

An Economist’s Guide to Epidemiology Models of Infectious Disease

**Christopher Avery, William Bossert, Adam Clark,
Glenn Ellison, and Sara Fisher Ellison**

Around mid-March 2020, as the United States and much of the rest of the world was facing an unprecedented health threat in the form of COVID-19, an abrupt shift in the tone and policies of the United States and United Kingdom occurred. In early March, Prime Minister Boris Johnson said that “we should all basically just go about our normal daily lives.” Likewise, on March 11, President Donald Trump reassured the American people that for “[t]he vast majority of Americans, the risk is very, very low.” Just five days later, the Trump administration recommended that “all Americans, including the young and healthy, work to engage in schooling from home when possible. Avoid gathering in groups of more than 10 people. Avoid discretionary travel. And avoid eating and drinking at bars, restaurants, and public food courts” (as reported by Keith 2020). The British government likewise markedly changed course, with a series of partial measures preceding a March 23 lockdown order. Although Trump and Johnson had been receiving briefings about COVID-19 for several weeks, the proximate cause of the

■ *Christopher Avery is Roy E. Larsen Professor of Public Policy, Harvard Kennedy School, Cambridge, Massachusetts. William Bossert is David B. Arnold, Jr., Professor of Science, Emeritus, Harvard John A. Paulson School of Engineering and Applied Sciences, Cambridge, Massachusetts. Adam Clark is Assistant Professor, Institute of Biology, University of Graz, Graz, Austria. Glenn Ellison is Gregory K. Palm Professor of Economics, Cambridge, Massachusetts. Sara Fisher Ellison is Senior Lecturer in Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. Their email addresses are christopher_avery@hks.harvard.edu, bossert@seas.harvard.edu, adam.tclark@gmail.com, gellison@mit.edu, and sellison@mit.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at <https://doi.org/10.1257/jep.34.4.79>.

shift in both countries appears to have been the March 16 release of a headline-grabbing epidemiological model produced by London's Imperial College, which predicted that there could be as many as 2,200,000 deaths in the United States and 510,000 in the United Kingdom" (as reported by Landler and Castle 2000).

The Imperial College model was not the only one to feature prominently in public policy. The Institute for Health Metrics and Evaluation (IHME) at the University of Washington released and frequently updated state-level estimates which garnered substantial attention. Its predictions contrasted markedly with (the most extreme) ones from Imperial College. Both sets of predictions turned out to be quite far off in important ways. This fact should not be surprising. There is, unavoidably, much uncertainty about key parameters early in an epidemic. It also takes longer to produce models that use frontier methods and incorporate data from multiple sources. Still, the models can be faulted for providing standard errors that did not accurately reflect the degree of uncertainty underlying the course of the epidemic.

Given the importance of the topic and the impact that these early models had, it is not surprising that many economists quickly became interested in applying their skills to improve understanding of the COVID-19 pandemic. One goal of this paper is to provide an overview of the extant epidemiological literature to facilitate the work of economists who wish to make incremental contributions. We begin by introducing the classic SIR (susceptible/infected/recovered) model, which serves as the basis of much of modern epidemiology of infectious disease, both theoretical and empirical. As we will discuss, the classic model is useful for building intuition about the possible paths of a pandemic. Researchers typically build on this model in a variety of ways, depending on the specific research question, the characteristics of the epidemic, and the available data. We then turn to methods and challenges of implementing these models in empirical epidemiology. With this background in place, we return to the two high-profile forecasting models, explain where they fit into the landscape of empirical epidemiology, discuss the policy imperatives which drove their prominence, and offer critiques. Finally, we consider the related economics papers, ones that expand on SIR-type models, leverage them to provide policy advice, and offer estimates that could help inform them.

The COVID-19 pandemic poses a wealth of policy challenges. We believe that there are fruitful synergies for economists who acquaint themselves with some basic epidemiology models and empirical techniques. We then consider how their economist's toolbox could dovetail with the existing epidemiology literature to produce useful insights.

Epidemiological Theory

Epidemiological theory has been rooted in empirical facts from the start. In 17th-century London, haberdasher turned statistician John Graunt kept weekly records of the causes of death in London parishes. He used these data to estimate the risks of dying from different diseases. His work was instrumental in the

development of biostatistics, demography, and epidemiology. After him, doctors and medical researchers started relying on statistics and then statistical models to help them predict the spread of infectious disease. In the 18th century, Daniel Bernoulli (1766) devised the first true epidemiological model to study the spread of smallpox. In 1906, W.H. Hamer suggested that the spread of infection should depend on the number of susceptible and infected people. He introduced the mass action law for the rate of new infections. Kermack and McKendrick (1927) leveraged these insights to create the SIR model—the workhorse model still the basis of much of modern epidemiology.

In the past century, the field of epidemiology has advanced along lines similar to those of economics. Theorists have developed more sophisticated models to bring out many insights. In recent years the field has taken an empirical turn, developing increasingly sophisticated models that leverage vast and detailed new data sources. It should be noted that just as a relatively small share of economists focus on real-time forecasting of the economy, a relatively small share of epidemiologists focus on real-time forecasting of new pandemics. Epidemiology is a much broader subject, encompassing the study of the distribution and determinants of health and disease outcomes across various populations. The particular niche of the epidemiology literature that is especially relevant for the current pandemic are the models that focus on the spread of an infectious disease. We will start with a discussion of the workhorse model in this class, the SIR model. We note that this classic model both offers basic insights and provides a tractable framework amenable to being built upon.

The Standard SIR Model

SIR is an acronym for the three states (sometimes referenced as “compartments”) in the model: Susceptible, Infected, and Recovered. At each time t , each member of the population is in one of these states, with proportions in these states given by $S(t)$, $I(t)$, and $R(t)$ where $S(t) + I(t) + R(t) = 1$ for a population of unit mass.

There are only two ways to move from one state to another. First, currently infected people may become noninfectious and move to the recovered state. Second, a susceptible person can contract the disease through contact with a currently infected person. People in the recovered state may still be sick (or even dead) but they share two key characteristics: they are not infectious and also not susceptible to future infection. Transition rates between states are governed by parameters γ and R_0 , which serve as summary statistics for (1) the recovery rate and (2) the number of people an infectious person would infect over the course of their disease in a fully susceptible population.

One way to motivate the model is to suppose that agents are uniformly randomly matched in continuous time. Assume that each meets on average $R_0\gamma$ others per unit time and that any susceptible agent matched with an infected agent becomes infected. As a result, new infections occur at a flow rate of $\gamma R_0 S(t) I(t)$ per unit time. Suppose also that each infectious agent recovers with probability γ per unit time, creating a flow of $\gamma I(t)$ individuals per unit time moving from the Infected to the

Recovered state. These dynamics can be summarized by the following continuous time dynamic equations for the values of $S(t)$, $I(t)$, and $R(t)$ given the two possible transitions from S to I for new infections and from I to R for sick people who become non-infectious:

$$\dot{S}(t) = -S(t)I(t)R_0\gamma,$$

$$\dot{I}(t) = S(t)I(t)R_0\gamma - \gamma I(t),$$

$$\dot{R}(t) = \gamma I(t).$$

The number of periods that an infected agent remains in the infected state follows an exponential distribution with parameter γ , so the expected amount of time in the infected state is $\frac{1}{\gamma}$. With $R_0\gamma$ contacts per person per unit time with others, each infected person has an expected number of R_0 contacts while infected. That is, the parameter R_0 can be thought of as the expected number of people that a newly infected person will directly infect in a population where everyone is susceptible.¹

The initial level of infection at time 0 is another exogenous parameter of the model and is typically assumed to be quite small (for example, one infection per 10 million people). If $R_0 > 1$, the number of infections is larger than the number of recoveries in early periods, while the proportion in the susceptible state remains close to 1. As a heuristic approximation, we would expect contacts with people infectious at time 0 to directly produce a total of $R_0 I(0) S(0)$ new infections, which is approximately $R_0 I(0)$ if $S(0)$ is close to 1. This set of new infections would produce approximately $R_0^2 I(0)$ subsequent new infections, and these would produce $R_0^3 I(0)$, and so on. For this reason, the initial growth rate of infections in an SIR model with $R_0 > 1$ is approximately exponential. Formally, one equilibrium of the system is $S(t) = 1$, $I(t) = 0$, $R(t) = 0$ for all t , but if this equilibrium is locally unstable if $R_0 > 1$, then adding a small number of infected agents leads to contagious growth of $I(t)$. By contrast, an equilibrium with $I(t) = 0$ is locally stable if $R_0 < 1$, as a small infection dies out in that case.

Over time, the growth rate of infections declines because the proportion of people in the susceptible state diminishes continuously as the infection spreads. Regardless of when the infection takes place, each infected person has an expected number R_0 of contacts with others while infectious, but as time passes, more and more of those contacts are with people who are not susceptible. The model has a “herd immunity” threshold of $\bar{S} \equiv 1/R_0$. When $S(t) = \bar{S}$, the expected number of people that a newly infected person will directly infect is equal to 1. The important implication of this property is that once the fraction of the population that is

¹A common alternative description of the SIR model defines $\dot{S}(t) = -S(t)I(t)\beta$ and $\dot{I}(t) = S(t)I(t)\beta - \gamma I(t)$, and then identifies R_0 separately as the ratio $R_0 = \frac{\beta}{\gamma}$. It is also equivalent to assume a proportionally higher probability $KR_0\gamma dt$ (where K is a known positive constant) that any pair of agents meet in combination with probability $\frac{1}{K}$ that a susceptible agent matched with an infected agent becomes infected.

susceptible is below the herd immunity threshold \bar{S} , a small infection introduced into the population will die out with the size of the infectious population never increasing.²

Importantly, note that reaching “herd immunity” does not mean that people will not continue to be infected. New infections continue to occur. They are just outnumbered by recoveries that are occurring. When R_0 is large, the number of people who are infectious when the herd immunity threshold is reached is large, so being limited by the number of recoveries is not comforting. Indeed, in these models there can be substantial “overshooting” with many more than $1 - \bar{S}$ people eventually infected. The number of people who escape the epidemic does not have as simple a formula, but is obviously very important practically. In an uncontrolled epidemic, it can be described as the solution to a simple implicit equation.³ Numerical examples indicate that overshooting can be dramatic with a significant fraction of the population getting infected after herd immunity is reached. For example, with $R_0 = 2$ we reach “herd immunity” when half the population has been infected, but the infection will not completely die out until another 30 percent of the population has been infected. With $R_0 = 2.5$, herd immunity is reached when 60 percent have been infected, but only 11 percent of the population will remain uninfected in an uncontrolled epidemic. In short, even with a moderate R_0 , few escape an uncontrolled epidemic. The “social distancing” policies that have been used to suppress COVID-19 infection rates are essentially an attempt to reduce R_0 .

One other noteworthy feature of SIR models is that for many values of R_0 , the time-path of new infections (and deaths) has a shape that is fairly symmetric about its peak and looks somewhat like a normal density. This provides a potential explanation for one of the earliest empirical observations in epidemiology: Farr (1840) noted that the time series of deaths in a smallpox epidemic and in four other epidemics “which have not yet been effectually controlled by medical science” were roughly symmetric and bell-shaped. Figure 1 below reproduces Figure 1A from Ferguson et al. (2020) illustrating the predictions of their SIR-like model for Great Britain and the United States.

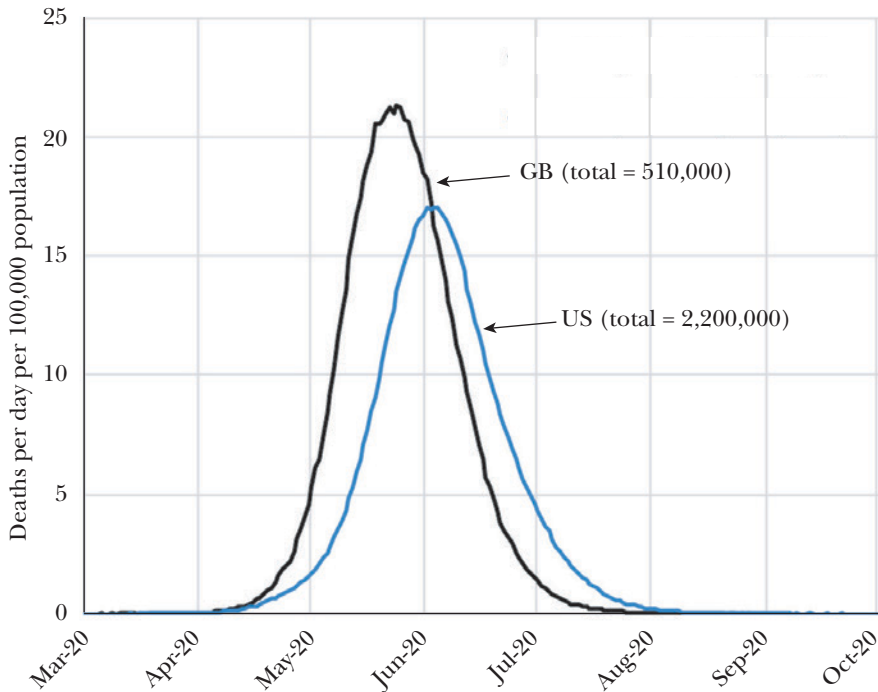
Some Conceptual Lessons from the Standard SIR Model

When a serious contagious disease becomes prevalent, two reactions will typically occur: people will modify their behavior to avoid getting sick and governments will enact policies aimed at slowing or stopping the spread. We can think of the original R_0 as a compound parameter, one that embodies both the underlying biological ability of the pathogen to jump from person to person in various types of

²Formally, the herd immunity threshold is such that $S(t) = S$, $I(t) = 0$, $R(t) = 1 - S$ is a stable equilibrium in the model for any $S \leq \bar{S}$.

³Formally, we can define the fraction who escape infection, $S(\infty)$, as $S(\infty) \equiv \lim_{t \rightarrow \infty} S(t)$. The equation that can be solved to find it is $S(\infty) = e^{-R_0(1-S(\infty))}$. Intuition for the formula is that $1 - S(\infty)$ agents are eventually infected. Each on average has R_0 interactions with others that would cause infection in someone who is susceptible. So the probability of escaping infection is the probability of zero events given a distribution that is Poisson with mean $R_0(1 - S(\infty))$.

Figure 1

Unmitigated Epidemic Scenarios from Imperial College Model

Source: Figure reproduced from Ferguson et al. (2020) Figure 1A: “Unmitigated epidemic scenarios for GB and the US. (A) Projected deaths per day per 100,000 population in GB and US.”

interactions as well as the number of interactions of each type that people have in the ordinary course of their daily lives.⁴ As self-interested behavior and government policies reduce interactions, it is as if the R_0 parameter in the equation describing how infections transmit is reduced to some time- and state-dependent variable R_0^t .⁵ It is important to remember that all the parameters of SIR models are simple encapsulations of more complex biological events. The cycle of infection involves the population biology of the pathogen outside the host, the behavior and population

⁴This approach has parallels to a classic predator-prey theory in biology, whose models have almost exactly the same form and dynamics as an SIR model. In that literature, there is a parameter governing transition from “freely roaming” to “prey,” which is a compound parameter with a fixed attack rate for a particular predator-prey combination as well as a contact rate between predator and prey, which can vary geographically and over time. See Gotelli (2008) for a description.

⁵See Chernozhukov, Kasaha, and Schrimpf (2020) and Goolsbee and Syverson (2020) for empirical evidence on the impact of endogenous behavioral changes and various government policies.

biology of the host, and the interaction of the pathogen and the host. Spatial, temporal, and between-host differences in the details of these events lead to the heterogeneity of the parameters that modelers now find important. While much of epidemiology is focused on understanding these details, they are typically absent from the models currently used to predict the course of diseases.

Policies that reduce the reproduction rate R_0 are often described as “flattening the curve,” referring to the graph that shows the rise of cumulative infections over time. A change in behavior that reduces R_0 to R_0^t at any time t affects the fraction of the population that permanently escapes infection. But the standard formula for the herd immunity threshold remains relevant to thinking about the possible long run outcomes: if we are not in the herd immunity region—that is, if $S(t) > \bar{S}$ —then the infection will once again spread if government restrictions are removed and people go back to their normal behaviors. If we are in the herd immunity region, then the infection will die out even if all restrictions are removed. Indeed, in this way the SIR model illustrates a clear intuition for how temporary policies can provide long-term benefits: implementing policies that reduce R_0^t at future times when we are approaching the herd immunity threshold will reduce overshooting.

In the case of COVID-19, reaching the herd-immunity threshold is widely believed to entail a devastating loss of life. The SIR model suggests that two other approaches may be appealing in such situations. First, we might put in place policies to reduce R_0^t with the intention of maintaining those policies until a vaccine is developed, thereby keeping the system from ever reaching the herd-immunity region. Second, we might enact more aggressive temporary measures for a period of time sufficient to drive prevalence to a level that is low enough so that less economically costly means of keeping $R_t \equiv R_0^t S(t)$ below 1 become feasible. For example, Hong Kong’s suppression of COVID-19 has involved, among other measures, hospitalizing everyone who tests positive to ensure isolation and conducting aggressive contact tracing. This is extremely expensive on a per-infected person basis but has cost trillions less than the US approach, not to mention limiting Hong Kong’s loss of life.

The SIR model is also helpful for thinking about vaccines. Vaccines are typically not perfect and neither available to nor willingly received by everyone. Suppose, for instance, that a vaccine was effective in preventing the disease completely and permanently in 60 percent of the people who received it and did nothing for the other 40 percent who received it. Administering such a vaccine to the entire population with, say, 10 percent infected or recovered would result in an additional $0.9 \times 0.6 = 54$ percent of the population immune, so that $S(t) = 1 - 0.1 - 0.54 = 0.36$. Depending on the value of R_0 , that number could be sufficient to achieve herd immunity. Achieving herd immunity via a vaccine rather than via infections is also advantageous in that it mitigates overshooting.

Variants of the SIR Model

There are many variants of the SIR model. As usual, the choice to add or subtract complexity from a model should depend on what one is studying. Common variants of the SIR model add additional disease states, referred to as “compartments,”

to provide a more realistic model of disease progression and transmission. The SEIR model includes an “exposed” state to account for individuals who have been infected with the disease but are not yet themselves infectious (Hethcote 2000; Li and Muldowney 1995). The SAIR variant includes an “asymptomatic” compartment for individuals who are infectious but may never develop symptoms. Because of the apparently strong contribution of asymptomatic and pre-symptomatic carriers to the spread of COVID-19, these variants, and particularly the SEIR model, have been quite common in recent epidemiological studies (for example, Kissler et al. 2020; Prem et al. 2020). Epidemiologists sometimes also introduce additional compartments not to reflect disease states but as a mathematical means of making the transmission time process more flexible as in Champredon et al. (2018), although this aim can be accomplished directly as in Zhigljavsky et al. (2020). These variants may be especially useful if one were interested in studying the impact of policies for which timing within the disease cycle is critical, like protocols for testing, contact tracing, and quarantining. For an excellent review of many of these extended forms, see Blackwood and Childs (2018).

A broader category of models divides compartments even further into dozens or even hundreds of different geographic and age states and then allows contact, infection, and recovery rates to vary across classes (Blackwood and Childs 2018; Hethcote 2000). Ebola, for example, is spread through contact with bodily fluids even after death, and one might capture this effect on disease dynamics by considering populations of health care and funeral workers (Champredon et al. 2018). Given the current understanding about how COVID-19 seems to be transmitted, it is easy to think of subpopulations who will have many more risky interactions than average: those living in crowded urban apartments, frequenting bars and nightclubs, using public transportation, attending crowded religious services, working in a nursing home, and so forth.

Models with heterogeneous subpopulations again behave much like the classic SIR model whereby the growth rate of a contagious disease is initially exponential then slows (and eventually dies out) over time (for example, Diekmann, Heesterbeek, and Metz 1990; Dushoff and Levin 1995; Lajmanovich and Yorke 1976). A common pattern in these models is that variations in within-class contact or transmission rates across subgroups produce a faster overall spread of infection than in a well-mixed SIR, with infections concentrated in certain high-risk subgroups. Thereafter, however, dynamics tend to slow down relative to a well-mixed model because contact rates between subgroups are typically lower than the average transmission rate (Bolker 1999). In general, these features tend to lead to less complete spread of diseases in age- and spatially structured models than an analogous homogeneous SIR model, although this is not always the case (Gomes et al. 2020; Hébert-Dufresne et al. 2020). Britton, Ball, and Trapman (2020) provide an illustration in which heterogeneity reduces the herd immunity threshold from 60 to 43 percent. In addition, heterogeneity can also lead to a longer overall persistence of diseases. For example, geographic structure can make it difficult to eradicate a disease fully, allowing periodic resurgences (Lloyd and May 1996).

The polio virus provides an example of the perverse impacts that can emerge from heterogeneity. Changes in hygiene practices in the United States around the middle of the twentieth century led to a decrease in infectiousness in polio, which in turn, led to an increase in its average age of onset. Because younger children typically experienced much milder cases of the virus, this increase in age of onset led to an overall increase in the mortality and morbidity associated with being infected with polio, which persisted until the widespread adoption of a vaccine (Melnick 1990).

Real-world disease states and processes are more complex than those assumed in all of these models, of course. For example, “infected” could be treated as a multidimensional continuum of states, instead of a single state. People can vary in the severity of their symptoms, their health outcomes, and the degree of infectiousness. Likewise, whether an exposure results in an infection can depend on the nature and dosage of the exposure. The extent to which people develop immunity will vary. All of these factors are subject to individual, spatial, and temporal heterogeneity.

Empirical Epidemiology

The field of epidemiology does not divide itself into theory and empirical work as neatly as does economics. There is more diversity in research styles and questions. It does appear, though, that like economics (as discussed in Angrist et al. 2020), the field of epidemiology has become more empirically oriented over time. Most relevant to economists, perhaps, are branches estimating parameters of disease processes, forecasting the courses of epidemics, and estimating policy effects. As noted above, forecasts by epidemiologists of the future course of COVID-19 received tremendous attention in the early days of the epidemic. These forecasts can combine theoretical modeling, calibration of some parameters, and estimation of others. Broadly speaking, forecasting models are often regarded as falling into two main styles. Those based on SIR-type models are in a class called “mechanistic,” which, like structural empirical models in economics, assume that a model is exactly correct and calibrate or estimate parameters to obtain a predictive model. There is another class of predictive models termed “phenomenological,” which may be motivated by theories of disease spread but are not derived directly from those theories. Instead, they posit a functional form for the evolution of cases or apply time-series methods to predict future outcomes based on available observations. This distinction is not a neat one, however, and forecasts can combine elements of both types.

In economics, choice of empirical model and technique is often driven by realities of data quality and availability. Economists interested in policy evaluation have, for instance, invested enormous effort into developing techniques for causal inference with observational data, which is what economists often have to work with. Something similar is true for epidemiologists interested in forecasts: their models are designed to leverage the data available on an epidemic in its earliest crucial

stages to greatest advantage. These early numbers tend to come from boots-on-the-ground efforts such as contact tracing or case counts, and they can be used to estimate parameters of either phenomenological or mechanistic models. To be clear, data from contact tracing differs from case counts in that it has information about the source of and the resulting infections from a particular infection, but it may not include most or all infections. Case counts attempt to document all infections, but not the tree of connections among them.

Mechanistic Forecasts

Even under ideal circumstances, reliably estimating parameters of mechanistic epidemiological models, such as the SIR, can be quite challenging due to their nonlinear and dynamic nature. The simplest idea for estimating R_0 —that is, making a list of initial infections, tracking down the number of additional infections that can be traced directly to each of those initial ones, and dividing to obtain an estimate of R_0 —is not an accepted practice due to the fact that incomplete contact tracing and asymptomatic cases would lead to downward-biased estimates. Instead, researchers often employ some more sophisticated variant of the following two-step method-of-moments approach: start with the log growth rate of the epidemic as implied by an SIR model, $\gamma(R_0 - 1)$, and equate that to an empirical log growth rate from the case counts. To identify γ and R_0 separately, then, one can use (potentially incomplete) contact tracing data to infer the distribution of length of time between infections, which helps tie down γ .

Most of us have internalized the notion that more data always lead to better estimates, but a counterintuitive situation can exist here. As the epidemic spreads and more data become available, the quality of (at least some of) the data can be compromised. First, contact tracing efforts will inevitably fall behind in a fast-growing epidemic, and the resulting data might be increasingly lower quality. Second, as an epidemic grows, behavioral responses can emerge, which could contaminate an estimate of R_0 . Third, increased testing can identify asymptomatic cases which could contaminate case growth rates, because cases which would not have been included in early case counts are included in later ones. In short, more data can lead to worse estimates, as discussed in Ferretti et al. (2020). There is a trade-off, though: these limited sample sizes early in an epidemic make capturing heterogeneity of many types problematic, to say nothing of capturing changes in parameters over time.

We should stress that epidemiologists have studied these issues in depth for many years. Asymptotic analyses of the properties of maximum likelihood estimation and other estimators of parameters in homogeneous and heterogeneous SIR models can be found in Rida (1991) and Britton (1998). Markov Chain Monte Carlo methods for the Bayesian estimation of heterogeneous SIR models are described in Demiris and O'Neill (2005). And modern applications of disease models typically involve parameterization approaches that are more sophisticated than those described above. For examples of work along these lines, useful starting points include Mills, Robins, and Lipsitch (2004), Massad et al. (2010), and Viboud et al. (2018).

Phenomenological Forecasts

In contrast to mechanistic methods, phenomenological approaches are often relatively straightforward to implement for the early stages of an epidemic. Early case data are used to fit the assumed growth curve (for instance, using maximum likelihood estimation). As additional case data come in, the parameter estimates are refined to reflect the new information. Information on the source of any particular case, typically provided by contact tracing, would not be necessary. With limited early data, it can be difficult to estimate as many parameters as one would want to estimate for a realistic compartmental (mechanistic) model, and this fact can make simple phenomenological approaches appealing. For example, Tuite and Fisman (2018) use a simple functional form with just three parameters, estimated by maximum likelihood, in which the way an epidemic declines is determined by one of the parameters. They note that they “are agnostic about the nature of factors that slow growth, but they could be postulated to include behavioural change, public health interventions, increased immunity in the population, or any other dynamic change that slows disease transmission.”

As epidemics progress, phenomenological approaches that use time-series techniques to predict changes remain well-suited to making near-term predictions. These models can be less useful, however, for other tasks. Observation error can rise as larger swaths of a population are infected and contact tracing becomes less reliable, and tightly parameterized models lack the flexibility to respond to qualitative changes in disease behavior that are inconsistent with earlier apparent patterns. For example, a model which posits a symmetric, bell-shaped evolution of cases over time cannot accommodate repeated changes in the rate of spread due to changing regulations, changing public perception, and “quarantine fatigue.” In a later section, we will see how early fits from the IHME model accurately characterized initial growth rates in case numbers across much of the United States, but its predictions of peak infection numbers and long-term dynamics have proven much less reliable.

Policies and Causal Inference

Epidemiologists and other health researchers have long been interested in the effects of healthcare interventions. The use of randomized controlled trials—often called the “gold standard” for causal inference—was pioneered by health researchers. During epidemics, however, the earliest data available are typically observational. Even in randomized trials, noncompliance raises concerns about selection biases. And, of course, the very nature of an infectious disease implies that a treatment applied to one agent may affect others. As a result, epidemiologists have recognized that the methods most commonly used in other medical fields for policy evaluation may be less appropriate for epidemiological applications (Halloran and Struchiner 1995; Hernán and Robins 2006). By now, however, epidemiologists have developed a variety of techniques to address field-specific concerns (for an extensive exposition, see Hernán and Robins 2020).

Analyses of Genomic Data

Analysis of the SARS-CoV-2 genome has revealed thousands of different strains of the virus circulating around the world (for the current phylogeny, consult nextstrain.org/ncov/global). The medical community has many reasons to be interested in these multiple strains. For instance, they may differ in communicability or virulence, or there could be less-than-perfect immunity across strains. Korber et al. (2020) present laboratory and epidemiological evidence suggesting that the COVID-19 variant which is now most common is more infectious than the strain that was dominant in Wuhan.

For the purposes of estimating epidemiological models, another immediately useful application of these techniques is to trace the spread of various mutations to determine where and when epidemics began in various regions. In fact, genomic data can serve as a type of substitute for contact tracing or detailed micro-level data on social networks and other human interactions, allowing researchers to trace the source of a particular group of infections without ever knowing anything about the agents' contacts. Researchers in Israel used genomic data, for instance, to produce the often-cited fact that 80 percent of all COVID infections there were caused by 1–10 percent of infected agents (Miller et al. 2020). In another genomic study, Worobey et al. (2020) note that although cases have been reported as early as January 2020 in the United States and Europe, genetic evidence suggests that these introductions failed to spread, and that it was only through later introductions at higher incidence that SARS-CoV-2 was able to establish in the general population. If these findings hold up in follow-up research, they may indicate that even if the virus cannot be fully eradicated, control measures may well prove to be effective if incidences can be brought low enough.

Early High Profile Models—What Went Wrong?

The introduction recounted how an early prediction model from Imperial College had a seemingly huge effect on policy decisions in the United States and the United Kingdom. In fact, one could argue that policy imperatives drove the prominence of that and another high-profile prediction model from IMHE early in the pandemic. Policy-makers were desperate for guidance on mask-wearing and social distancing, predictions on the number of intensive care hospital beds necessary in a particular city, likely timing of peak infections, and so forth. Those two models were up and running early in the pandemic and provided those numbers that policy-makers needed. It is instructive to take a closer look to understand how their predictions were produced and what ultimately went wrong.

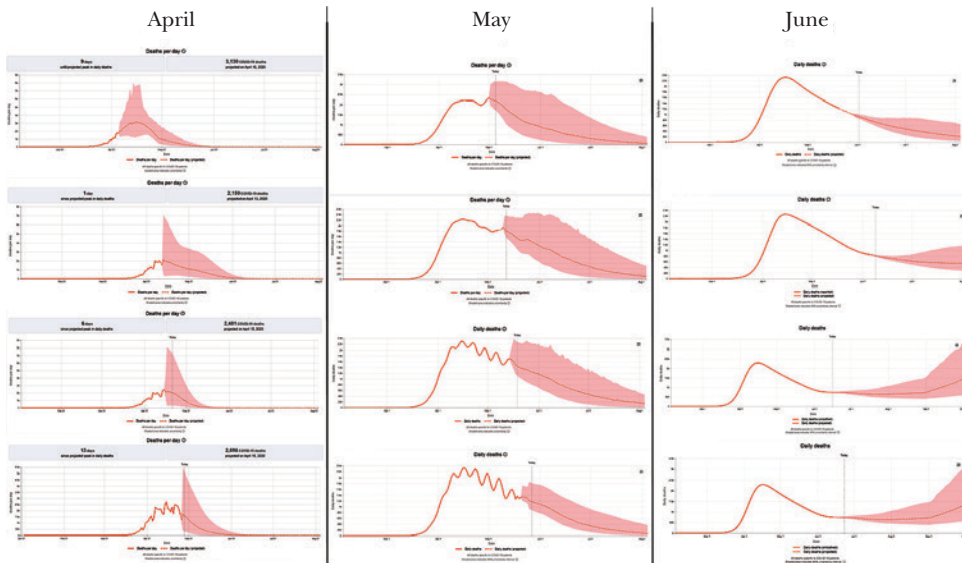
The headline-grabbing figures from the Imperial College model were the most extreme predictions out of many that they produced. They arose from assumptions that governments would not mandate any mitigation strategies, such as mask-wearing or social distancing, and indeed that people would not choose to engage in any of those strategies themselves. Those assumptions were often omitted from the initial

reporting and public discussion of the predictions. Much of the Imperial College report, however, consisted of discussions of the potential impact of such policies, along the lines of an earlier policy discussion on mitigating pandemic influenza in Ferguson et al. (2006). Some information about the details of the Imperial College model were given, but initially the source code was not public. The early reports made certain details clear: the model was based on the familiar SIR framework and that extreme predictions were derived assuming that neither official actions nor individual choices would be taken to slow the spread of the virus. The R_0 term was taken as a single, fixed parameter, with a value of 2.4. Their estimated death rate for those infected was 0.9 percent. Both estimates were based on early experience with COVID-19 in places such as China and Italy but obviously associated with significant uncertainty. The source code for the model was eventually released at the end of April, and researchers were able to reproduce its results from its assumptions by early June (as reported by Chawla 2020). Although this delay is understandable, it was also arguably a contributor to confusion surrounding predictions early in the pandemic.

Meanwhile, as the number of COVID-19 cases was ramping up in the United States, alternative predictions were being offered by IHME at the University of Washington. Their phenomenological model began by assuming a particular functional form for how the number of cases in a locality would rise and then fall over time, with location-specific parameters estimated to fit early case numbers. The model could easily be fit separately to data on each state, and predictions were refined as new data came in. The intention was that local officials could then use these location-specific and daily predictions to plan extra hospital capacity and procure medical equipment, which many of them did. The notion, however, of a common functional form—that is, that the basic shape of increase, peak, and decline of infections would be the same in all locations, from Italy to India, from Wuhan province to Topeka, Kansas—seems to ignore crucial information about how mitigation strategies varied across locations and changed over time. More recent versions of the IHME model have taken an alternative approach, as we discuss in a moment. Roughly speaking, the originally publicized IHME model was assuming a bell shape for the daily deaths and trying to find the parameters governing that bell shape based on the early observations. In a model of this form, once growth has started to slow, there will be limited uncertainty about the size or timing of the peak. Also, the bell-shape symmetry implies that deaths will start falling as rapidly as they grew.

Figure 2 shows a series of screen grabs from the IHME model predicting daily US deaths (from the Internet Archive), at approximately one-week intervals starting in early April 2020. The first four predictions, going down the first column and through the end of April, have several common features resulting from the bell-shape assumptions: the predicted shape of deaths over time is symmetric; the predicted number of deaths goes to zero quickly, around June 1; and the error bands are large in the short run and go to zero around the time that the predicted number of deaths goes to zero. Note that in these first four panels, estimates of the parameters are being updated regularly as new data come in.

Figure 2
Weekly Screenshots of the IHME US Deaths Predictions



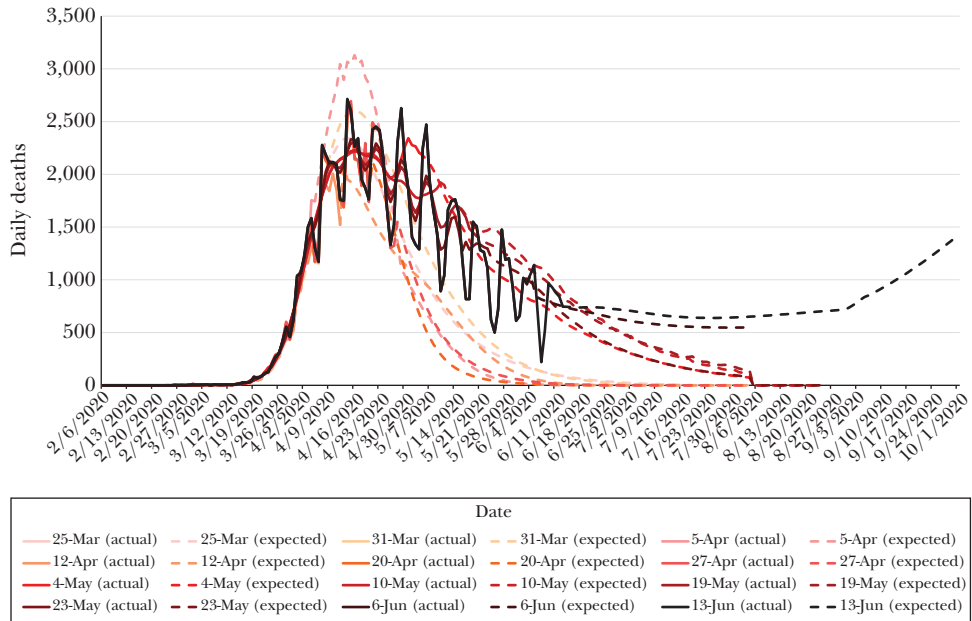
Source: covid19.healthdata.org.

Note: This figure was constructed with a series of screen grabs from covid19.healthdata.org, IHME's website, from the Internet Archive located at archive.org. The screen grabs are at approximate one-week intervals throughout April, May, and June 2020.

In early May 2020, IHME switched away from the curve-fitting approach to a more mechanistic SIR-type framework. The model predicted roughly the deaths in the next few days in a phenomenological way and then fit an SIR-based model to the past and short-term future predictions to generate long-run predictions. The middle column shows that starting in May, the model allowed for asymmetry. It also started using a smoothing algorithm on the existing case data. The way error bands were calculated changed, but error bands still shrunk eventually instead of growing, reflecting that declining deaths implied that epidemics in SIR models in which R_t falls to less than 1 die out in an exponential manner. As a result of these changes, predictions of positive numbers of deaths stretched into summer 2020. Starting in June 2020, the final column of the figure, another substantial change was made to the calculation of error bands, whereby they start small and increase as time proceeds, reflecting increasing, not decreasing, uncertainty in predictions further into the future.⁶

⁶For additional discussion of how the models did not reflect the degree of uncertainty early in the epidemic, see Avery et al. (2020). Stock (2020) also notes the importance of uncertainties that existed early in the pandemic.

Figure 3
IHME US Daily Deaths Predictions Overlaid



Source: covid19.healthdata.org.

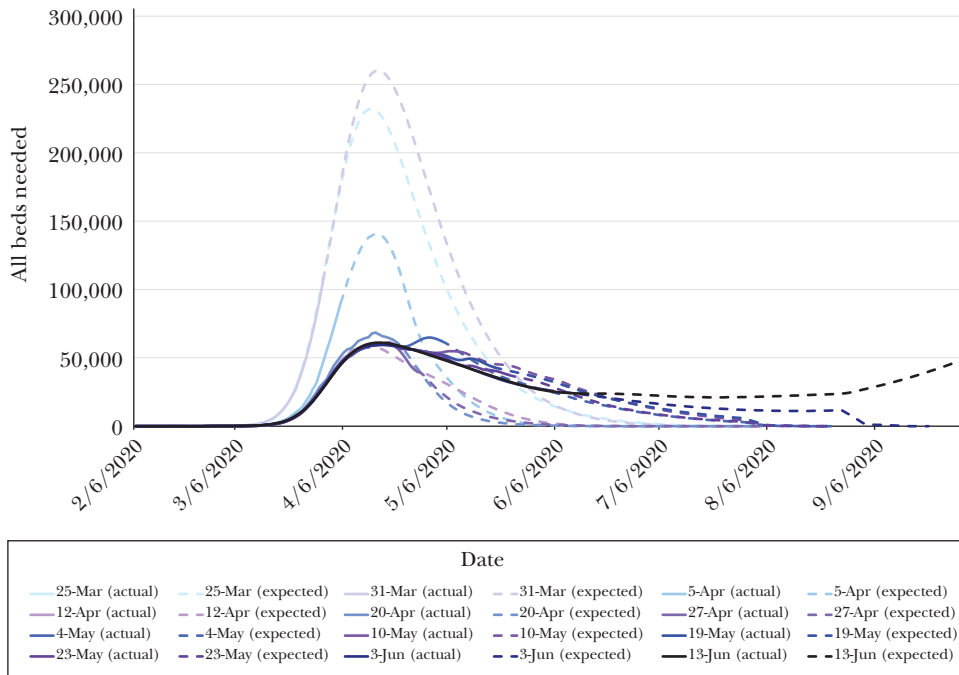
Note: This figure was constructed from data downloaded directly from IHME’s website at the University of Washington.

In Figure 3, we overlay these same predictions on a common scale, color-coded so that earlier predictions are lighter. For readability, we do not include error bands. Clearly, the IMHE predictions of US deaths over time change as it becomes clear that the pandemic will not die out at the beginning of the summer, and a symmetric model of US deaths is inaccurate. Even so, the initial predictions of the size and location of the (first) peak were fairly accurate.

Figure 4 shows a different output of the IHME model: predictions of hospital utilization. With this outcome, the initial predictions are starkly different from later ones. Not coincidentally, many locations prepared for much greater hospital utilization during the first “surge” than was needed. We should note that IHME does publish their source code and is forthcoming about changes. That being said, the model is complicated enough that reading through the source code and documented changes is difficult and time-consuming, certainly for us, but also, one would imagine for most researchers.

The Imperial College and IMHE models filled a void early on for policy-makers scrambling to understand the pandemic, to decide how strongly to react, to

Figure 4
IHME US Hospital Use Predictions Overlaid



Source: covid19.healthdata.org.

Note: This figure was constructed from data downloaded directly from IHME’s website at the University of Washington.

convey policies to constituents, and to allocate resources. But many other predictive models are now available, some with well-designed online dashboards where users can insert different assumptions, some backed by state-of-the-art epidemiology theory, and some leveraging empirical innovations and new information. We cannot hope to survey all of the predictive models here, but both the Centers for Disease Control and Prevention (CDC) and the website FiveThirtyEight.com highlight and compare several of the most well-known and well-received ones.⁷ Table 1 shows 15 models highlighted by FiveThirtyEight.com (including IHME), with a few words about their basic approaches and some details about their implementation. These models largely agree in their short-run predictions, but divergence appears at forecasting horizons of six weeks or more. We have organized them by predicted

⁷The Centers for Disease Control has come under criticism from many quarters for allowing political considerations to influence how they present and describe predictive models.

Table 1
Some Predictive Epidemiological Models

<i>Source</i>	<i>Approach</i>	<i>Details</i>
High Predicted Mortality Level (by Sept. 5th)		
The University of Texas COVID-19 Modeling Consortium, University of Texas https://covid-19.tacc.utexas.edu/projections/	Model 1 uses a curve fitting approach, and Model 2 is an SEIR model with compartment "D" (dead)	Uses anonymized mobile phone data and daily reported deaths to make predictions for three weeks ahead
COVID Scenario Pipeline, Johns Hopkins University https://github.com/HopkinsIDD/COVIDScenarioPipeline	SEIR model	Projects the spread of the epidemic and impacts on health care for different interventions
ERDC SEIR Model, U.S. Army Engineer Research and Development Center https://github.com/reichlab/covid19-forecast-hub/blob/master/data-processed/USACE-ERDC_SEIR/metadata-USACE-ERDC_SEIR.txt	SEIR model with compartments for unrecorded infections and isolated individuals	Uses Bayesian inference to choose parameters
EpiGro, University of Arizona https://www.sciencedirect.com/science/article/pii/S1755436516300329	Curve fitting model	Based on properties of curves implied by SIR model
Medium Predicted Mortality Level (by Sept. 5th)		
DeepCOVID Model, Georgia Tech https://www.cc.gatech.edu/~badityap/covid.html	Deep learning model	Assumes that the effect of interventions is implicitly captured in mobility data
IHME COVID-19 Projections, IHME, University of Washington https://covid19.healthdata.org/united-states-of-america	Hybrid model that incorporates statistical and disease transmission models	Uses social distancing information and mobile phone data to estimate contact between people
COVID-19 Projections Using Machine Learning, Youyang Gu https://covid19-projections.com/	SEIR with machine learning to choose parameters	Estimates incorporate all infected individuals of SARS-CoV-2 virus, not only individuals who tested positive from a COVID-19 test
Columbia University COVID-19 Projections, Shaman Group https://github.com/shaman-lab/COVID-19Projection	Metapopulation SEIR with filtering to determine parameters	Includes projections for daily cases, infections, mortality, and cumulative hospital usage
Global Epidemic and Mobility Model (GLEAM), Northeastern University https://covid19.gleamproject.org/	SEIR model with mobility data	Region-level model with several types of human mobility between regions
Low Predicted Mortality Level (by Sept. 5th)		
COVID-19 Simulator, MGH, Harvard Medical School, Georgia Tech, Boston Medical Center https://www.covid19sim.org/	SEIR model	Includes state-level variations in mobility and tracks hospital usage
Bayesian SEIRD Model, University of Massachusetts https://github.com/dsheldon/covid	SEIR model with additional compartments "D" (death) and "H" (hospitalized-and-will-die)	Employs Bayesian inference and time-varying dynamics
UCLA-SuEIR Model, UCLA https://covid19.uclaml.org/	SuEIR model	Has compartment for unobserved infections
A Shiny App, Iowa State http://www.covid19dashboard.us/	New spatiotemporal epidemic modeling (STEM) framework	Nonparametric model emphasizing 7-day forward projections down to county level
DELPHI Epidemiological Model, MIT https://www.covidanalytics.io/	SEIR with under-detection, hospitalization, and government interventions	Varies effective contact rate and societal/government response by state
LANL Model, Los Alamos https://covid-19.bsvgateway.org/	Dynamic model that forecasts future cases and deaths	Allows for a variety of interventions, resulting in a wide prediction interval

mortality levels. In part, this divergence may reflect different assumptions about how social distancing and government policies will evolve.

As this article was being completed in late summer 2020, it seemed that predictive models about the future of the epidemic had faded from popular discourse. Discussions of reported cases, deaths, and trends seemed, by mid-July, to be getting more attention than forecasts from epidemiological models. Google Trends indicates that searches for “IHME Model” peaked in mid-April and had fallen by 90 percent by early July. Attention by academics also seems to have fallen: Google Scholar indicates that Ferguson et al. (2020), released on March 16, had already been cited 828 times in early July, while the later May 21 report by the Imperial group (Unwin et al. 2020) providing more sophisticated estimates of R_0 for US states had been cited just three times.

One likely reason for the initial surge and subsequent decline of interest in predictive models is that earlier in 2020 they were seen as relevant to policy choices: whether to require businesses to close and people to stay home and how much to invest in hospital bed capacity. By contrast, predictive models appear to have much less relevance to the pressing decisions of fall 2020 such as when to reopen in-person schools. In addition, predictive models have likely lost popular credibility. The initial IHME forecast predicted that the epidemic would all but die out in the United States by early June. The Imperial College model was often linked to its most extreme predictions. Finally, the waning interest may also reflect that the future course of the disease is not readily predictable by any model, but rather, will depend to a considerable extent on how individuals behave and what policies are enacted.

Epidemiology-Related Research in Economics

Economists have responded enthusiastically to demands for COVID-related research and analysis. We cannot attempt to cover this burgeoning literature in its entirety. Rather, our focus will be tighter: on research that leverages SIR-type models, expands upon them, or offers estimates that could help inform them. We chose this sub-literature as our focus because we feel that it is an area where cross-discipline knowledge and the use of complementary models and tools have already continued and will continue to yield real insights.

It is useful to organize much of this sub-literature into three strands. These strands represent salient features of this pandemic as opposed to previous ones, and we feel that economists are well-positioned to make contributions in those three areas. First, economists have recognized the potential endogeneity of parameters such as R_0^t , as the precautions taken could be a function of disease prevalence or current cases. Second, several economics papers have focused on the effects of allowing various types of heterogeneity in SIR-type models. Third, economists have taken the political economy issues involving endogenous social distancing and government policies seriously—issues which could also greatly influence the pattern of R_t over time. We will discuss each of the three strands in turn.

Endogeneity

The R_0^t parameter in an SIR model is a potentially endogenous parameter, which reflects both how easily communicable a particular pathogen is as well as how people behave and interact given the current state of an epidemic. It is natural that economists would recognize this endogeneity and model it theoretically and allow for it in empirical analyses. Applying traditional economics approaches to incorporating behavioral responses into epidemiological models is not new and dates back at least to work on the AIDS epidemic in the 1990s (Kremer 1996; Philipson and Posner 1993). Recently, a strand of COVID-related literature accommodating and studying an endogenous reproduction number has emerged. Toxvaerd (2020) and Kudlyak, Smith, and Wilson (2020) develop models that endogenize the social distancing as reflecting a cost and benefit of avoiding infection and discuss impacts on the time path of infections. Farboodi, Jarosch, and Shimer (2020) develop a tractable model of forward-looking individual distancing in which they can compare equilibrium and social optimizing distancing. They calibrate to epidemiological estimates of R_0 from early in the pandemic. They then show that, given a particular choice for the disutility of social distancing, the laissez-faire equilibrium, where social distancing is the result of endogenous individual choices, roughly matches the degree of distancing in the United States as measured by cell-phone mobility data. They find that the optimal government policy in the United States, taking externalities into account, is immediate—but not particularly restrictive—social distancing of long duration. Eichenbaum, Rebelo, and Trabandt (2020) develop another model in which the primary channel for distancing is to reduce consumption of social goods, which is restrictive as a model of distancing activities but creates clean connections to macroeconomic activity.

Goolsbee and Syverson (2020) study endogenous social distancing from an empirical perspective. They provide an estimate of how important endogenous individual actions are relative to government policies designed to lower R_t . Using county-level mobility data in a border discontinuity design, they find that of the 60 percent decrease in US activity observed, only about 7 percentage points can be explained by government regulations across different states and municipalities. Their research suggests that ignoring endogeneity in these models could be problematic and could, in particular, lead researchers to mistakenly attribute effects on disease dynamics to government policies. Chernozhukov, Kasaha, and Schrimpf (2020) find substantial causal effects of government policies using a more sophisticated dynamic model of consumer choices, while still finding that providing information on risks is also quite important.

The endogeneity of R_0^t has also been recognized and addressed by epidemiologists. Reluga (2010) is most similar to how some economists have set up the problem—it develops a differential game version of the SIR model in which agents can, at each instant, take a costly social distancing action that reduces their instantaneous probability of infection. It computes equilibria for several sets of parameter values covering scenarios in which the disease spreads at different rates and a vaccine is closer or farther off, and compares equilibrium payoffs to the social optimum.

Reluga (2010) also provides references to earlier literature, much of which is less utility-focused. A recent example of work of this style is Eksin, Paarporn, and Weitz (2019), which discusses variants of the SIR model that make how people distance in response to current or cumulative cases as a primitive (instead of deriving this from a utility function) and notes that distancing could make the long-run fraction infected much lower than would be predicted by an SIR model calibrated in early stages of the epidemic. While economists' first inclination will be to regard it as a drawback that distancing behavior is a primitive rather than derived from dynamic optimization given an assumed utility function, a skeptic could easily note that there is quite limited evidence on the utility-consistency of the ways in which people socially distance over the course of an epidemic, and that models with utility functions calibrated to rationalize how people have distanced in past epidemics may not provide better predictions than would models in which behavior itself is calibrated to behavior in past epidemics.

Heterogeneity

In many branches of economics, it has become standard to incorporate heterogeneous consumer preferences and/or firm profit functions. Given this norm, it is not surprising that economists are also increasingly incorporating heterogeneity into their COVID-related work.

One of the most striking features of COVID-19 is how fatality rates vary with age. The calibrations in Ferguson et al. (2020), for example, assume an infection fatality rate of 9.3 percent for those over 80, 2.2 percent for those 60–69, 0.15 percent for those 40–49, and 0.03 percent for those 20–29. Economic activities also vary with age, of course. Therefore, it is natural to assess the potentially disparate impact that policies may have on different age groups, consider explicitly age-varying policies, or both.

Several recent papers use calibrated multi-population SIR models where subpopulations are interpreted as age groups to discuss the economic and health consequences of lockdown and reopening policies. Rampini (2020) considers a two-population model calibrated to reflect those under and over age 55 and notes that a two-phase reopening in which the young are released before the old can reduce hospital overcrowding, mortality, and economic losses. Favero, Ichino, and Rustichini (2020) and Baqaee et al. (2020) make finer distinctions of subpopulations. The former considers a 15-population model corresponding to subsets defined by five age groups and three occupation types. The latter uses a five population model corresponding to age groups but calibrates interactions between age groups using contact survey data,⁸ data on activity differences across occupations, and industry-specific worker age distributions. In other words, they take an estimate of the average R_0 from the epidemiology literature and choose a matrix of subgroup-to-subgroup infection rates that is consistent both with that R_0 and with

⁸ "Contact surveys" are distinct from contact tracing. The former simply obtains data on typical daily contacts of randomly-selected people, both within and across various subgroups.

the differences across groups in the contact surveys and mobility data. The results of Baqaee et al. (2020) are sobering: even slow reopening policies that prioritize industries on a GDP-to-risk basis tend to produce conditions that require subsequent reversals of policy with new shutdowns if individuals relax their levels of social distancing. Acemoglu et al. (2020) analyze a much broader class of time- and age-varying policies and provide estimates of the Pareto frontier of optimal policies that minimize economic losses and deaths. They note that age-dependent policies can provide substantial gains relative to uniform policies, with the greatest improvement coming from doing as much as one can to protect those in the oldest group when prevalence is high among those in younger age groups.

Ellison (2020) builds on models in the epidemiology literature that take a broader view of heterogeneity—reflecting that those who ride public transportation or frequent bars will have many more contacts than others in their age group, for example—and discusses their implications for an analysis of COVID-19. Jackson and López-Pintado (2013) is an example within economics. One cautionary observation is that these models have more parameters that need to be calibrated, and long run outcomes can be sensitive to activity levels of the less active, particularly when we are considering relaxing restrictions. It is difficult to calibrate these parameters early in an epidemic, and predictions that do not allow for heterogeneity may be overconfident.

Ellison (2020) also notes that conclusions drawn from applying homogeneous SIR models to a world that is more like a heterogeneous SIR model would be biased in a number of ways. As noted earlier, homogeneous SIR models may substantially overstate the fraction of the population that must be infected in order to achieve herd immunity. A related observation is that (targeted) lockdown policies can also be more cost-effective in heterogeneous populations. There can be substantial gains either from taking permanent measures to reduce spread among the highly active or from temporarily locking down less active groups to minimize overshooting of herd immunity thresholds. We look forward to seeing such heterogeneities incorporated into more policy analyses.⁹

Political Economy

An extraordinary characteristic of this health crisis in the United States is the degree to which it has been politicized, even to the extent that simple precautions like wearing a mask have become freighted with political meaning. Evidence suggests that social distancing and mask-wearing are very important weapons in combating COVID-19 (Abaluck et al. 2020; Chernozhukov, Kasaha, and Schrimpf 2020), so understanding political obstacles to improving, or simply variation in, these behaviors is quite important. A trio of papers attempt to address this issue by looking specifically at the role of the media. They have found evidence of correlation or

⁹Given the substantial fraction of deaths which have occurred in nursing homes, one such extension that seems very natural would be to incorporate a nursing home sector. This would allow one to model impacts of policies like those discussed in Chen, Chevalier, and Long (2020).

causal effects of media consumption on knowledge about COVID-19 and behavior regarding it. Jamieson and Albarracín (2020) find that, controlling for party affiliation and other demographics, use of conservative media was associated with significantly lower levels of knowledge about the virus and the disease characteristics associated with it. Simonov et al. (2020) exploit quasi-random assignments of channel positions in a cable lineup to estimate the effect of full Fox News viewership on non-compliance with stay-at-home orders, finding an increase of 12–25 percent noncompliance. Finally, Bursztyjn et al. (2020), also interested in the effect of Fox viewership, exploit a different instrument, the broadcast time of Hannity and Tucker Carlson Tonight relative to sunset in a particular location. They document a much different tone to the COVID-related content on the two shows early in the epidemic and find that areas with greater exposure to Hannity—more dismissive of the risks—experienced significantly more cases and deaths.

Barrios and Hochberg (2020) use data on internet searches to document that Republican-dominated areas perceive less risk from the virus than do Democratic-dominated areas. Finally, Ajzenman, Cavalcanti, and Da Mata (2020) find similar political effects in Brazil, another country struggling with high caseloads and deaths and with a president dismissive of the severity of the pandemic. They find differential effects on behavior following presidential speeches disparaging social distancing, based on the level of political support for the president by location. Additional papers documenting the political divide and its effects on behavior and health outcomes during the pandemic are cited in these papers as well.

Although none of these papers use epidemiological models or methods, their estimates are useful for understanding how the parameters in the epidemiological models might vary over time and by geographic location. In fact, their specifications and results suggest ways in which R_0^t might be parameterized in an empirical model with a variety in covariates.

Conclusion

A symbiotic relationship between academic research and government policy-making existed long before the spring of 2020. Many researchers aim to produce research that is topical, useful, and policy-relevant. In turn, policymakers seek out expert advice and prediction, often in the form of theoretical or empirical models. Our current crisis, however, has put the structure and the mechanics of this relationship in stark relief.

We think that it is important to draw a distinction between two roles that models have served during the pandemic. Models can help us predict, and they can help us understand, and policy-makers have demanded both types. For instance, models can help us predict timing and magnitude of infections and hospitalizations as well as the need for equipment and other resources. The ability to generate detailed predictions for specific localities is important, especially for local decision-makers who have to set policy and allocate resources. Ultimately, though, the test

of the usefulness of these models is typically empirical in nature, whether that be using retrospective data to judge various models after the fact or using previous and contemporary data from similar settings. The opacity of such models may not be entirely unimportant, but it could be second-order: as long as a “black box” works, we may not care what is in it.

Alternatively, models can help us understand. They can help us understand, for instance, an important interaction of factors, or a mechanism which can indirectly affect the spread of a disease. Such models need not be able to generate location- and day-specific predictions of the number of hospital beds needed, but they are no less important in informing policy-making and resource allocation in different ways.

Understanding the process by which these models’ predictions and insights can be accessed by policymakers has also gained importance. The normal process of writing, vetting, and publishing scientific and economic research is being stretched to its limits given the urgency of the pandemic. Direct and wide dissemination can work for certain types of knowledge: detailed predictions from empirical models lend themselves to the now ubiquitous COVID “dashboards” that make those predictions available to policy-makers and others with just a click or two. There is no reason to believe that the models which have the best designed websites and interfaces are the ones producing the most careful and accurate predictions, though. Conveying more subtle insights, such as how government policies might interact with endogenous social distancing, seems substantially more difficult but no less important. One would hope that robust lines of communication and established respectful relationships between experts and policy-makers could facilitate such dialogues.

We wrote this paper in hopes of spurring interesting and important research by economists on epidemics and COVID-19, in particular. If this extraordinary period in time also spurs a rethinking of the complicated relationship between research and policy-making, the dialog between experts and non-experts, and the practical uses of both theoretical and empirical modeling, we will all reap the benefits.

■ *We are grateful to Marcy Alsan, Ben Bolker, Amitabh Chandra, Bill Clark, Jonathan Dushoff, Michael Kremer, Nolan Miller, Ziad Obermeyer, Elizabeth Rourke, Bruce Sacerdote, Doug Staiger, Jim Stock, and Richard Zeckhauser for useful conversations and Eva Demsky for outstanding research assistance.*

References

- Abaluck, Jason, Judith A. Chevalier, Nicholas A. Christakis, Howard Paul Forman, Edward H. Kaplan, Albert Ko, and Sten H. Vermund. 2020. "The Case for Universal Cloth Mask Adoption and Policies to Increase the Supply of Medical Masks for Health Workers." *Covid Economics* 5 (2020): 147–59.
- Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston. 2020. "Optimal Targeted Lockdowns in a Multi-Group SIR Model." NBER Working Paper 27102.
- Ajzenman, Nicolas, Tiago Cavalcanti, and Daniel Da Mata. 2020. "More Than Words: Leaders Speech and Risky Behavior during a Pandemic." Unpublished.
- Angrist, Joshua, Pierre Azoulay, Glenn Ellison, Ryan Hill, and Susan Feng Lu. 2017. "Economics Research Evolves: Fields and Styles." *American Economic Review Papers and Proceedings* 107 (5): 293–97.
- Avery, Christopher, William Bossert, Adam Thomas Clark, Glenn Ellison, and Sara Fisher Ellison. 2020. "Policy Implications of Models of the Spread of Coronavirus: Perspectives and Opportunities for Economists." *Covid Economics* 12: 21–68.
- Baqee, David, Emmanuel Farhi, Michael J. Mina, and James H. Stock. 2020. "Reopening Scenarios." NBER Working Paper 27244.
- Barrios, John M., and Yael V. Hochberg. 2020. "Risk Perception through the Lens of Politics in the Time of the COVID-19 Pandemic." Becker Friedman Institute Working Paper 2020-32.
- Bernoulli, Daniel. 1766. "Essai d'une nouvelle analyse de la mortalité causée par la petite vérole." In *Mémoires de Mathématique et de Physique*, edited by Académie royale sciences.
- Blackwood, Julie C., and Lauren M. Childs. 2018. "An Introduction to Compartmental Modeling for the Budding Infectious Disease Modeler." *Letters in Biomathematics* 5 (1): 195–221.
- Bolker, Benjamin M. 1999. "Analytic Models for the Patchy Spread of Plant Disease." *Bulletin of Mathematical Biology* 61 (849).
- Britton, T. 1998. "Estimation in Multitype Epidemics." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 60 (4): 663–79.
- Britton, Tom, Frank Ball, and Pieter Trapman. 2020. "A Mathematical Model Reveals the Influence of Population Heterogeneity on Herd Immunity to SARS-CoV-2." *Science* 369 (6505): 846–49.
- Bursztyn, Leonardo, Aakaash Rao, Christopher P. Roth, and David H. Yanagizawa-Drott. 2020. "Misinformation during a Pandemic." NBER Working Paper 27417.
- Champredon, David, Michael Li, Benjamin M. Bolker, and Jonathan Dushoff. 2018. "Two Approaches to Forecast Ebola Synthetic Epidemics." *Epidemics* 22: 36–42.
- Chawla, Dalmeet Singh. 2020. "Influential Pandemic Simulation Verified by Code Checkers." *Nature*, June 18. <https://media.nature.com/original/magazine-assets/d41586-020-01685-y/d41586-020-01685-y.pdf>.
- Chen, M. Keith, Judith A. Chevalier, and Elisa F. Long. 2020. "Nursing Home Staff Networks and COVID-19." NBER Working Paper 27608.
- Chernozhukov, Victor, Hiroyuki Kasaha, and Paul Schrimpf. 2020. "Causal Impact of Masks, Policies, Behavior on Early Covid-19 Pandemic in the US." *arXiv*.
- Demiris, Nikolaos, and Philip D. O'Neill. 2005. "Bayesian Inference for Stochastic Multitype Epidemics in Structured Populations via Random Graphs." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67 (5): 731–45.
- Diekmann, Odo, Johan Andre Peter Heesterbeek, and Johan A.J. Metz. 1990. "On the Definition and the Computation of the Basic Reproduction Ratio R_0 in Models for Infectious Diseases in Heterogeneous Populations." *Journal of Mathematical Biology* 28 (4): 365–82.
- Dushoff, Jonathan, and Simon Levin. 1995. "The Effects of Population Heterogeneity on Disease Invasion." *Mathematical Biosciences* 12 (1–2): 25–40.
- Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt. 2020. "The Macroeconomics of Epidemics." NBER Working Paper 26882.
- Eksin, Ceyhan, Keith Paarporn, and Joshua S. Weitz. 2019. "Systematic Biases in Disease Forecasting—The Role of Behavior Change." *Epidemiology* 27: 96–105.
- Ellison, Glenn. 2020. "Implications of Heterogeneous SIR Models for Analyses of COVID-19." NBER Working Paper 27373.
- Farboodi, Maryam, Gregor Jarosch, and Robert Shimer. 2020. "Internal and External Effects of Social Distancing in a Pandemic." NBER Working Paper 27059.

- Farr, William. 1840. "Causes of Death in England and Wales." In *Second Annual Report of the Registrar General of Births, Deaths and Marriages in England*, 100–53. London: W. Clowes and Sons.
- Favero, Carlo A., Andrea Ichino, and Aldo Rustichini. 2020. "Restarting the Economy while Saving Lives under Covid-19." CEPR Discussion Paper DP14664.
- Ferguson, Neil M., Derek A.T. Cummings, Christophe Fraser, James C. Cajka, Philip C. Cooley, and Donald S. Burke. 2006. "Strategies for Mitigating and Influenza Pandemic." *Nature* 442: 448–52.
- Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, et al. 2020. *Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand*. Swindon, United Kingdom: Medical Research Council: The Royal Society.
- Ferretti, Luca, Chris Wymant, Michelle Kendall, Lele Zhao, Anel Nurtay, Lucie Abeler-Dörner, Michael Parker, David Bonsall, and Christophe Fraser. 2020. "Quantifying SARS-CoV-2 Transmission Suggests Epidemic Control with Digital Contact Tracing." *Science* 368 (6491).
- Gomes, M. Gabriela M., Rodrigo M. Corder, Jessica G. King, Kate E. Langwig, Caetano Souto-Maior, Jorge Carneiro, Guilherme Gonçalves, Carlos Penha-Goncalves, Marcelo U. Ferreira and Ricardo Aguas. 2020. "Individual Variation in Susceptibility or Exposure to SARS-CoV-2 Lowers the Herd Immunity Threshold." <https://doi.org/10.1101/2020.04.27.20081893>.
- Goolsbee, Austan, and Chad Syverson. 2020. "Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020." Becker-Friedman Working Paper 2020-80.
- Gotelli, Nicholas J. 2008. *A Primer on Ecology*. Oxford: Oxford University Press.
- Halloran, M. Elizabeth, and Claudio J. Struchiner. 1995. "Causal Inference in Infectious Diseases." *Epidemiology* 6 (2): 142–51.
- Hébert-Dufresne, Laurent, Benjamin M. Althouse, Samuel V. Scarpino, and Antoine Allard. 2020. "Beyond R0: Heterogeneity in Secondary Infections and Probabilistic Epidemic Forecasting." <https://doi.org/10.1101/2020.02.10.20021725>.
- Hernán, Miguel A., and James M. Robins. 2006. "Instruments for Causal Inference: An Epidemiologists Dream?" *Epidemiology* 17 (4): 360–72.
- Hernán, Miguel A., and James M. Robins. 2020. *Causal Inference: What If*. Boca Raton, FL: CRC Press.
- Hethcote, Herbert W. 2000. "The Mathematics of Infectious Diseases." *SIAM Review* 42 (4): 599–653.
- Jackson, Matthew O., and Dunia López-Pintado. 2013. "Diffusion and Contagion in Networks with Heterogeneous Agents and Homophily." *Network Science* 1 (1): 49–67.
- Jamieson, Kathleen Hall, and Dolores Albarracín. 2020. "The Relation between Media Consumption and Misinformation at the Outset of the SARS-CoV-2 Pandemic in the US." *The Harvard Kennedy School Misinformation Review* <https://doi.org/10.37016/mr-2020-012>.
- Keith, Tamara. 2020. "Timeline: What Trump Has Said and Done about the Coronavirus." *NPR*, April 21. <https://www.npr.org/2020/4/21/837348551/timeline-what-trump-has-said-and-done-about-the-coronavirus>.
- Keppo, Jussi, Elena Quercioli, Kudlyak, Marianna, Lones Smith, and Andrea Wilson. 2020. "For Whom the Bell Tolls: Avoidance Behavior at Breakout in COVID19." Virtual Macro Seminar, 1:25:07. April 2020.
- Kermack, William Ogilvy, and Anderson G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society of London A* 115 (772): 700–21.
- Kissler, Stephen M., Christine Tedijanto, Edward Goldstein, Yonatan H. Grad, and Marc Lipsitch. 2020. "Projecting the Transmission Dynamics of SARS-CoV-2 through the Postpandemic Period." *Science* 368 (6493): 860–68.
- Korber, Bette, Will M. Fischer, Sandrasegaram Gnanakaran, Hyejin Yoon, James Theiler, Werner Abfalterer, Nick Hengartner, et al. 2020. "Tracking Changes in the SARS-CoV-2 Spike: Evidence that D614G Increases Infectivity of the COVID-19 Virus." *Cell* 182 (4): 812–27.
- Kremer, Michael. 1996. "Integrating Behavioral Choice into Epidemiological Models of AIDS." *Quarterly Journal of Economics* 111 (2): 549–73.
- Lajmanovich, Ana, and James A. Yorke. 1976. "A Deterministic Model for Gonorrhoea in a Nonhomogeneous Population." *Mathematical Biosciences* 28 (3–4): 221–36.
- Landler, Mark, and Stephen Castle. 2020. "Behind the Virus Report that Jarred the U.S. and the U.K. to Action." *New York Times*, March 17. <https://www.nytimes.com/2020/03/17/world/europe/coronavirus-imperial-college-johnson.html>.
- Li, Michael Y., and James S. Muldowney. 1995. "Global Stability for the SEIR Model in Epidemiology." *Mathematical Biosciences* 125 (2): 155–64.

- Lloyd, Alun L., and Robert M. May. 1996. "Spatial Heterogeneity in Epidemic Models." *Journal of Theoretical Biology* 179 (1): 1–11.
- Massad, E., F.A.B. Coutinho, M.N. Burattini, and M. Amaku. 2010. "Estimation of R0 from the Initial Phase of an Outbreak of a Vector-Borne Infection." *Tropical medicine & International Health* 15 (1): 120–26.
- Melnick, J.L. 1990. "Poliomyelitis." In *Tropical and Geographical Medicine*, edited by Kenneth S. Warren and Adel A.F. Mahmoud, 558–76. New York: McGraw Hill.
- Miller, Danielle, Michael A. Martin, Noam Harel, Talia Kustin, Omer Tirosh, Moran Meir, Nadav Sorek, et al. 2020. "Full Genome Viral Sequences Inform Patterns of SARS-CoV-2 Spread into and within Israel." <https://www.medrxiv.org/content/10.1101/2020.05.21.20104521v1.full.pdf>.
- Mills, Christina E., James M. Robins, and Marc Lipsitch. 2004. "Transmissibility of 1918 Pandemic Influenza." *Nature* 432: 904–06.
- Philipson, Tomas J., and Richard A. Posner. 1993. *Private Choices and Public Health: The AIDS Epidemic in an Economic Perspective*. Cambridge, MA: Harvard University Press.
- Prem, Kiesha, Yang Liu, Timothy W. Russell, Adam J. Kucharski, Rosalind M. Eggo, Nicholas Davies, Stefan Flasche, et al. 2020. "The Effect of Control Strategies to Reduce Social Mixing on Epidemic in Wuhan, China: A Modelling Study." *The Lancet Public Health* 5 (5): 261–70.
- Rampini, Adriano A. 2020. "Sequential Lifting of Covid-19 Interventions with Population Heterogeneity." NBER Working Paper 27063.
- Reluga, Timothy C. 2010. "Game Theory of Social Distancing in Response to an Epidemic." *PLOS Computational Biology* 6 (5): 1–9.
- Rida, Wasima N. 1991. "Asymptotic Properties of Some Estimators for the Infection Rate in the General Stochastic Epidemic Model." *Journal of the Royal Statistical Society: Series B* 53 (1): 269–83.
- Simonov, Andrey, Szymon K. Sacher, Jean-Pierre H. Dubé, and Shirsho Biswas. 2020. "The Persuasive Effect of Fox News: Non-compliance with Social Distancing during the COVID-19 Pandemic." NBER Working Paper 27237.
- Stock, James H. 2020. "Data Gaps and the Policy Response to the Novel Coronavirus." NBER Working Paper 26902.
- Toxvaerd, Flavio. 2020. "Equilibrium Social Distancing." Cambridge-INET Working Paper Series 2020/08.
- Tuite, Ashleigh R., and David N. Fisman. 2018. "The IDEA Model: A Single Equation Approach to the Ebola Forecasting Challenge." *Epidemics* 22: 71–77.
- Unwin, H. Juliette T., Swapnil Mishra, Valerie C. Bradley, Axel Gandy, Thomas A. Mellan, Helen Coupland, Jonathan Ish-Horowicz, et al. 2020. "State-Level Tracking of COVID-19 in the United States." <https://www.medrxiv.org/content/10.1101/2020.07.13.20152355v1.full.pdf>.
- Viboud, Cécile, Kaiyuan Sun, Robert Gaffey, Marco Ajelli, Laura Fumanelli, Stefano Merler, Qian Zhang, Gerardo Chowell, Lone Simonsen, Alessandro Vespignani. 2018. "The RAPIDD Ebola Forecasting Challenge: Synthesis and Lessons Learnt." *Epidemics* 22: 13–21.
- Worobey, Michael, Jonathan Pekar, Brendan B. Larsen, Martha I. Nelson, Verity Hill, Jeffrey B. Joy, Andrew Rambaut, Marc A Suchard, Joel O. Wertheim, and Philippe Lemey. 2020. "The Emergence of SARS-CoV-2 in Europe and the US." Unpublished.
- Zhigljavsky, Anatoly, Ivan Fesenko, Henry Wynn, Kobi Kremnitzer, Jack Noonan, Jonathan Gillard, and Roger Whitaker. 2020. "A Prototype for Decision Support Tool to Help Decision-Makers with the Strategy of Handling the COVID-19 UK Epidemic." <https://doi.org/10.1101/2020.04.24.20077818>.