mismodeling

If the large-scale behaviors of a system are correctly described, specific details can be understood in terms of their effects on these behaviors. But if a model's assumptions do not yield the same general large-scale behaviors of the system being modeled, adding additional details to the model will serve only to create a false sense of confidence.

For example, models using continuous variables to represent fractions or probabilities of individuals being infected may predict that although a lockdown can produce an exponential decline in cases, the number of cases will inevitably rebound once the lockdown is lifted. However, the assumption of approximately continuous behavior breaks down for small numbers of infections.

We know that zero cases cannot grow back—and even a few may not grow back. If a total case number that is fractional appears in the final output of the model, human judgment can correct for the error (e.g., by interpreting a fraction of a case in the model as the virus having been eliminated in reality). But if these small numbers arise as intermediate values in the model, the model will mistakenly predict exponential growth once the lockdown is lifted, despite the fact that the model is no longer valid in this regime because there may in fact be zero cases.[†]

A rebound in infections after lockdown measures are lifted is a potential large-scale behavior of the system, but it is not inevitable (as predicted by continuous models). Rather, it depends on our actions: If interventions strong enough to create an exponential decay in the number of active infections are held in place for a sufficient amount of time, the virus will be eliminated.[‡]

Some might object that even if the fraction of the population infected becomes very small, if the size of the population being considered is large enough, then the number of cases will nonetheless be large enough to be approximated as continuous. However, whereas models often consider the entire population of a country together, disease transmission is far more local in reality (and can be made even more so with lockdowns and travel reductions). Thus, the sizes of the populations for which the models apply will be far smaller than that of an entire country. The locality in the dynamics (the degree of

which can be increased by travel restrictions) makes it more likely that a small fraction of the population infected in the model corresponds to the virus being eliminated in reality, and it also allows for the lockdown to be lifted region by region, rather than remaining in all regions until the entire country is cleared of the virus. The specific detailed assumptions of particular models may differ from this example; ultimately, what matters

The question of predicting the disease trajectory is less important than questions related to what's necessary to 1) cause an exponential decrease rather than increase in new infections and 2) cause this decrease to occur as quickly as possible.

is whether or not they appropriately characterize the large-scale behaviors of the disease.

Modeling for Policy

Finally, "What will happen?" is a different question than "What should we do?" For COVID-19 the latter question is far easier to answer than the former. In the absence of a full understanding of a system's details, answering the latter question involves understanding how our potential actions impact the relevant large-scale parameters of the system, which for COVID-19 are the rate of growth or decay in each region and the probabilities of transmission between regions.

Even if we cannot precisely predict the impact of any given intervention, we know of many interventions that will reduce the rates of transmission within and between regions. And based on the empirical understanding of COVID-19 transmission and the fact that many countries have eliminated or nearly eliminated the virus, we know that combining enough interventions together will reduce the rate of transmission sufficiently to achieve exponential decline and stop the outbreak (7). This, in and of itself, is a simple but powerful formal model that captures the large-scale behaviors of interest.

The question of predicting the disease trajectory is less important than questions related to what's necessary to 1) cause an exponential decrease rather than increase in new infections and 2) cause this decrease to occur as quickly as possible. The point is not the specific predictions for each intervention, such as social distancing, mask wearing, isolation in and outside of homes, testing/contact tracing/quarantines, and travel restrictions. The point is that if enacted in concert they can eliminate the virus.

This distinction is of particular importance because scientists often make predictions based on the assumption that societies are unwilling or unable to eliminate the virus. It's an assumption that's been invalidated by the actual actions and outcomes in countries such as Australia, Belize, China, Estonia, Greece, New Zealand, Norway, Slovakia, Switzerland, Thailand, Vietnam, and many others, as well as U.S.

[†]The virus may still be reimported, but if elimination is a stable fixed point of a collection of regions (Fig. 1), the number of regions with nonzero infections will decrease to zero over time.

[‡]The elimination of the virus can be hastened by testing, contact tracing, and quarantine, which may become more feasible and/or effective once the number of infections has been sufficiently reduced.

states such as Montana and Vermont. These regions controlled their outbreaks and have had little or no community transmission at the time of writing (1). Furthermore, the assumption that the virus cannot be beaten without a vaccine can become a self-fulfilling prophecy; policymakers may not take doable steps because they are discouraged by purported scientific predictions.

More generally, the use of models in pandemic response showcases a key difference between academically relevant research and policy-relevant analysis. The former can tolerate assumptions and models that are exploratory in nature, increasing our knowledge of the wide range of conditions that

might happen at some time in the future or some location—or even in an alternative reality—thereby increasing the scope of our understanding. The latter must focus on validated assumptions and real-world risk, including uncertainty in both our data and our understanding. Policy actions must be guided by only sound assumptions, because mistaken assumptions may cost millions of lives. Instead of assuming that we fundamentally differ from all of the countries that have achieved or are nearing elimination (an assumption that becomes increasingly implausible as the number of these countries grows), we should instead focus on how we can replicate their common success.

-
- 1 Johns Hopkins Center for Systems Science and Engineering, COVID-19 dashboard. <https://coronavirus.jhu.edu/map.html>. Accessed 17 June 2020.
 - 2 Y. Bar-Yam, From big data to important information. *Complexity* **21** (S2), 73–98 (2016).
 - 3 N. M. Ferguson *et al.*, Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Imperial College COVID-19 Response Team (2020).
 - 4 F. Ball, D. Mollison, G. Scalia-Tomba, Epidemics with two levels of mixing. *Ann. Appl. Probab.* **7**, 46–89 (1997).
 - 5 V. Colizza, A. Vespignani, Invasion threshold in heterogeneous metapopulation networks. *Phys. Rev. Lett.* **99**, 148701 (2007).
 - 6 A. F. Siegenfeld, Y. Bar-Yam, Eliminating COVID-19: The impact of travel and timing. arXiv:2003.10086 (2020).
 - 7 H. V. Fineberg, Ten weeks to crush the curve. *N. Engl. J. Med.* **382**, e37 (2020).