Policy Implications of Models of the Spread of Coronavirus: Perspectives and Opportunities for Economists

Christopher Avery, William Bossert, Adam Thomas Clark, Glenn Ellison,
Sara Fisher Ellison

CESifo

Policy Implications of Models of the Spread of Coronavirus: Perspectives and Opportunities for Economists

Abstract

This paper provides a critical review of models of the spread of the coronavirus (SARS-CoV-2) that have been influential in recent policy discussions. It notes potentially important features of the real- world environment that the standard models do not incorporate and discusses reasons why estimating critical parameters is difficult. These limitations may bias forecasts and lead forecasters to overstate confidence in their predictions. They also provide social scientists with opportunities to advance the literature and enable improved policies. This paper also discusses how optimal policies might depend on what is learned from new data and models.

JEL-Codes: C020, I100, I180.

Christopher Avery
Harvard Kennedy School
Cambridge / MA / USA
christopher_avery@hks.harvard.edu

William Bossert
Harvard University
Cambridge / MA / USA
bossert@seas.harvard.edu

Adam Thomas Clark
German Centre for Integrative Biodiversity
Research / Leipzig / Germany
adam.tclark@gmail.com

Glenn Ellison
Massachusetts Institute of Technology
Cambridge / MA / USA
gellison@mit.edu

Sara Ellison
Massachusetts Institute of Technology
Cambridge / MA / USA
sellison@mit.edu

I. Introduction

Predictive models have taken on newfound importance in response to the spread of the COVID-19 illness, and the SARS-CoV-2 virus that causes it. The fast pace of reports, predictions, and policy changes that have gripped much of the world since March have highlighted both the political relevance and the great uncertainty associated with real-time forecasting in complex systems. On Thursday March 12, Boris Johnson announced that the U.K. would not limit large gatherings or close schools because "The scientific advice is that this could do more harm than good at this time;"[1] Patrick Vallance, the chief science advisor to the government, described a containment plan based on herd immunity, noting that this would likely require 60% of the population to contract the virus (Robinson and Blanchard (2020)). Four days later, Johnson reversed course and stated "now is the time for everyone to stop non-essential contact with others"[2], coordinating his announcement with a related briefing on a new study from Imperial College to indicate that the results of that study had influenced the government's change in policy (Booth (2020)). The United States also changed its guidelines that same day, suggesting that gatherings be limited to ten people (after the CDC had suggested a different limit of 50 people the day before)[3], with Dr. Deborah Birx, Coronavirus Response Coordinator for the White House Coronavirus Task Force, also referencing the Imperial College report in support of the policy change: "What had the biggest impact in the model is social distancing, small groups, not going in public in large groups" (Fink (2020)). From a modeling perspective, Birx's comment is telling because those predictions about the relative impact of different distancing procedures are completely driven by *ad hoc* assumptions about how each affects contacts.

[1] https://www.gov.uk/government/speeches/pm-statement-on-coronavirus-12-march-2020
[2] https://www.gov.uk/government/speeches/pm-statement-on-coronavirus-16-march-2020
[3] There is some suggestion that the United States previously more relaxed policy had been influenced by a blog post by law professor Richard Epstein with back of the envelope calculations suggesting no more than 500 deaths due to the epidemic in the United States (Dawsey et al., (2020)). We discuss Epstein's estimate in Section V.

One perplexing element of these references to the Imperial College study was that government officials cited predictions that varied by surprising orders of magnitude. Donald Trump has emphasized the prediction of 500,000 deaths in the UK and 2.2 million deaths in the U.S. from uncontrolled spread of the epidemic, whereas Patrick Vallance focused on a prediction of 20,000 deaths in the UK from an extreme response of "suppression" (Ferguson et al., 2020) which requires substantive changes in behavior until the production of a vaccine (Johnston, 2020). On March 31, Deborah Birx reported that the White House consulted "five or six international or domestic modelers," and concluded that full mitigation would reduce total deaths in the United States to 100,000 to 200,000. At the same time, Anthony Fauci, Director of the National Institute of Allergy and Infectious Diseases, cast doubt on these predictions: noting that "… I know my modeling colleagues are going to not be happy with me, but models are as good as the assumptions you put into them. And as we get more data, then you put it in and that might change."[4]

At present, after several weeks of "social distancing," deaths that have been officially linked to SARS-CoV-2 in the UK and US have only amounted to a fraction of those projected in the most pessimistic scenarios for the Imperial College model, but are also already above the predictions of other widely publicized projections. Since it is still an early phase of the pandemic, there should be a role for well-founded models to provide information and guide critical policy decisions. This paper provides a critical review of the literature from several fields with the goal of providing a foundation for economists and other social scientists to conduct original research that may both save lives and bolster economic outcomes.

---

[4] https://www.whitehouse.gov/briefings-statements/remarks-president-trump-vice-president-pence-members-coronavirus-task-force-press-briefing-15/

Economists have responded energetically and have produced a considerable amount of policy-oriented research related to the spread of the pandemic in a short period of time. Of most relevance to this paper, Atkeson (2020a) and Berger, Herkenhoff, and Mongey (2020) apply and propose extensions to the standard ("SIR") epidemiological model, while Stock (2020) argues that policy analysis is limited for the moment given the limitations of existing data. We discuss several other recent papers by economists in the text and conclusion.

The paper proceeds as follows. Section 2 outlines the types of models generally used in the ecology/epidemiology literature to study the spread of disease. Section 3 describes a number of heterogeneities that are important in practice but are not incorporated in the baseline versions of these models. Section 4 describes how existing data is frequently used to parameterize existing models. Section 5 provides detail of five estimates of the spread reported to have influenced government policy and critiques the models used to provide those estimates. Section 6 discusses outstanding policy questions and suggests pressing research questions for economists to take on. Section 7 concludes.

II. Types of Models and Their General Properties

There are two primary approaches for modeling the spread of disease: (1) "Mechanistic" and (2) "Phenomenological".[5] The distinction between these models is generally analogous to the distinction between structural and reduced form models in economics. Just as proponents of structural work tout the ability to extend their models to conduct counterfactual analysis, advocates for mechanistic approaches to disease modeling highlight the importance of out of sample predictions: *"All other things being equal, mechanistic models are more powerful since they*

---

[5] See Hurford (2012) for further discussion.

tell you about the underlying processes driving patterns. They are more likely to work correctly when

extrapolating beyond the observed conditions." (Bolker (2008), p. 7).

As a biological process, disease spread is usually much better characterized by models

that focus on rates of change, rather than models that focus on the states of individual variables

themselves. Thus, models usually end up with the same general components, e.g. representing

the rates at which new individuals become infected, at which infected individuals recover or die,

etc. The primary distinction between phenomenological and mechanistic models in

epidemiology, therefore, tends to be more directly related to how models have been

parameterized than on the functional forms themselves. Models fit based on a priori biological

assumptions, or boots-on-the-ground efforts to identify infected individuals and trace their

contacts and resulting infections, tend to be labeled as mechanistic. In this sense, mechanistic

models in this literature may be seen as analogous to macroeconomic models that fit some

parameters to data and then calibrate other parameters to match external evidence. Models that

are parameterized through curve-fitting based on reported case or mortality data, tend to be

labeled as phenomenological.

The most common approach that has been used to model the spread of SARS-CoV-2 is

the "Susceptible / Infectious / Recovered" (SIR) model.[6] In essence, SIR models can be viewed

as continuous-time Markov chain models where only a limited number of transitions between

states are possible. In its most basic form, the SIR model considers dynamics within a single

homogeneous population (i.e. what physicists would call a "mean-field model"), following the

form

---

[6] Kermack and McKendrick published seminal articles studying this model from 1927 to 1933. These articles were republished in 1991 (Kermack and McKendrick (1991a), (1991b), (1991c)). Sattenspiel (1990) provides a detailed discussion of the history of models of the spread of infectious disease.

$$\mathrm{d}S/\mathrm{d}t \quad = \quad -\beta IS/N$$

$$\mathrm{d}I/\mathrm{d}t \quad = \quad \beta IS/N - \gamma I$$

$$\mathrm{d}R/\mathrm{d}t \quad = \quad \gamma I$$

where S, I, and R are the abundance of the three model states, β combines information about encounter rates and infectivity, and γ describes the combined rates of recovery and mortality from the disease.[7]

In this model, the disease will increase in incidence until $\beta IS/N = \gamma I$ (i.e. $S/N = \gamma/\beta$), at which point new infections can no longer keep up with the recovery rate (n.b. this point will always be reached eventually, since recovery results from a linear rate process that is directly proportional to I, whereas infection rates are non-linear and depend jointly on I and S). Thereafter, the number of new infections will begin to drop, and the disease will die out. This threshold (often referred to as the "*herd immunity*" threshold) describes the minimum size of the susceptible pool required for the infection to spread through a well-mixed population. Importantly, individuals continue to become infected (especially if the infected pool is large), but on average, the number of new infections per time step will begin to decrease once the herd immunity threshold has been reached.

One potentially confusing aspect of the basic SIR model is that commonly used methods for identifying and analyzing point equilibria are not necessarily indicative of states that will be realized as time goes to infinity. This is because of the "inertia" in the model caused by existing infections. Although the rate at which new individuals become infected drops after the herd immunity threshold is reached (and, recall that this threshold represents the "equilibrium" where

---

[7] Chapter 15 of Lehman, Loberg, and Clark (2019) provides a textbook description of the SIR and other models.

dI/dt = 0 and dS/dt = -dR/dt), the fraction of the population that will ultimately be infected depends on a number of other factors, including the current number of infections, the rate at which new individuals are infected, and the recovery rate. There are some specialized forms of the SIR model for which such equilibria have been derived (Harko, Lob, and Mak (2014)).

Note that individuals in the SIR model are assumed to be infinitely divisible (i.e. there is no "minimum" step size for transitions among states), and that all state transitions are modeled as exponential processes which can be characterized by their half-lives. Additional compartments can be used to incorporate time-lags into the model. For example, since the incubation period for SARS-CoV-2 is estimated to be about five to six days, the extension of SIR to SEIR to allow for a fourth state ("Exposed") has been especially common for studying SARS-CoV-2 (Hethcote & Driessche (1991) and Li & Muldowney (1995)), although many other variants of the model exist (e.g. "SAIR" models, which include an additional "Asymptomatic" infected state). Collectively, these models are known among epidemiologists as "compartmental models".[8]

The structure of transfer rates and states, and of resulting model predictions can vary enormously even within a specific type of compartmental model. Some factors that differ across SIR models include modeling of frequency vs. density-dependent transmission (Rudolf & Antonovics (2005)), separation of β into terms describing contacts vs. infection events (Blackwood & Childs (2018)), separation of γ and the "Recovered" bin to distinguish between recoveries and deaths (Gallos & Fefferman (2015)), inclusion of geographic (Bolker (1999)), and demographic (Hethcote (2000)) structure, the application of reproduction and death rates across population bins (Harko et al. (2014)), and loss of immunity (Hethcote & Driessche (1991)). Although these many variants of the model allow it to capture different kinds of biological

---

[8] See Brauer (2008) and in Blackwood & Childs (2018) for general reviews of this term and associated concepts.

nuance, it can also inhibit comparisons of predictions across models, e.g. if parameters describe fundamentally different processes or relate to very different types of individuals.

Bolker has compiled a list of well over 40 open source models of the SARS-CoV-2 epidemic.[9] These models are mostly SIR models with considerable overlap in mathematical structure. Some primary differences among them involve: (i) the number and identity of compartments in the model; (ii) whether time-steps are continuous or discrete; (iii) whether dynamics are entirely deterministic or include stochastic aspects; (iv) how any incorporated stochasticity is introduced into the model (e.g. discrete time-steps, Gillespie's method, tau-leaping)[10]; (v) whether methods are employed to quantify or reduce observation error; (vi) the inclusion of demographic and/or spatial structure in the model; and (vii) the types of scenarios and outcomes that they are able to consider (e.g. hospital capacity, number of ICU beds, social distancing strategies, etc.). Jointly, these models and methods provide a summary of the "state of the art" methods currently available for simulating and parameterizing SIR-like models.

There is a distinct literature that models the spread of disease in a network context. These models typically take bilateral links between people as exogenous primitives and study the probabilistic spread of infection from an initial set of infected people through the rest of the network. It is possible to solve for steady state levels of disease in some versions of a network model with SIR structure, where this steady state level varies with the geometry of the network and the underlying properties of the infection.[11] The epidemiological models that we discuss in Section 5 do not seem to make use of network techniques.

---

[9] https://docs.google.com/spreadsheets/d/1hUZlVDPfa5C8KgURoP_3dAiUQgI6rdb7A5e_g8NcPaY/edit#gid=0)
[10] See Wilkinson (2011) and Gillespie's notes (https://www.slideshare.net/csgillespie/the-27257701) for details.
[11] See Acemoglu and Ozdaglar (2009) and Easley and Kleinberg (ch. 21, 2010) for class and textbook summaries of the SIR model in a network context.

A notable shortcoming of the basic SIR model is that it does not allow for heterogeneity in state frequencies and rate constants.[12] We discuss several different sources of heterogeneity in more detail in Section 2.

The most important and challenging heterogeneity in practice is that individual behavior varies over time. In particular, the spread of disease likely induces individuals to make private decisions to limit contacts with other people. Thus, estimates from scenarios that assume unchecked exponential spread of disease, such as the reported figures from the Imperial College model of 500,000 deaths in the UK and 2.2 million in the United States, do not correspond to the behavioral responses one expects in practice. Further, these gradual increases in "*social-distancing*" that can be expected over the courses of an epidemic change dynamics in a continuous fashion and thus blur the distinctions between mechanistic and phenomenological models.[13] Each type of model can be reasonably well calibrated to an initial period of spread of disease, but further assumptions, often necessarily ad hoc in nature, are needed to extend either type of model to later phases of an epidemic.

In addition to problems related to heterogeneity, there are three additional challenging aspects of SIR models that make their dynamics especially difficult to predict. First, because the model is nonlinear, small changes in parameter values and initial states have large effects on dynamics. For example, given the reported doubling-times of SARS-CoV-2 early in the epidemic, comparing two otherwise identical regions that begin with 1 vs. 100 infected individuals can lead to more than a three-week lag-time in case numbers. This challenge is especially pernicious because of the high uncertainty that is usually associated with all of these

---

[12] Murray et al. (2020) specifically argues that a phenomenological approach has an advantage over SEIR models that base transition probabilities on an assumption of "random mixing".

[13] See Wood (2001) for an example of a "semi-mechanistic" model that incorporates phenomenological components in an otherwise mechanistic model.

values, meaning that predictions that fully incorporate uncertainty often span many orders of magnitude. Second, because dynamics in these models tends to be both complex and non-monotonic, classic model diagnostics and fitting tools may not be good indicators of whether a model will produce good extrapolations. For example, many models have high predictive ability when fit to the early stages of an epidemic, where growth in the number of infected individuals is approximately exponential. However, comparatively few models are able to predict the saturation point at which the number of new infections begins to decline, much less the expected number of infected individuals at the peak of the epidemic. Third, these problems are compounded by the fact that disease transmission involves substantial time lags. Not only can these lags confound models that fail to incorporate them correctly, but they also make it difficult to identify the effects of interventions on disease spread, since changes in observed case numbers lag infection events by at least several days.

Sadly, all three of these challenges seem to be particularly acute for SARS-CoV-2. Uncertainty in rates and states are especially high for the current pandemic, both because forecasts must be made in real-time, and because there are unprecedented stresses at a global scale on the institutions that are typically responsible for collecting and reporting data. Due to the unusually high asymptomatic rate of infection, it is also difficult to distinguish changes in reported case numbers that are associated with actual infections vs. reporting biases, so nonlinearities and saturation points are particularly difficult to identify. Because SARS-CoV-2 appears to have a particularly long latent stage in comparison to similar diseases, lags between policy interventions and observable changes in dynamics have often taken a week or longer.

III. Heterogeneities and Modeling

Several sources of heterogeneity complicate any modeling approach to SARS-CoV-2; it would be natural to attempt to incorporate one or more of them in future models.

Heterogeneous Exposure

The standard SIR model assumes a single transmission rate for the entire population, but systematic differences in the routine of daily life may result in different patterns of social interactions. Some of these differences are predictable functions of population density, and, thus, it is natural to expect different levels of social contact in urban than in rural areas. But daily routines may also be quite different across major cities: more than half of daily commuters take public transit in New York City, whereas only ten percent do so in Los Angeles.[14]

Household and building structure may also promote or inhibit the spread of disease. Nursing homes and long-term care facilities may be particularly susceptible because of the heavy degree of interactions between residents and staff that take place on a moment-to-moment basis. Some early statistics are consistent with the hypothesis that residents in these facilities are at unusual risk of infection: as of April 14, residents and workers at long-term care facilities accounted for 13.8% of all positive tests (3,907 of 28,163) in Massachusetts.[15] Similarly, 36% of residents at Pine Street Inn, a homeless shelter in Boston, tested positive for SARS-CoV-2.[16]

It also appears that some types of social contacts may be unusually likely to transmit SARS-CoV-2. A single "Beer Pong" party where participants shared drink glasses at an Austrian ski resort is credited with producing hundreds of infections in Denmark, Germany, and Norway (Hruby, 2020). A soccer game with record attendance may help to explain why

---

[14] https://en.wikipedia.org/wiki/List_of_U.S._cities_with_high_transit_ridership
[15] https://www.mass.gov/doc/covid-19-cases-in-massachusetts-as-of-april-14-2020/
[16] https://www.wbur.org/commonhealth/2020/04/14/coronavirus-boston-homeless-testing

Bergamo is an epicenter of the pandemic, perhaps exacerbated by the outcome of the game, as fans hugged and kissed each time Atalanta (the team from Bergamo) scored. Atalanta won the game 4 to 1 (Azzoni and Dampf (2020)). A March 10[th] choir practice in Washington State with 60 attendees resulted in 45 infections and two deaths.[17] The phenomenon of "superspreaders" is well known in the epidemiology literature: as Stein (2011) summarizes, "The minority of individuals who infect disproportionately more susceptible contacts, as compared to most individuals who infect few or no others, became known as super-spreaders, and their existence is deeply rooted in history." (See Hethcote and Van Ark (1987) for a model that incorporates some of these heterogeneities.)

Early studies of the spread of SARS-CoV-2 highlight the importance of networks and the advantages of flexible jobs that enable some people to work from home. Geographic concentrations of infections in New York and Italy can be traced to the strength of network ties to early cases (Kuchler et al. (2020), see also Borjas (2020)), while people have taken greatest steps for social distancing in regions with relatively high median incomes and internet access (Chiou and Tucker, 2020). These observations suggest that there may be systematically different trajectories for the spread of disease across incomes and regions, and perhaps even across regions according to political predilections (Allcott et al. (2020)). Case counts to date suggest possible differences for infection rates by race and ethnicity (Kendi (2020)) but testing rates may vary across these groups and data for the race of infected individuals is very spotty.[18]

Heterogeneous response across people

---

[17] https://www.cnn.com/2020/04/01/us/washington-choir-practice-coronavirus-deaths/index.html

[18] In Massachusetts, for example, "Hispanic" (22.4%) and "Non-Hispanic Black/African American" (15.7%) groups are disproportionately represented among people with a positive test and a reported category for Race/Ethnicity but Race/Ethnicity is listed as "Unknown" for 45% and as "Missing" for 18% of people with positive tests. (https://www.mass.gov/doc/covid-19-cases-in-massachusetts-as-of-april-14-2020)

Observational data suggests strong correlations between age, sex, existing medical conditions, and risk from SARS-CoV-2 infection. On March 26, the Center for Disease Control (CDC) reported that people age 65 and over accounted for approximately half of hospitalizations and ICU admissions and 80% of deaths and those age 85 and over were particularly affected.[19] The standard SIR model assumes equal mortality rates for all infected people; the SIR models that we discuss typically allow for differential risks by age, but often do not incorporate other factors that appear to be strongly correlated with risk.

Though the United States is not tracking deaths by gender (Gupta (2020)), the observed mortality rate has been at least 50% higher for men than for women in 12 of the 14 countries for which data is available.[20] In addition, the CDC warns that "people of any age who have *serious underlying medical conditions* might be at higher risk for severe illness from COVID-19."[21] These guidelines are consistent with retrospective analysis of multi-center data from Wuhan, China, which found that hypertension, diabetes, and heat disease were the most common comorbidities and that Sequential Organ Failure Assessment (SOFA) and d-dimer scores were significantly associated with death for hospitalized patients (Zhou et al. (2020)). The CDC includes "severe obesity" (BMI $\geq$ 40) as an underlying medical condition that increases risk in this case. David Ludwig and Richard Malley suggest that the high rate of adult obesity may increase the mortality rate from SARS-CoV-2 in the United States, as "Only 12 percent of Americans over age 20 are considered metabolically healthy." (Ludwig and Malley (2020)). Anthony Fauci noted that "comorbidities … are, unfortunately, disproportionately prevalent in the African American population," suggesting another reason that there might be variations in

---

[19] https://www.cdc.gov/mmwr/volumes/69/wr/mm6912e2.htm
[20] http://globalhealth5050.org/covid19/, retrieved April 7, 2020.
[21] https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/groups-at-higher-risk.html

outcomes by race and ethnicity.[22] It may be possible to account for these variations in risk

across cities, countries, and subgroups by post-stratifying the final size of the epidemic by the

estimated proportions falling within specific risk groups, but that would require information on

the interactions between underlying medical conditions and mortality rates from SARS-CoV-2

that is not yet available.

Medical Capacity by Location

The standard SIR model also assumes consistent medical response independent of one's

geographic location. Yet, there may be important variation in medical capacity across locations.

The United States ranks 32nd in the world with a bit less than 3 hospital beds per thousand

people. By comparison, South Korea, Japan, and Germany each have at least 8 hospital beds per

thousand people, though the United States has greater per-capita ICU capacity than any of those

countries.[23] Hospital capacity may be of most importance at moments of peak local infection.

To avoid the strain of going overcapacity in some places, France (Chang (2020)) and Germany

(Ellyat (2020)) have been moving patients to under-utilized hospitals. By contrast, in the United

States, there has only been some movement of medical personnel to New York City and other

hot spots of disease.

Dosage and the Nature of Exposure

The standard SIR model does not adjust the risk of infection or mortality rates conditional

on infection to account for the dose of exposure, or more generally, for the nature of social

contact. Some researchers have conjectured that exposure to a higher "viral load" can result in

---

[22] https://www.whitehouse.gov/briefings-statements/remarks-president-trump-vice-president-pence-members-coronavirus-task-force-press-briefing-april-7-2020/

[23] https://en.wikipedia.org/wiki/List_of_countries_by_hospital_beds

more severe illness, perhaps motivated in part by the death of Li Wenliang, the 33 year old

ophthalmologist who worked to publicize the disease in December, 2019. As American doctors

Rabinowitz and Bartman comment, "Dose sensitivity has been observed for every common acute

viral infection that has been studied in lab animals, including coronaviruses". (Rabinowitz and

Bartman (2020)). Early evidence from China and Italy only finds a strong association between

viral load and the likelihood of infecting others, not an association between viral load and the

outcome for the person with that viral load. (See Cereda et al. (2020), He et al. (2020), Geddes

(2020), and Heneghan, Brassey, and Jefferson (2020) for discussion of this point.)

Multiple strains of virus

At present, eight strains of SARS-CoV-2 have been identified by researchers, though it is

not clear how many other strains might exist. These original eight strains are sufficiently similar

that it seems that the virus is not mutating quickly enough to produce a version that is more

deadly to humans (Weise (2020)). While anti-viral immunity can be long lasting, it is not certain

that a person who has been exposed to one strain of SARS-CoV-2 will be immune to exposure to

other strains.   Contrary to some statements in the popular media, it remains unclear whether less

virulent (i.e. less deadly) SARS-CoV-2 strains will emerge on time scales that are relevant for

this epidemic. First, although these sorts of changes are well-supported by theory[24], in practice

they seem to occur very slowly, especially for diseases with relatively low mortality rates.

Moreover, the same geographic heterogeneity that may work to slow the spread of the virus will

almost certainly slow the spread of less virulent strains, even if they have a considerable fitness

advantage over other strains.

IV. Data

---

[24] See Read (1994) and Lively (1999) for discussions of evolutionary models of virulence.

The World Health Organization (WHO) defines a pandemic as "the worldwide spread of a new disease".[25] Data collection in the early stages of a pandemic is both especially important and especially challenging because a pandemic typically involves a relatively novel disease, and by definition involves a globally occurring stress on public institutions which can make information gathering and dissemination difficult.

Early aggregate statistics for infections and hospitalizations in an epidemic are notoriously unreliable. Many regions will be slow to identify and test the first people who contract a new disease, and those who are infected might not be inclined to go to a hospital for treatment and might not be admitted given the seemingly common nature of their symptoms. By contrast, death rates have been seen as more reliable measures for tracking the initial spread of diseases such as Ebola and SARS.

Several factors limit the value of death rates for short-run analysis in modeling the spread of SARS-CoV-2. Deaths seem to lag exposures by at least two weeks, so any analysis based on death rates requires a substantial time lag and reduces one's ability to assess the effect of a policy change. There may also be an additional time lag before deaths are incorporated in administrative data; official revisions have resulted in upward adjustments of previously reported deaths by as much as 78% in the United Kingdom (Giles (2020)). Further, there is often ambiguity in the cause of death from SARS-CoV-2 because it interacts so frequently with an underlying condition and because pneumonia is frequently the proximate cause of death. There may even be systemic differences across countries - for instance, doctors in the United States have much more discretion in the choice of cause of death than doctors in the UK (Henriques, 2020). The difficulty of assessing death rates from historical data is underscored by the ongoing

---

[25] https://www.who.int/csr/disease/swineflu/frequently_asked_questions/pandemic/en/

debate about the effects of the 1918 Flu,[26] and similar uncertainty exists even for more recent widespread influenza outbreaks.

An alternative method for assessing death rates is to use year-over-year difference-in-difference comparisons of all deaths at a regional level, comparing areas with known infections to other areas, though that approach is often not possible in real time. Using this method, there is growing evidence that the number of deaths attributed to SARS-CoV-2 is a substantial underestimate of the actual effect of the disease.[27] The difference-in-difference method also likely includes the effect of negative externalities of the epidemic, such as the loss of resources for doctors to treat other medical conditions, in the death rates ascribed to SARS-CoV-2. While it is essential to account for negative externalities in cost-benefit analyses, they likely distort estimates of the spread of disease that are based on death rates in overwhelmed areas.

There is now an abundance of data at a local and daily level on test results and hospitalizations. Unfortunately, these data are incomplete and inconsistent across cities and countries because testing is not universal and because there is no universal standard for hospitalization. Asymptomatic individuals may make up a substantial share of the infected, and are now believed to be contagious, but are infrequently tested. (See Stock, Aspelund, Droste, and Walker (2020) for estimates from Iceland, which has tested 6% of its citizens). Since we don't know how the tested population is related to the general population in almost any location, it is not clear what we learn from observing aggregate test results by region or country.

Sero-surveys

---

[26] For example, Johnson and Mueller (2002) and Spreeuwenberg, Kroneman, and Paget (2018) provide dramatically different estimates of these death rates.

[27] https://www.economist.com/graphic-detail/2020/04/03/covid-19s-death-toll-appears-higher-than-official-figures-suggest

Properties of the standard "polymerase chain reaction" (PCR) test for SARS-CoV-2 are not known with certainty. The proportion of false negative tests may vary by site (Cummins, 2020) and may be as large as 30%. Given the absence of reliable test information for representative samples of the population, Dushoff (2020) and others suggest that we should try to conduct *retrospective* blood-based tests known as sero-surveys to determine the population infection rate. Sero-surveys test for the existence of antibodies in the blood as evidence of past infection. The Johns Hopkins Center for Health Security estimates that the sensitivity—the fraction of infected people who test positive—is 93.8% and the specificity—the fraction of non-infected who test negative—is 95.6% for the RDT serology test, the first test approved for use in the U.S.[28] The CDC plans to conduct national sero-surveys, but will not start them until summer 2020 (Branswell (2020)). One initial study provides serology results for a sample of 3,330 people in Santa Clara County, California (Bendavid et al. (2020)), but the interpretation of those results are complicated by uncertainty about the specificity of the test that was used (Gelman (2020)).

V. Estimating Parameter Values for the SIR Model

The basic reproduction number R_0 describes the expected number of new infections that will be produced by a single infected individual in a "naive" population (i.e. where all other individuals are susceptible). It is a very helpful statistic for comparing different models and tends to be one of the rate parameters that is easiest to accurately measure. In general, there are two methods that can be used to estimate R_0.

The first, more mechanistic, approach is to use contact-tracing data from early in an epidemic. Health organizations often deploy contact tracing to attempt to slow or eradicate the

---

[28] http://www.centerforhealthsecurity.org/resources/COVID-19/Serology-based-tests-for-COVID-19.html

spread of disease – identifying and quarantining all people who come in contact with the earliest-known patients.   Because very few individuals are infected, and because institutions are comparatively unstressed, it is ideally possible to identify and test all individuals who came into contact with an infected individual, and from that to accurately compute R_0 (Ferretti et al., (2020)). However, this approach is less effective as the number of infected individuals grows large, as it can be difficult to definitively identify the source of new infections, and because comprehensive testing of contacts may no longer be feasible. Additionally, this method of estimation may fail to accurately capture heterogeneity in infectiousness among individuals or regions, especially if the distribution has a long tail (e.g. due to "super-spreaders"). The second, more phenomenological, approach is often straightforward to implement for the early stages of an epidemic, e.g. fitting the rate of growth to a known functional form. However, as case numbers grow, this approach becomes less effective, and because it is parametrized based on the number of confirmed cases, biased or temporally inconsistent testing will reduce the reliability of this estimate. Nevertheless, there has been high similarity in phenomenologically estimated R_0 values across different countries, suggesting that there is some underlying biological reality to these estimates.[29]

Even under ideal circumstances, it is important to note that parameterizing an SIR model with reported case data is necessarily a tricky endeavor. A classic method used in many textbook exercises would be to fit the basic model we introduce at the beginning of this text by setting the recovery rate $\gamma = 1/$(mean infectious period) and $R_0 = \beta/\gamma$. This approach is especially well-suited for some types of models, e.g. individual-based stochastic simulations, as it treats the rate terms as waiting times drawn from an exponential distribution. Thus, if one were to simulate an initial

---

[29] Ridenhour, Kowalik, and Shay (2014) estimate somewhat different values of R_0 across countries for the 2009 flu epidemic.

cohort of infected individuals within an arbitrarily large susceptible population, and track them until all of the original infected individuals recovered, one would find that the average duration of infection across those original individuals was $1/\gamma$, and that the total number of new infections directly caused by the original individuals, divided by the total number of original individuals, would be R_0. However, this definition is somewhat problematic in practice, as it requires the ability to fully track outcomes of all individuals in the initial cohort, until the entire cohort has recovered. If any individuals are omitted – e.g. because they had especially mild cases and recovered before they could be identified or infect others, or because they had especially severe cases and did not recover before the end of the study, the resulting rate coefficients will be biased. One partial solution is to fit rates based on "typical" rather than "average" outcomes. For example, if we know that 50% of patients recover within $t_{1/2}$ days, then we can estimate γ from the half-life of the infection – i.e. - $\gamma \approx \log(1/2)/t_{1/2}$. Similarly, we can roughly approximate R_0 in terms of the average number of new infections per half-life of the infection – i.e. $R_0 + 1 \approx \exp(\beta t_{1/2})$, and thus $\beta \approx \log(R_0 + 1)/t_{1/2}$, where "$R_0 + 1$" signifies R_0 new infections over $t_{1/2}$ days, plus the original infection. The advantage of this approach is that it can be applied even when only partial data from a cohort is available, and does not require information about individual-level outcomes, which can be difficult to acquire for legal and logistical reasons. We stress, however, that in modern applications of disease models, parameterization approaches are generally much more sophisticated than those that we present here – see Massaud et al. (2010) for examples of some suitable methods.

Economists might consider other approaches for estimating R_0 or might develop models for which this parameter is not fundamental, as the similarity in existing estimates of R_0 across countries can be viewed as suspicious. A structural economic model would provide

microfoundations so that R_0 emerges endogenously in equilibrium from a model of social interactions. Further, economists might focus on a more specific parameter, such as the probability of transmission per interaction (or possibly per unit of time of interactions) to facilitate analysis of policies that limit or penalize interactions. An alternative reduced form model might estimate R_0 as a function of city characteristics, policies adopted, and the degree to which people self-modify behavior as a function of prevalence.

While R_0 for SARS-CoV-2 can be estimated with comparatively high accuracy relative to other rate coefficients, there remains considerable uncertainty even about the (average) value of this parameter. For example, a collation of studies by the Germany's Robert Koch Institute[30] suggests that in the absence of control measures, R_0 falls between about 2.4 and 3.3 (although they acknowledge excluding several studies with especially high estimates). If we assume that half of infected individuals recover within 5-6 days, these rates would imply a doubling-time of about 2.5-3.5 days, which closely matches initial growth rates of confirmed case numbers in most countries (though again, all of these estimates are highly uncertain). Although this is a deceptively small range, it is important to remember that small changes in rates have major consequences. After a month of uncontrolled growth starting from a single infection, a doubling time of 2.5 days would yield about 4000 infections, whereas a doubling time of 3.5 would lead to about 400. Note, however, that much of this uncertainty is ultimately related to the generation time of the infection, rather than the rate at which it spreads in a population. Thus, phenomenological estimates of overall growth rates in confirmed case numbers have generally been less variable than estimates of R_0, and, correspondingly, tend to provide more accurate short-term forecasts of case numbers.

---

[30] See https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Steckbrief.html for summary statistics and related citations (in German):

There is minimal data for estimating parameters other than R_0, and in many cases uncertainty spans a wide range of possible values. Using the summary statistics of the Robert Koch Institute, as reported on 7 April 2020, as an example: incubation time has been reported to be about 5-6 days, but with observed values ranging from 1 to 14 days; reported case mortality rates (i.e. deaths divided by confirmed cases) range from 0.1% to 22%; the fraction of confirmed cases relative to actual cases (i.e. including undetected cases) has been roughly estimated at 5-9.2% based on early reports from China; the duration of contagiousness is so uncertain that no range is provided. Because each of these ranges of values contributes additional uncertainty to model forecasts, it seems likely that predictions which fully incorporate all of these sources of uncertainty will likewise span at least an order of magnitude, and likely much more.

Estimates for the fraction of infections that have been detected include at least two sources of uncertainty: first, that in most locations, not all symptomatic people can be tested, and second, not all people infected with SARS-CoV-2 are symptomatic (or at least, their infections are "sub-clinical", in that they are not sick enough to seek medical care). Stock (2020) defines the "asymptomatic rate" as the fraction of the infected who are not recorded as infected. Early estimates of asymptomatic rates ranged from 18% to 31% (Stock, 2020), but more recent data suggests values in a much higher range, with Stock et al. (2020) estimating a value of 90% based on more representative sampling in Iceland. Stock (2020) also observes that different reports of the asymptomatic rate may not be comparable because they do not also use the same definition of "asymptomatic". Uncertainty about the asymptomatic rate is closely linked with uncertainty about the mortality rate for SARS-CoV-2; Atkeson (2020b) shows that existing data are consistent with a wide range of mortality rates.

VI. Existing Models and Predictions

A handful of models that have received attention in the media and yield a fairly disparate set of results; we review five of them here.[31]

<u>Epstein Estimates</u>

Richard Epstein predicted that the "adaptive response" and other factors would dramatically limit the spread of disease: "Even if there is some undercounting, it is highly unlikely, given the relatively short (two-week) incubation period, that the number of current cases will more than double or triple." (Epstein (2020)) This model is largely mechanistic, in that it hypothesizes specific processes that drive dynamics, although the specific rates chosen for the model have been heavily criticized. Epstein initially predicted 500 deaths or fewer in the United States, then later revised that estimate to 5,000 deaths. Needless to say, these predictions have already been falsified.

<u>Oxford Study</u>

Lourenco et al (2020) consider three scenarios with a range of point estimates from 36% to 68% of UK adults infected as of March 19. The twin implications of these estimates was that the country was already likely close to the herd immunity point and – dividing observed deaths by the implied number of infected people – that the mortality rate for the disease must be quite low and thus that the Imperial College projections were dramatic overestimates. But the paper essentially assumes these conclusions: "Our ov*erall approach rests on the assumption that only a very small proportion of the population is at risk of hospitalisable illness." (Their MCMC

---

[31] We chose to review the Epstein, Ferguson, and Murray models because they have been reported to be used by governments, the Lourenco model because it gained attention for providing a prediction contrary to that of Ferguson, and the Kissler model because one of the co-authors of that paper, Marc Lipsitch, has been publicly identified as one of the modelers consulted by the U.S. government. Rogers and Molteni (2020) provide a very detailed popular account of the relationship between policy choices and the projections of various disease models.

estimation starts from a prior that the case fatality rate is about 1/700 in some versions and about 1/7000 in others.) This model is largely mechanistic, and follows a relatively standard SIR form.

IHME

The IHME paper (Murray et al. (2020)) is designed to estimate the demand for hospital beds, and thus to guide local-level policy decisions. At the same time, it yields estimates for caseload and deaths, and those estimates have been widely publicized and used by state governments for planning purposes.

This model takes as data the observed death rates on a day-by-day basis for each U.S. state and fits a cumulative Gaussian cdf to produce the predicted future course of the log of death rates in each state after accounting for state-level variations in age distribution. It then works backwards from future predicted death rates to project the demand for hospital beds and ICU admissions per capita for each state through the course of the epidemic.[32]

These results rely on two critical modeling assumptions. First, the authors assume a single change in behavior (and thus the nature of the fitted curve) at the time of the formal announcement of social distancing measures in each state. Second, the authors assume that social distancing induces changes in death rates that are analogous to the rises and falls indicated by the official statistics for Wuhan. As the paper comments, "*Modeling for US states based on one completed epidemic, at least for the first wave, and many incomplete epidemics is intrinsically challenging. The consequent main limitation of our study is that observed epidemic curves for COVID-19 deaths define the likely trajectory for US states*." Over time, IHME has

---

[32] The general method of fitting a deterministic curve to a cumulative set of observed cases tends to produce biased results because it does not adequately account for uncertainty. King et al. (2015) use simulation results to demonstrate this bias for predictions of the spread of Ebola.

been using Bayesian methods to incorporate data from the U.S. as it is released, while still putting heavy weight on the entire cycle of results observed in Wuhan.

The lockdown in Wuhan was more strict and more strictly enforced than any in the US—even after the recent reopening there are limits on how long residents can spend out of their residential compounds and checkpoints at which authorities take residents' temperatures and check a phone app that indicates if they are considered at risk of being infected. Accordingly, it seems likely that US infection rates will not decline as quickly as in Wuhan given our less stringent practices of social distancing. For this reason, the IMHE model may underestimate the fraction of deaths which will occur after the peak.

The IHME model is probably the most phenomenological of all of the models that we consider here, in that its predictions are primarily based on fitting curves to historical data, rather than on *a priori* biological assumptions about disease dynamics. This property has allowed the model to provide real-time location-specific forecasts. However, it has also led to intense criticism based on the suggestion that its predictions are more strongly driven by statistical and functional form assumptions than they are indicative of actual likely trends.

Kissler Seasonality Model

Kissler et al., (2020) use historical data from five years of US hospital admissions for two strains of coronavirus (HCoV-HKU1 and HCoV-OC43, which typically induce common colds) to fit parameters for an mechanistic SIR model that vary on a seasonal basis and allows for periods of immunity and cross-immunity after recovery from a given infection. They then expand the model to include SARS-CoV-2 as a third coronavirus under the assumption that it has similar seasonality to the two milder strains, which have been shown to spread in any season,

though with different peaks by season. Focusing on scenarios where policies do not prevent the herd immunity threshold from being reached, they note that substantial uncertainty about the strength and duration of immunity following an infection and levels of cross-immunity makes many future patterns possible. SARS-CoV-2 outbreaks might recur in regular annual or biennial intervals, or the virus could seemingly be mostly eliminated and then resurge five years later.

<u>Ferguson / Imperial College Model</u>

Ferguson et al. (2020) is likely the closest of the five we examine to a structural economics model. It uses Census data to model the spread of disease in a mechanistic SIR model at a fairly local level, considering four sources of social interactions that could produce "transmission events": within household, at work, in school, or in the community at large. The model is parameterized so that initial transmission events occur approximately equally at home, at school/work, or in the community, matching a stylized fact from previous studies suggesting that approximately one-third of transmissions occur in each of these places.

The model accounts for social distancing by applying particular rules for the effect of five different interventions: "Case Isolation", "Voluntary Home Quarantine", "Social Distancing of those over age 70", "Social Distancing of entire population", "Closure of schools and universities". For example, it assumes that closing schools increases contact rates within affected families by 50% and also increases contact rates by 25% in the community in general. Any combination of these five interventions is assumed in the model to induce specific changes in social contacts, and in turn determines the transmission rate R_0 in one's local area at time t.

The changes in contact rates assumed in this model are never justified and, in fact, appear to be entirely arbitrary and in some cases clearly inaccurate.[33] Economists would likely prefer the transparency of Atkeson (2020a), which characterizes the behavioral and institutional response to the spread of infection in reduced form with a change in the value of R_0.

Comparison of the Five Models

One critical distinction between the predictions of IMHE and Ferguson is that IMHE predicts a much earlier peak and steeper decline in the rate of infection, presumably because IMHE puts so much weight on the experience of Wuhan, where an extreme form of social distancing was enforced.[34] Much as this modeling choice for IMHE seems to have a huge influence on the qualitative nature of the results, the contrarian predictions of Epstein and Lourenco appear to follow directly from their assumptions. The conclusion of Kissler et al. (2020) that SARS-CoV-2 could have peaks in any season appears to follow similarly from their modeling choice to use patterns from other coronaviruses to model the spread of SARS-CoV-2.

Confidence Intervals and Scenarios

Stock (2020) argues that there is too little information to calibrate SEIR models of SARS-CoV-2 at this time, and in particular, that due to selection into testing, we know neither the true number of people who are infected in any country, nor the mortality rate conditional on infection.[35] Stock then demonstrates that two distinct plausible values of the asymptomatic rate (the proportion of infected people who have no symptoms if illness) suggested by data from different subsamples in Wuhan yield dramatically different predictions for the effects of three

---

[33] The scenario for the closing of schools and universities assumes that 25% of universities remain open, though in the United States, it seems that all but Liberty University will be closed from mid-March through at least May 1.
[34] Bergstrom suggests that this feature of the model has led to systemic errors in its predictions for Italy and Spain. https://twitter.com/CT_Bergstrom/status/1250304069119275009.
[35] See also Bendavid and Bhattacharya (2020).

social distancing polices. These simple simulations cast profound doubt on the precision of any existing model that provides predictions of the future course of the spread of SARS-CoV-2.

Unfortunately, there seems to be a general tendency for researchers to report a greater degree of confidence than is warranted for an existing model, in part because it is not straightforward to quantify parameter uncertainty or to trace the effect of those uncertainties in a non-linear model. Realistic confidence intervals in this context would also be so wide as to seem vacuous. As Jonathan Dushoff commented in an e-mail communication, "*Retrospectively, the most successful looking model is also likely to be a model with narrow confidence intervals, where there was some luck involved in making the model forecasts look good. In cases where there's a lot of models, and not so many realizations, this is very likely to happen.*"

Among the five models we discuss in the previous section, only Murray's (2020) IMHE model provides explicit confidence intervals for any predicted results. For example, that paper reports a 95% confidence interval of (38,242, 162,106) for total deaths in the United States. While this is a reasonably wide range, it is not clear from the paper how this interval was computed, and it still seems likely to be too narrow.[36] In particular, the paper seems to make no attempt to allow for uncertainty over the degree of reduction in contacts in the United States from mid-March on: presumably this confidence interval is the range of possible values conditional on the assumed functional form for deaths, which maintains that deaths rise and fall

---

[36] The paper explains that uncertainty in the forecast is driven primarily by variance in the "fixed" and "random" effects in the model. It does not provide further details of the computation of any confidence intervals. In a standard hierarchical modelling framework, this version of forecasting errors would imply that their estimates include both variability attributable to differences in mortality rates (of reported deaths) among localities, modelled as "random" parameters drawn from a hyperprior (typically a Gaussian distribution), and uncertainty in the mean mortality rate computed across all localities, modelled as a "fixed" effect. One caveat for such an approach is that this estimate of uncertainty is predicted on two assumption: (1) the data used to fit the model represents an unbiased sample of the true underlying process; and (2) that the true process generating dynamics matches the analytical function hypothesized in their model. See also Bergstrom's critique of this approach. https://twitter.com/CT_Bergstrom/status/1243838213086539776

symmetrically about some peak. The functional form is said to have been chosen because others "did not fit the data as well." This is implicitly relying heavily on reported data from the two epidemics that have risen and fallen, Wuhan and South Korea, and this dependence is presumably not taken into account.

The four other papers report mean estimates for a range of scenarios, suggesting a confidence interval interpretation without using that language. For example, Lourenco (2020) states that the results from analysis of various scenarios of the Oxford Model indicate "*significant population level immunity accruing by mid March in the UK as ρ is decreased to plausible values*", where ρ is the proportion of the population susceptible to hospitalisable illness. Kissler et al (2020) report five qualitative implications of their analysis, justifying them with statements such as "*In all modeled scenarios*", "*many scenarios lead to*". Similarly, Ferguson et al (2020) states, "*Such policies are robust to uncertainty in both the reproduction number, R_0 (Table 4) and in the severity of the virus (not shown),*" but allowing for a small degree of variation in a single parameter is hardly a reasonable test of robustness.[37]

In sum, the language of these papers suggests a degree of certainty that is simply not justified. Even if the parameter values are representative of a wide range of cases within the context of the given model, none of these authors attempts to quantify uncertainty about the validity of their broader modeling choices.

---

[37] The paper only provides the results for variations in the value of R_0 and only appears to consider variation in this parameter in a range from 2.0 to 2.6, which appears to be a much narrower range than a 95% confidence interval for R_0.

VII. "Flattening the Curve" and Future Policy Considerations

The policy response to SARS-CoV-2 in the US (and several other nations) has been a tragic failure. By the end of this month the US will have lost tens of thousands of lives. The two trillion dollar relief package will ameliorate only a portion of economic suffering that has resulted from the outbreak. And the most rosy realistic forecast for the future is that after another month of social distancing we might be roughly back to a state of the world that is about as promising as where we could have been in mid-March given better policy. Even in that scenario, the number of infections will likely be much higher in mid-May than they were in mid-March. The infectious will be more dispersed. Production of personal protective equipment will have increased, but this will likely be offset by depletion of the existing stockpile. At present, testing capacity in the United States is dramatically higher than it was, but still lags what South Korea achieved.

With this sobering realization in mind, we still need to address the question of policy going forward. A starting point for thinking about optimal policies is to work backwards from the end. In the long run we will presumably reach one of three endgames.

First, we may work to hold infections to a moderate level until an effective vaccine becomes available, perhaps in 12 to 24 months. An optimal policy reaching this endpoint would make tradeoffs between holding down infection levels and sustaining economic activity until the vaccine is available.

Second, we may get some good news that leads us to choose to allow the infection to spread at a more rapid, but controlled, rate. This news could involve learning that some antiviral treatment or some practice like universal wearing of N95 masks reduces the health consequences

of infection or slows its spread. Or it could simply involve learning that the virus was less harmful than is currently believed, e.g because asymptomatic cases are more common than we know.

Third, we may reach the point where we cannot or choose not to adopt policies that keep transmission rates in check, and infection proceeds to the point where herd immunity stops the spread. Current simulations suggest that the health consequences of reaching this endpoint rapidly would be very, very bad (Greenstone and Nigam (2020)). And there would likely be severe economic disruption even if we abandoned all government-mandated distancing efforts– few people would want to go to work, send their children to school, or eat in a restaurant if we reached a peak where fully 10% of the US population was currently infectious.

Flattening the rate at which the disease expands has the obvious advantage of reducing the possibility of overloading hospital systems and thereby reducing negative externalities. It also buys time, potentially allowing a larger fraction of those who will eventually become infected to benefit from more effective treatments that may be developed. "Flattening the curve" will also reduce the extent to which a disease is able to over-shoot the herd immunity threshold (since the incidence of the disease at the point that the threshold is reached determines the number of additional cases that will still occur before the disease dies out). For example, if we take the simple three-compartment model described at the beginning of the paper with R_0 of 2.6 and a 5.5 day half-life for the infectious period (i.e. $\beta = \log(R_0+1)/5.5 = 0.23$, $-\gamma = \log(1/2)/5.5 = -0.13$) and 0.0013% of the population (i.e. 100,000 out of 7.5 billion people) infected at the start of the simulation, then we would expect an "uncontrolled" spread of the disease to ultimately infect about 75% of the population, a peak incidence of about 13%, and for 99% of all infections to occur within about 6 months. In contrast, if control methods are put into place to prevent the

incidence from ever exceeding 1% of the population, only about 56% of the population will ultimately be infected (roughly equivalent to the herd immunity threshold in this case), but it would take more than one and a half years for 99% of the infections to occur.

That is, "flattening the curve" is no panacea, as it lengthens the time required to reach the herd immunity threshold and still is projected to produce a large number of deaths. It is possible that we will learn that asymptomatic infections are currently much larger than is known and this endpoint is not so bad in less densely populated areas. Some countries, especially in the developing world may find that they cannot avoid it even if it is as bad as feared. The same could be true in the US given the incomplete social distancing in place and rising discord regarding even these incomplete measures. But we focus here on thinking about optimal policies that reach one of the first two endpoints.

We think it is useful to think separately about two types of policies that may be implemented to alleviate health-related and economic suffering while waiting for an effective vaccine/treatment.

First, we may pursue an aggressive policy of testing and contact tracing. Such policies are thought to be a key part of how South Korea, Hong Kong, and Singapore have mostly kept infection rates in check. Absent failures to build testing capacity in the US, we could perhaps have achieved comparable success. Effective contract tracing, however, requires a capacity for testing and public health outreach that is far beyond what could be implemented today in the US. And once contract tracing alone cannot effectively mitigate the spread, and strong social distancing policies are adopted, the incremental value from contact tracing is reduced.

If, in the future, active cases have been reduced to a point where contact tracing is again feasible, it could become a central part of our approach to SARS-CoV-2. Controlling the disease while distancing only those who are known to be infected is extremely attractive from an economic viewpoint. Whether the US could hope to keep R_0 below one mostly by these means, however, is unclear. Many countries practicing contact tracing have done so in part by making extensive mandated use of phone-based and other geolocation data that the US has not yet contemplated. And even with such aggressive use of technology, reports from Hong Kong suggest that it has been challenging to keep R_0 below one without also implementing some social distancing measures. Much of the US is much less densely populated than is Hong Kong, so less aggressive contact tracing might suffice. It is even conceivable that testing-based policies may allow us to simultaneously reduce infections and sustain economic activity even if distancing cannot first reduce infections to a low level. For example, the cost and availability of home-pregnancy-style tests could fall to such a degree that a substantial fraction of the population could take one every day.[38] But we are now very far from that level, and have not been making the massive investments in testing research and developing the public health infrastructure that such an aspiration would seem to require.

Second, we will pursue some set of social distancing measures. Many variants of social distancing have been implemented. We find it useful to think of individual restrictions as lying on a continuum in terms of the ratio of economic costs incurred per unit of virus-slowing. At one extreme, there are some policies with a negative ratio – society would be better in both dimensions if long lines and crowded waiting rooms at the department of motor vehicles were replaced with expanded online services and scheduled appointments with minimal waits.

---

[38] Silver (2020) provides a detailed analysis of the implications of future testing scenarios given the possibilities of both false positives and negatives. See also https://twitter.com/zbinney_NFLinj/status/1245789672833417217.

Consensus seems to be shifting that universal wearing of facemasks may also provide substantial

benefits – the slower spread of the virus and low number of deaths in Japan is intriguing and

suggests that universal adoption of high-grade facemasks could be incredibly valuable. Costs are

certainly minimal compared to most other interventions, so any substantial benefit would make

the ratio attractive. Work-from-home and distance education have more substantial costs, but

also presumably provide substantial benefits. Shutting construction sites, at least relative to

continuing with modified work practices that increase social distancing, mask usage, and

handwashing, seems like a yet larger ratio of economic costs to disease slowing benefits. Many

other activities have not been suspended presumably because the economic costs are seen as

prohibitive. Recent reductions in new hospitalizations are encouraging, but to some extent we

have been taking it on faith that R_0 can be reduced below one without taking the more extreme

measures used in China.[39]

Absent successful contact tracing, reductions in social distancing would likely induce a

transmission rate $R_0 > 1$ and could lead to recurrence of disease. Cyclical peaks are well known

from prior diseases: measles outbreaks took place every few years for decades in England and

the U.S. (Dalziel et al., 2016). There is reason to believe that subsequent outbreaks of the

disease will be less intense, since people who previously contracted it may remain immune for

some time, and will grow more slowly than the initial outbreak (and presumably, that they will

therefore be easier to control). Nevertheless, in biological models there are rarely any

guarantees, and a well-known counterexample to this phenomenon is the 1918 influenza

pandemic, for which the second wave of infections was more severe than the first in many

locations. Theoretically, we would also not expect the effect to be very large unless the

---

[39] Flaxman et al. (2020) examining European data up through March 28, 2020 estimate confidence intervals for R_0
which include a fairly wide range of values above and below one.

susceptible fraction has moved substantially away from one. If our current distancing efforts limit infections to just a few percent of the population, as we hope, then this effect would not be large.

Serology testing will be a third component of policy response. Having at least some random serology testing seems very important – the information it provides on how widespread infection has been is tremendously important for thinking about how the virus will spread and how severe the health consequences will be. It can also provide economic benefits in that those who are immune can safely return to work. But this economic benefit may not be large in practice, both because we hope that we can avoid a situation where a large share of the population is infected, and because there will be uncertainty about how strongly and for how long prior infection is protective.

An optimal policy for the period until vaccines/treatments become salient will presumably involve choosing some point on the social distancing spectrum and whether to attempt aggressive contact tracing at each point in time.[40] The costs and benefits of policies will naturally vary with the current prevalence of active infections, both because the number of future infections prevented by a unit decrease in R_0 is proportional to the number currently affected and because economic impacts will vary. For example, the loss of learning when schools move online will be smaller if the alternative was to have teachers calling in sick and leaving classes in the hands of substitutes. Hence, the optimal policy will be an adaptive one with measures that adjust to prevailing conditions.

---

[40] It is unclear a priori whether the optimal control policy should be stationary. (We are grateful to Michael Kremer for suggesting this point.) Morris et al. (2020) and Alvarez, Argente, and Lippi (2020) provide separate analyses for SIR models where the optimal control policy is not stationary.

If the optimal policy keeps society in the region where the susceptible fraction S is close to one, then the optimal control problem of choosing an optimal policy for each (I, S) might be thought of as mostly just a one-dimensional problem where one must choose an optimal policy for each I. Depending on what we learn about fatality rates, rates of spread, and the effects of various policies on these characteristics, it is plausible that the optimal policy might take on one of several forms.

In an optimistic scenario, we may find that social distancing policies similar to those we have adopted, plus more face masks, can drive R_0 far below one. In that case, continuing these policies for a moderate period of time may get us to the point where we can switch to primary reliance on contact tracing combined with just the most efficient distancing policies on the spectrum. To be prepared to take advantage of such an opportunity, should it present itself, it might be sensible to start practicing for it right away by attempting to immediately institute aggressive contact tracing in areas of the country that are relatively isolated and where prevalence is low.

It is also plausible that the combination of contact tracing and highly efficient social distancing measures will not be enough to keep cases from expanding. In that case, it may be that we need to also adopt social distancing practices intermediate on the spectrum alongside contact tracing for quite some time. For example, it may be that schools will not open for 1-2 years, and we should start preparing for such eventualities.

There may also be inherent nonconvexities in the set of social distancing policies available to us. For example, it may be that the combination of shutting schools, shutting bars, restaurants, and retail stores, encouraging work from home, and limiting gathering is more efficient than an intermediate policy involving a subset of these elements. An optimal control

framework might say that the best thing to do in such a situation is to adopt a policy where we mix in continuous time between the aggressive distancing and testing plus moderate distancing alternatives.

But in practice, there are presumably transition costs incurred whenever stores open and close, firms shift to work from home, and so on. In this case, the optimum could potentially be an alternation between more and less strict shutdown policies, with new shutdowns triggered every time the disease prevalence reaches some threshold. An appealing side-benefit of extensive testing and contact tracing is that it can make it easier to recognize when prevalence has moved into the region where testing hits its capacity constraint and more aggressive distancing is needed. Alternation between periods of low and high prevalence of disease could also be optimal in a world where occasional returns to more normal activity can help to maintain worker-firm connections and allow firms to occasionally clear bottlenecks in production and shipping that arise from social distancing.

A limitation of any such strategy is that periods with $R_0 > 1$ must be balanced with periods with $R_0 < 1$ to keep prevalence in check. Given that R_0 is well above one under minimal social distancing, we will be unable to spend much time in such a regime if the R_0 in the social distancing regime is only a little below one. There will also be a temptation to loosen restrictions in the summer if seasonal differences slow the spread at that time. Whether that is advisable depends on whether the effect of a restriction is smaller in proportional terms.

New policy discussions will need to occur whenever the best available treatment changes or when knowledge about the epidemiology changes. Most consequentially, as the set of available treatments improves we will need to think about whether social distancing restrictions then in force are still warranted.

An important consideration for all policy discussions is that we must recognize that we are making decisions under substantial uncertainty. For all the reasons we noted earlier, we still know very little about the epidemiology of SARS-CoV-2. And we are also following an economic path that is unprecedented in modern times and should recognize that there is great uncertainty about how quickly and well modern economies will recover from shorter vs. longer and more vs. less extensive shutdowns. Some heuristics that we may want to keep in mind for dealing with such uncertainty are that it may be useful to try to keep the economy and disease situation as close as we can to situations with which we have some experience, and that we should think about worst-case consequences of policies. Adaptive policies that implement stricter policies whenever prevalence is increasing are appealing in part because they are well adapted to uncertainty in these ways.

VIII. Conclusion and Opportunities for Further Research

There are some ways in which we are now better positioned than when we first faced the SARS-CoV-2 threat. Most importantly, we have more information about what we are facing. Current and future research is extremely valuable because we still face great uncertainty and there are momentous decisions to come.

In the past month, economists and political scientists have produced new and valuable research with papers that may guide the way for future efforts. There is an urgent need for better data (Stock (2020), Stock, Aspenlund, Droste, and Walker (2020)), and for creative and entrepreneurial methods of interpreting the limited data that is available (Fang et al (2020), Harris (2020)). It is not clear that models of the spread of disease have played a wholly positive

role in shaping policy in the first three months of 2020. Similarly, some reports suggest that a misguided understanding of behavioral science (the fear of "behavioral fatigue") delayed the response of the UK government to the spread of the pandemic.[41] But there have already been clear contributions in recent weeks for both behavioral approaches (Barari et al. (2020), Bricese et al. (2020)) and theoretical modeling (Atkeson (2020), Berger, Herkenhoff, and Mongey (2020)). We include an appendix listing standard and novel sources of data about the pandemic and hope it will be a useful resource for researchers.

After an initial policy debate about whether and when to close down businesses and schools in February and March 2020, we can anticipate the next debate about how and when to reopen them in May and June 2020 and beyond. Using the information we have and will gain to make better policy choices in our second opportunity is critically important.

---

[41] https://behavioralscientist.org/why-a-group-of-behavioural-scientists-penned-an-open-letter-to-the-uk-government-questioning-its-coronavirus-response-covid-19-social-distancing/

References:

Acemoglu, D. & Ozdaglar, A. (Fall, 2009). 6.207/14.15: Networks Lecture 8: Diffusion through Networks. MIT Open. https://economics.mit.edu/files/4626

Allcott, H., et. al. (2020). "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic". NBER Working Paper #26946.

Alvarez, F. E., Argente, D., & Lippi, F. (2020). "A Simple Planning Problem for COVID-19 Lockdown". NBER Working Paper #26981.

Atkeson, A., (2020b). "How Deadly Is COVID-19? Understanding the Difficulties with Estimation of Its Fatality Rate". NBER Working Paper #26965.

Atkeson, A., (2020a). "What Will Be the Economic Impact of Covid-19 in the US? Rough Estimates of Disease Scenarios," NBER Working Paper 26867.

Azzoni, T, and A. Dampf, (2020, March). "Game Zero?' Soccer Game Attended by 40,000 Fans Likely Made This Italian City a Coronavirus Epicenter" *Time.* https://time.com/5809848/game-zero-soccer-game-italy/

Barari, S. et. al., (2020). "Evaluating COVID-19 Public Health Messaging in Italy: Self-Reported Compliance and Growing Mental Health Concerns,". Working Paper. https://j.mp/39btyT2

Bendavid E., & Bhattacharya, J. "Is the Coronavirus Epidemic as Deadly as They Say?" *Wall Street Journal*, March 24, 2020.

Bendavid E., et al. "COVID-19 Antibody Seroprevalence in Santa Clara County, California," *Medrxiv*, https://www.medrxiv.org/content/10.1101/2020.04.14.20062463v1.full.pdf

Berger, D., Herkenhoff, K., & Mongey, S. (2020). "*An SEIR Infectious Disease Model with Testing and Conditional Quarantine*," NBER Working Paper 26901.

Blackwood, J., & Childs, L. (2018). An introduction to compartmental modeling for the budding infectious disease modeler. *Letters in Biomathematics*, 5(1), pp.195-221.

Bolker, B. (1999). Analytic models for the patchy spread of plant disease. *Bulletin of Mathematical Biology*, 61(5), 849-874.

Booth, W. (2020). "A chilling scientific paper helped upend U.S. and U.K. coronavirus strategies", *Washington Post.* March 17, 2020.

Borjas, G. (2020). "Demographic Determinants of Testing Incidence and COVID-19 Infections in New York City Neighborhoods". NBER Working Paper #26952.

Branswell, H., (2020). CDC launches studies to get more precise count of undetected Covid-19 cases, *STAT.* April 4, 2020.

Brauer, F., Van den Driessche, P., & Wu, J. (2008). *Mathematical epidemiology* (Lecture notes in mathematics (Springer-Verlag). 1945. Berlin: Springer.

Bricese, G., Lacetra, N., Macis, M., & Tonin, M. (2020). "Compliance with COVID-19 Social-Distancing Measures in Italy: The Role of Expectations and Duration," NBER Working Paper 26916.

Chang, B. (2020). "France has converted its high-speed TGV trains into ambulances to transport critically ill coronavirus patients across the country," *Business Insider.* April 1, 2020.

Cereda, D., et al., (2020). The early phase of the COVID-19 outbreak in Lombardy, Italy. *ArXiv.org.* Mar 20, 2020. https://arxiv.org/ftp/arxiv/papers/2003/2003.09320.pdf

Chiou, L. & Tucker, C. (2020) "Social Distancing, Internet Access, and Inequality". NBER Working Paper #26982.

Cummins, E. (2020). "Why the Coronavirus Test Gives So Many False Negatives," Slate.com, April 6, 2020.

Dalziel, B., et. al. (2016). Persistent Chaos of Measles Epidemics in the Prevaccination United States Caused by a Small Change in Seasonal Transmission Patterns. *Plos Computational Biology, 12*(2), E1004655.

Dawsey, J., Abutaleb, Y., Stein, J., Wagner, J. (2020). "Trump weighs restarting economy despite warnings from U.S. public health officials," *Washington Post*, March 23, 2020

Dushoff, J., (2020). "Preparing for the Storm," March 23, 2020, http://dushoff.github.io/notebook/covid.html

Easley, D., & Kleinberg, J. (2010). *Networks, crowds, and markets: Reasoning about a highly connected world*. New York: Cambridge University Press.

Ellyat, H. (2020). "Germany has a low coronavirus mortality rate: Here's why", *CNBC.com*. April 3, 2020.

Epstein, R., "Coronavirus Perspective," working paper. *Hoover Institution*. Stanford University, March 16, 2020. https://www.hoover.org/research/coronavirus-pandemic

Fang, H., Wang, L. & Yang, Y. (2020), "Human Mobility Restrictions and the Spread of the Novel Coronavirus (2019-NCOV) in China," NBER Working Paper 26906.

Ferguson, N., et al., (2020). "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand," Working Paper, Imperial College London. https://spiral.imperial.ac.uk:8443/handle/10044/1/77482

Ferretti, L. et al., (2020). "Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing," *Science.* March 31, 2020.

Fink, S., (2020). "White House Takes New Line After Dire Report on Death Toll," *New York Times*, March 16, 2020.

Flaxman, S. et al. (2020). "Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries," Working Paper, Imperial College London. https://www.imperial.ac.uk/mrc-global-infectious-disease-analysis/covid-19/report-13-europe-npi-impact/

Gallos, L., & Fefferman, N. (2015). The Effect of Disease-Induced Mortality on Structural Network Properties. *PLoS One, 10*(8), E0136704.

Geddes, L., (2020). "Does a high viral load or infectious dose make covid-19 worse?" *New Scientist.* March 27, 2020

Gelman, A. (2020). "Concerns with that Stanford study of coronavirus prevalence,**"** https://statmodeling.stat.columbia.edu/2020/04/19/fatal-flaws-in-stanford-study-of-coronavirus-prevalence/

Giles, C., "Deaths from coronavirus far higher in England than first reported" *Financial Times*. April 7, 2020.

Greenstone, M. & Nigam, V. (2020). "Does Social Distancing Matter?" *Becker Friedman Institute for Economics Working Paper Series* 2020-26.

Gupta, A. (2020). "Does Covid-19 Hit Women and Men Differently? U.S. Isn't Keeping Track." *New York Times*. April 3, 2020.

Harko, T., Lobo, F., & Mak, M. (2014). Exact analytical solutions of the Susceptible-Infected-Recovered (SIR) epidemic model and of the SIR model with equal death and birth rates. *Applied Mathematics and Computation, 236*(C), 184-194.

Harris, J. (2020). "The Coronavirus Epidemic Curve is Already Flattening in New York City," NBER Working Paper No. 26917.

He, X., et al., (2020). "Temporal dynamics in viral shedding and transmissibility of COVID-19," *Medrxiv.* https://www.medrxiv.org/content/10.1101/2020.03.15.20036707v2.full.pdf

Hengehan, C., Brassey, J., & Jefferson, T. (2020). "SARS-CoV-2 viral load and the severity of COVID-19," Working Paper, *Centre for Evidence-Based Medicine*, Oxford University, March 26, 2020.

Henriques, M., (2020). "Coronavirus: Why death and mortality rates differ". BBC.com. April 1, 2020.

Hethcote, H. (2000). The Mathematics of Infectious Diseases. *SIAM Review, 42*(4), 599-653.

Hethcote, H., & Driessche, W. (1991). Some epidemiological models with nonlinear incidence. *Journal of Mathematical Biology, 29*(3), 271-287.

Hethcote, H., & Van Ark, J. (1987). Epidemiological models for heterogeneous populations: Proportionate mixing, parameter estimation, and immunization programs. *Mathematical Biosciences, 84*(1), 85-118.

Hruby, D. (2020). "How an Austrian ski resort helped coronavirus spread across Europe". CNN.com, March 24, 2020

Hurford, A. (2012). "Mechanistic Models – What is the Value of Understanding?" https://theartofmodelling.wordpress.com/2012/02/19/mechanistic-models-what-is-the-value-of-understanding/

Johnson, N., & Mueller, J. (2002). Updating the Accounts: Global Mortality of the 1918-1920 "Spanish" Influenza Pandemic. *Bulletin of the History of Medicine, 76*(1), 105-115.

Johnston, J., (2020). "Patrick Vallance says keeping coronavirus death toll below 20,000 would be 'good outcome" *Politics Home*. March 17, 2020.

Kendi, I. (2020). "What the Racial Data Show. The pandemic seems to be hitting people of color the hardest," The Atlantic, April 6, 2020

Kermack, W., & McKendrick, O. (1991a). Contributions to the mathematical theory of epidemics—I. *Bulletin of Mathematical Biology, 53*(1-2), 33-55.

Kermack, W., & McKendrick, O. (1991b). Contributions to the mathematical theory of epidemics—II. The problem of endemicity. *Bulletin of Mathematical Biology, 53*(1-2), 57-87.

Kermack, W., & McKendrick, O. (1991c). Contributions to the mathematical theory of epidemics—III. Further studies of the problem of endemicity. *Bulletin of Mathematical Biology, 53*(1-2), 89-118.

King, A., Magpantay, F., & Rohani, P. (2015). Avoidable errors in the modeling of outbreaks of emerging pathogens, with special reference to Ebola. *ArXiv.org, 282*(1806), 20150347.

Kissler, S., et. al., (2020). "Projecting the transmission dynamics of SARS-CoV-2 through the post-pandemic period," *MedRxiv*, https://doi.org/10.1101/2020.03.04.20031112.

Kuchler, T., Russel, R., & Stroebel, J. (2020). "The Geographic Spread of COVID-19 Correlates with Structure of Social Networks as Measured by Facebook". NBER Working Paper #26990.

Lazar, K., Ryan, A. (2020) "How Accurate Are Coronavirus Tests?" *Boston Globe*, April 2, 2020.

Lehman, C., Loberg, Clark, & Open Textbook Library. (2019). *Quantitative Ecology: A New Unified Approach*. University of Minnesota Libraries Publishing.

Li, M., & Muldowney, J. (1995). Global stability for the SEIR model in epidemiology. *Mathematical biosciences*, 125(2), pp.155-164.

Lively, C. (1999). Migration, Virulence, and the Geographic Mosaic of Adaptation by Parasites. *The American Naturalist, 153*(S5), S34-S47.

Lourenco, J. et al., (2020). "Fundamental principles of epidemic spread highlight the immediate need for large-scale serological surveys to assess the stage of the SARS-CoV-2 epidemic," *MedRxiv*. https://doi.org/10.1101/2020.03.24.20042291.

Ludwig, D., & Malley, R. (2020). "Americans Are Already Too Diseased to Go Back to Work Right Now, *New York Times*. March 30, 2020.

Massad, E., Coutinho, F., Burattini, M., & Amaku, M. (2010). Estimation of R 0 from the initial phase of an outbreak of a vector-borne infection. *Tropical Medicine & International Health, 15*(1), 120-126.

Morris, D., Rossine, F., Plotkin, J., & Levin, S. (2020). "Optimal, near-optimal, and robust epidemic control". Working paper. 2020. https://osf.io/rq5ct/

Murray, C. et. al., (2020). "Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months. *Medrxiv*. March 30, 2020. https://www.medrxiv.org/content/10.1101/2020.03.27.20043752v1

Rabinowitz, J. & Bartman, C. (2020). "These Coronavirus Exposures Might Be the Most Dangerous," *New York Times.* April 1, 2020.

Read, A.F. (1994). The evolution of virulence. *Trends in microbiology*. 2(3), pp.73-76.

Ridenhour, B., Kowalik, J.M., & Shay, D. K. (2014). Unraveling R0: Considerations for public health applications. *The American Journal of Public Health, 104*(2), 32.

Robinson, M, & Blanchard, S. (2020). "Chief scientific adviser wants 40million Britons - 60 percent of the entire population - to catch coronavirus for herd immunity to take effect". *Daily Mail*. March 13, 2020

Rogers, A. & Molteni, M. (2020). "The Mathematics of Predicting the Course of the Coronavirus". *Wired*. March 30, 2020.

Rudolf, V., & Antonovics, J. (n.d.). Species Coexistence and Pathogens with Frequency-Dependent Transmission. *The American Naturalist, 166*(1), 112-118.

Sattenspiel, L. (1990). Modeling the spread of infectious disease in human populations. *American Journal of Physical Anthropology, 33*(S11), 245-276.

Silver, N. (2020). "Coronavius Case Counts are Meaningless," *538.com.* April 4, 2020. https://fivethirtyeight.com/features/coronavirus-case-counts-are-meaningless/

Spreeuwenberg, P., Kroneman, M., & Paget, J. (2018). Reassessing the Global Mortality Burden of the 1918 Influenza Pandemic. *American Journal of Epidemiology, 187*(12), 2561-2567.

Stein, R. (2011). Super-spreaders in infectious diseases. *International Journal of Infectious Diseases, 15*(8), E510-E513.

Stock, J. (2020). "Data Gaps and the Policy Response to the Novel Coronavirus" NBER Working Paper No. 26902.

Stock, J., Aspelund, K., Droste, M., & Walker, C. (2020). "Estimates of the Undetected Rate among the SARS-CoV-2 Infected using Testing Data from Iceland," Working Paper.

Weise, E. (2020). "8 strains of the coronavirus are circling the globe. Here's what clues they're giving scientists" *USA Today*. March 27, 2020.

Wood, S. (2001). Partially Specified Ecological Models. *Ecological Monographs. 71*(1), 1-25.

Zhou, F., et al., (2020). "Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study". *The Lancet*, 395:10229.1054-1062. March 28, 2020.

We have compiled a list of data sources, mostly publicly available, that researchers may find useful. Some are just repackaging of official government statistics on infections, hospitalizations, deaths, and so forth, but, depending on the particular website, they might be from different governments and at different levels of aggregation. Others are entirely novel data sources, often shared by private firms who have collected them in the course of their business practice or as a public service. With the exception of the data on the 1918 Flu Pandemic, all data sets cover a period of time relevant for study of the current pandemic.

We offer the list with the caveat that new data dumps, curated repositories, and official statistics are becoming available on a daily basis. Furthermore, links might become stale as websites are moved or removed. (We try to offer searchable descriptions to help researchers track down the data sources.)

1. Statistics on infections, hospitalizations, recoveries, and deaths at the **country** level
   a. https://data.humdata.org/dataset/novel-coronavirus-2019-ncov-cases
   b. European CDC data https://www.ecdc.europa.eu/en/covid-19-pandemic
   c. Institute for Health Metrics and Evaluation data http://www.healthdata.org/covid
2. Statistics on infections, hospitalizations, recoveries, and deaths at the US **state** level, as well as data on ventilator, hospital, and ICU usage and number of tests
   a. https://ourworldindata.org/covid-testing
   b. https://covidtracking.com/data
3. Statistics on infections, hospitalizations, recoveries, and deaths at the US **county** level
   a. https://usafacts.org/visualizations/coronavirus-covid-19-spread-map/ (where it says Download data (Jan. 22- April 6))
4. NY Times data at the US **county** level of confirmed cases and deaths by day
   a. https://www.nytimes.com/article/coronavirus-county-data-us.html, available at https://github.com/nytimes/covid-19-data
5. New York City-specific data
   a. https://www.nytimes.com/article/coronavirus-county-data-us.html, available at https://github.com/nychealth/coronavirus-data
6. Historic data for 43 US cities on the 1918 Flu Epidemic, including deaths, death rates, and non-pharmaceutical interventions (such as social distancing), from US Census data---used in Markel, Lipman, and Navarro (2007), but not clear if and where the data are available for download
   a. https://jamanetwork.com/journals/jama/fullarticle/208354
7. Data on the timing of non-pharmaceutical interventions (social distancing, closure of schools and universities, closure of nonessential business, gathering size limitations, etc.) in the US
   a. www.keystonestrategy.com/covid-19/, available at https://github.com/Keystone-Strategy/covid19-intervention-data/
8. Data on distances traveled using cell phone movements from Unacast---can view and interact with their mobility map, but not download the data. (They may be offering it to universities and non-profits for free.)
   a. https://www.unacast.com/covid19/social-distancing-scoreboard
9. Mobility data is becoming available from other sources as well, such as Google, CityMapper, and SafeGraph

a. https://www.google.com/covid19/mobility/
    b. https://citymapper.com/cmi
    c. https://safegraph.com
10. Restaurant reservation data by city and day for Australia, Canada, Germany, Ireland, Mexico, United Kingdom, and the United States from Open Table
    a. https://www.opentable.com/state-of-industry
11. Data on airport usage---can view map but not available for download
    a. https://wanderlog.com/coronavirus-airports-effect
12. Data on fever incidence from Kinsa, a firm that makes smart thermometers---can view infographics, but not available for download
    a. https://www.kinsahealth.co/images-and-infographics-from-kinsas-health-weather-map-and-data/
13. Data on air pollution (at weekly or daily frequency)
    a. For European countries: https://www.eea.europa.eu/themes/air/air-quality-and-covid19/monitoring-covid-19-impacts-on
    b. Worldwide: https://www.covidexplore.com/PM25 (https://github.com/mayukh18/covidexplore)
    c. India and China (though they might have data on more countries): https://energyandcleanair.org/blog/
    d. https://aqicn.org/data-platform/covid19/
14. Satellite data on air pollution data in Europe
    a. https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-5P/Coronavirus_lockdown_leading_to_drop_in_pollution_across_Europe (https://scihub.copernicus.eu)
15. Data from Foursquare on consumer location, derives measures of impacts on flights, stocking up behavior, sit-down restaurants, fast food, etc.---can view article, but data not available for download
    a. https://enterprise.foursquare.com/intersections/article/understanding-the-impact-of-covid-19/
16. Data on planned, ongoing, and completed trials for medical interventions
    a. https://covid-evidence.org
17. Google Trends, a free tool offered to allow researchers to download (normalized) data on search volumes over time and across geographic regions
    a. https://trends.google.com/trends/
    b. Example article using these data https://www.nytimes.com/2020/04/05/opinion/coronavirus-google-searches.html