ANDREW ATKESON
University of California, Los Angeles

STEPHEN KISSLER
University of Colorado, Boulder

The Impact of Vaccines and Behavior on US Cumulative Deaths from COVID-19

ABSTRACT We estimate that the combination of changes in behavior to slow the spread of COVID-19 and the delivery of vaccines to a substantial majority of the American population by mid-2021 saved close to 800,000 American lives relative to what would have occurred had vaccines not been developed. We argue that the duration and magnitude of this behavioral response—and thus its overall success in delaying infections—came as a surprise, relative to both our historical experience with pandemic influenza and to model-based projections based on that experience. Thus, we take from our experience with COVID-19 over the past four years the important public health lesson that behavior change can be a powerful force for slowing the spread of a dangerous infectious respiratory disease for a long time. At the same time, these behavioral changes to slow the spread of COVID-19 came at a tremendous economic, social, and human cost. To avoid similar pain from mitigation in the next pandemic, we argue that we need to make investments now not only in vaccine development, but also in data infrastructure so that we can precisely target behavior-oriented mitigation efforts to minimize their economic and social impacts in the next pandemic.

Conflict of Interest Disclosure: Stephen Kissler has served as a consultant for ModernaTx, advising on infectious disease modeling for vaccine strategy. The authors did not receive financial support from any firm or person for this paper or, other than the aforementioned, from any firm or person with a financial or political interest in this paper. The authors are not currently an officer, director, or board member of any organization with a financial or political interest in this paper.

Brookings Papers on Economic Activity, Spring 2024: 67–112

Starting in March 2020, the American public undertook massive changes in behavior in response to the threat from COVID-19.[1] These behavioral changes arose partly in response to public mandates and partly as a spontaneous private reaction to this new disease threat. These public and private disease mitigation efforts succeeded in slowing the transmission of SARS-CoV-2 to a remarkable extent through 2020 and well into 2021, by which time effective vaccines had been developed and delivered to much of the American population.

As a result of these mitigation efforts, a large majority of Americans were able to get vaccinated for COVID-19 before experiencing their first infection. We document this using nationwide serology data, which lets us estimate the cumulative number of infections and vaccinations over time. Population-level data on vaccine efficacy indicate that this success in delivering vaccines to many Americans prior to their first SARS-CoV-2 infection substantially reduced the infection fatality rate (IFR) these Americans suffered when they did contract COVID-19.

In this paper, we use these observations, together with a structural epidemiological model, to argue that the combined success in slowing SARS-CoV-2 transmission through behavior change and the widespread delivery of vaccines saved close to 800,000 American lives.

We argue that relative to historical experience with pandemic influenza and modeling based on this experience, this public health success was a surprise. As of March 2020, it was not at all clear that it would be possible to slow the spread of SARS-CoV-2 long enough to develop vaccines and deliver them to the American population in time to save lives. We see the success of behavior-based mitigation of SARS-CoV-2 transmission as one of the most important public health lessons of this pandemic—it is, in fact, possible to slow the spread of a dangerous respiratory disease for quite a long time.

But, at the same time, these mitigation efforts came at a tremendous economic, social, and human cost. To avoid similar pain from mitigation in the next pandemic, we argue that we need to make investments now not only in vaccine development, but also in data infrastructure so that we can precisely target mitigation efforts to minimize the economic and social impacts of mitigation with the next pathogen. One might think of these investments in data infrastructure as similar in spirit to the huge investments

1. In what follows, we use the term "COVID-19" to refer to the disease caused by the SARS-CoV-2 virus.

made in the infrastructure to gather economic data after World War II to better guide economic policy. For population-level infectious disease mitigation policies to be effective at low economic and social cost, they need to be guided by detailed real-time epidemiological, demographic, and behavioral data, which are only available in a crisis if one is prepared in advance to gather such data.

I. Our Estimate of Lives Saved

We estimate that behavioral mitigation and vaccination together saved close to 800,000 American lives between February 15, 2020 and February 15, 2024. This estimate is based on three data sources: serology data capturing immunity derived from SARS-CoV-2 infection and COVID-19 vaccination in the American population, data on the dynamics of COVID-19-associated deaths, and linked vaccine and COVID-19 mortality data from thirty US states. We describe the construction and interpretation of these data in section II.

Our estimate rests on two central premises. First, due to the immune evasion capabilities of SARS-CoV-2, the overwhelming majority of Americans would have become infected with SARS-CoV-2 by February 2024 under any realistic vaccination and behavioral mitigation scenario. Second, the health risk of a person's first infection, when unvaccinated, is vastly higher than one's risk after having been vaccinated or previously infected. Thus, the benefit of behavioral mitigation and vaccination came principally from vaccinating individuals before their first SARS-CoV-2 infection. The serology data indicate that slightly more than two-thirds of the US population were vaccinated prior to their first infection with SARS-CoV-2; it is this group that principally contributes to our estimate of lives saved.

In support of these premises, an estimated 94 percent of Americans had been infected with SARS-CoV-2 by late 2022, despite the behavioral mitigation and vaccine uptake in the preceding years (Klaassen and others 2023). Population-level data on COVID-19 mortality for those who had been vaccinated versus those who had not been vaccinated gathered from thirty US states with linked mortality and vaccine data are consistent with the view that COVID-19 was extremely dangerous for those who contracted it for the first time without protection from vaccines. For those contracting COVID-19 after vaccination or prior infection, the disease is much less dangerous.

Table 1. Model-Implied Cumulative COVID-19 Deaths

Baseline and alternative scenarios	
Baseline behavior and vaccines	1,180,000
Baseline behavior, no vaccines	1,979,000
No mitigation with vaccines	3,345,000

Source: Authors' calculations.

Based on these premises, we construct both a back-of-the-envelope calculation of the lives saved by mitigating behavior and vaccines and an estimate from a structural epidemiological model that considers behavior, decline in the COVID-19 IFR over time, and waning immunity against both reinfection and severe disease. The back-of-the-envelope calculation, which conjectures that the 68 percent of Americans that managed to get vaccinated prior to their first SARS-CoV-2 infection would have suffered an IFR four times higher had they not been vaccinated, leads us to an estimate of 845,000 lives saved.

We develop a full structural model to delve a bit deeper into this calculation and set ourselves up for conducting counterfactual exercises. The model combines a fairly detailed epidemiological description of the various variants of SARS-CoV-2 that have appeared over the past four years with a simple model of how mitigating behavior reacts to the rise and fall of daily deaths from the disease as well as parameters governing the administration of vaccines.[2] We argue that this model fits both the dynamics of the data on COVID-19 deaths and the dynamics of the serology data on infections and vaccinations quite well. Our model's implications for cumulative COVID-19 deaths from February 15, 2020 through February 15, 2024 are shown in the first row of table 1.

We simulate the model with vaccines turned off to arrive at a counterfactual prediction for the dynamics of COVID-19 deaths in the absence of vaccines, with results for cumulative mortality in this counterfactual reported in line 2 of table 1. The use of the full structural model with its added detail delivers our preferred estimate of just under 800,000 lives saved as the difference between cumulative deaths reported on line 2 and line 1.

We then draw out four lessons for future pandemics from these data and our counterfactual modeling exercises.

2. We have presented versions of this model in earlier work, including Atkeson (2021a, 2021b, 2023b).

I.A. Lesson 1

First, we argue that it was the combination of mitigating behavior and vaccines together that saved lives.

To illustrate this point, we conduct two counterfactual model simulations. We simulate our model with its baseline specification of mitigating behavior but without vaccines. Without vaccines, behavior alone would have postponed infections, but in the end, nearly everyone would have been infected and subject to a high IFR from that first infection.

We then simulate our model with vaccines distributed starting at the end of December 2020 but with no mitigating behavior before that time. We report our model-implied cumulative death toll for this scenario in the third row of table 1. In this counterfactual simulation, we see that, without a behavioral response, vaccines would have come too late to save lives. Our model implies that cumulative COVID-19 deaths would also have been substantially higher in this scenario without mitigation because our serology and deaths data imply that COVID-19 was substantially more dangerous in 2020 than in 2021, and most infections in this scenario would have occurred in 2020.

One might be tempted to use this scenario of an unmitigated epidemic as a benchmark against which to argue that the combination of vaccines and behavior together saved over 2 million lives. We argue that such a comparison would be an overstatement as it seems highly implausible that there would be no private efforts to avoid transmission even in the absence of any public mitigation policies. The model simulation of an unmitigated epidemic has the daily death toll peaking at over 60,000 deaths per day. It seems highly likely that people would have reacted on their own to such an outcome even in the absence of any public policies toward the epidemic.

I.B. Lesson 2

This success of delaying infections for many months through changes in behavior was a surprise relative to historical experience and modeling of pandemic influenza.

We take as the strongest piece of evidence in favor of this claim the conclusion of Ferguson and others (2006), a prominent study of mitigation options for a pandemic influenza in the United States, regarding the timing of administration of vaccines that these vaccines would have "almost no effect" (451) if started after 120 days after the first worldwide case because at this time horizon, they would be too late to save lives. Clearly,

mitigation of COVID-19 bought us many more than 120 days for vaccines to have a significant impact on COVID-19 cumulative mortality.

I.C. Lesson 3

To a remarkable extent, this strong behavioral response to COVID-19 through 2020 and 2021 was universal across all fifty states.

Certainly, there are significant differences in cumulative mortality from COVID-19 across states, but we argue that the outcomes across US states have much more in common than any of them (except New York City) have with the predicted impact of an unmitigated epidemic. We take these common dynamics of COVID-19 across states as strong evidence of the importance of an endogenous behavioral reaction to current disease incidence as predicted by many economic models.[3]

And yet, this observation leads us to our fourth lesson.

I.D. Lesson 4

It is unclear what behavioral reaction to expect in response to the next epidemic.

Epidemiologists have noted the impact of changes in behavior on the dynamics of prior epidemics, particularly in attenuating the initial phase of exponential growth of infections predicted by simple epidemiological models.[4] But figuring out how to predict the quantitative impact of such changes in behavior and how private behavior will respond to public health measures has proved an unsolved challenge.[5]

We see any successful theory of behavior as having to confront a wide range of data across different epidemics. For example, as noted above, the success of public and private changes in behavior in slowing the spread of COVID-19 came as a surprise relative to historical experience. And yet New York City suffered a terrible first wave of deaths from COVID-19 early in the pandemic largely due to a delayed reaction to the disease despite clear warnings from the Italian experience a few weeks earlier. Somehow the evidence of COVID-19 deaths in New York City seemed to have a much bigger impact on behavior elsewhere in the United States than did the European experience despite objective evidence that air travel links were likely to spread the disease across the globe.

3. See, for example Atkeson (2021b), Gans (2022), and Atkeson, Kopecky, and Zha (2024) and the papers cited therein.

4. See, for example, Chowell and others (2016) and Eksin, Paarporn, and Weitz (2019).

5. See, for example, Ferguson (2007) and Funk and others (2015).

Of particular concern is the question of how our collective experience with COVID-19 over the past four years will influence behavioral responses to the next pandemic for perhaps a generation or more.

The remainder of our paper is organized as follows. In section II, we review the data used in our study. In section III, we summarize the main features of our structural epidemiological model. In section IV, we present our main results and the four main lessons we take away from these results. In section V, we lay out more specifically the types of investments in data infrastructure that we believe would be useful in preparing to do more targeted mitigation in the next pandemic. Finally, in section VI, we conclude.

In section A of the online appendix, we compare the implications of our model to other estimates of lives saved in the literature. In section B of the online appendix, we look more closely at the cross-section of outcomes for COVID-19 cases, vaccinations, and deaths across US states. We use our model to argue that the range of outcomes observed are consistent with plausible variation either in the strength of the behavioral reaction or in state-specific structural factors having an impact on transmission rates. Disentangling the importance of these factors as well as state-level variation in IFRs is something we leave for future research. In section C of the online appendix, we give a full description of our model and its parameters.

II. Serology and Mortality Data

In this section, we review the serology data and the data on mortality from COVID-19 that we use in choosing parameters for our model and constructing our estimate of the impact of behavior and vaccines on cumulative mortality from this disease.

II.A. Serology Data

The serology data we use are drawn from two surveys.

As described in Jones and others (2021), the Centers for Disease Control and Prevention (CDC) measured SARS-CoV-2 seroprevalence (the population-level prevalence of immune markers in the blood) from 2020–2022 by testing for antibodies against two distinct viral antigens in samples from blood donors. One of these antibody types (against type S antigen) is generated in response to either a prior infection or vaccination. The other antibody type (against type N antigen) is generated only in response to prior infection. Thus, with some caveats, the pair of positive or negative results for each sample allows one to measure whether the individual making the blood donation had been previously infected (with

or without vaccination), vaccinated without prior infection, or neither vaccinated nor previously infected.[6] We refer to this survey as the Blood Donor Survey.[7]

As described in Bajema and others (2021), serology data were also collected from samples from commercial blood testing labs. These data measure only whether the person giving the sample had previously been infected. We refer to this as the Commercial Lab Survey.[8]

We note that these serology surveys were drawn from different convenience samples—one a sample of blood donors and the other a sample of those having blood drawn as part of their medical checkups or care. We check for consistency of the measure of those infected across these two sources. Unfortunately, no serology data from a sample designed to be representative of the population are available.

In figure 1, we show the results of the Blood Donor and Commercial Lab serology surveys at the national level for the overall population. The crosses show estimates from the Blood Donor Survey of the cumulative percentage of the population that had experienced infection by the survey date (showing a response to the N antigen). The dots show estimates from the Commercial Lab Survey of the cumulative percentage of the population that had experienced infection by the survey date. We see that the two serology surveys give consistent estimates for the percentage of the population infected at least through the first Omicron wave in early 2022.

The circles in figure 1 show estimates from the Blood Donor Survey of combined seroprevalence. That is, it adds to the percentage showing a response of the N antigen, the percentage of those showing a response to the S antigen but not the N antigen. This additional group is presumed to be vaccinated but not yet infected; these circles show the sum of those with measurable antibodies from infection (whether or not they also have

6. Such caveats include waning immunity, which can cause a previously infected or vaccinated person to test negative on one or both of the antigen tests, and heterogeneity in the immune response, in which a person can mount an abnormally low immune response to the N antigen despite being infected. See also Ong and others (2021).

7. CDC, "2020–2021 Nationwide Blood Donor Seroprevalence Survey Infection-Induced Seroprevalence Estimates," https://data.cdc.gov/Laboratory-Surveillance/2020-2021-Nationwide-Blood-Donor-Seroprevalence-Su/mtc3-kq6r/about_data; "2022–2023 Nationwide Blood Donor Seroprevalence Survey Combined Infection- and Vaccination-Induced Seroprevalence Estimates," https://data.cdc.gov/Laboratory-Surveillance/2022-2023-Nationwide-Blood-Donor-Seroprevalence-Su/ar8q-3jhn/about_data.

8. CDC, "Nationwide Commercial Laboratory Seroprevalence Survey," https://data.cdc.gov/Laboratory-Surveillance/Nationwide-Commercial-Laboratory-Seroprevalence-Su/d2tw-32xv/about_data.

Figure 1. National-Level Results of the Blood Donor and Commercial Lab Serology Surveys for the Overall Population

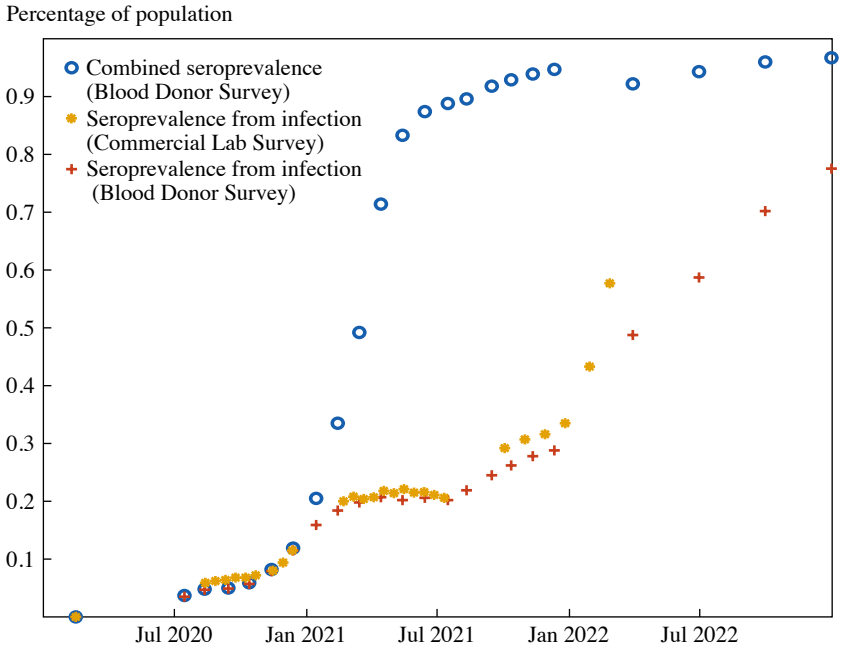

Source: CDC.

been vaccinated) and those with antibodies from vaccination but not from infection.

Several features of these serology survey data stand out. First, we see that the estimated cumulative percentage of the population infected as of January 2021 was quite low—well below 20 percent for the overall population. That is, efforts to slow the spread of SARS-CoV-2 through 2020 appear to have succeeded.

We see from the gap between the circles and the crosses/dots that the rapid deployment of vaccines succeeded in vaccinating a large portion of the population prior to first infection by the late summer of 2021. Consistent with this rapid deployment of vaccines in the first half of 2021, we see slow growth in the estimate of cumulative infections between January 2021 and July 2021. From the start of 2021 through the summer of that year, the combination of behavior and vaccinations appeared to be on a path of ending the epidemic.

Unfortunately, due to the combination of new variants (Delta and Omicron) and waning of the immunity provided by vaccines and prior infection, in the fall of 2021, we see infections continue to rise, particularly so in 2022. Given that variants of Omicron have continued to show the ability to infect those who had previously been vaccinated (and reinfect those with prior infections), it is likely that by early 2024, an overwhelming majority of the population has experienced a COVID-19 infection. Considerations of herd immunity that were prominently discussed early in the pandemic have turned out ex post not to be relevant due to a combination of immune evasion by new variants and waning immunity.

II.B. Mortality Data

We now turn to our data on mortality from COVID-19. We draw these data from the CDC's COVID Data Tracker website.[9] This data set counts deaths from COVID-19 at both the national and state levels, with deaths for New York City broken out separately.

Figure 2 shows cumulative and weekly COVID-19 deaths at a daily rate for the United States from February 2020 to February 2024 in panels A and B, respectively. The mortality data are shown as dotted lines. The outcomes predicted by our baseline model simulation are shown as solid lines.

As shown in this figure, the COVID-19 epidemic in the United States has played out in a series of waves, particularly over the first two years of the epidemic. While these waves garnered considerable attention at the time, what we find most striking about this pattern is that, from very early on in the epidemic, cumulative COVID-19 deaths grew roughly linearly. This linear growth of cumulative deaths is clearly faster in the first two years of the epidemic (from February 2020 through February 2022) than in the second two years of the epidemic.

Why do we find the linear growth of cumulative COVID-19 deaths over the past four years striking? What is missing in figure 2 is any substantial initial period of exponential growth of cumulative deaths as would be predicted by standard epidemiological models for a novel pathogen. To our minds, this observation of linear growth in cumulative deaths sustained over a four-year period is one of the most remarkable features of the COVID-19 epidemic in contrast with historical experience with influenza

9. The data can be downloaded from CDC, "Provisional COVID-19 Death Counts, Rates, and Percent of Total Deaths, by Jurisdiction of Residence," https://data.cdc.gov/NCHS/Provisional-COVID-19-death-counts-rates-and-percen/mpx5-t7tu/about_data.

Figure 2. Baseline Model COVID-19 Deaths for the United States (February 2020 to February 2024)

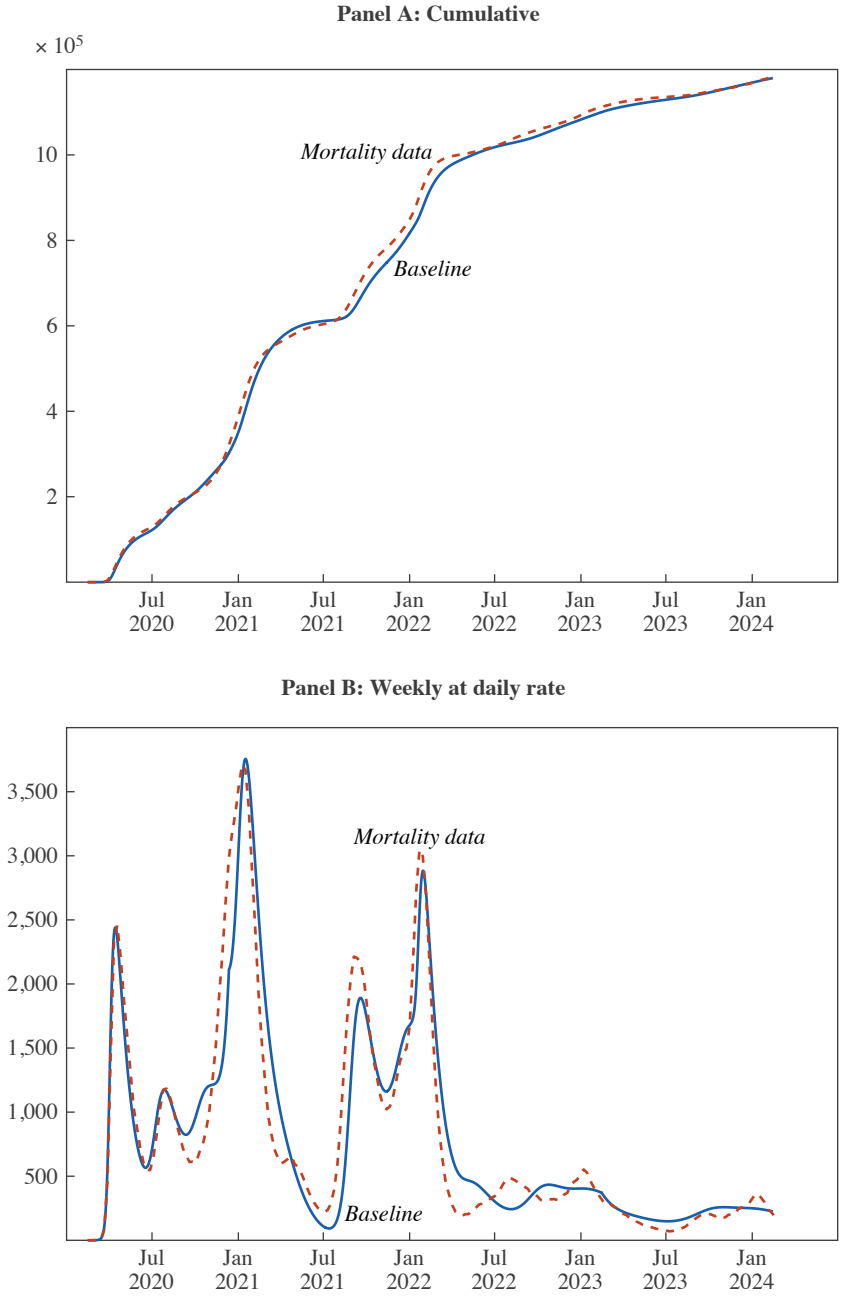

Source: CDC COVID Data Tracker and authors' calculations.

and the predictions of many epidemiological models. In our model, this outcome is attributed to the strength of the public and private behavioral responses to mitigate transmission of SARS-CoV-2.

II.C. Implied Infection Fatality Rates (IFRs)

Note that these serology and deaths data together imply that the IFR for COVID-19 declined over the course of 2020 and 2021 and then again with Omicron. In particular, the Blood Donor and Commercial Lab Surveys give identical estimates that 11.5 percent of the US population had been infected by December 2020. The cumulative COVID-19 death toll by the end of 2020 was close to 390,000. Given a US population of 332 million, this would imply an overall IFR close to 1 percent.

Looking at the same numbers prior to the first big Omicron wave, in November 2021, the Blood Donor Survey estimated that 27.8 percent of the population had been infected and the Commercial Lab Survey estimated that 31.6 percent of the population had been infected, while the CDC estimates that just over 800,000 Americans had died of COVID-19 by the end of November 2021, implying that close to 61 million Americans were infected with SARS-CoV-2 between January 1 and November 2021 (if we take 30 percent infection-induced seroprevalence as our estimate for November 2021). These numbers imply an IFR closer to 0.66 percent over the course of 2021 prior to Omicron. The equivalent numbers after the first large Omicron wave show a substantial further decline in the implied IFR. We use these estimates as a guide for parameterizing the IFR in our model.

II.D. Mortality by Vaccine Status

We make use of population-level data on the realized COVID-19 mortality rates of the vaccinated and unvaccinated. As discussed in Jia and others (2023), thirty states of the United States integrated their vaccine databases with their reporting of mortality data. Thus, for these states, on a weekly basis, one can measure the number of COVID-19-related deaths among those who had received the two doses of the primary series of vaccines at least fourteen days before death and COVID-19-related deaths among those who had not received these primary vaccines. The CDC also estimates the number of people in these states in these two groups, and thus one can construct a weekly COVID-19 mortality rate for the vaccinated and unvaccinated populations.

In panel A of figure 3, we show data on the weekly age-adjusted COVID-19 mortality rates for those with two doses of a primary vaccine

Figure 3. Weekly COVID-19 Mortality Rates for the Vaccinated and Unvaccinated Populations

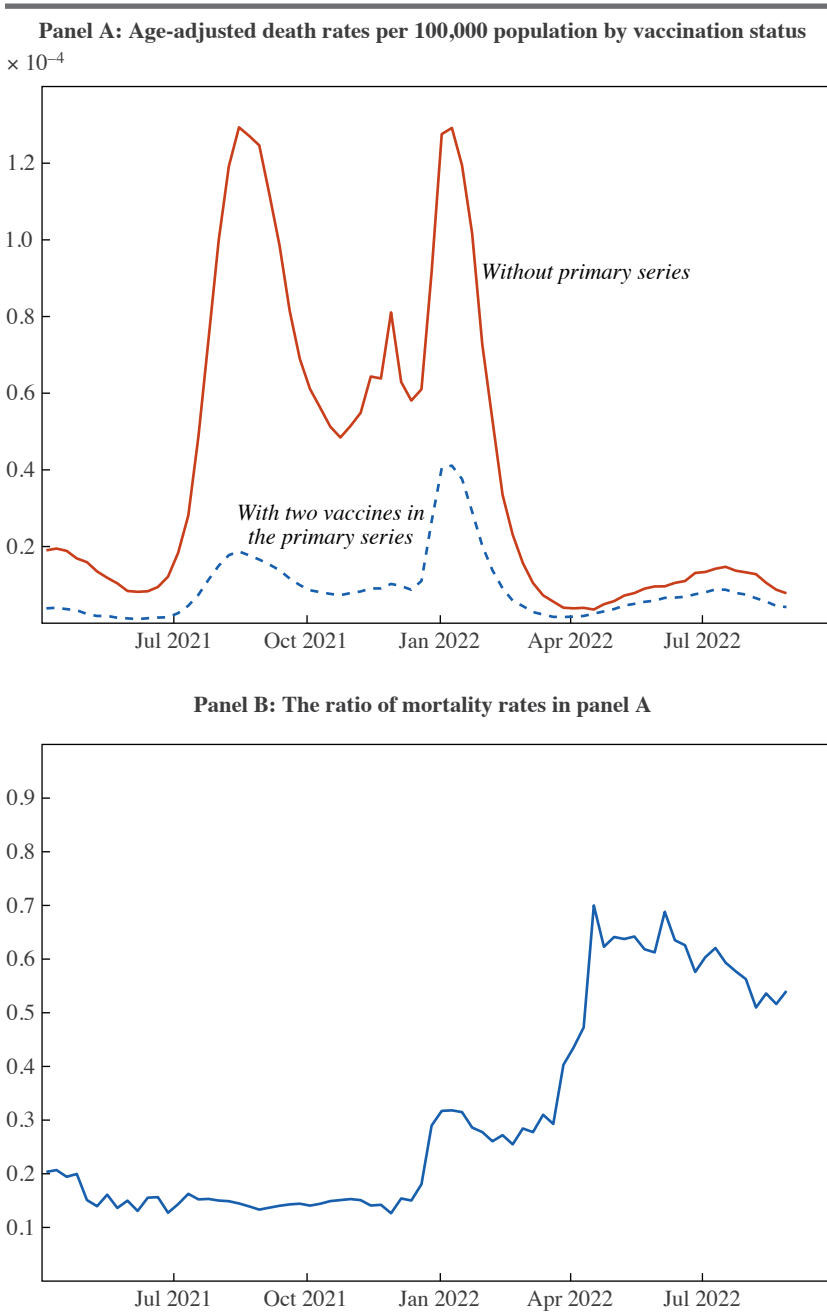

Source: CDC.

(from the first half of 2021) at least fourteen days prior to death (dashed line) and those without this protection from vaccines (solid line).[10] The dates given on the *x* axis are the year and week number used in the CDC's *Morbidity and Mortality Weekly Report* (*MMWR*). We see in this figure that the weekly mortality rate for the unvaccinated was much higher than for the vaccinated in 2021. After the first big Omicron wave, the weekly COVID-19 mortality rate for the unvaccinated falls to meet the low mortality rate for the vaccinated.

In panel B of figure 3, we show the ratio of these two mortality rates. We see in the panels of this figure that vaccination reduced the COVID-19 mortality rate on the order of 85 percent until the first Omicron wave. After that first Omicron wave, we see that the difference in mortality rates by vaccination status was much smaller. We conjecture, based on the serology data, that this outcome arose as the majority of the unvaccinated had come to have the protection of a prior SARS-CoV-2 infection. Thus, the mortality rates for the unvaccinated fell to a level much closer to that for the vaccinated as both groups were largely protected after this first Omicron wave.

II.E. Waning Immunity and the Long Tail of COVID-19 Deaths

With the emergence of Omicron variants, we have seen that both vaccines and prior infection provide only temporary protection against new infections. As a result, the prevalence of SARS-CoV-2 infections has remained high over the past two years despite the fact that by the end of the first quarter of 2022, the overwhelming majority of the US population had already been vaccinated or experienced a prior SARS-CoV-2 infection or both.

This outcome is a result of two factors. One is that the protection offered by vaccines and prior infection against reinfection wanes over time. The other is that the continual evolution of the virus allows new versions of it to evade immune defenses. After two years of Omicron and three years of experience with mRNA vaccines, it is clear that both processes are at work with COVID-19, but their relative importance is difficult to disentangle.[11]

10. We use the data available at CDC, "Rates of COVID-19 Cases or Deaths by Age Group and Vaccination Status and Booster Dose," https://data.cdc.gov/Public-Health-Surveillance/Rates-of-COVID-19-Cases-or-Deaths-by-Age-Group-and/d6p8-wqjm/about_data.

11. See, for example, Jung and others (2024).

III. Summary Description of the Model

We now present our structural model of the impact of behavior and vaccines on cumulative mortality from COVID-19 in the United States over the period from February 15, 2020 to February 15, 2024. This model extends that in Atkeson (2023b). A full description of this model is given in the online appendix that accompanies this paper.

III.A. Purpose and Fit of the Model

Recall that our estimate of the impact of behavior and vaccines on cumulative COVID-19 mortality is based primarily on an accounting of the number of Americans who were able to get vaccinated prior to their first SARS-CoV-2 infection, and to a lesser degree on an estimate of the benefits of delaying infections due to a decline over time in the IFR of COVID-19. We see this model as a formal accounting device to account for the dynamics of the COVID-19 infection fatality ratio implied by the serology data and the transition of the epidemic toward an endemic steady state.

Thus, while we do not formally estimate the parameters of this model, we do evaluate it as an accounting device on the basis of its fit to the dynamics of SARS-CoV-2 infections and COVID-19 vaccinations at the national level as measured by the serology data in figure 1 as well as the dynamics of deaths from COVID-19 at the national level as shown in figure 2.

We show the model fit to the serology data in figure 4. Panel A compares the model estimate of the fraction of the population with protection from severe disease due to prior infection, taking into account waning immunity as described below (this fraction can be either vaccinated or not), to the serology data on the fraction of the population showing antibodies from prior infection.

Panel B of figure 4 compares the model estimate of the fraction of the population showing antibodies from vaccination but not prior infection, again taking into account waning immunity as described below, to the serology data on the fraction of the population showing antibodies from vaccination but not prior infection.

In this figure, we see that the fit of the model with its baseline parameters to the dynamics of infections, vaccinations, and deaths is quite good.

We then use this model to assess several counterfactuals to estimate the impact of behavior and vaccines on cumulative mortality from COVID-19 in the United States over the past four years. It is here that the structure of the model is harder to assess as we do not observe these counterfactual outcomes in the data. As we describe the model, we aim to describe what

Figure 4. Fit of the Model to Serology Data

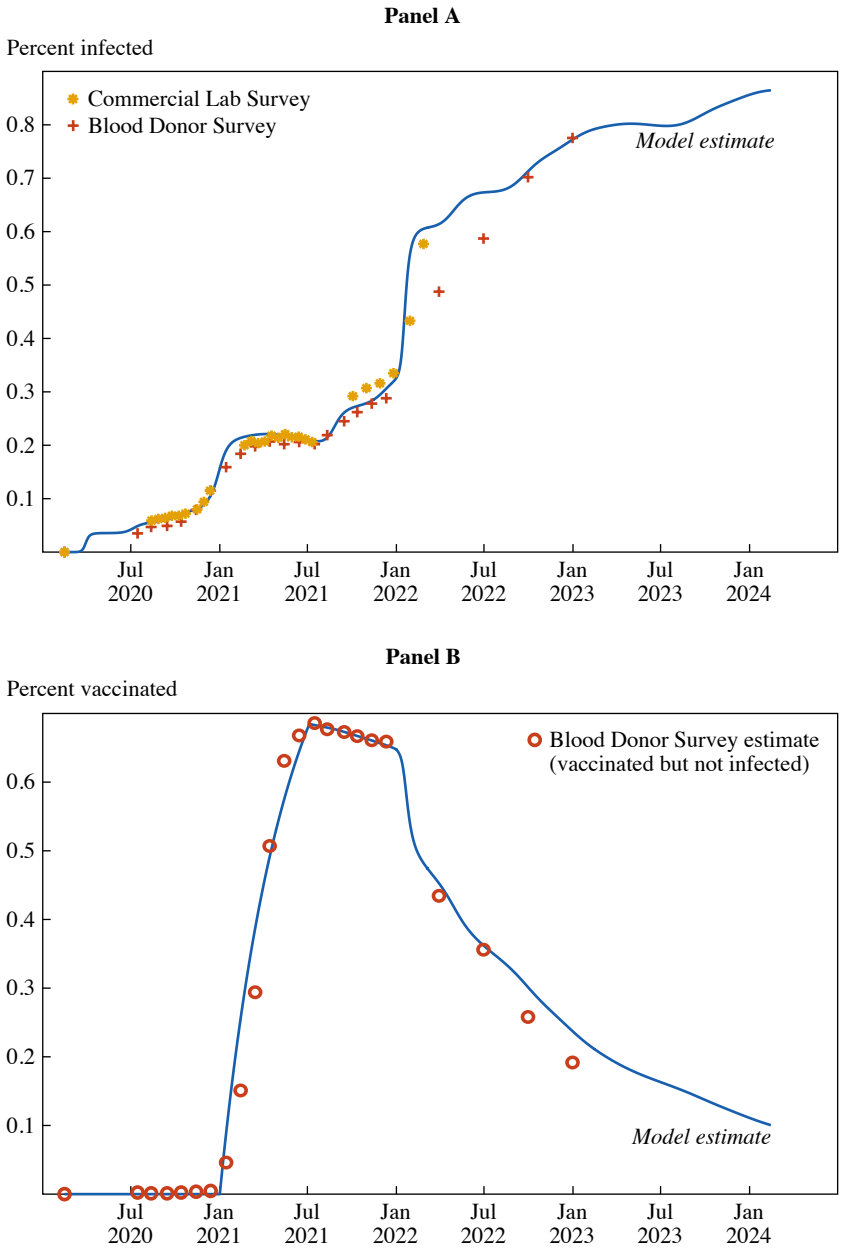

Source: CDC and authors' calculations.

Note: In panel B, the model-implied percentage vaccinated equals to $V(t)/0.75$ where $V(t)$ is the portion of the population with effective protection after vaccination in the model.

features of the data that we do observe allow us to identify the key parameters driving our model's implications for these counterfactuals. In particular, we focus on describing why we have some confidence in our choices for the parameters governing the nature and strength of the behavioral response in the model.

III.B. Model Structure

The model is a susceptible-exposed-infectious-hospitalized-resistant-susceptible (SEIHRS) model with waning immunity and introduction of the Alpha, Delta, and Omicron variants as the epidemic progresses. This model extends the workhorse susceptible-infectious-recovered (SIR) epidemiological model in several dimensions. We explain the reasons for these extensions after reviewing some basic epidemiological concepts.

To begin, a standard SIR model of an epidemic views the population at any point in time as being divided into three categories: susceptible to infection $S(t)$, currently infected and capable of spreading the disease $I(t)$, and resistant to the disease $R(t)$ either from natural immunity (including that induced by prior infection) or from vaccination.

This distribution of characteristics across the population is assumed to evolve over time as follows. Those that are currently infectious, $I(t)$, are assumed to stop being infectious at rate γ per unit time. A fraction η of those who stop being infectious do so because they die. We thus refer to η as the IFR.

Those currently infectious encounter other agents in the population at random and transmit their disease to those agents met at a rate $\beta(t)$ per unit time. We refer to $\beta(t)$ as the *transmission rate*. We allow the transmission rate to depend on factors inherent to the pathogen and the environmental location as indicated by a parameter $\bar{\beta}$ as well as time-dependent factors such as seasonality and behavioral responses.

Since the expected length of time that an infectious agent is expected to be in this state is $1/\gamma$, the average number of agents that an infectious person will transmit their disease to is given by $\beta(t)/\gamma$. Since only fraction $S(t)$ of those agents are actually susceptible to the disease, the expected number of new infections caused by a single infectious agent is given by what is called the *effective reproduction number*:

$$\mathcal{R}_{eff}(t) = \frac{\beta(t)}{\gamma} S(t).$$

Note that the average length of time that an infectious agent remains infectious (here $1/\gamma$) in this model also corresponds to the average length of time between one individual becoming infectious and subsequent infections caused by that individual. This length of time is referred to as the *generation interval*.

The effective reproduction number is related to the SIR model implied growth rate of the fraction of the population that is infectious by

$$\frac{\dot{I}(t)}{I(t)} = \left(\mathcal{R}_{eff}(t) - 1\right)\gamma$$

where $\dot{I}(t)$ denotes the derivative of infections with respect to time.

We note two points from this formula. First, we see that the question of whether the epidemic—in terms of $I(t)$—is growing or shrinking over time is determined by whether the effective reproduction number is above or below one.

Second, the speed of growth of infections per unit time is determined both by the effective reproduction number and the generation interval. Thus, to match data on the growth rate of infections (or deaths) per unit time, one must take a stand on these two parameters. In our model, we hold the generation interval fixed across variants and aim to match the dynamics of weekly deaths in the data with differences in inherent transmissibility of different variants, a seasonal influence on transmissibility, and a behavioral response to the current level of deaths.

We now explain the dimensions in which we extend this simple model and why we do so. We then review our choices for parameter values, with a focus on the generation interval, IFRs, transmission rates, and the impact of behavior on these transmission rates.

We add compartments to the simple SIR model as follows. We add both an exposed state E and the hospitalized state H. Agents in the exposed state have contracted the disease but are not yet infectious. This is a common modification of the SIR framework. Inclusion of this state enriches the dynamics of initial growth of the epidemic. We describe below the purpose of the hospitalized state H. We also add a vaccinated state V to count those who have been vaccinated prior to their first SARS-CoV-2 infection. In terms of protection against infection and severe disease, this state is equivalent to the R state counting those with immunity from prior infection.

To allow for different SARS-CoV-2 variants to have different transmission rates and different IFRs, the compartments E and I are further broken

down by variant i, where i indexes the original variant, and the Alpha, Delta, and Omicron variants.

The rate at which agents leave the E_i compartment for both the normal and more transmissible variants is σ and the rate at which agents leave the I_i compartments for all variants is γ. We also include compartments E_i and I_i corresponding to those experiencing breakthrough Omicron infections. These individuals are modeled as having immunity to previous variants but not to Omicron. The purpose of these additional states is to allow the IFR for breakthrough infections to differ from that of other infections.

With these assumptions, the mean generation time for the model is then $1/\sigma + 1/\gamma$. We set this generation time in line with estimates from the CDC.[12] As mentioned above, this generation time sets the time scale of the epidemic implied by the model.

III.C. The Model of Behavior and Disease Transmission

We use an ad hoc model of the impact of behavior on transmission rates. Specifically, the reduced form for the behavioral response of the transmission rate to the level of daily deaths is given by

$$\beta_i(t) = \bar{\beta}_i \exp\left(-\kappa(t)\frac{dD(t)}{dt} + \psi(t)\right)$$

where the parameters $\bar{\beta}_i$ control the inherent transmissibility of the original and subsequent variants of SARS-CoV-2, the parameter $\psi(t)$ is used to introduce seasonality in transmission, and $\kappa(t)$ is the semi-elasticity of transmission with respect to the level of daily deaths. Thus, public and private behavior having an impact on transmission is assumed to respond only to the current level of daily deaths.

Five comments regarding this model of behavior are in order.

First, we have assumed that behavior reacts to the current level of daily deaths. As described in Atkeson (2021b), this form of behavior serves to regulate the effective reproduction number and drive it down to one in the initial phase of the epidemic and then keep it close to one for the remaining course of the epidemic. More specifically, such behavior regulates the model-implied growth rate of cumulative deaths to remain roughly

12. See CDC, "COVID-19 Pandemic Planning Scenarios," https://archive.cdc.gov/www_cdc_gov/coronavirus/2019-ncov/hcp/planning-scenarios.html. On that web page, the CDC notes a mean time of approximately six days between symptom onset in one person to symptom onset in another person infected by that individual.

constant over time. We argue throughout this paper that this outcome of roughly linear growth of cumulative COVID-19 deaths is one of the most striking features of the data on COVID-19 deaths, not only in the United States but around much of the world. While this outcome might be predicted by economic theory, it is not universally observed across epidemics. For example, as we discuss in the online appendix, mitigating behavior seems to have taken a different and more persistent form in the recent mpox epidemic. Thus, it is not clear that behavior will take the same form in the next epidemic.

Our second comment concerns the role of the dual assumptions that behavior responds to the daily death rate and not the level of infections and that, due to the presence of the H compartment, daily deaths are essentially a distributed lag of past levels of $I(t)$. As discussed in the appendix of Atkeson (2021b), these assumptions appear to be remarkably successful in allowing the model to match the size of the waves of COVID-19 deaths with each new variant over the past four years. Models in which behavior reacts to the level of infections directly or that do not include this lag have difficulties in matching the size of these waves as, in these cases, behavior is too successful at keeping the effective reproduction number close to one. That is, mitigating behavior reacts so quickly to changes in the level of infections that waves are cut off.[13]

Third, we see the introduction of new variants as exogenous shocks to transmission rates that allow us to identify the strength and timing of the behavioral response of the model. We thus take the observation that the model can match the size and shape of the waves of deaths associated with the introduction of the Alpha, Delta, and Omicron variants as validation of the parameter choices governing the behavioral response in the model, including the delay induced by the H compartment. Moreover, we take from the Omicron wave in which new infections spiked much higher than in previous waves and much higher than deaths did as validation of the assumption that behavior responds to deaths and not infections.

Fourth, the waves of COVID-19 deaths appear to have a seasonal pattern, with summertime lows, which we match with our seasonal factor $\psi(t)$ chosen to follow a sine wave.

Fifth, ideally, one would want to build a model in which agents are fully rational and make decisions about their mitigation behavior, particularly for understanding behavioral responses in the counterfactuals that we

13. On this point, see Droste and Stock (2021) and Atkeson, Kopecky, and Zha (2021).

consider. To build such a model, however, one must take a stand on what agents believe about the risks that they face from the disease, and this can be hard to do in real time. Moreover, it might also be difficult to incorporate the delayed responses of behavior that appear to be critical in reproducing the dynamics of the epidemic that we have observed. We leave these challenges to future research.

Fitting this model to the data has been an ongoing project starting with a first version in Atkeson (2021a). The goal has been to explore whether one could account for the dynamics of the COVID-19 epidemic with a simple model with a stable formulation of behavior. To that end, in previous work and in this model, we find that the strength of the response of public and private behavior having an impact on transmission to the level of daily deaths as indexed by $\kappa(t)$ appears to have relaxed in the late fall of 2020 and remained consistent since then. Specifically, we choose an initially high value of $\kappa(t)$ for the period February 15, 2020 until November 2020, and then $\kappa(t)$ declines to a new level equal to 35 percent of its initial value. We refer to this apparent relaxation of behavior in the face of the level of daily COVID-19 deaths as "fatigue." We find that this onetime change in behavior in our model is required for the model to match the height of the waves of COVID-19 in late 2020 and beyond.[14] This formulation of behavior was chosen early on in this modeling process starting with the first version of this model in February 2021 and has been kept constant since that time.

III.D. Key Parameters

We set the IFRs for the SARS-CoV-2 variants prior to Omicron to be a declining function of time. As discussed above, the serology data estimates for the percentage of the population infected as of the end of 2020 and the data on cumulative deaths at that time imply an IFR of 1 percent for 2020.

We use that value for the IFR for 2020. The corresponding IFR implied by the serology and deaths data for 2021 for the period prior to Omicron is 0.5 percent.

To match the big jump in infections with the first Omicron wave with an increase in deaths that is modest in comparison to what would have happened if Omicron was as deadly as prior variants, we use a lower IFR of 0.15 percent for those infected with Omicron out of the S compartment. We also allow Omicron to infect those in the R compartment (those with

14. Andersson and others (2021) and De Gaetano and others (2023) argue that the impending arrival of effective vaccines may have caused such a relaxation of behavior.

protection from prior infection or vaccination) with a very low IFR. We refer to such infections as *breakthrough infections*.

Having chosen these parameters, we choose the parameters for inherent transmissibility $\overline{\beta}_t$ to match the dynamics of the waves of deaths associated with each of them. As described in the online appendix, these parameters imply a relative transmissibility across variants indexed by the ratio of these parameters that is in line with established estimates.

In modeling the transmissibility of Omicron, one must set two parameters—the constant $\overline{\beta}_t$ reflecting its inherent transmissibility and a parameter governing the probability that a vaccinated or recovered individual suffers a breakthrough infection. These two parameters combine to give Omicron a growth advantage over Delta. There is considerable uncertainty regarding the relative importance of these two parameters. We choose them to match data from South Africa that Omicron had a growth advantage of a factor of three relative to Delta in a population with 85 percent protected by prior immunity as well as our serology and deaths data in that first Omicron wave.[15] We find that our model's implications for the first wave of Omicron deaths are largely invariant to the particular choice of $\overline{\beta}_t$ for Omicron. What does vary as this parameter is varied (and the probability of a breakthrough infection modified to keep the growth advantage of Omicron over Delta at three times) is the size of the wave of initial Omicron infections. We have chosen a pair of parameters to match this growth advantage for Omicron and this wave of infections as indicated in the serology data.

To model the impact of vaccines, we set the rate at which susceptible agents are moved from the S compartment directly to the V compartment equal to $\lambda(t) = 0.0065$ starting on January 1, 2021, and zero before that date. Vaccines are administered at this rate for the first 185 days of 2021. The rate of vaccination then drops to $\lambda(t) = 0.0065/5$ until the end of 2022 and then $\lambda(t)$ is set to zero after that. In the model, the V compartment is equivalent to the R compartment and is simply used to count vaccinations prior to infection.

In our model, agents in compartment $V(t)$ enjoy full protection from infection by the Alpha and Delta variants and substantial protection against death from Omicron in the same way as agents with prior infection (in the R compartment). Thus, we regard the number of agents in this compartment

15. See, for example, Raquel Viana, Sikhulile Moyo, Daniel G. Amoako, Houriiyah Tegally, Cathrine Scheepers, Christian L. Althaus, and others, "Rapid Epidemic Expansion of the SARS-CoV-2 Omicron Variant in Southern Africa," *Nature* 603 (2022): 679–86.

as representing the population that is both vaccinated prior to a first SARS-CoV-2 infection and that gained protection from that vaccination. To model that vaccines are not 100 percent effective, we assume that the portion of those who arrive in the V compartment is 75 percent of the total vaccinated. Thus, when we compare the model implications for $V(t)$ to the measures from the serology data on those vaccinated but not infected in figure 4, we plot $V(t)/0.75$ as a measure of the total population vaccinated.

We assume that agents flow out of the V and R compartments back to the S compartment and thus become susceptible again to severe disease at a rate corresponding to expected duration of protection against severe disease of three years. Because Omicron can also infect those in the R and V compartments with breakthrough infections (but with a much lower IFR), our model allows protection against reinfection to wane much faster than protection against severe disease. It is this second process that largely accounts in the model for the long tail of COVID-19 deaths that we see over the past two years. Both estimates for the speed of waning are subject to considerable uncertainty.

IV. Main Model Results: Four Lessons

We now use the model to conduct counterfactual experiments to explore the impact of behavior and vaccines on cumulative mortality from COVID-19 in the United States over the past four years. We focus on drawing four lessons from the model.

IV.A. Lesson 1: Behavior and Vaccines Together

As discussed above, we show our model's baseline implications for the dynamics of COVID-19 deaths, infections, and vaccinations in figures 2 and 4. We show the model's baseline implications for cumulative COVID-19 mortality over the four-year period from February 15, 2020 to February 15, 2024 in the first row of table 1 above.

We show the model implications for COVID-19 deaths with the baseline parameters governing behavior but with no vaccines in figure 5 and in the second row of table 1. As indicated in the second row of table 1, the model implies that absent vaccines, the cumulative death toll over the past four years would have been 1,979,000. That is, our model implies that, given baseline behavior, vaccines saved 799,000 lives. We take this as the headline result of this paper.

We see in figure 5 that most of these additional deaths would have occurred in 2021. After the first big Omicron wave in early 2022, the model

Figure 5. Baseline Model Behavior but No Vaccines

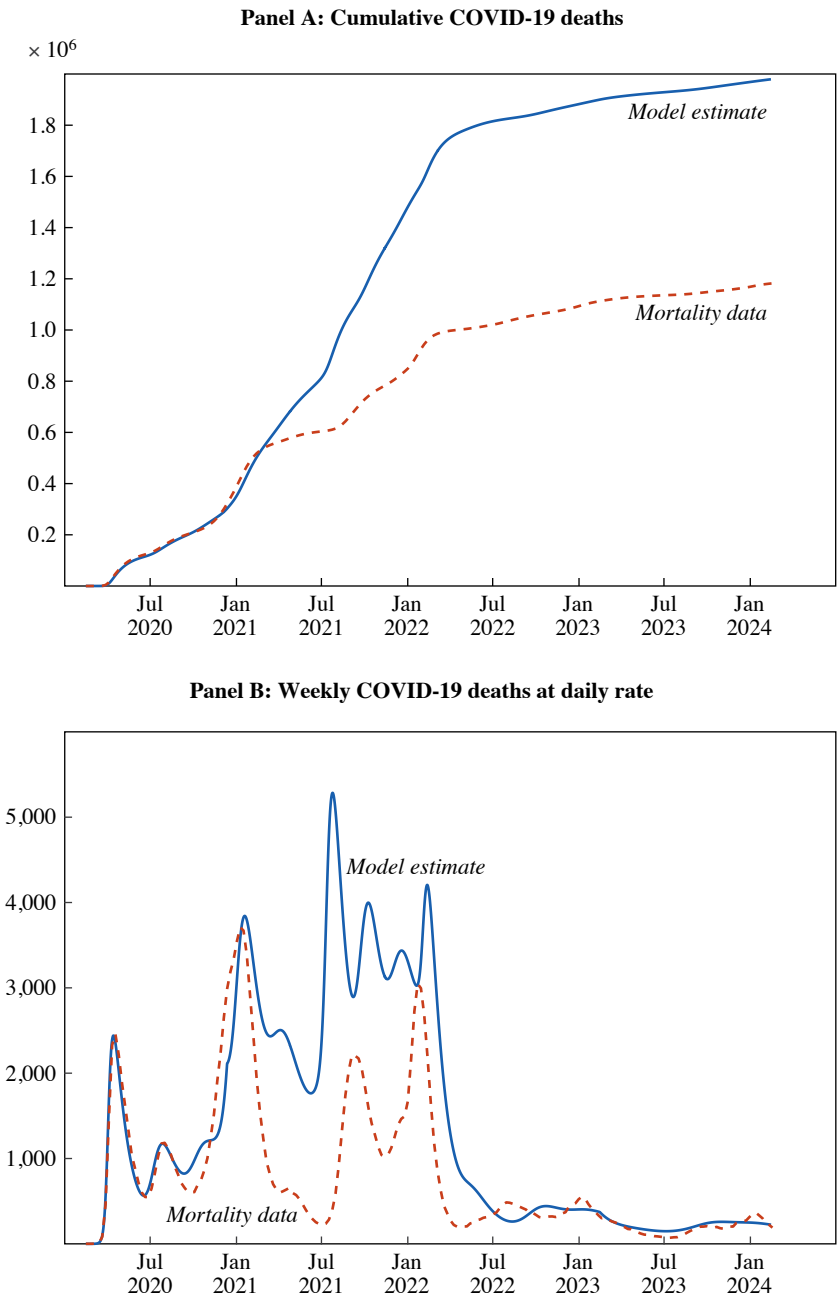

Source: CDC COVID Data Tracker and authors' calculations.

Figure 6. Dynamics of Infections in the Absence of Vaccines

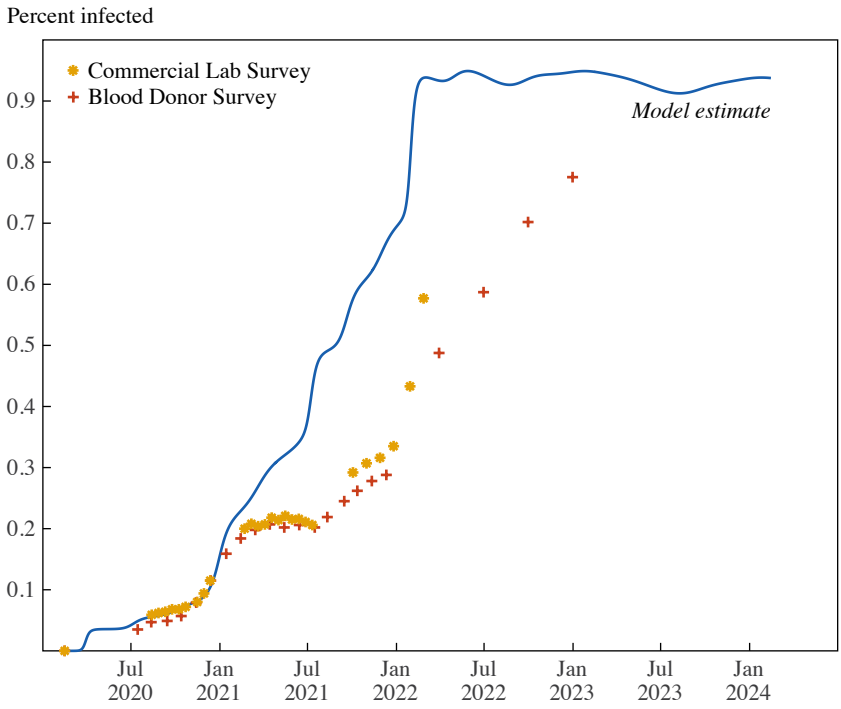

Source: CDC and authors' calculations.

implications for COVID-19 deaths with and without vaccines are nearly the same. This is because, in the absence of vaccines, the model implies that close to 95 percent of the population would have experienced their first SARS-CoV-2 infection by the end of that first Omicron wave and thus the level of population protection against severe disease after that point would have been similar with or without vaccines. This prediction of our model for the dynamics of infections in the absence of vaccines is shown in figure 6.

In our model, we assume that vaccination reduced the IFR from first infection with SARS-CoV-2 in 2021 from 0.005 to 0.0013 (or 25 percent of the IFR for the naive unvaccinated).[16] Thus, to understand our counterfactual

16. Recall that we assume that 25 percent of those who receive a vaccine do not end up with protection, while the other 75 percent gain complete protection until either their immunity wanes or they suffer a breakthrough infection with Omicron.

estimate of the COVID-19 death toll in the absence of vaccines, imagine that this 68 percent of the population had instead been infected without the protection of vaccines and had, as a result, suffered the full IFR of 0.005 rather than 0.0013. Had this occurred, the counterfactual death toll from COVID-19 in the absence of vaccines would have been 847,000 higher than the baseline with vaccines.[17] Our full model delivers a slightly lower estimate of lives saved due to the arrival of Omicron in late 2021, which had a lower IFR than prior variants, and the assumption that the protection against severe disease offered by vaccines (and prior infection) wanes over time. But one can clearly see from this calculation the simple logic underlying our estimate of the impact of behavior and vaccines on cumulative mortality from COVID-19 in the United States.

These blood serology data highlight how important the interaction of behavior change and vaccine development and deployment were in saving lives. Had SARS-CoV-2 swept through the US population in an unmitigated epidemic, it is likely that the overwhelming majority of the US population would have been infected by early fall of 2020, leaving much less room for people to benefit from being vaccinated prior to their first infection.

We illustrate this point by simulating the model with the behavioral parameter $\kappa(t) = 0$. As shown in figure 7, in this simulation, the vast majority of the US population gets infected by the late summer of 2020.[18] We report the implied cumulative death toll in the third row of table 1. Here we see an extraordinary model-implied death toll, consistent with an IFR of 1 percent applied to nearly the entire US population in 2020 together with subsequent deaths in later years due to waning immunity.

Clearly, any estimate of lives saved depends on the assumed counterfactual. What impact should this have on our thinking about the next pandemic? From an ex ante perspective as of March 2020, the premises on which our ex post estimation is based would have been hard to predict. Was it going to be possible to delay transmission for the time required to develop and deliver effective vaccines? If vaccines had taken much longer to arrive or had offered less protection against severe disease, would the whole exercise of slowing transmission have been a wasted effort?

17. The calculation is $0.68 \times 0.75 \times 0.005 \times 332{,}000{,}000$ where the last term is the US population.

18. Our results for the cumulative mortality of an unmitigated epidemic during 2020 are worse than those in Ferguson and others (2020) in part because our estimate of the basic reproduction number of the original variant is higher (we assume 3 while they assumed 2.5) and thus an unmitigated epidemic infects more of the population and in part because our estimated IFR at the start of the epidemic is slightly higher (1 percent versus 0.9 percent).

Figure 7. Model with No Behavioral Response

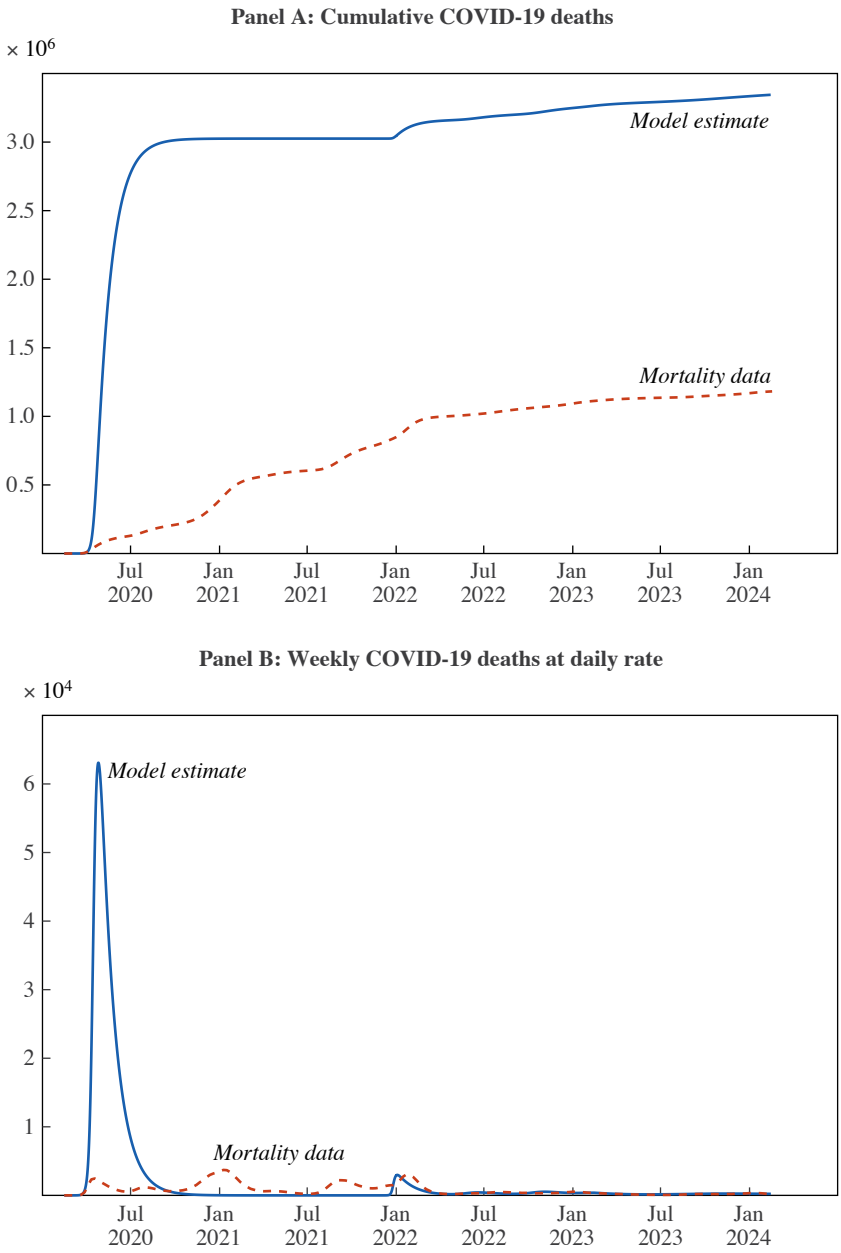

Source: CDC COVID Data Tracker and authors' calculations.

Based on these simulations comparing the death toll with baseline behavior and no vaccines to that with no mitigating behavior, we argue, in short, no—as such, mitigation efforts would have still helped to reduce strain on a severely overburdened health care system and bought critical time to learn how to better care for patients with severe disease even in the absence of vaccines. Such considerations are important to bear in mind when considering which behavioral interventions should be adopted. Without behavioral responses to the epidemic, an unmitigated epidemic would have been much more severe than even our counterfactual with behavior but without vaccines.

IV.B. Lesson 2: Strength and Duration of the Behavioral Response Was a Surprise

We argue now that the success in slowing the spread of COVID-19 during 2020 and 2021 evident in the serology data came as a surprise relative to both historical experience with pandemic influenza and model-based estimates of the impact of mitigation measures on transmission based on that historical evidence.

In many ways, pandemic influenza was the closest historical and epidemiological parallel to the COVID-19 epidemic. Both diseases are fast-moving respiratory diseases with potentially high IFRs. The case of the 1918–1919 "Spanish Flu" epidemic was viewed as particularly relevant, but the epidemics of 1957, 1968, and 2009 also served as examples.

The risk of a new pandemic influenza has been viewed as a substantial threat for a long time. See, for example, the disease and economic scenarios laid out by the President's Council of Economic Advisers in September 2019 (CEA 2019), which foresaw the potential for hundreds of thousands of deaths and trillions of dollars of economic disruption from a pandemic influenza.

In response to this threat from pandemic influenza, epidemiologists have invested considerable effort into studying historical experiences and modeling the impact of various mitigation options on influenza transmission.[19]

19. For examples of studies of transmission during the 1918–1919 pandemic influenza, see Mills, Robins, and Lipsitch (2004), Fraser and others (2011), and Eggo, Cauchemez, and Ferguson (2011). For studies of the impact of mitigation on transmission during the 1918–1919 pandemic influenza, see, for example, Bootsma and Ferguson (2007), Hatchett, Mecher, and Lipsitch (2007), Correia, Luck, and Verner (2022), and Velde (2022).

Of particular interest in this regard is figure 1 in Hollingsworth and others (2011), which shows the duration (in weeks) and effectiveness (in terms of percentage reduction in transmission rates) of historical interventions to slow the spread of the 1918–1919 influenza and SARS-CoV-1. That figure estimates that interventions in the 1918–1919 influenza pandemic reduced transmission rates by less than 50 percent in all cases and much less than that amount in many cases. Moreover, these interventions were sustained for less than fifteen weeks. As shown in that figure, mitigation efforts for SARS-CoV-1 were estimated to be much more effective, but these were also sustained for less than fifteen weeks. In comparison, with COVID-19, we see from the serology data that efforts to slow disease spread substantially had an impact for many months through late 2021.

Prominent studies of the possibilities for using public health interventions to contain a new influenza strain at its source include Longini and others (2005) and Ferguson and others (2005). Prominent modeling studies of the use of broader public health measures, including school closures and social distancing, to slow the spread of a pandemic influenza that broke through efforts to contain it at the source include Ferguson and others (2006) and Germann and others (2006). Universally, these studies predict short periods of very rapid spread of disease even in the modeled presence of intense public health efforts to slow disease spread so that available flu vaccines can be administered.

Particularly telling in this regard is the caption of figure 4 in Ferguson and others (2006, 451) that notes, regarding the timing of administration of vaccines, these vaccines would have "almost no effect" if started after 120 days after the first worldwide case.[20] This conclusion is clearly too pessimistic about the possibility of controlling the spread of a respiratory pathogen through behavioral mitigation, as COVID-19 vaccines still had a major benefit despite arriving more than a year after the first worldwide case. The COVID-19 pandemic fundamentally changed our conception of what is possible with respect to behavioral mitigation.

This contrast between the anticipated and observed impact of behavioral change on slowing the transmission of COVID-19 is even more remarkable given that the original strain of SARS-CoV-2 was more contagious than a pandemic influenza strain was expected to be, had the ability to spread prior to the onset (or in the absence) of symptoms, and ultimately

20. See also figure 2 in Germann and others (2006).

generated new variants with substantially increased transmissibility. The cards were stacked against us, even relative to the modeled scenarios for pandemic influenza that served as our basis for our earliest understanding of SARS-CoV-2.[21]

To see this point, consider the scenarios for pandemic influenza expected by modelers as laid out in Meltzer and others (2015). Table 1 in that paper lays out the range of scenarios for transmissibility and clinical severity of potential new pandemic influenzas typically considered, and figure 1 in that paper places historical pandemics in this space of transmissibility and clinical severity. The original strain of SARS-CoV-2 had higher transmissibility than the worst-case scenario and was near to the worst-case scenario in terms of its clinical severity.

The fact that SARS-CoV-2 could be transmitted prior to showing symptoms made epidemiologists (including ourselves) pessimistic that its spread could be effectively controlled. As described in Fraser and others (2004, 6146), "the success of . . . control measures is determined as much by the proportion of transmission occurring prior to the onset of overt clinical symptoms (or via asymptomatic infection) as the inherent transmissibility of the etiological agent (measured by the reproductive number R_0)." Likewise, early in the COVID-19 epidemic, Hellewell and others (2020) pointed to the pre- and asymptomatic transmission of COVID-19 as a reason to be pessimistic about our ability to contain its spread.

These features of COVID-19, together with the hazy prospects as of March 2020 for developing an effective vaccine in time to be useful, meant that despite all the planning for a pandemic influenza the set of actionable targeted mitigation policies available to slow the spread of COVID-19 in a cost-effective manner was very small. In fact, in an early and highly cited article from March 9, 2020, giving broad outlines of options for mitigating the coming pandemic, Anderson and others (2020, 934, emphasis added) remarked that "it is easy to suggest a 60 percent reduction in transmission will do it or quarantining within 1 day from symptom onset will control transmission, *but it is unclear what communication strategies or social distancing actions individuals and governments must put in place to achieve these desired outcomes*."

We argue that one of the main lessons of our experience with COVID-19 is there are far greater possibilities for slowing transmission of a deadly respiratory virus than previously thought. Given that new knowledge, we

21. See, for example, Davies and others (2020).

should work urgently to determine how to achieve similar behavioral mitigation in the next pandemic but at far lower cost.

IV.C. Lesson 3: Behavior and State-Level Outcomes

There has been great interest in comparing the impact of COVID-19 across states of the United States in the press and in some academic work.[22] Certainly the outcomes for cumulative mortality for COVID-19 vary widely across the states of the United States. What accounts for these differences? We address this question in greater detail in section B of our online appendix.

Here we make the argument that, relative to the historical and modeling benchmarks for pandemic influenza discussed above, residents of all fifty states made surprisingly strong and lasting efforts to slow the spread of SARS-CoV-2 so that vaccines came in time to save a considerable number of lives.

To illustrate this point, in online appendix figure B.6, taken from Chitwood and others (2022), we show the dynamics of the effective reproduction number for SARS-CoV-2 for each of the fifty states of the United States. In this figure, we observe that behavior in all fifty states changed rapidly and dramatically so as to drive the effective reproduction number of COVID-19 in the state down to one very early on in the epidemic. Moreover, this behavior was sufficiently sustained to keep this effective reproduction number close to one throughout 2020. Atkeson, Kopecky, and Zha (2024) find similar results for both US states and many countries.

As we have discussed above, if the effective reproduction number of a disease remains close to one, then the growth rate of current infections is close to zero. Equivalently, the growth rate of cumulative infections and deaths is then roughly constant. This is precisely the dynamics we observe in cumulative COVID-19 mortality at the state level.

To illustrate this point, in online appendix figure B.7, we show the dynamics of cumulative COVID-19 deaths as an age-adjusted death rate per 100,000 of the population for selected states. In the left panel of this figure, we show the dynamics of cumulative COVID-19 deaths for California, Florida, New York (excluding New York City), and Texas. We see that New York State had a very rapid growth of cumulative deaths in the initial phase of the pandemic and then settled into a lower growth rate. Texas had

22. See, for example, Barro (2022), Bollyky and others (2023), and Kerpen, Moore, and Mulligan (2022).

a high growth rate of cumulative deaths throughout the first two years of the pandemic. Given the rhetoric surrounding this topic, we find it striking how similar the age-adjusted outcomes for COVID-19 deaths have been for California and Florida over the past four years.

In the right panel of online appendix figure B.7, we show the dynamics of cumulative COVID-19 deaths as an age-adjusted death rate per 100,000 of the population for New York City and seven other states representing extreme high and low mortality outcomes across states. With the exception of New York City, we see largely linear growth in cumulative deaths over the first two years of the COVID-19 epidemic for all of these locations. As evident in the figure, New York City suffered exceptionally rapid initial growth of cumulative COVID-19 deaths in the first wave of the epidemic, likely due to the surprise introduction of a large number of hidden cases from Europe in early 2020.

For further evidence of this commonality of responses across US states, in online appendix figure B.8, we show estimates from the Commercial Lab and Blood Donor serology surveys of cumulative infections and combined seroprevalence for the fifty states of the United States. While these surveys show considerable variation in the estimated percentage infected across states, we see in this figure that all of the states followed similar dynamics of slow growth in infections in the first two years of the pandemic and rapid deployment of vaccines in the first half of 2021.[23] We discuss these state-level serology data in greater detail in online appendix section B.

Based on this evidence, we argue that the most important feature of the outcomes across US states (and even countries around the world) is how much they have in common relative to outcomes that were expected given prior epidemiological modeling of and past experiences with pandemic influenza. To a large extent, residents of every state in the United States outside of New York City reacted very strongly to COVID-19 very early on and took significant actions to slow its spread all through 2020 and 2021. We regard the observation that this could be done and done nearly universally across different states of the United States, as a great surprise.

To expand further on this point, observe that the model-based forecast in Ferguson and others (2020) for peak deaths with unmitigated spread of COVID-19 was over sixteen deaths per day per 100,000 population (implying over 50,000 deaths per day in the United States as a whole) with

23. Chitwood and others (2022) argue that the serology data underestimate the true portion of the population ever infected for a variety of reasons. This paper presents alternative estimates of the state-level portion of the population infected through 2020 in its figure 7.

75 percent of the population being infected by late summer of 2020. This forecast was not out of line with what was experienced in locations that did little to mitigate the spread of SARS-CoV-2. For example, we note that seroprevalence studies in Manaus, Brazil indicated an attack rate of 75 percent in the first wave of the pandemic (Buss and others 2021). We see nothing like this rapid spread of COVID-19 in the serology data across US states.

In March and April 2020, New York City experienced the worst wave of COVID-19 cases and mortality of anywhere in the United States over the past four years. Its peak weekly mortality rate was sixty per 100,000 population (less than ten per 100,000 per day)—in the range of one-half that predicted in Ferguson and others (2020) for peak deaths with unmitigated spread. Seroprevalence estimates for New York City indicate up to 20 percent of that population of 8 million people was infected in the first wave in the spring of 2020 (Stadlbauer and others 2021).

We illustrate the extent to which the first wave of COVID-19 deaths in New York City was an outlier in online appendix figure B.9. In that figure, we show the dynamics of weekly COVID-19 deaths for the fifty states at an age-adjusted rate per 100,000 of population. As is clear from the figure, the first wave of COVID-19 deaths in New York City was much larger than any other wave experienced in any state in the United States. That is, the response to flatten the curve and dramatically slow the transmission of COVID-19 was universal across the fifty states of the United States.

We now turn to our fourth lesson regarding the prospects for a similar behavioral response next time.

IV.D. Lesson 4: Unclear If Behavior Will Be the Same Next Time

From our perspective, the success of this sustained and fairly uniform behavioral response to slow transmission for this length of time to allow for the deployment of vaccines and improved medical care is perhaps the biggest surprise of the COVID-19 pandemic. Clearly, a strategy of slowing transmission for eight to fifteen months as needed to develop and deliver an effective vaccine is based on the premise that people can be persuaded to go along with that plan. To an extent that seems well outside historical experience with pandemic influenza and predictions based on that historical experience, Americans did go along with that plan, with or without mandates from state governments.

Many economists, one of us included, have argued ex post that this pattern of adjusting behavior to keep the growth rate of new infections and deaths relatively close to zero, observed nearly universally in the United

States and across many countries, is precisely the response that economic theory would predict.[24]

But this argument then raises the puzzle of why did we not see a quantitatively similar response to pandemic influenzas, in particular the 1918 Spanish Flu? And comparison of these different outcomes for pandemic influenza and COVID-19 raises the question of which behavioral response should we expect to see in the next pandemic? Will it be a short, sharp wave as for COVID-19 in New York City in March and April 2020 and in most cities for which we have data from 1918? Or will it be a long, drawn-out affair as for COVID-19 in the rest of the United States? The answer to this question will have a big impact on the range of mitigation strategies available in the face of the next pandemic and is a great challenge in epidemiological modeling (Funk and others 2015).

The world has already experienced an outbreak of another emerging pathogen. Starting in May 2022, mpox, formerly known as monkeypox, began to spread rapidly primarily through sexual contact between men, with this spread being particularly alarming since it showed up in a large number of countries in a short period of time. Mpox is an example of a known pathogen endemic to a relatively small area (in West Africa) suddenly spreading rapidly well outside that region.

Simple examination of the exponential growth of cases in the United States between May and August 2022 indicated that this disease had the potential to spread quite broadly, at least within a subset of the US population. Instead, the number of cases began to die out rapidly in late August, and new cases in the United States have been held at a low level throughout 2023. What explains this path of this epidemic? It appears that through a combination of a sustained change in private sexual behavior and the targeted application of vaccines, it was possible to dramatically reduce the number of cumulative cases relative to what would be predicted for an unmitigated epidemic. In other words, the behavioral response to mpox appears to be a remarkable success.

Zhang and others (2024) quantify the impact of behavior and vaccines on the spread of mpox with an epidemiological model using data on the mpox outbreak in the United Kingdom, which exhibited an epidemic curve similar to that in the United States. These authors argue that changes

24. See Atkeson (2021b, 2023a) and Atkeson, Kopecky, and Zha (2024). See Gans (2022) for a broader survey of the economics papers on this topic.

in behavior and vaccination together played an important role in shaping this epidemic.

We find it interesting to note, however, that their accounting of the impact of behavior and vaccines on the trajectory of this epidemic is quite different from our accounting of the impact of these factors on the trajectory of the COVID-19 epidemic. In particular, they find that the response of behavior (in terms of men reducing the number of their sexual partners) was strong and persistent enough to drive the effective reproduction number of mpox below one on a sufficiently sustained basis to drive the number of new cases to a very low level. This never happened with COVID-19. They then estimate that the use of pre-exposure vaccines for susceptible men limited the threat of resurgence.

The estimated combined impact of these interventions was then very substantial in limiting the size of the outbreak: the United Kingdom had 3,250 observed cases over the study period relative to an estimated final size of an uncontrolled epidemic of 169,400 cases. We see these estimates, together with the discussion in Daskalakis, Romanik, and Jha (2024), as driving home the message that targeted interventions in combination with vaccination can have a powerful impact on outcomes of an epidemic.

Another factor to consider going forward is the extent to which our experience with COVID-19 will shape reactions to new epidemics going forward for decades to come. Will the public be more skeptical of public health warnings about new pathogens? Or will our collective experience with significant mortality from an infectious disease outbreak lead us to take future threats more seriously? Addressing such questions seems of first-order importance for research going forward.

V. What Is Needed to Make Mitigation Less Painful Next Time?

The behavioral mitigation measures undertaken during the COVID-19 pandemic helped to save many thousands of lives, but they came at a high social and economic cost. Uncertainty about key features of COVID-19 and about the human behaviors that had an impact on its spread forced us to take stronger, more widespread, and longer-lasting behavioral mitigation measures than might have been necessary in a more information-rich setting. Likewise, individuals largely lacked the tools they needed to make informed assessments about their risk of becoming infected or transmitting disease. For example, widespread and cheaply available diagnostic tests—along with clear guidance on how to report and interpret

them—could have helped alleviate the need for general physical distancing measures like school and workplace closures that lasted for many months into the pandemic.[25]

The next pandemic may look very different from COVID-19, but it will nevertheless be critical to find ways to rapidly reduce our uncertainty about the pathogen's characteristics and the human behaviors that underlie its spread, and likewise to rapidly develop and deploy the tools that will empower individuals to make informed behavioral choices. This will require developing off-the-shelf research protocols for learning about transmission routes, the natural history of infection, and the dynamics of immunity for an emerging pathogen soon after it is first detected. In the meantime, we must also invest in ongoing data collection efforts to provide baseline measurements against which data on an emerging infectious disease can be meaningfully compared. A detailed discussion of the steps needed to effectively prepare for the next pandemic is provided by Lipsitch and others (2023). Here, we outline a few key considerations.

V.A. Assessing Transmission Routes

When an emerging outbreak is detected, a critical first task is to determine the pathogen's routes of transmission. Beyond the most basic information on transmission route (e.g., sexual versus vector-borne versus respiratory, and [if respiratory] droplet versus aerosol versus fomite), it is also important to identify the venues and behaviors that are most conducive to spread. For example, it became evident early in the COVID-19 pandemic that outdoor transmission was far less common than indoor transmission (Bulfone and others 2021) and that singing was a particularly high-risk activity (Hamner and others 2020). Preapproved study designs, backed with funding for rapid deployment, would help to more rapidly clarify how and where the bulk of transmission occurs in the event of an emerging outbreak.

To place these studies in the proper context, we also require detailed studies on interpersonal contact patterns, both at baseline and as they evolve over the course of an outbreak, much like the CoMix study did in the context of COVID-19 (Gimma and others 2022). Such studies recruit representative cohorts and ask questions about their behaviors (e.g., conversational or sexual contacts) that may be relevant to the spread of disease. Mobility data—gathered, for example, using mobile phones—can also be useful (Buckee and others 2020), though such data must be interpreted with

25. See, for example, the discussion in Atkeson and others (2020).

care since the owners of mobile devices or the users of a given app may not be representative of the broader population (Wesolowski and others 2016). Data access and privacy issues should also be proactively addressed well in advance of a public health crisis.

Detailed contact tracing data can be useful for determining the level of risk associated with various types of contact. For respiratory infections, household transmission studies like the Office for National Statistics (ONS) Coronavirus (COVID-19) Infection Survey in the United Kingdom (Pouwels and others 2021) can be helpful for assessing the level of risk associated with close contact. For sexually transmitted infections, partnership surveys can serve the same purpose.[26] The value of such studies can be greatly enhanced by collecting pathogen genomic information, allowing researchers to distinguish direct within-household (or within-partnership) transmission from new introductions from the community.

The production and distribution of nonpharmaceutical interventions (NPIs) should be rapidly scaled up in the event of an emerging outbreak. In the early stages of an outbreak, plausible effectiveness should be enough to justify the use of sufficiently low-impact NPIs—for example, plausible effectiveness would justify the widespread use of masks against the early spread of SARS-CoV-2 or condoms to prevent transmission of mpox, even in the absence of direct studies assessing the efficacy of those interventions for those specific pathogens. In tandem, the effectiveness of these NPIs should be continuously monitored so that their use can be founded on more direct, pathogen-specific evidence or, if no effectiveness is found, their use can be phased out.

V.B. Describing the Course of Infection

Once infection occurs, it is critical to understand the risk of various health outcomes. Key statistics like the IFR are subject to bias that can affect early estimates in both directions: early in an epidemic, the most severe cases are the ones that are most likely to be detected, thus skewing the IFR upward; yet, if the epidemic is spreading rapidly, a simple division of mortality by cumulative prevalence can skew the IFR downward, since recently infected individuals have not yet had time for their cases to worsen. This underscores the need for principled studies to track the range, timing, and probability of potential health outcomes in an emerging epidemic. An understanding of the IFR and related risks of various health outcomes helps to set the appropriate level of behavioral response.

26. For example, Ueda and others (2020).

Similarly, it is important to rapidly assess how a person's infectiousness varies over time. Again, household or partnership studies can be helpful, especially when coupled with frequent, quantitative diagnostic testing (e.g., RT-qPCR tests to assess pathogen load) and detailed symptom reporting. A critical piece of information to gather early in an epidemic is how the timing of symptoms relates to infectiousness, as this relationship plays a major role in determining how difficult it is to ultimately control a pathogen's spread (Fraser and others 2004). If infectiousness precedes symptoms, the need to develop and deploy rapid diagnostic tests becomes paramount.

V.C. Tracking Incidence and Immunity

Public health response in the United States is largely coordinated at the state level, which poses major challenges for data sharing and standardization. The need for improved data collection, standardization, and dissemination is a major focus area of the new Center for Forecasting and Outbreak Analytics (CFA) based at the CDC. The CFA has taken many cues from the National Weather Service (George and others 2019), and indeed a digital infrastructure for providing information on current epidemiological conditions and a near-term forecast would go a long way toward informing more targeted behavioral responses in the event of another public health crisis.

Alongside information on disease incidence, well-designed serological studies can be invaluable both for reconstructing what has happened after an outbreak ends (as we have tried to do in this report) and for informing on the dynamics of immunity. It is important to conduct ongoing serological studies so that proper baselines can be set, especially because serological tests can cross-react.[27] Serological studies can inform on the duration of immunity to infection, thus helping individuals to calibrate their behavior to better match their risk of infection.

VI. Conclusion

The behavioral response to COVID-19 in 2020–2022 was highly—and unexpectedly—effective in reducing cumulative COVID-19-related mortality in the United States. We estimate that the combination of behaviorally driven transmission reduction and vaccination resulted in roughly 800,000 lives saved during that time period, in line with other estimates.

27. For example, serological tests for SARS-CoV-1 can turn positive based on exposure to a related common coronavirus; see Patrick and others (2006).

Critically, we see that both of these factors—a strong behavioral response and the relatively fast development of an effective vaccine—were needed to yield a substantial reduction in mortality. Had a vaccine not been developed or had behavior not changed, we anticipate that much of the US population would have received their first immunological exposure to SARS-CoV-2 from infection rather than vaccination, and thus the total mortality from the pandemic would have been much higher.

We had three main goals in writing this report: (1) we sought to provide an evidence-based estimate of the value of behavior change during the COVID-19 pandemic in terms of reduced mortality; (2) we sought to describe a straightforward modeling framework that can be adapted to assess counterfactual scenarios for COVID-19 and for infectious diseases more generally; and (3) we sought to discuss the lessons of the COVID-19 pandemic from three hypothetical perspectives: the ex ante perspective of a public health planner in March 2020, with knowledge of basic parameters of the virus but no certainty about its future evolution; the ex post perspective, where we are today, performing an assessment of how we actually performed given our knowledge of how the pandemic actually unfolded; and the perspective of future public health planners, who will be responsible for responding to new, possibly very different, emerging infectious diseases. We now discuss each of these goals in turn.

There are many ways to estimate lives saved during the COVID-19 pandemic, some relying on sophisticated models of transmission and immunological dynamics. We pursue a simpler tack, estimating the total mortality in the scenarios with no behavioral change prior to rollout of a vaccine in January 2021 and with behavior change but no vaccine. Based on serology data, we estimate that less than 20 percent of the US population—and a substantially smaller fraction of individuals over age 65—had been infected with SARS-CoV-2 before the introduction of vaccines. Yet other areas of the world that experienced an impact early on, before an effective behavioral response could be mounted, saw estimated attack rates of up to 75 percent within a short few months, which also aligns with epidemiological models for an unmitigated epidemic with transmissibility similar to the ancestral strains of SARS-CoV-2. As such, we can attribute a mortality reduction in the roughly 55 percent of the population who were able to be vaccinated prior to their first infection to the transmission-reducing behavioral response. Had a successful vaccine not been developed, however, it is unclear whether behavioral response would have had a substantial impact on cumulative mortality through the present day, since immune evasion and increasingly contagious variants of the virus have rendered herd immunity moot.

By introducing a modeling framework, we were able to compare more nuanced counterfactual scenarios and to better separate the impact of behavior from that of vaccination. The framework we discuss here is completely standard, perhaps with the exception of the form of the behavior term, which reduces the transmissibility parameter β_j proportionally to an exponentially decaying function of a parameter κ_j that captures the strength of the behavioral response relative to some disease metric (e.g., total infections or the rate of increase in mortality). It is possible to compare many counterfactual scenarios using this framework, but the main takeaway is that behavioral transmission reduction and vaccination have a powerful positive synergy, where the timing of both is paramount—that is, early behavior change, coupled with the rapid development of an effective vaccine, can pay dividends in reduced mortality.

The ex ante perspective of the public health planner in March 2020 is one lacking in many critical details about the pandemic's ultimate course, and yet it is perhaps the most informative perspective to consider when assessing the best course of action in future pandemics. Early modeling work during the COVID-19 pandemic, including our own, anticipated that SARS-CoV-2 would become endemic (Kissler and others 2020; Shaman and Galanti 2020; Murray and Piot 2021) but failed to anticipate both the ratcheting transmissibility of the virus with successive variants and the relatively swift development of an effective vaccine. The ex post perspective is useful for determining what we might have done differently, but this has limited application for future pandemics.

Instead, rather than thinking about what we should have done differently in hindsight from a management perspective and doing that going forward, we should instead ask what types of information we would have wanted during the early days of the pandemic to make more informed, ex post–like decisions and determine how best to put mechanisms in place now to collect that data. For example, key elements of the natural history of infection and the route of infection—such as the frequency of asymptomatic infections, the role of presymptomatic transmission, and the importance of aerosols in transmission—were unclear for far longer than they should have been.

Developing protocols for rapidly identifying cases, charting their course, and determining likely routes of transmission through prospective household and contact surveys, like the ONS Coronavirus Infection Survey in the United Kingdom (Pouwels and others 2021) and the European CoMix Survey (Gimma and others 2022), are critical for future pandemics. Likewise, it is clear that behavior can change spontaneously in response to a

perceived infectious threat. It is less important to have an exact model for how behavior changes in response to threat than it is to have a robust framework for measuring the relevant changes in behavior when they actually happen. This will require a robust survey-taking machinery to be rapidly deployed in the event of an emerging pathogen. Such work may be augmented by the development of secure contact tracing technologies, like the ones developed for contact notification during the COVID-19 pandemic. Regardless, we must avoid the trap of "fighting the last pandemic," recognizing that while another coronavirus pandemic could occur within our lifetimes, there are many other threats that should be carefully thought through and incorporated into the data-collecting mechanisms discussed here. That said, the experience with mpox, and the fact that behavioral mitigation measures during the COVID-19 pandemic strongly suppressed the spread of various other pathogens (Koutsakos and others 2021), suggests that behavioral mitigation can be an important tool for addressing a wide range of infectious disease threats.

Our findings are limited by a substantial degree of uncertainty in the actual number of infections that occurred during the pandemic and a lack of reliable data capturing the dynamics of behavioral change during the pandemic. Regarding the lack of behavioral data, it is unclear even what an ideal data set would look like, given that we do not have a solid grasp on what types of interactions are necessary and sufficient for the transmission of a respiratory pathogen.[28] Conversational encounters are often used as a proxy, but the precise dynamics of interpersonal transmission in real settings remain poorly understood. The models we use are intentionally simplified and so gloss over much important variation in baseline risk factors, population structure, and viral attributes that can, and do, have a major impact on transmission patterns. Our goal here is to provide a scaffold to guide thinking about behavior-modulated disease transmission, rather than to faithfully recapitulate the dynamics of a particular outbreak—though we note that, under reasonable assumptions, a fairly faithful recapitulation of those dynamics is possible with a model like the one presented here.

ACKNOWLEDGMENTS We would like to acknowledge many helpful conversations with Ben Moll, Michael Droste, Karen Kopecky, James Stock, and Tao Zha.

28. See, for example, Ferretti and others (2023).

References

Anderson, Roy M., Hans Heesterbeek, Don Klinkenberg, and T. Déirdre Hollingsworth. 2020. "How Will Country-Based Mitigation Measures Influence the Course of the COVID-19 Epidemic?" *The Lancet* 395, no. 10228: 931–34.

Andersson, Ola, Pol Campos-Mercade, Armando N. Meier, and Erik Wengström. 2021. "Anticipation of COVID-19 Vaccines Reduces Willingness to Socially Distance." *Journal of Health Economics* 80: 102530.

Atkeson, Andrew G. 2021a. "A Parsimonious Behavioral SEIR Model of the 2020 COVID Epidemic in the United States and the United Kingdom." Working Paper 28434. Cambridge, Mass.: National Bureau of Economic Research.

Atkeson, Andrew G. 2021b. "Behavior and the Dynamics of Epidemics." Working Paper 28760. Cambridge, Mass.: National Bureau of Economic Research.

Atkeson, Andrew G. 2023a. "COVID-19: Epidemiological Models." *Annual Review of Financial Economics* 15: 7–27.

Atkeson, Andrew G. 2023b. "The Impact of Vaccines and Behavior on U.S. Cumulative Deaths from COVID-19." Working Paper 31525. Cambridge, Mass.: National Bureau of Economic Research.

Atkeson, Andrew G., Michael C. Droste, Michael Mina, and James H. Stock. 2020. "Economic Benefits of COVID-19 Screening Tests." Working Paper 28031. Cambridge, Mass.: National Bureau of Economic Research.

Atkeson, Andrew G., Karen A. Kopecky, and Tao Zha. 2021. "Behavior and the Transmission of COVID-19." *American Economic Association Papers and Proceedings* 111: 356–60.

Atkeson, Andrew G., Karen A. Kopecky, and Tao Zha. 2024. "Four Stylized Facts about COVID-19." *International Economic Review* 65, no. 1: 3–42.

Bajema, Kristina L., Ryan E. Wiegand, Kendra Cuffe, Sadhna V. Patel, Ronaldo Iachan, Travis Lim, and others. 2021. "Estimated SARS-CoV-2 Seroprevalence in the US as of September 2020." *JAMA Internal Medicine* 181, no. 4: 450–60.

Barro, Robert J. 2022. "Vaccination Rates and COVID Outcomes across U.S. States." *Economics and Human Biology* 47: 101201.

Bollyky, Thomas J., Emma Castro, Aleksandr Y. Aravkin, Kayleigh Bhangdia, Jeremy Dalos, Erin N. Hulland, and others. 2023. "Assessing COVID-19 Pandemic Policies and Behaviours and Their Economic and Educational Trade-offs across US States from Jan 1, 2020, to July 31, 2022: An Observational Analysis." *The Lancet* 401, no. 10385: 1341–60.

Bootsma, Martin C. J., and Neil M. Ferguson. 2007. "The Effect of Public Health Measures on the 1918 Influenza Pandemic in U.S. Cities." *Proceedings of the National Academy of Sciences* 104, no. 18: 7588–93.

Buckee, Caroline O., Satchit Balsari, Jennifer Chan, Mercè Crosas, Francesca Dominici, Urs Gasser, and others. 2020. "Aggregated Mobility Data Could Help Fight COVID-19." *Science* 368, no. 6487: 145–46.

Bulfone, Tommaso Celeste, Mohsen Malekinejad, George W. Rutherford, and Nooshin Razani. 2021. "Outdoor Transmission of SARS-CoV-2 and Other

Respiratory Viruses: A Systematic Review." *Journal of Infectious Diseases* 223, no. 4: 550–61.

Buss, Lewis F., Carlos A. Prete Jr., Claudia M. M. Abrahim, Alfredo Mendrone Jr., Tassila Salomon, Cesar de Almeida-Neto, and others. 2021. "Three-Quarters Attack Rate of SARS-CoV-2 in the Brazilian Amazon during a Largely Unmitigated Epidemic." *Science* 371, no. 6526: 288–92.

Chitwood, Melanie H., Marcus Russi, Kenneth Gunasekera, Joshua Havumaki, Fayette Klaassen, Virginia E. Pitzer, and others. 2022. "Reconstructing the Course of the COVID-19 Epidemic over 2020 for US States and Counties: Results of a Bayesian Evidence Synthesis Model." *PLOS Computational Biology* 18, no. 8: e1010465.

Chowell, Gerardo, Lisa Sattenspiel, Shweta Bansal, and Cécile Viboud. 2016. "Mathematical Models to Characterize Early Epidemic Growth: A Review." *Physics of Life Reviews* 18: 66–97.

Correia, Sergio, Stephan Luck, and Emil Verner. 2022. "Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu." *Journal of Economic History* 82, no. 4: 917–57.

Council of Economic Advisers (CEA). 2019. *Mitigating the Impact of Pandemic Influenza through Vaccine Innovation*. Washington: Executive Office of the President.

Daskalakis, Demetre, Nikki Romanik, and Ashish K. Jha. 2024. "Lessons from the Mpox Response." *JAMA* 331, no. 5: 387–88.

Davies, Nicholas G., Adam J. Kucharski, Rosalind M. Eggo, Amy Gimma, and W. John Edmunds. 2020. "Effects of Non-pharmaceutical Interventions on COVID-19 Cases, Deaths, and Demand for Hospital Services in the UK: A Modelling Study." *The Lancet Public Health* 5, no. 7: e375–85.

De Gaetano, Alessandro, Paolo Bajardi, Nicolò Gozzi, Nicola Perra, Daniela Perrotta, and Daniela Paolotti. 2023. "Behavioral Changes Associated with COVID-19 Vaccination: Cross-National Online Survey." *Journal of Medical Internet Research* 25: e47563.

Droste, Michael, and James H. Stock. 2021. "Adapting to the COVID-19 Pandemic." *American Economic Association Papers and Proceedings* 111: 351–55.

Eggo, Rosalind M., Simon Cauchemez, and Neil M. Ferguson. 2011. "Spatial Dynamics of the 1918 Influenza Pandemic in England, Wales and the United States." *Journal of the Royal Society Interface* 8, no. 55: 233–43.

Eksin, Ceyhun, Keith Paarporn, and Joshua S. Weitz. 2019. "Systematic Biases in Disease Forecasting—The Role of Behavior Change." *Epidemics* 27: 96–105.

Ferguson, Neil M. 2007. "Capturing Human Behaviour." *Nature* 446, no. 7137: 733.

Ferguson, Neil M., Derek A. T. Cummings, Simon Cauchemez, Christophe Fraser, Steven Riley, Aronrag Meeyai, Sopon Iamsirithaworn, and Donald S. Burke. 2005. "Strategies for Containing an Emerging Influenza Pandemic in Southeast Asia." *Nature* 437, no. 7056: 209–14.

Ferguson, Neil M., Derek A. T. Cummings, Christophe Fraser, James C. Cajka, Philip C. Cooley, and Donald S. Burke. 2006. "Strategies for Mitigating an Influenza Pandemic." *Nature* 442, no. 7101: 448–52.

Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, and others. 2020. *Impact of Non-pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand*. London: Imperial College.

Ferretti, Luca, Chris Wymant, James Petrie, Daphne Tsallis, Michelle Kendall, Alice Ledda, and others. 2023. "Digital Measurement of SARS-CoV-2 Transmission Risk from 7 Million Contacts." *Nature* 626, no. 7997: 145–50.

Fraser, Christophe, Derek A. T. Cummings, Don Klinkenberg, Donald S. Burke, and Neil M. Ferguson. 2011. "Influenza Transmission in Households during the 1918 Pandemic." *American Journal of Epidemiology* 174, no. 5: 505–14.

Fraser, Christophe, Steven Riley, Roy M. Anderson, and Neil M. Ferguson. 2004. "Factors That Make an Infectious Disease Outbreak Controllable." *Proceedings of the National Academy of Sciences* 101, no. 16: 6146–51.

Funk, Sebastian, Shweta Bansal, Chris T. Bauch, Ken T. D. Eames, W. John Edmunds, Alison P. Galvani, and Petra Klepac. 2015. "Nine Challenges in Incorporating the Dynamics of Behaviour in Infectious Diseases Models." *Epidemics* 10: 21–25.

Gans, Joshua S. 2022. "The Economic Consequences of R = 1: Towards a Workable Behavioural Epidemiological Model of Pandemics." *Review of Economic Analysis* 14, no. 1: 3–25.

George, Dylan B., Wendy Taylor, Jeffrey Shaman, Caitlin Rivers, Brooke Paul, Tara O'Toole, and others. 2019. "Technology to Advance Infectious Disease Forecasting for Outbreak Management." *Nature Communications* 10, article no. 3932.

Germann, Timothy C., Kai Kadau, Ira M. Longini Jr., and Catherine A. Macken. 2006. "Mitigation Strategies for Pandemic Influenza in the United States." *Proceedings of the National Academy of Sciences* 103, no. 15: 5935–40.

Gimma, Amy, James D. Munday, Kerry L. M. Wong, Pietro Coletti, Kevin van Zandvoort, Kiesha Prem, and others. 2022. "Changes in Social Contacts in England during the COVID-19 Pandemic between March 2020 and March 2021 as Measured by the CoMix Survey: A Repeated Cross-Sectional Study." *PLoS Medicine* 19, no. 3: e1003907.

Hamner, Lea, Polly Dubbel, Ian Capron, Andy Ross, Amber Jordan, Jaxon Lee, and others. 2020. "High SARS-CoV-2 Attack Rate following Exposure at a Choir Practice—Skagit County, Washington, March 2020." *Morbidity and Mortality Weekly Report* 69, no. 19: 606–10.

Hatchett, Richard J., Carter E. Mecher, and Marc Lipsitch. 2007. "Public Health Interventions and Epidemic Intensity during the 1918 Influenza Pandemic." *Proceedings of the National Academy of Sciences* 104, no. 18: 7582–87.

Hellewell, Joel, Sam Abbott, Amy Gimma, Nikos I. Bosse, Christopher I. Jarvis, Timothy W. Russell, and others. 2020. "Feasibility of Controlling COVID-19 Outbreaks by Isolation of Cases and Contacts." *The Lancet Global Health* 8, no. 4: e488–96.

Hollingsworth, T. Déirdre, Don Klinkenberg, Hans Heesterbeek, and Roy M. Anderson. 2011. "Mitigation Strategies for Pandemic Influenza A: Balancing Conflicting Policy Objectives." *PLoS Computational Biology* 7, no. 2: e1001076.

Jia, Katherine M., William P. Hanage, Marc Lipsitch, Amelia G. Johnson, Avnika B. Amin, Akilah R. Ali, Heather M. Scobie, and David L. Swerdlow. 2023. "Estimated Preventable COVID-19-Associated Deaths due to Non-vaccination in the United States." *European Journal of Epidemiology* 38, no. 11: 1125–28.

Jones, Jefferson M., Mars Stone, Hasan Sulaeman, Rebecca V. Fink, Honey Dave, Matthew E. Levy, and others. 2021. "Estimated US Infection- and Vaccine-Induced SARS-CoV-2 Seroprevalence Based on Blood Donations, July 2020–May 2021." *JAMA* 326, no. 14: 1400–1409.

Jung, Sung-mok, Sara L. Loo, Emily Howerton, Lucie Contamin, Claire P. Smith, Erica C. Carcelén, and others. 2024. "Potential Impact of Annual Vaccination with Reformulated COVID-19 Vaccines: Lessons from the U.S. COVID-19 Scenario Modeling Hub." *PLoS Medicine* 21, no. 4: e1004387.

Kerpen, Phil, Stephen Moore, and Casey B. Mulligan. 2022. "A Final Report Card on the States' Response to COVID-19." *International Journal of the Economics of Business* 30, no. 2: 139–58.

Kissler, Stephen M., Christine Tedijanto, Edward Goldstein, Yonatan H. Grad, and Marc Lipsitch. 2020. "Projecting the Transmission Dynamics of SARS-CoV-2 through the Postpandemic Period." *Science* 368, no. 6493: 860–68.

Klaassen, Fayette, Melanie H. Chitwood, Ted Cohen, Virginia E. Pitzer, Marcus Russi, Nicole A. Swartwood, Joshua A. Salomon, and Nicolas A. Menzies. 2023. "Changes in Population Immunity against Infection and Severe Disease from Severe Acute Respiratory Syndrome Coronavirus 2 Omicron Variants in the United States between December 2021 and November 2022." *Clinical Infectious Diseases* 77, no. 3: 355–61.

Koutsakos, Marios, Adam K. Wheatley, Karen Laurie, Stephen J. Kent, and Steve Rockman. 2021. "Influenza Lineage Extinction during the COVID-19 Pandemic?" *Nature Reviews Microbiology* 19, no. 12: 741–42.

Lipsitch, Marc, Mary T. Bassett, John S. Brownstein, Paul Elliott, David Eyre, M. Kate Grabowski, and others. 2023. "Infectious Disease Surveillance Needs for the United States: Lessons from COVID-19." ArXiv:2311.13724. Ithaca, N.Y.: Cornell University.

Longini, Ira M., Jr., Azhar Nizam, Shufu Xu, Kumnuan Ungchusak, Wanna Hanshaoworakul, Derek A. T. Cummings, and M. Elizabeth Halloran. 2005. "Containing Pandemic Influenza at the Source." *Science* 309, no. 5737: 1083–87.

Meltzer, Martin I., Manoj Gambhir, Charisma Y. Atkins, and David L. Swerdlow. 2015. "Standardizing Scenarios to Assess the Need to Respond to an Influenza Pandemic." *Clinical Infectious Diseases* 60, no. suppl. 1: S1–S8.

Mills, Christina E., James M. Robins, and Marc Lipsitch. 2004. "Transmissibility of 1918 Pandemic Influenza." *Nature* 432, no. 7019: 904–6.

Murray, Christopher J. L., and Peter Piot. 2021. "The Potential Future of the COVID-19 Pandemic: Will SARS-CoV-2 Become a Recurrent Seasonal Infection?" *JAMA* 325, no. 13: 1249–50.

Ong, David S. Y., Paraskevi C. Fragkou, Valentijn A. Schweitzer, Roy F. Chemaly, Charalampos D. Moschopoulos, and Chrysanthi Skevaki. 2021. "How to Interpret

and Use COVID-19 Serology and Immunology Tests." *Clinical Microbiology and Infection* 27, no. 7: 981–86.

Patrick, David M., Martin Petric, Danuta M. Skowronski, Roland Guasparini, Timothy F. Booth, Mel Krajden, and others. 2006. "An Outbreak of Human Coronavirus OC43 Infection and Serological Cross-Reactivity with SARS Coronavirus." *Canadian Journal of Infectious Diseases and Medical Microbiology* 17, no. 6: 330–36.

Pouwels, Koen B., Thomas House, Emma Pritchard, Julie V. Robotham, Paul J. Birrell, Andrew Gelman, and others. 2021. "Community Prevalence of SARS-CoV-2 in England from April to November 2020: Results from the ONS Coronavirus Infection Survey." *The Lancet Public Health* 6, no. 1: e30–e38.

Shaman, Jeffrey, and Marta Galanti. 2020. "Will SARS-CoV-2 Become Endemic?" *Science* 370, no. 6516: 527–29.

Stadlbauer, Daniel, Jessica Tan, Kaijun Jiang, Matthew M. Hernandez, Shelcie Fabre, Fatima Amanat, and others. 2021. "Repeated Cross-Sectional Sero-Monitoring of SARS-CoV-2 in New York City." *Nature* 590, no. 7844: 146–50.

Ueda, Peter, Catherine H. Mercer, Cyrus Ghaznavi, and Debby Herbenick. 2020. "Trends in Frequency of Sexual Activity and Number of Sexual Partners among Adults Aged 18 to 44 Years in the US, 2000–2018." *JAMA Network Open* 3, no. 6: e203833.

Velde, François R. 2022. "A Model of Economic Activity in San Francisco during the 1918 Influenza Epidemic." Working Paper 2022-04. Federal Reserve Bank of Chicago. https://www.chicagofed.org/publications/working-papers/2022/2022-04.

Wesolowski, Amy, Caroline O. Buckee, Kenth Engø-Monsen, and Charlotte Jessica Eland Metcalf. 2016. "Connecting Mobility to Infectious Diseases: The Promise and Limits of Mobile Phone Data." *Journal of Infectious Diseases* 214, suppl. 4: S414–20.

Zhang, Xu-Sheng, Sema Mandal, Hamish Mohammed, Charlie Turner, Isaac Florence, Josephine Walker, and others. 2024. "Transmission Dynamics and Effect of Control Measures on the 2022 Outbreak of Mpox among Gay, Bisexual, and Other Men Who Have Sex with Men in England: A Mathematical Modelling Study." *The Lancet Infectious Diseases* 24, no. 1: 65–74.

Comments and Discussion

COMMENT BY

THOMAS PHILIPPON Atkeson and Kissler provide an important analysis of how vaccines and behavioral changes saved lives during the COVID-19 epidemic. They estimate that the vaccines and associated behavioral responses saved 800,000 lives compared to a counterfactual with no vaccine.

This is a large effect. The lives saved represent approximately 0.25 percent of the US population. As a comparison, the H1N1 epidemic of 1918–1919 killed about 0.65 percent of the US population.

A key result of the paper is that there is a strong complementarity between vaccines and behavioral responses. Without any change in behavior, the death toll would have been close to 1 percent of the population, in large part because many people would have become infected before vaccines became available. By contrast, Atkeson and Kissler estimate that roughly two-thirds of the population got vaccinated before their first infection and that vaccines were very efficient, lowering fatality by a factor of five at the peak of the epidemic in 2021.

THE ROLE OF VACCINES Atkeson and Kissler use granular data to document some key facts and develop a structural epidemiological model to interpret the facts and compute counterfactuals.

Their data come from three sources: serology data to keep track of immunity following infections and vaccinations; deaths associated with COVID-19 over time and across regions; and data linking vaccines and mortality from thirty US states.

The first key estimate is vaccines saved about 800,000 lives, as can be seen from the difference between the first and second lines of table 1 in the

Brookings Papers on Economic Activity, Spring 2024: 113–132

paper. It contains the estimates from the structural model, but a simple back-of-the-envelope calculation is also possible. The serology data suggest that essentially everyone has been infected at least once by now, but that 68 percent of the population received a vaccine before their first infection. The US population in 2020 was around 330 million, and the authors estimate that 1.2 million people have died from COVID-19.

The infection fatality rate (IFR) changes a lot over time for a variety of reasons: new variants, better treatments, vaccines, and so on. In 2020, the virus infected 11.5 percent of the population and killed 390,000 people, which implies an IFR around 1 percent. In 2021, the IFR decreased to 0.66 percent. After 2021, it was around 0.2 percent. Vaccines reduced the IFR by a factor of more than five during 2021 but much less afterward.

Atkeson and Kissler capture these changes in several ways. In the pre-Omicron period, they assume that the IFR falls over time, starting from a high value of 1 percent in early 2020 and eventually decreasing to 0.5 percent, as indicated by the serology data. When the first wave of Omicron arrives with its large increase in infections but a smaller increase in deaths, the implied IFR decreases further to 0.15 percent.[1]

In the authors' baseline calibration, the average IFR over the sample period is then 0.5 percent for the "naive unvaccinated" and 0.13 percent for vaccinated people. A rough estimate of lives saved is then the difference in IFR applied to the population that was vaccinated before the first infection: $(0.5\% - 0.13\%) \times 0.68 \times 330$ million = 830,000. This is in the ballpark of the more precise estimate from the structural model.

THE "SURPRISINGLY" LARGE IMPACT OF BEHAVIORAL CHANGES The second key takeaway from the serology survey data is that mitigation played a crucial role during 2020. By January 2021, less than 20 percent of the population had been infected. By contrast, model simulations predict that essentially the entire population would have been infected without a behavioral response. Together with the 1 percent IFR discussed earlier, this would have led to more than 3 million deaths.

The large role of behavioral changes provides a strong motivation for the development of the structural model. The structural model allows Atkeson and Kissler to study counterfactual experiments that would otherwise be unknowable.

The model is quite advanced and granular. It features waning immunity and takes into account the appearance of variants (Alpha, Delta, and

1. The 0.15 percent applies to susceptible agents (S). The authors also allow breakthrough infections from Omicron (in the R population) but with a very low IFR.

Omicron). A standard susceptible-infectious-recovered (SIR) model splits the population into three groups: susceptible, infectious, and resistant. The authors add two more groups: exposed but not yet infectious (E) and hospitalized (H). They also account for people who are vaccinated prior to their first infection. The E and I groups are indexed by the variant of the virus, and the R group is subject to (rare) breakthrough infections from the Omicron variant.

With this model Atkeson and Kissler obtain reliable estimates of the likely death rates under alternative scenarios. They then argue convincingly that the behavioral response was much stronger and longer lasting than in previous epidemics.

Should we then call it a surprise? I suppose it all depends on the relevant information set. A surprise is, by definition, the difference between an outcome and its expectation based on the prior information set. If we take as our information set the average strength and duration of behavioral responses in previous epidemics, as illustrated by the surveys published before the pandemic, then we must agree with the authors.

I would argue, however, that one should include both the existence of the internet and of the welfare state in our information set. Jones, Philippon, and Venkateswaran (2021) show that the possibility to work and shop remotely had a large impact on COVID-19 mitigation and saved approximately 200,000 lives. These options did not exist in the past, but they were (at least partly) predictable.

Similarly, the 1918 influenza occurred before the expansion of the welfare state. For many households, not working to slow down the spread of the virus would have meant extreme hardship. The $5 trillion fiscal response (Romer 2021) would have been simply unimaginable at the time. The social insurance and public health components of the fiscal response to COVID-19 afforded households the possibility to reduce in-person labor supply, even though the rest of the spending was arguably superfluous (Romer 2021).

PREPARING AGAINST FUTURE CRISES Atkeson and Kissler argue that, to prepare for the next epidemic, we must improve data collection and analytics. We must determine the pathogen's transmission routes and keep track of incidence and immunity.

I fully agree with these points, but I would add several nonpharmaceutical interventions to the list. A striking feature of the COVID-19 epidemic is its unequal impact across groups and locations. As is well known, the virus was ten times more dangerous for old people than for young people. Similarly, Jones, Philippon, and Venkateswaran (2021) show that the risk of

infection varied by a factor of five across occupations. Finally, we see from the current paper that outcomes also differed by a factor of five across states: in terms of fatality rates, New Hampshire and Vermont look like Denmark, while Arizona and Mississippi look like Russia.

These large differences across demographic groups, occupations, and locations imply that we can reduce the severity of future pandemics with targeted interventions. An obvious one is to improve options for remote work and schooling, starting with universally available broadband internet. While some tasks cannot be done remotely, the large differences in exposures across occupations suggest that significant improvements are possible. Similarly, regarding schooling, while we know that in-person teaching is preferable, it seems likely that some form of remote learning will be needed in future crises, and it is therefore important to ensure equal access to computers and reliable internet connections for all students.

Differences across locations are harder to interpret since they reflect differences in preferences as well as governance choices for given preferences. The scale of differences in fatality rates, however, suggests that differences in preferences are unlikely to account for all the variation that we observe across states. Learning and emulating best practices can therefore improve the policy trade-off between mitigation, individual freedoms, and economic damages.

REFERENCES FOR THE PHILIPPON COMMENT

Jones, Callum, Thomas Philippon, and Venky Venkateswaran. 2021. "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home." *Review of Financial Studies* 34, no. 11: 5188–223.

Romer, Christina D. 2021. "The Fiscal Policy Response to the Pandemic." *Brookings Papers on Economic Activity*, Spring: 89–110.

COMMENT BY

COADY WING The COVID-19 pandemic upended life in the United States. People changed their behaviors to mitigate the risk of infection and mortality. Governments imposed new regulations designed to encourage further reductions in the transmission of the virus and to promote vaccine take-up once the vaccine became available. By now, there is a large body of literature in social and health sciences that documents these behavioral responses and tries to evaluate the intended and unintended consequences of various policy initiatives (Gupta and others 2021b; Autor and others 2022; Chetty, Friedman, and Stepner 2024).

In their paper, Atkeson and Kissler evaluate the high-level effect of the pandemic response on the level and time series dynamics of COVID-19 mortality. They use a structural epidemiological model to decompose the way that mortality is determined by transmission-related behaviors, vaccine take-up, and shifts in transmissibility and virulence of the virus. This is a compartmental model in which a population of susceptible people transitions through a collection of health states. Transitions are governed by model parameters that define—at each point in time—the viral transmission rate, the duration of infectiousness, the take-up and effectiveness of the vaccine, and the infection fatality rate. The parameters of the model are not estimated from the data. In some cases, the authors draw on epidemiological studies to guide the choice of parameters that represent infection fatality rates of different strains. But for the most part, Atkeson and Kissler choose parameterizations that seem plausible, most likely using some amount of trial and error. The main evidence that the chosen parameters are sensible is that the model does an excellent job of reproducing the observed time series of COVID-19 mortality in the United States. It also fits the time series estimates of the cumulative share of the population that had been infected by COVID-19 based on convenience samples from the Blood Donor Survey and Commercial Lab Survey.

Treating the model as correct, Atkeson and Kissler examine counterfactual scenarios to measure the role of specific mechanisms in causing mortality. For example, in one scenario, behavioral changes are maintained but the vaccine never arrives. In another simulation, people do not engage in major behavioral changes, but the vaccine becomes available on schedule. In both cases, they use the model to compute the number of COVID-19 deaths that would have occurred if the parameters of the model are correct but certain events played out differently. The simulations suggest that the combination of behavioral changes and the eventual availability of the vaccine led to substantial reductions in mortality. Without behavioral changes, the vaccine would have arrived too late to matter. Without the vaccine, the behavioral changes would mostly have reallocated deaths over time. One way to see it is that behavioral changes that mitigate transmission early in the pandemic increase the marginal health benefits of vaccines later in the pandemic. Behavior and vaccines are complements in an aggregate health production function.

In my discussion, I focus on three main topics. First, I try to provide an intuitive account of the type of model that Atkeson and Kissler use in their analysis and to point out some of the key assumptions involved in such models. Second, I discuss some of the ways we might judge the credibility

Figure 1. The Classic SIR Model

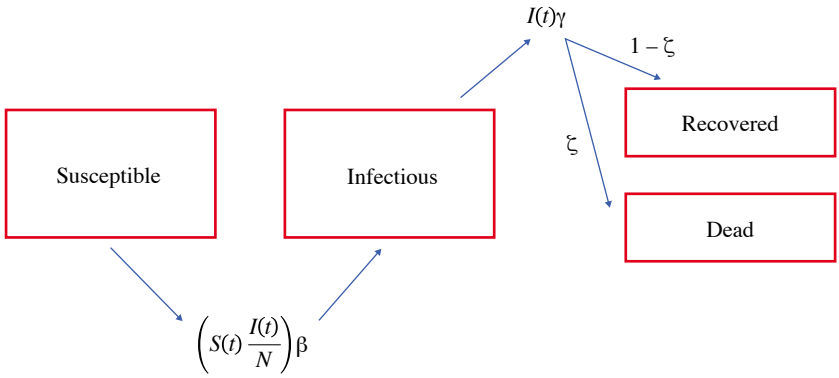

Source: Author's illustration.

of structural epidemiological models and suggest some ideas for incorporating quasi-experimental study designs. Third, I shift focus to questions about the behavioral determinants of vaccine take-up with particular attention to take-up coming out of the COVID-19 pandemic. I present some early research examining the way that vaccine take-up may be undermined by breakthrough infections.

STRUCTURAL EPIDEMIOLOGICAL MODELS The most famous model of an epidemic is the susceptible-infectious-recovered (SIR) model, which was developed by Kermack and McKendrick (1927). Figure 1 is a visual representation of a basic SIR model. In this setup, the whole population starts out susceptible to the disease except for a single index patient who is infected. The people in the susceptible and infectious compartments mix at random. When an infectious person and a susceptible person come into contact, the susceptible person is infected at rate β. Infected people recover and leave the infectious compartment according to the recovery rate, γ. The infection mortality rate is ζ.

The share of the population in each compartment at a given point in time is determined by these transmission, recovery, and mortality parameters. In a setting where the whole population is susceptible, an infected person generates $R = \frac{\beta}{\gamma}$, new infections. When $R < 1$, the disease dies out over time. When $R > 1$, there is an exponential outbreak in the number of new infections.

Real-world conditions are more complicated than the simple model implies. For example, the population of susceptible and infectious people

Figure 2. A Stylized Version of Atkeson and Kissler's SEIHRS Model

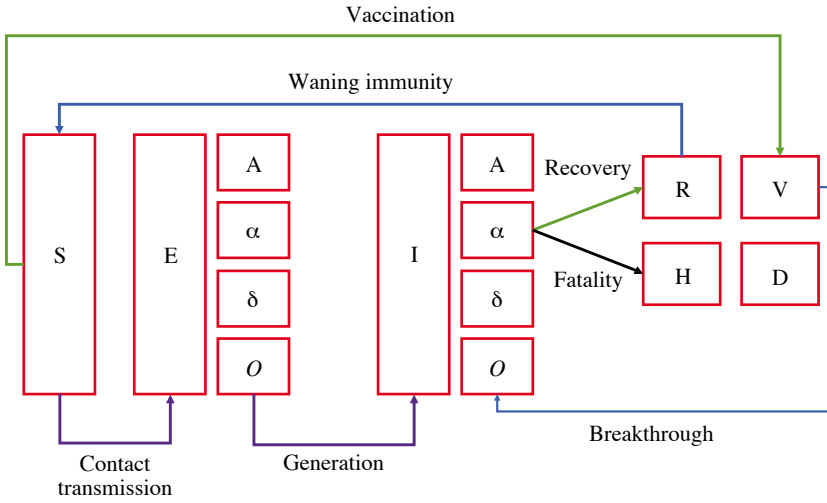

Source: Author's illustration.

might be structured so that some people have higher mixing rates than others, violating the random mixing assumption of the basic model. The basic SIR model is not capable of representing epidemics that exhibit repeated waves of infections and in which the properties of the pathogen itself might change over time. And—of course—economists often point out that the simple SIR framework does not allow for people who change their behaviors in response to prevailing epidemiological conditions, new public policies, or expectations about future technologies such as vaccines or cures.

The contemporary literature has elaborated and complexified the basic SIR model to make it more realistic. A visual representation of Atkeson and Kissler's model would look something like figure 2. The first thing to notice is that there are many more compartments. At a high level, the model is organized around six main states: susceptible, exposed, infectious, hospitalized, resistant, and dead. But the exposed and infectious compartments are subdivided further by the viral strain—there is a box each for the original ancestral strain as well as the Alpha, Delta, and Omicron viral variants. Each of the strains has a different transmission rate and infection mortality rate. The resistant compartment includes space for people who have recovered from an actual infection and for people who have been vaccinated. The flow of people through the compartments is no longer in a

single direction. Immunity from infection may wane over time, creating a flow of people from the resistant to the susceptible. Vaccinated people may experience breakthrough infections.

Simulating the model involves specifying the timing of certain shocks to the system. Some pathways open up at particular calendar times because a new viral strain emerges. The parameters governing transmission and mortality change over time too. Taken together, the model simulations involve a schedule of epidemiological shocks (new viral variants), technology shocks (vaccine availability), seasonality, behavioral relationships (transmission-mortality elasticity), and structural changes in behavioral relationships (fatigue).

To get a feel for how the model works, it helps to consider the block of the model that shapes the transmission rate at a point in time. In Atkeson and Kissler's model, the transmission rate at a point in time follows a seasonal component and a behavioral component in which the transmission rate responds to the number of daily deaths. The idea is that when death rates get high, people change behavior to reduce infection risk. After a time, fatigue sets in and the responsiveness of transmission to fatalities shrinks. Formally, the transmission relationship in the model is:

$$\beta_j(t) = \bar{\beta}_j \times \exp\left[-\kappa(t)\frac{dD(t)}{dt} + \psi(t)\right].$$

In this expression, $\bar{\beta}_j$ is the inherent transmissibility of viral strain j, $\psi(t)$ is the effect of the season on transmission, $\frac{dD(t)}{dt}$ is the level of daily deaths, and $\kappa(t)$ is the semi-elasticity of transmission with respect to deaths. In practice, the seasonal function $\psi(t)$ is a cosine wave parameterized to match fall-winter versus spring-summer patterns. For the original COVID-19 strain, the authors set $\bar{\beta}_j = 1.2$. To represent the behavioral component of transmission, they set the baseline semi-elasticity to be $\bar{\kappa} = 250{,}000$. To incorporate the idea of behavioral fatigue, they smoothly shrink the behavioral response down to $.35 \times \bar{\kappa}$ by late November 2020. Once a set of parameters has been chosen, Atkeson and Kissler run their model, pushing an initial population through the various compartments and keeping track of how many people are dead, infectious, and vaccinated at each point in time.

IDENTIFICATION PROBLEMS AND EPIDEMIOLOGICAL MODELS Constructing epidemiological models is challenging for some of the same reasons that constructing macroeconomic models is challenging. The parameters of the

model are hard to cleanly identify because the variation generated during real-world outbreaks is not randomized across places, time periods, and people. And the outcomes realized during an epidemic seem very contingent on a high-dimensional set of conditions and constraints. This makes it hard to accumulate knowledge across different settings or to discriminate between one hypothesized model and another.

How should we judge the credibility of Atkeson and Kissler's model? One natural strategy is to compare the outputs of the model with observed outcomes from the real world. That is the approach that Atkeson and Kissler take. Figures 2 and 4 from their paper show a tight correspondence between COVID-19 mortality and infections as generated by the model and the actual time series data on COVID-19 mortality and COVID-19 infections. Comparing the model-predicted mortality and infection series with their real-world counterparts is essentially the same idea macroeconomists use when they form judgments based on how well a specific model is able to match moments observed in the real world.

Although the close fit shown in Atkeson and Kissler's figures 2 and 4 is impressive, a fundamental question is whether this collection of compartments and parameters is really a good representation of the process that generated those deaths and infections. Are there other combinations of parameters and compartments that might also fit the data very well but could generate quite different counterfactual simulations? To what extent is the good fit of the model akin to a regression that fits the data well in sample but performs badly at out-of-sample forecasts?

The model makes strong assumptions about the mechanism and even the specific numerical values of key causal parameters. The payoff from these kinds of assumptions is substantial: you can use the model to simulate the pandemic under alternative conditions, which is just the type of thing you need to do to study alternative policy options. But the credibility of the counterfactual simulations depends on the plausibility of the underlying modeling choices. For example, is $\overline{\kappa} = 250{,}000$ a plausible value for the behavioral response to mortality in the early pandemic? Does this choice undershoot the degree to which transmission responded to mortality? It seems hard to decide something like this through intuition.

One idea is to combine the structural epidemiological methods with research strategies that are common in empirical microeconomics, which focus on identifying causal effects using plausibly exogenous variation. Nakamura and Steinsson (2018) discuss identification problems in macroeconomics, pointing out that macroeconomists often judge the credibility of specific models by their ability to match moments observed in the real

world. Often the moments used in this type of work are simple aggregate means and variances. But Nakamura and Steinsson (2018) highlight that more recent research involves efforts to match *identified* moments. Identified moments are causal effects identified using the methods popular in empirical microeconomics: regression discontinuity designs, difference-in-differences designs, or instrumental variable designs. Taking advantage of identified moments to judge the performance of a more complex structural model or to pin down the value of a class of parameters from a structural model is a strategy that may be useful for future work on epidemiological models.

To take one very small step in this direction, we could compare estimates from Atkeson and Kissler's model with identified moments from related quasi-experimental studies. For example, Gupta and others (2021a) use a generalized difference-in-differences regression to make reduced-form estimates of the effects of the early vaccination campaign on cumulative COVID-19 mortality over the first five months of the vaccination campaign. Their estimates come from quasi-likelihood Poisson regression models with the following basic form:

$$M_{st} = \exp\left[\sum_{k=0\ldots4} \delta_k V_{st-k} + a_s + b_t\right] + e_{st}.$$

In the regression, M_{st} is the cumulative number of COVID-19 deaths per 100 adults in state s as of week t, and V_{st} is the cumulative number of doses administered per 100 adults in state s by week t. The model includes state and week fixed effects and is intended to measure the effects of the vaccine rollout by exploiting variation in the speed of vaccine distribution across states. The estimated parameters from this two-way fixed effects specification are used to estimate the counterfactual cumulative COVID-19 mortality rate in the absence of the vaccination campaign. The results imply that by the second week of May in 2021 the vaccination campaign had already averted about 139,393 COVID-19 deaths. How do these reduced-form estimates line up with the simulations from Atkeson and Kissler's model? The numbers underlying the left panel of their figure 5 imply that by the second week of May the vaccination campaign had averted 126,664 COVID-19 deaths. This is well inside the confidence interval of the estimate by Gupta and others (2021a), perhaps suggesting that Atkeson and Kissler's model fares pretty well at matching an identified moment that is a bit more removed from the analysis than the

mortality time series itself.[1] Atkeson and Kissler's model-based estimates extend beyond the first few months of the campaign, and they suggest that the impact of the vaccination campaign continued to grow over the course of 2021 and then diminished in 2022. The model-based estimates imply that most of the additional deaths that would have occurred in the absence of the vaccine would have happened by the end of 2021. This is because without the vaccine nearly everyone would have been infected during the Omicron wave in early 2022.

Atkeson and Kissler are surely correct that people respond to epidemiological conditions and the availability of vaccines. And public policies designed to control an epidemic are premised on the idea that behavior is both malleable and an important determinant of the path of the epidemic. Their model provides an excellent example of how to integrate these policy relevant relationships into an epidemiological model. But these models would be more compelling if there were more quasi-experimental studies trying to pin down the specific ways that people respond to changing conditions and how those changes affect downstream population health outcomes. In particular, economists could be useful by developing the identification strategies and data sources needed to estimate things like: (a) the causal effect of mortality on disease transmission (behavioral responses); (b) the causal determinants of vaccine take-up and behavioral fatigue; and (c) the role of differentiated contact patterns on disease outcomes.

VACCINE TAKE-UP One of the main lessons that Atkeson and Kissler draw from their analysis is that behavioral adaptations that reduced transmission rates during the first year of the pandemic allowed the COVID-19 vaccine to substantially reduce overall mortality from COVID-19. Specifically, in simulations where they keep the behavioral parameters fixed but never turn on the availability and take-up of the vaccine, there would have been almost 800,000 additional COVID-19 deaths.

The implications of the model are quite encouraging in certain ways. They suggest that it is possible to use behavioral modifications to suppress a pandemic long enough to develop and distribute a vaccine soon enough for the vaccine to actually save lives. At the same time, as Atkeson and Kissler are careful to point out, there is a lot of historical contingency involved in this analysis. If the highly transmissible Omicron variant had

1. Not that removed, of course: the two-way fixed effects regression in Gupta and others (2021a) is based on state *x* week-level mortality data. Atkeson and Kissler are working with a national mortality time series rather than a state *x* week panel.

arrived sooner, then the vaccine likely would not have saved many lives: by the time the vaccine arrived, it would have been too late. On the other hand, if the variant had not appeared or had appeared even later, then the vaccine would have had even more of an impact. Governments have little influence on the appearance and characteristics of new viral strains and so the strategy of behavior-induced transmission reduction followed by vaccination is somewhat inherently risky. However, the distribution and take-up of the vaccine itself is something that may deserve more attention.

In particular, it would make sense for economists to develop a better understanding of the determinants of vaccine take-up both during an epidemic and during regular conditions. Neoclassical economics suggests that vaccine take-up may be *too low* from a social welfare point of view because vaccines may produce positive externalities. Acton and others (2022) studied college vaccine mandates and found some evidence that mandates led to lower rates of COVID-19 spread in nearby communities. Freedman and others (2023) use linked micro data on COVID-19 tests, vaccinations, and health care records to study vaccine spillovers in middle schools and households with children in Indiana. They find little evidence of spillovers in schools, but they do find vaccine spillovers in households. Since households might plausibly internalize these vaccine spillovers, it is not obvious that free-riding on positive externalities is a major determinant of low vaccine take-up.

Another explanation for low vaccine take-up is that people's assessment of the private net benefits of the vaccine is somewhat lower than expected. Recent work by Carlin and others (2022) used discrete choice survey experiments to measure people's willingness to pay to be vaccinated for COVID-19 during early 2021. They found that median willingness to pay was around $50. Back-of-the-envelope calculations based on estimates of the value of statistical life from other contexts imply people should have been willing to pay around $2,700 to be vaccinated, given the mortality effects of the vaccine and prevailing caseloads. Thus, people seem to undervalue the COVID-19 vaccine. This could be one explanation for relatively low vaccine take-up in the United States. People might undervalue vaccines for many different reasons, including concerns about the safety of the vaccine or the perceived costs of vaccine side effects, needle aversion, or the political symbolism of the vaccine.

Another possibility is that people's demand for a vaccine is partly derived from their own experience of the vaccine and the underlying illness. For example, Jin and Koch (2021) study the relationship between influenza vaccination and influenza infection at the individual level over time. They

find that contracting influenza in one year affects people's take-up of the vaccine in future years, suggesting that people learn from suffering. However, they also find that what people learn depends on their vaccination history. People who were unvaccinated and infected in the baseline year are more likely to be vaccinated in the future. But this learning-induced demand for the vaccine is offset for people who were vaccinated and experienced a breakthrough infection. One interpretation is that breakthrough infections undermine people's assessment of the usefulness of a preventive vaccine. This is almost certainly the wrong conclusion for people to draw: breakthrough infections can and do occur even when the vaccine is effective. Nevertheless, this type of misguided behavioral response is not difficult to understand.

LEARNING BY SUFFERING IN INDIANA A similar dynamic may hold for the COVID-19 vaccine as people make choices about boosters and vaccine take-up in nonepidemic conditions. To shed some light on the issue, I used linked administrative data from Indiana to study the relationship between vaccine take-up and prior vaccination and COVID-19 infection experiences. There are three main data sources: (1) Indiana COVID-19 vaccination registry; (2) Indiana COVID-19 lab test registry; and (3) Indiana Network for Patient Care (INPC) research database. INPC is a database of electronic medical records contributed by most of the hospitals and clinics in Indiana. I constructed a study sample of people who had at least one health care encounter in the INPC system between 2018 and 2019. Then I linked these individual records with COVID-19 vaccination records and COVID-19 lab tests and test results from 2020 to 2022.

With the data in hand, I estimate simple cross-sectional regressions with the following form:

$$Vax_i^{2022} = \alpha_0 + \alpha_1 Vax_i^{2021} + \alpha_2 Covid_i^{2021} + \alpha_3 \left(Vax_i^{2021} \times Covid_i^{2021} \right) + X_i \beta + e_i.$$

In the regression, Vax_i^{2022} is a binary variable indicating whether person i received the COVID-19 vaccine in 2022. Vax_i^{2021} indicates whether the person was vaccinated in 2021, and $Covid_i^{2021}$ indicates whether the person had a lab-confirmed COVID-19 infection in 2021. X_i is a covariate vector that adjusts for gender, race, and age fixed effects. I fit an overall regression to the full sample as well as separate regressions for younger, middle aged, and older people. The estimated regression coefficeints in table 1 show that the 2022 vaccination rate is much lower than the vaccination rate during the main pandemic. The take-up rate in 2022 was about 12 percent among

Table 1. Regressions of COVID-19 Vaccine Take-Up on Prior Season COVID-19 Infection and Vaccination

	Age 18–39	*Age 40–64*	*Age 65+*
Intercept	–0.011***	0.008***	0.033***
	(0.002)	(0.001)	(0.001)
2021 infection	0.015***	0.023***	0.024***
	(0.001)	(0.001)	(0.001)
2021 vaccine	0.167***	0.119***	0.089***
	(0.001)	(0.001)	(0.001)
2021 infection × 2021 vaccine	–0.017***	–0.012***	–0.013***
	(0.002)	(0.002)	(0.002)
N	820,124	1,107,246	876,854
R^2	0.06	0.04	0.03
Mean of outcome	0.061	0.091	0.118

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Source: Author's calculations.

Note: The regressions adjust for gender-, race-, and age-fixed effects. Standard errors are estimated using a heteroskedasticity robust variance matrix.

people over age 65, 9 percent among middle-aged adults age 40–64, and 6 percent among younger adults age 18–39.

Vaccination in 2022 was much higher among people who were vaccinated in 2021, suggesting preferences for the vaccine in the past are a strong predictor of vaccination in subsequent seasons. The coefficient on the prior COVID-19 infection indicator is also positive and quite large. Among those age 40–64 and age 65 and older who did not get vaccinated in 2021, the 2022 vaccination rate was about 2.4 percentage points higher among people who contracted COVID-19 in 2021 than among people who did not contract COVID-19. This learning by suffering effect is a bit smaller—only 1.5 percentage points—among younger adults age 18–39. However, the coefficient on the interaction term is negative, suggesting that experiencing a breakthrough infection offsets the learning by suffering effect perhaps because it undermines confidence in the vaccine. For middle-aged and older adults, the breakthrough effect offsets the learning by suffering effect by about 50 percent. For younger adults the breakthrough effect offsets the learning by suffering effect by over 100 percent, completely undoing any induced demand from prior infection.

This analysis suggests that the dynamics of individual COVID-19 vaccination exhibit some of the same patterns reported by Jin and Koch (2021) for influenza. In particular, take-up of the vaccine is partly determined by firsthand experience of the disease, and breakthrough infections seem to reduce subsequent demand for the vaccine. This type of response to health shocks seems undesirable from a public health point of view. It would

probably be better if people did not lower their opinion of the efficacy of the vaccine on the basis of their own recent health experiences. But it is not at all hard to understand how a breakthrough infection might be a salient event that does motivate behavioral changes. Understanding how people interpret and change their behavior in response to salient health events in their own lives or in the lives of other people in their family may be an important way to develop more realistic models of epidemics.

REFERENCES FOR THE WING COMMENT

Acton, Riley K., Wenjia Cao, Emily E. Cook, Scott A. Imberman, and Michael F. Lovenheim. 2022. "The Effect of Vaccine Mandates on Disease Spread: Evidence from College COVID-19 Mandates." Working Paper 30303. Cambridge, Mass.: National Bureau of Economic Research.

Autor, David, David Cho, Leland D. Crane, Mita Goldar, Byron Lutz, Joshua Montes, William B. Peterman, David Ratner, Daniel Villar, and Ahu Yildirmaz. 2022. "The $800 Billion Paycheck Protection Program: Where Did the Money Go and Why Did It Go There?" *Journal of Economic Perspectives* 36, no. 2: 55–80.

Carlin, Patrick, Brian E. Dixon, Kosali I. Simon, Ryan Sullivan, and Coady Wing. 2022. "How Undervalued Is the Covid-19 Vaccine? Evidence from Discrete Choice Experiments and VSL Benchmarks." Working Paper 30118. Cambridge, Mass.: National Bureau of Economic Research.

Chetty, Raj, John N. Friedman, and Michael Stepner. 2024. "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data." *Quarterly Journal of Economics* 139, no. 2: 829–89.

Freedman, Seth M., Daniel W. Sacks, Kosali I. Simon, and Coady Wing. 2023. "Direct and Indirect Effects of Vaccines: Evidence from COVID-19." Working Paper 30550. Cambridge, Mass.: National Bureau of Economic Research.

Gupta, Sumedha, Jonathan Cantor, Kosali I. Simon, Ana I. Bento, Coady Wing, and Christopher M. Whaley. 2021a. "Vaccinations against COVID-19 May Have Averted up to 140,000 Deaths in the United States." *Health Affairs* 40, no. 9: 1465–72.

Gupta, Sumedha, Thuy D. Nguyen, Shyam Raman, Byungkyu Lee, Felipe Lozano-Rojas, Ana I. Bento, Kosali I. Simon, and Coady Wing. 2021b. "Tracking Public and Private Responses to the COVID-19 Epidemic." *American Journal of Health Economics* 7, no. 4: 361–404.

Jin, Ginger Zhe, and Thomas G. Koch. 2021. "Learning by Suffering? Patterns in Flu Vaccination Take-Up." *American Journal of Health Economics* 7, no. 1: 68–94.

Kermack, William Ogilvy, and A. G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 115, no. 772: 700–721.

Nakamura, Emi, and Jón Steinsson. 2018. "Identification in Macroeconomics." *Journal of Economic Perspectives* 32, no. 3: 59–86.

GENERAL DISCUSSION James Stock noted an important contribution of the paper is a fairly simple model given the complexity of the task. On the interaction of behavior and vaccination, Stock reflected on the importance of waiting: first, and rather obvious, it gives you time to get the vaccine; second, given the high fatality rate in the first wave, those who took self-protective measures and waited and then became infected during later waves were better off. Stock further pointed to the notable result in figure 10 of the paper's conference draft, which shows the convergence of the effective reproduction number to one for every state, despite significant state differences in political views and COVID-19-related interventions.[1] This, he argued, points to the extensive self-protection measures taken by individuals across all states. Finally, Stock underscored the importance of continued work in this area, including better data collection, to better prepare us for similar events in the future.

Speaking to the cross-state variation in COVID-19 incidence, Louise Sheiner agreed with the authors' emphasis on the importance of behavior in explaining these differences. She argued that politicization was a major contributor, pointing to data that show that the variation in the labor force participation rate, unemployment, and consumption can be explained by political affiliation—the share of the state population that voted for Joe Biden. The same goes for state variation in vaccination rates. Sheiner concluded that this underscores people's attitudes as the primary driver of differences, rather than state lockdowns and other mandates during this time.

On the behavioral response, Carol Graham suggested that the authors further explore the great variation within states in vaccination rates—that observed between counties. Looking at the standard deviation within states, for example, would be useful in trying to better understand the outcome that seems to show there is convergence across states in the aggregate.

Stefanie Stantcheva asked the authors whether they had explicitly taken the multidose nature of the COVID-19 vaccine into account in their model, and if not, how such dynamics may alter the findings. Stantcheva then raised the role of trust in government as a salient issue to consider when contemplating future responses to a pandemic.

Tying together issues of data collection and trust in public officials, Steven Davis argued that there is a need to make data both transparent and credible to avoid politicization. Davis then turned to discuss the economic

1. This figure is included as figure B.6 in the online appendix of the paper's final version published in this *BPEA* volume.

resilience value of working from home. In his paper with Barrero and Bloom, survey results where respondents were asked if their internet connectivity affected their ability to work productively when working from home showed that while 75 percent said their internet connection was perfectly adequate, the other 25 percent reported that their productivity was negatively affected by the poor quality of their internet connection.[2] The paper estimates that subpar internet connectivity lowered US labor productivity by 3 percent in the period from May 2020 to April 2021. Davis also suggested that people may also be more willing to undertake voluntary efforts to slow the transmission of a virus if they can productively work from home, further underscoring the increased economic resilience that widespread access to high-quality internet would provide.

Kenneth Rogoff wondered how the model from the paper could be used to learn more about the economic cost of different choices. He observed that the learning losses for students during COVID-19 were stunning. What kind of information can we gather to help us better, and more quickly, determine what the right mitigation efforts are while we wait for a vaccine the next time around?

Maurice Obstfeld appreciated the contribution of the authors' work in successfully pinning down the role that delay played in the spread of COVID-19, particularly in the current politicized environment. However, Obstfeld worried that the next pandemic could be very different and called for better preparedness in general. We were lucky this time, Obstfeld reasoned, that mRNA technology was available, speeding up the rollout of vaccinations notably. He added that future pandemics are likely to be just as politicized, raising issues such as what the optimal approach to school closures would be if younger demographics were more vulnerable, as was the case with the 1918 influenza epidemic. Obstfeld brought up previous efforts, including two panels in 2021 (the Independent Panel for Pandemic Preparedness and Response and the High-Level Independent Panel) focused on funding and global cooperation on pandemic surveillance and response, as well as a book by Bill Gates, *How to Prevent the Next Pandemic*, on the topic. But he noted that few of the recommendations had been implemented since, and he cautioned that international attention to this issue had dwindled significantly with access to antivirals and the slowing fatality of

2. Jose Maria Barrero, Nicholas Bloom, and Steven J. Davis, "Internet Access and Its Implications for Productivity, Inequality, and Resilience," In *Rebuilding the Post-Pandemic Economy*, edited by Melissa S. Kearney and Amy Ganz (Washington: Aspen Institute Press, 2021). https://www.economicstrategygroup.org/publication/barrero-bloom-davis/.

COVID-19. Obstfeld stressed the need for funding and international cooperation to continue the work, including surveillance of animal reservoirs and thinking about what types of vaccines may be necessary depending on specific viruses of future pandemics.

Andrew Atkeson recalled how economists had sprung into action at the onset of the pandemic, thinking about sectoral-level interventions and testing, among other things. He firmly believed that it was not for lack of ideas that things did not happen. He pointed out the sometimes lackluster response of the public health sector and epidemiologists to the ideas of economists during this time. Atkeson stated that, just like the military, we need to plan for all types of contingencies and be ready with a response no matter the circumstances, noting that the willingness to spend given the economic cost in these situations is very high.

To Obstfeld's point, Şebnem Kalemli-Özcan agreed that the next pandemic could be very different but said there are still lessons to be drawn from COVID-19. She called for better financial targeting of funds in general. Kalemli-Özcan acknowledged that while mandated lockdowns may be a second-best option, relying on behavioral responses to mitigate the spread of a virus may not be feasible—especially in countries where a greater share of the labor force is informal. She suggested that to make lockdowns more efficient and financially sustainable, funds should be targeted to specific sectors: in the case of COVID-19, contact-intensive sectors. She further highlighted some of her own work on the topic and agreed that more data collection was needed.[3]

Stan Veuger cautioned that some of the policy suggestions in the paper might be difficult to implement. Mandating testing, he believed, would likely face opposition in the United States as well as in Western Europe. On data collection, Veuger was skeptical of the role the Centers for Disease Control and Prevention (CDC) would be able to play. He pointed to the absence of a US-wide representative survey of COVID-19 incidence despite a significant amount of additional funding to CDC during the pandemic.[4]

3. Cem Çakmaklı, Selva Demiralp, Şebnem Kalemli-Özcan, Sevcan Yeşiltaş, and Muhammed A. Yıldırım, "The Economic Case for Global Vaccinations: An Epidemiological Model with International Production Networks," working paper 28395 (Cambridge, Mass.: National Bureau of Economic Research, 2024).

4. Congressional Research Service, *US Public Health Service: COVID-19 Supplemental Appropriations in the 116th Congress* (Washington: Author, 2021), https://crsreports.congress.gov/product/pdf/R/R46711/3; and *American Rescue Plan Act of 2021 (P.L. 117-2): Public Health, Medical Supply Chain, Health Services, and Related Provisions* (Washington: Author, 2021), https://crsreports.congress.gov/product/pdf/R/R46834.

Atkeson said that they had purposefully tried to stay clear of the politics but agreed that we may not want federal mandates. He argued that we probably ought to accept the different decisions of democratically elected state officials, some of whom, during the COVID-19 pandemic, chose to accept higher death rates in their states.

Laurence Ball asked how events might have transpired if things had been a lot better than they were. How many fewer deaths could there have been had we pursued the optimal policy? Atkeson responded by suggesting that in the absence of vaccines being made available earlier, there was probably not much we could have done better.

Tristan Reed asked about the external validity of the authors' findings for developing countries, saying that most of the delay in vaccine delivery to developing countries during the pandemic could be attributed to them ordering later.[5] This may seem highly irrational at first, Reed conceded, and justifying not buying a vaccine seems to suggest a very low statistical value of life. However, Reed explained, the authors point to results in the paper that could suggest that in the absence of work-from-home technology, purchasing a vaccine is not worth it anymore after 120 days. He pondered whether when developing countries today are asked to financially contribute toward being able to purchase vaccines on day zero of a future pandemic and interest is muted, it reflects their belief that they do not have work-from-home technology.

On the low uptake of vaccines in emerging markets and its interaction with the need for better data, Raghuram Rajan talked about the specific case of India. He explained that India severely undercounted the deaths resulting from COVID-19. Initially, there were false stories spread about natural immunity against COVID-19 in India. Rajan said that the actual number of fatalities was suppressed, as the initial miscalculations would have negatively reflected on the capability of the public health system in each state. The second COVID-19 wave hit India hard because of the lack of immunization. In the official statistics, death rates in India are low, but taking undercounting into account raises the death toll significantly. In conclusion, Rajan said that this highlights the dangers of working with poor data and the policies made based on such data.

Hoyt Bleakley expressed his preference for adding standard tools of public economics: weighing marginal cost versus marginal benefit, including

5. Ruchir Agarwal and Tristan Reed, "Financing Vaccine Equity: Funding for Day-Zero of the Next Pandemic," *Oxford Review of Economic Policy* 38, no. 4 (2022): 833–50.

external cost and benefits. He agreed with the discussants' point that the benefits of getting vaccinated are largely internal, even though measuring the external benefits would be key for the design of a subsidy. Bleakley suggested using the model to measure not just the externality on the next person who could get infected, but also on the social and private incentives to delay infection before a vaccine becomes available. He also suggested looking at a much faster transmission rate in the authors' model, which would resemble that of Omicron at the onset of the pandemic, to see how that would alter the effect described in the paper.

Gerald Cohen noted that the discussion on the current paper and the paper by Stantcheva on inflation (also included in this *BPEA* volume) both focused on information available to the public: the current paper on the extent to which the public understood the propagation mechanisms for disease transmission during COVID-19 and Stantcheva's paper on whether the public are able to think correctly about inflation in general. Cohen suggested that economists should think about the importance of how to optimize people's information about the benefits of vaccines or private efforts of mitigation versus the benefits of the public better understanding inflation.

Speaking to the monetary losses during COVID-19, Robert Hall observed that a great number of workers were on temporary layoff around April 2020, and that output losses were large.[6] This came with a partially offsetting increase in leisure. He noted that there was no material decline in consumption. How do we put prices on these developments? Hall proposed that we carefully consider these different pieces to the puzzle in trying to measure the net loss of well-being from COVID-19.

6. According to Bureau of Labor Statistics (BLS), the number of unemployed people on temporary layoff reached 18 million in April 2020. BLS, "Temporary Layoffs Remain High following Unprecedented Surge in Early 2020," February 10, 2021, https://www.bls.gov/opub/ted/2021/temporary-layoffs-remain-high-following-unprecedented-surge-in-early-2020.htm.

Online Appendix for The Impact of Vaccines and Behavior on US Cumulative Deaths from COVID-19

Andrew Atkeson[1] and Stephen Kissler[2]

[1]Department of Economics, UCLA
[2]Department of Computer Science, University of Colorado Boulder

April 28, 2024

Contents

A Comparison to Other Estimates in the Literature

We have taken an approach of estimating the impact of behavior and vaccinations on cumulative COVID mortality based on an estimate of the fraction of the U.S. population that were vaccinated before they were first infected with COVID derived from serology data, an estimate of the fraction of those individuals who would have been infected with COVID absent vaccines based on an estimate of the

final size of the COVID epidemic, and an estimate of the difference between the infection fatality rate for those first infected by COVID before being vaccinated and those first infected by COVID after being vaccinated.

A.1 Key Sources of Uncertainty Regarding our Estimate of Lives Saved

Our estimate of lives saved is subject to uncertainty due to uncertainty in each of the numbers entering this calculation.

Our model may lead to an overestimate of lives saved because blood serology data may undercount infections prior to vaccination or overstate the protection offered by vaccination either because it is not representative of the entire population or because of waning antibodies or both. See, for example, Chitwood et al. (2022). See Eales et al. (2023) for alternative estimates of the evolution of the infection fatality rate from SARS-COV-2 using data from England. We do not consider this concern as having a significant impact on our estimate of lives saved by vaccines. In particular, Atkeson (2023) conducts a similar exercise estimating the impact of vaccines on lives saved using an estimate of cumulative infections over the course of 2020 closer to that in Chitwood et al. (2022) and arrives at an estimate of lives saved close to 750,000, not far from the updated estimate in this paper.

Our model may lead to an underestimate of lives saved because old people got vaccinated early and did more to avoid getting infected early. In Figure A.1 we show serology data from the Blood Donor and Commercial Lab surveys for the overall population in the left panel (reproducing Figure 1 from the main text) and the population 65 years old and older in the left panel. By comparing the two panels, we see that infections spread more slowly through this older population in 2020 and 2021 and that vaccinations were deployed more rapidly in this population in early 2021 than was the case for the general population.

Absent vaccines, this elderly population would have likely become infected, and once infected they would have died at a rate higher than our assumed infection fatality rate. To illustrate this point, we see in the serology data shown in Figure A.1 that for those over 65 years of age, 80% had been vaccinated before they were infected as of July 2021. In Figure A.2, we see that those over 65 years of age were slower to get infected than the general population consistently across states. Clearly, given the high infection fatality rate of COVID-19 for the elderly, the counterfactual of a much higher infection rate among this age group would imply a very high death toll and thus a much larger estimate of lives saved.

We note that some argued for an approach to COVID mitigation based on focused protection of the elderly. See “The Great Barrington Declaration" at `https://gbdeclaration.org/`. The serology data both at the national level

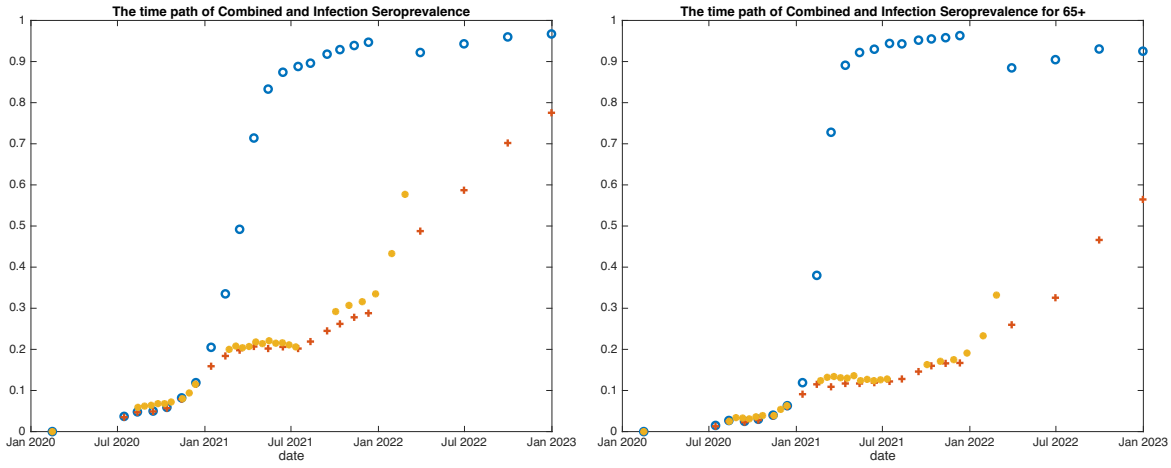

Figure A.1: Combined Seroprevalence from Blood Donor Surveys in blue and Seroprevalence from Infection (Blood Donor Survey in red, Commercial Lab Survey in Yellow). Left Panel: Overall Population. Right Panel: Population 65 and over

shown in Figure 1 in the main text and at the state level discussed below are consistent with the view that those over 65 did manage to get infected at a substantially slower rate than the general population. Would this greater mitigation by the elderly have made a difference for overall mortality absent vaccines? Given the arrival of the Delta and Omicron variants, we think not. It is unclear whether this success in avoiding infection by the elderly would have been more than temporary absent the development of vaccines. Of course, our statement here is based on the benefit of our ex-post knowledge of the emergence of these highly transmissible variants. Perhaps focused protection might have greater effect in the context of a different disease. Importantly, focused protection is most effective when overall infection rates are kept low, as targeted protection of just a single group – like the elderly in the context of COVID-19 – would likely fail in the absence of substantial mitigation efforts in the rest of the population as well.

A.2 Alternative Estimation Approaches

There are a number of alternative approaches to estimating the impact of vaccination on COVID mortality available in the literature that we review here.

Jia et al. (2023) apply an interesting methodology to construct an estimate of a different measure of the impact of vaccines on COVID mortality. They aim to measure the marginal impact on mortality if vaccination (here corresponding to the first two shots in the primary series) had happened faster or had been pushed

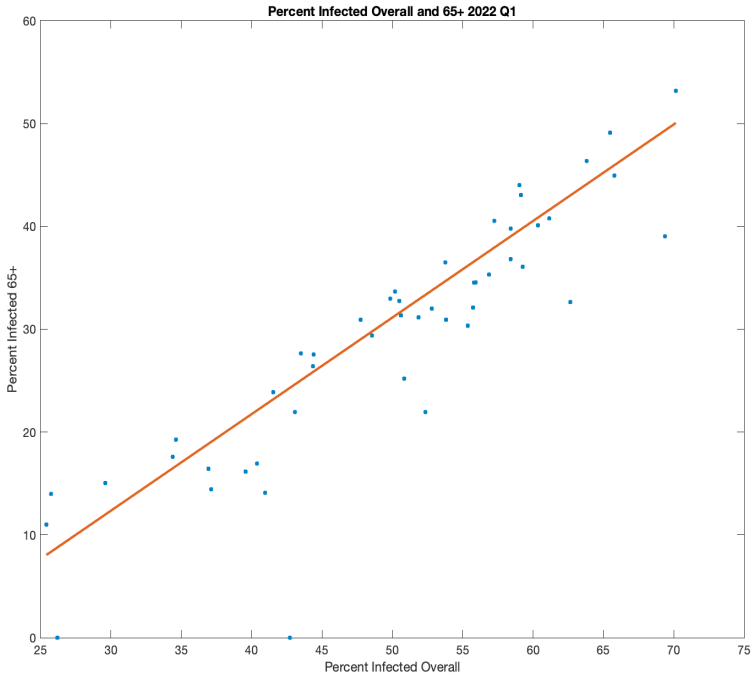

Figure A.2: Scatter plot of Blood Donor serology data on percent of those over 65 infected by state in the first quarter of 2022 vs. overall percent infected in the same survey. These data show that the elderly consistently avoided infection relative to the general population across states

further than it actually was. They focus on the time period from the end of May 2021 (when vaccination was well underway) and the beginning of September of 2022.

Their estimate of this marginal impact is based on the use of two data inputs. The first data input is a count of weekly COVID deaths by vaccination status over the study time period. That is, it consists of a weekly count of the number of COVID deaths among those who had not received their primary series of vaccinations 14 days prior to death and those that had received those vaccinations. These counts of COVID deaths by vaccination status are stratified by age and are available for 30 U.S. states that integrated their state-level vaccine records with their reporting of vital records statistics reported to the CDC. The second data input is a measure of the size of the vaccinated and unvaccinated populations in these 30 states derived from their state-level vaccine records.

These counts of COVID deaths by vaccination status can be combined with

measures of the size of the vaccinated and unvaccinated populations to estimate the weekly COVID population mortality rate by age and vaccination status. We reproduce the results from this calculation by Jia et al. (2023) reported in their figure S3 here in Figure A.3. The associated weekly measures of the population by vaccination status in these age groups used to derive the COVID mortality rates shown in Figure A.3 are reported by Jia et al. (2023) in their figure S2 reproduced here in Figure A.4

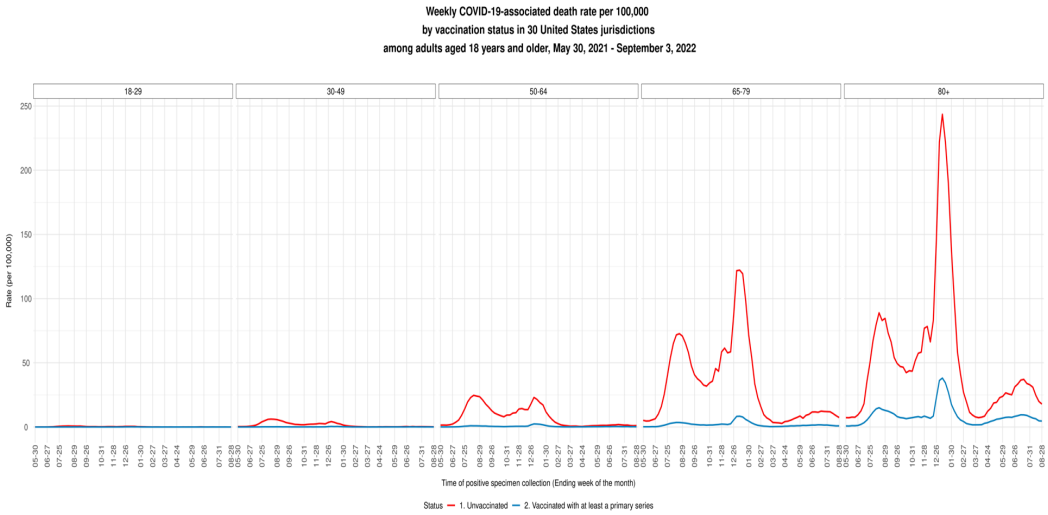

Figure A.3: Weekly Mortality by Vaccine Status and Age May 2021 to Sept 2022 from Jia et al. (2023) figure S3. Mortality rates for the unvaccinated are shown in red. Those for the vaccinated in blue.

Jia et al. (2023) construct an estimate of additional lives saved had vaccination been extended to the entire population by constructing a counterfactual estimate of the number of COVID deaths that would have occurred had everyone been vaccinated by the end of May 2021. To construct this counterfactual COVID death toll, they apply the weekly COVID mortality rate for the vaccinated population shown in Figure A.3 to the unvaccinated population shown in Figure A.4. The difference between the observed death toll in the unvaccinated population and this counterfactual death toll is their estimate of additional lives saved. Using this methodology, they find that, had the entire population received the primary round of two shots by the end of May 2021, an additional 232,000 lives could have been saved.[1]

[1]Zhong et al. (2023) conduct a related analysis of the mortality impact of marginal increases in the rate of vaccination on COVID mortality over the period January 2021 through April

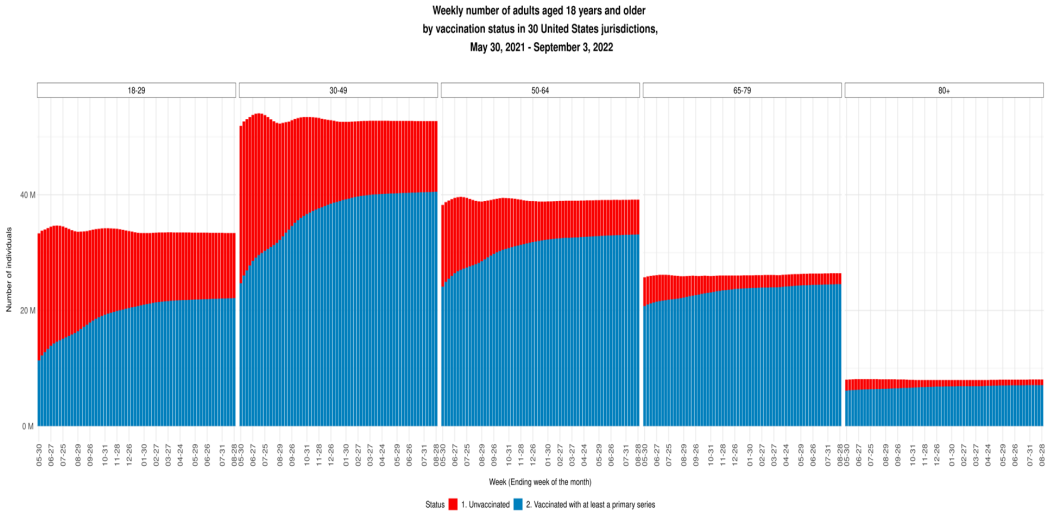

Figure A.4: Population by Vaccine Status and Age May 2021 to Sept 2022 from Jia et al. (2023) figure S2, The number of unvaccinated people in each age category is shown in red and the number of vaccinated in blue.

We note several striking features of these data shown in Figures A.3 and A.4.

First, in Figure A.4 we see that the administration of vaccines in 2021 covered a large portion of the population by the late summer of 2021, with this level of coverage being particularly high for the elderly population. Thus, these administrative data on vaccination are consistent with the serology data we discussed above. These observations of high levels of vaccine coverage by late summer of 2021 are also consistent with the hypothesis that the marginal impact on COVID mortality of raising vaccination coverage from observed levels to 100% of the population is significantly smaller than the impact on COVID mortality of raising vaccine coverage from 0% to observed levels that we aim to estimate.

Second, in Figure A.3, we see that the weekly COVID mortality rate for the unvaccinated was much higher than that for the vaccinated in the Delta and first Omicron waves in the fall of 2021 and early 2022. We see this evidence as consistent with the hypotheses that a first infection with COVID Delta or Omicron was quite dangerous for the unvaccinated and that vaccines provided considerable protection against severe disease during those episodes.

2022. They report that increasing vaccination coverage to 85% of the adult population over this time period would have saved an additional 178,000 lives and increasing it to 100% of the adult population would have saved 319,000 lives. Again these estimates are of the benefit of increasing vaccination rates above those actually achieved. We focus on estimating the impact of achieved vaccination rates.

Table 1: Model Implied Cumulative COVID Deaths:
Alternative Scenarios For Comparison with Other Estimates

Baseline behavior and vaccines	1,180,000
Baseline behavior and faster vaccines	904,000
Baseline behavior and vaccines through Sep 15 2021	659,000
Baseline behavior no vaccines through Sep 15 2021	1,107,000

We note in Figure A.3, the COVID mortality rate for the unvaccinated falls markedly even for the elderly after late March or early April of 2022. We see this decline in the mortality rate for the unvaccinated as consistent with the hypothesis that, after the first Omicron wave, the large majority of the unvaccinated population had likely already been exposed to COVID and hence had some protection from severe disease due to that prior infection. Note that this hypothesis is also consistent with the finding in Barro (2022) from cross-state regressions that the impact of vaccination rates on subsequent COVID mortality appears to decline substantially after the first COVID wave.

In our model, and in the data, the pace of vaccinations slows considerably in the second half of 2021. One might ask, as do Jia et al. (2023), how many lives might have been saved if the initially rapid pace of administering vaccines had been continued through the second half of 2021 and continued through into 2024.

We show our model results for this scenario in the first row of Table 1. Here we find that with continued rapid vaccination, the cumulative death toll would have been 907,000 over the four year duration of the epidemic. In comparison with our baseline death toll of 1,180,000, this counterfactual simulation suggests that more rapid vaccination continuing through the fall of 2021 would have saved an additional 273,000 lives.

As shown in Figure A.5, nearly all of this benefit is realized from the model's implication that more rapid vaccination would have very much reduced the wave of deaths associated with the Delta variant in the fall of 2021 and would have substantially mitigated the initial waves of deaths associated with the Omicron variant in early 2022. Continued rapid vaccination after early 2022 would have had a much smaller marginal impact on the model implied COVID death toll. We see these findings as in line with the results of Jia et al. (2023).

Steele, Couture, and et al (2022) develop an estimate of the impact of vaccines on COVID mortality over the time period December 2020 to September 2021 based on an alternative approach to calculating this impact. Their approach is based on an estimate of flows of new potential infections over the study time period. That is, they develop a structural model of the epidemic similar to an SIR model and attempt to estimate the rate per unit time at which infected individuals would

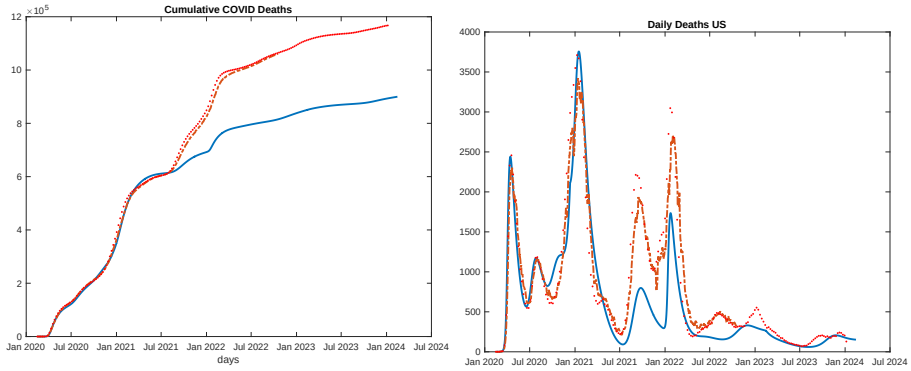

Figure A.5: Model baseline behavior with faster vaccines. Left Panel: Cumulative COVID deaths in blue and data in red. Right Panel: Weekly COVID deaths in blue and data in red

have transmitted their infection to vaccinated individuals had those individuals not been vaccinated together with the associated COVID deaths that would have followed those counterfactual infections.[2] They arrive at an estimate of 235,000 lives saved from December 2020 to September 2021.

This estimate is substantially lower than our estimate primarily because the estimation methodology is different. In particular, if we apply an incremental IFR of 0.05% to their estimate of 235,000 lives saved, we see that this corresponds to an estimate of 47 million additional infections in their counterfactual without vaccines over the period from December 2020 to September 2021, equal to 14% of the U.S. population. What their estimate leaves out relative to ours is an estimate of the additional infections and deaths that would have occurred as the COVID epidemic continued after September 2021 in the counterfactual without vaccines. It is these additional infections and deaths that we aim to count by incorporating an estimate of the final size of the epidemic in our estimation procedure.

To draw a comparison between our model results and those in Steele, Couture, and et al (2022), we simulate our model with baseline behavior with and without vaccines through September 15, 2021. We report these results in the third and fourth rows of Table 1. In this table, we see that we obtain an estimate that vaccines saved 448,000 lives through that date. Note, in our counterfactual, vaccines averted substantially more infections during this time period than estimated in Steele, Couture, and et al (2022).

The Commonwealth Fund has produced a series of estimates of lives saved by

[2]The model is based on one used to estimate the impact of vaccines on seasonal influenza here Tokars et al. (2018). See also Jones, Khader, and Branch-Elliman (2022) for an editorial commentary on this article.

vaccines.[3] The authors of these estimates use an agent-based model based on the model in their earlier paper Moghadas et al. (2021).

They report an estimate that, through March 2022, "Without vaccinations, there would have been an estimated 66 million additional COVID infections". That corresponds to an estimate that an additional 20% of the U.S. population would have been infected absent vaccines. This is an underestimate of counterfactual infections relative to our calculation of averted infections prior to Omicron based on a simple count of those vaccinated before infected as implied by the serology data together with our estimate of the final size of the COVID epidemic.

They then report an estimate of 2,265,222 deaths averted. This estimate of deaths averted seems very high in light of their estimate of infections averted. Using their estimate of infections averted, this estimate of deaths averted implies that those 66 million additional infections would have had an infection fatality rate of 3.4%. It is unclear what evidence this estimate of the incremental infection fatality rate is based on.

B What Drove Differences in State Level Outcomes?

In this section, we examine State level outcomes for cumulative mortality from COVID-19 in greater detail.

To start, we make the argument that, relative to the historical and modeling benchmarks for pandemic influenza discussed above, residents of all 50 states made surprisingly strong and lasting efforts to slow the spread of COVID-19 so that vaccines came in time to save a considerable number of lives. To illustrate this point, in Figure B.6, taken from Chitwood et al. (2022),we show the dynamics of the effective reproduction number for COVID-19 for each of the 50 states of the United States during 2020. In this figure, we observe that behavior in all 50 states changed rapidly and dramatically so as to drive the effective reproduction number of COVID-19 in the state down to one very early on in the epidemic. Moreover, this behavior was sustained so as to keep this effective reproduction number close to one throughout 2020. Atkeson, Kopecky, and Zha (2023) find similar results for both U.S. States and many countries.

As we have discussed above, if the effective reproduction number of a disease remains close to one, then the growth rate of current infections is close to zero.

[3] See, for example, `https://www.commonwealthfund.org/blog/2022/impact-us-covid-19-vaccination-efforts-march-update`. This estimate is an update of a previous estimate from December of 2021at `https://www.commonwealthfund.org/publications/issue-briefs/2021/dec/us-covid-19-vaccination-program-one-year-how-many-deaths-and`. That estimate was for 36 million infections averted and 1.1 million lives saved, again implying an incremental IFR over 3%.

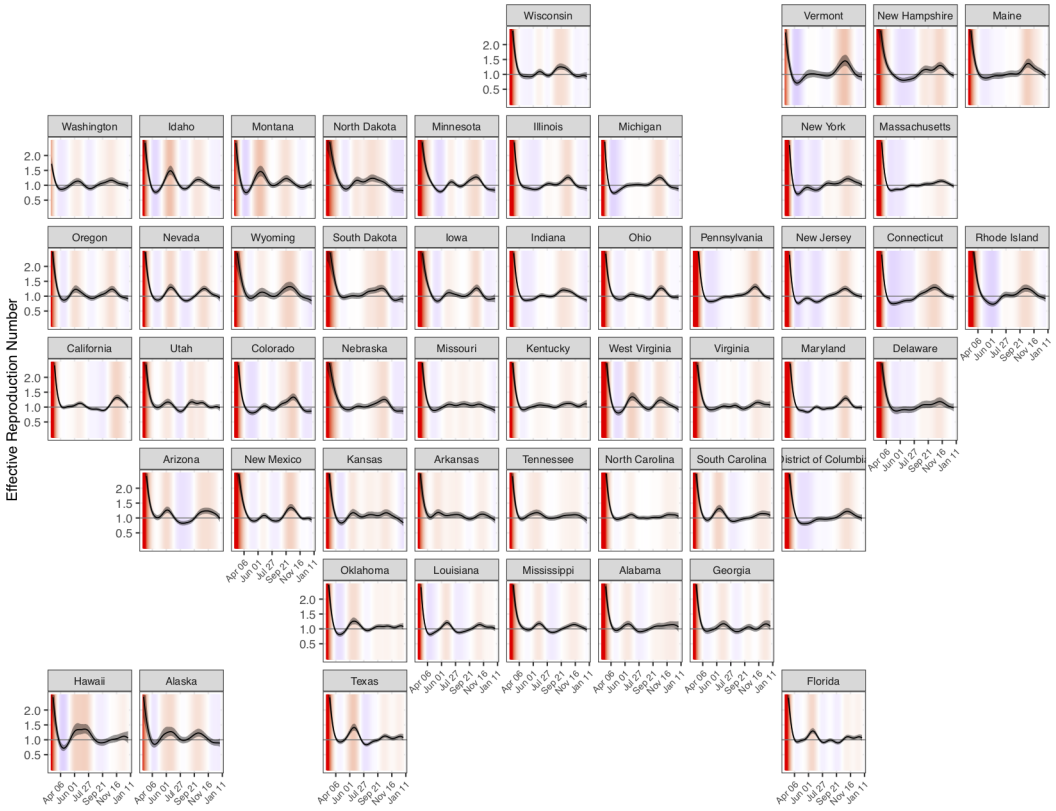

Figure B.6: Dynamics of the Effective Reproduction Number by State. This is Figure 4 from Chitwood et al. (2022). $\mathcal{R}_t$ estimates for each US state from March 1, 2020 to January 1, 2021. Background colors indicate whether $\mathcal{R}_t$ is substantially greater than 1 (red), close to 1 (white), or substantially less than 1 (blue). Grey line indicates $\mathcal{R}_t = 1$. Shaded areas represent 95% credible intervals.

Equivalently, the growth rate of cumulative infections and deaths is then roughly constant. This is precisely the dynamics we observe in cumulative COVID-19 mortality at the state level. In Figure B.7, we show the dynamics of cumulative COVID-19 deaths as an age-adjusted death rate per 100K of the population for selected states. In the left panel of this figure, we show the dynamics of cumulative COVID-19 deaths for California, Florida, New York (excluding New York City), and Texas. We see that New York had a very rapid growth of cumulative deaths in the initial phase of the epidemic, and then settled in to a lower growth rate. Texas had a high growth rate of cumulative deaths throughout the first two years of the epidemic. Given the rhetoric surrounding this topic, we find it striking how similar the age-adjusted outcomes for COVID-19 deaths have been for California

and Florida over the past four years.

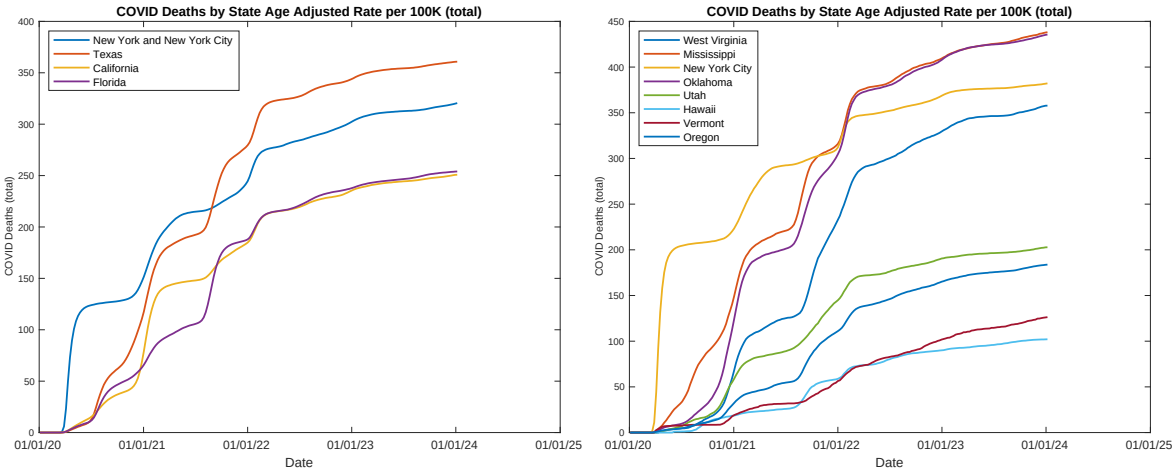

Figure B.7: Dynamics of Cumulative COVID Deaths by State Rate per 100K population age-adjusted. Left Panel: Four big states Right Panel: High and Low Outcomes

In the right panel of Figure B.7, we show the dynamics of cumulative COVID-19 deaths as an age-adjusted death rate per 100K of the population for New York City and seven other states representing extreme high and low mortality outcomes across states. With the exception of New York City, we see largely linear growth in cumulative deaths over the first two years of the COVID-19 epidemic for all of these locations. As evident in the figure, New York City suffered exceptionally rapid initial growth of cumulative COVID-19 deaths in the first wave of the epidemic, likely due to the surprise introduction of a large number of hidden cases from Europe in early 2020.

For futher evidence of this commonality of responses across U.S. states, in Figure B.8, we show estimates from the Commercial Lab and Blood Donor serology surveys of cumulative infections (in red) and combined seroprevalence (in black) for the 50 states of the United States. While these surveys show considerable variation in the estimated percent infected across states, we see in this figure that all of the states followed similar dynamics of slow growth in infections in the first two years of the epidemic and rapid deployment of vaccines in the first half of 2021.[4]

[4]Chitwood et al. (2022) argues that the serology data underestimates the true portion of the population ever infected for a variety of reasons. This paper presents alternative estimates of the state-level portion of the population infected through 2020 in its Figure 7.

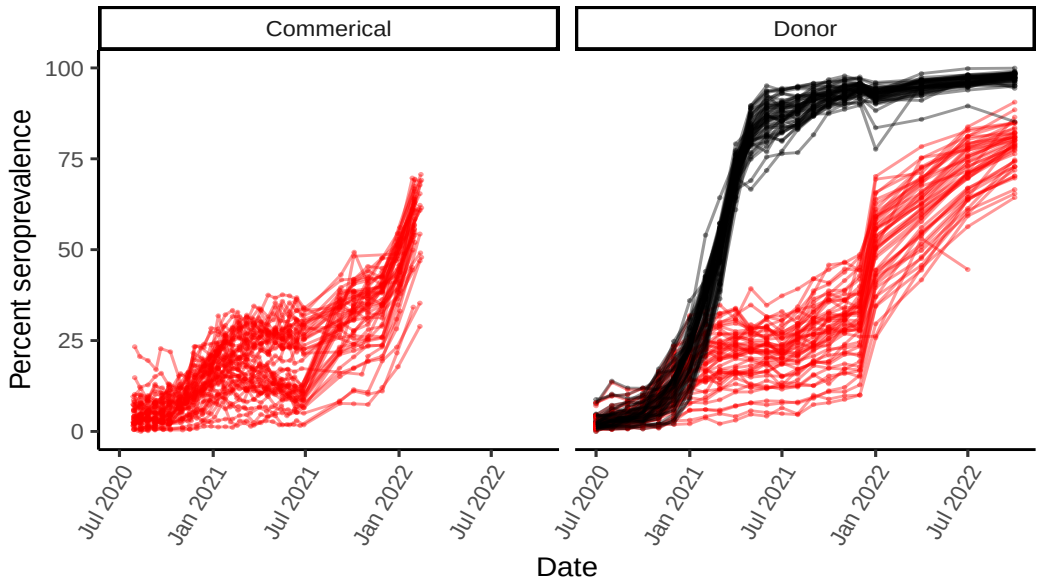

Figure B.8: Dynamics of Blood Serology Estimates of Cumulative Infections (red) and Combined Vaccinated and/or Infected (black). Left Panel: Commercial Lab Survey Right Panel: Blood Donor Surveys

Based on this evidence, we argue that the most important feature of the outcomes across U.S. states (and even countries around the world) is how much they have in common relative to outcomes that were expected given prior epidemiological modeling of and past experiences with pandemic influenza. To a large extent, residents of every state in the United States outside of New York City reacted very strongly to COVID-19 very early on and took significant actions to slow its spread all through 2020 and 2021. We regard the observation that this could be done, and done nearly universally across different states of the U.S., as a great surprise.

To expand further on this point, observe that the model-based forecast in Ferguson and et al (2020) for peak deaths with unmitigated spread of COVID-19 was over 16 deaths per day per 100K population (implying over 50,000 deaths per day in the U.S. as a whole) with 75% of the population being infected by late summer of 2020. This forecast was not out of line with what was experienced in locations that did little to mitigate the spread of SARS-CoV-2. For example, we note that seroprevalence studies in Manaus, Brazil indicated an attack rate of 75% in their first wave of the pandemic.[5] We see nothing like this rapid spread of

[5]See Buss et al. (2021).

COVID-19 in the serology data across U.S. states.

In March and April of 2020, New York City experienced the worst wave of COVID-19 infections and mortality of anywhere in the U.S. over the past four years. Its peak weekly mortality rate was 60 per 100K population (less than 10 per 100K per day) — in the range of one half that predicted in Ferguson and et al (2020) for peak deaths with unmitigated spread. Seroprevalence estimates for New York City indicate up 20% of that population of 8 million people was infected in the first wave in Spring of 2020.[6]

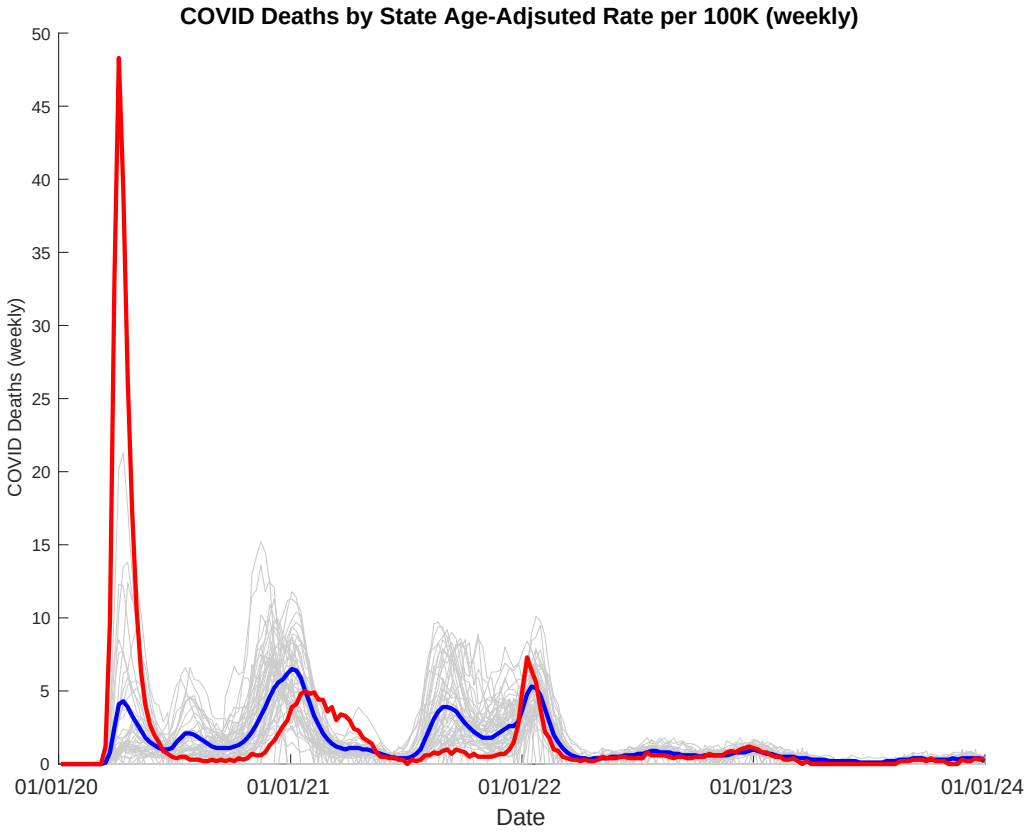

Figure B.9: Dynamics of Weekly COVID Deaths by State Rate per 100K population age-adjusted. Red: New York City Blue: United States

We illustrate the extent to which the first wave of COVID-19 deaths in New York City was an outlier in Figure B.9. In that figure, we show the dynamics of weekly COVID-19 deaths for the 50 states at an age-adjusted rate per 100K of population in gray. We show these dynamics for New York City in red and for the

[6]See Stadlbauer, Tan, and et al (2021).

United States overall in blue. As is clear from the figure, the first wave of COVID-19 deaths in New York City was much larger than any other wave experienced in any state in the United States. That is, the response to flatten the curve and dramatically slow the transmission of COVID-19 was universal across the 50 states of the United States.

As we have discussed above, there has been great interest in comparing the impact of COVID-19 across states of the U.S. in the press and in some academic work. We have argued that, in terms of broad strokes, the dynamics of the COVID-19 epidemic have more in common across states than would be expected from historical experience or projections from pandemic influenza models. We now examine the cross-section outcomes across states and argue that these outcomes are consistent with reasonable variation in either structural factors impacting virus transmissibility that might vary across states (such as weather, density, etc.) and/or reasonable variation in the strength of the behavioral response across states.

B.1 Cross-State Cumulative COVID Mortality and Infections

We start with a review of cross-state outcomes for cumulative deaths shown in Figure B.10. The different linear growth rates of cumulative COVID deaths across states of the United States, sustained over time, led to significant differences in cumulative outcomes through the first two years of the epidemic.

In the left panel of Figure B.10, we show cumulative deaths by state from the beginning of the epidemic through April 4, 2022 as an age-adjusted rate per 100K of population. This dispersion in cumulative COVID death rates is quite wide. In the right panel of Figure B.10, we show the cumulative deaths by state for the period April 4, 2022 through December 30, 2023 on the same scale. Here we see that the growth in cumulative deaths over the past two years has been much slower and more uniform across states than during the first two years.

In Figure B.11 we show results for seroprevalence measured in the Blood Donor Survey from the first quarter of 2022 (at the end of the first big Omicron wave) broken down at the state level. In the left panel of this figure, we see that there is considerable dispersion in the measure of the cumulative percentage of the population infected across states by the end of the first big Omicron wave, with some locations showing only 25-30% of the population infected and others showing roughly 70% of the population infected. In the right panel of this figure, we see, in contrast, that the population in nearly all states had achieved a high level of combined protection from either prior infection or vaccines by the end of the first Omicron wave.[7]

[7]Klaassen et al. (2023) Figures 1, 2 and 3 present alternative estimates of the portion of the population at the state level with effective protection from severe disease from either prior infection or vaccination over the course of 2021. These estimates are also consistent with the

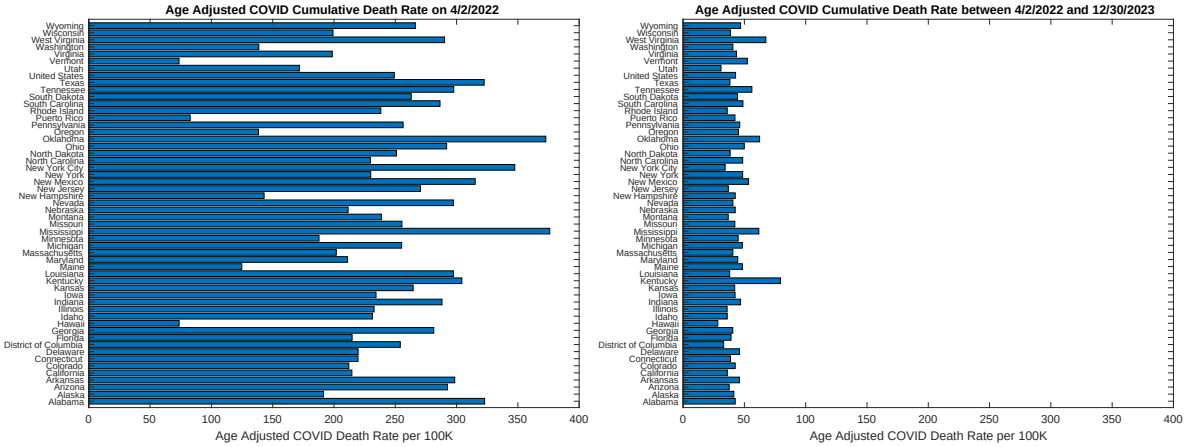

Figure B.10: Cumulative COVID Deaths by State Rate per 100K population age-adjusted. Left Panel: January 2020 - April 2, 2022 Right Panel: April 2, 2020 - December 30, 2023

At a mechanical level, the different outcomes for cumulative deaths across states through the first quarter of 2022 shown in the left panel of Figure B.10 are largely accounted for by differences in cumulative infections in the serology data with moderate variation in the implied infection fatality rates across states. Thus, it appears that the serology data are giving a meaningful measure of cumulative infections.

We illustrate this point in Figure B.12. The left panel of that figure shows a scatter plot of the state-level measure of the cumulative percent of the population infected from the Blood Donor Survey in the first quarter of 2022 on the x-axis and state-level cumulative COVID mortality at an age-adjusted rate per 100K population as of April 4, 2022 on the y-axis. The red line in that figure is a regression line with the intercept set to zero. The slope of that line is consistent with an estimated cumulative infection fatality rate very close to 0.5% based on the cross-state variation in measured infections and deaths by the end of the first quarter of 2022.

The right panel of that figure shows the implied cumulative infection fatality rate at the state level obtained by taking the ratio of the state-level cumulative COVID age-adjusted mortality rate as a percentage of the population and the state-level measure of the cumulative percentage of the population infected as of the 2022 Q1 Blood Donor Survey. (Note that this variation in state-level infection

view that the population across U.S. states had attained high levels of protection from severe disease by the end of 2021.

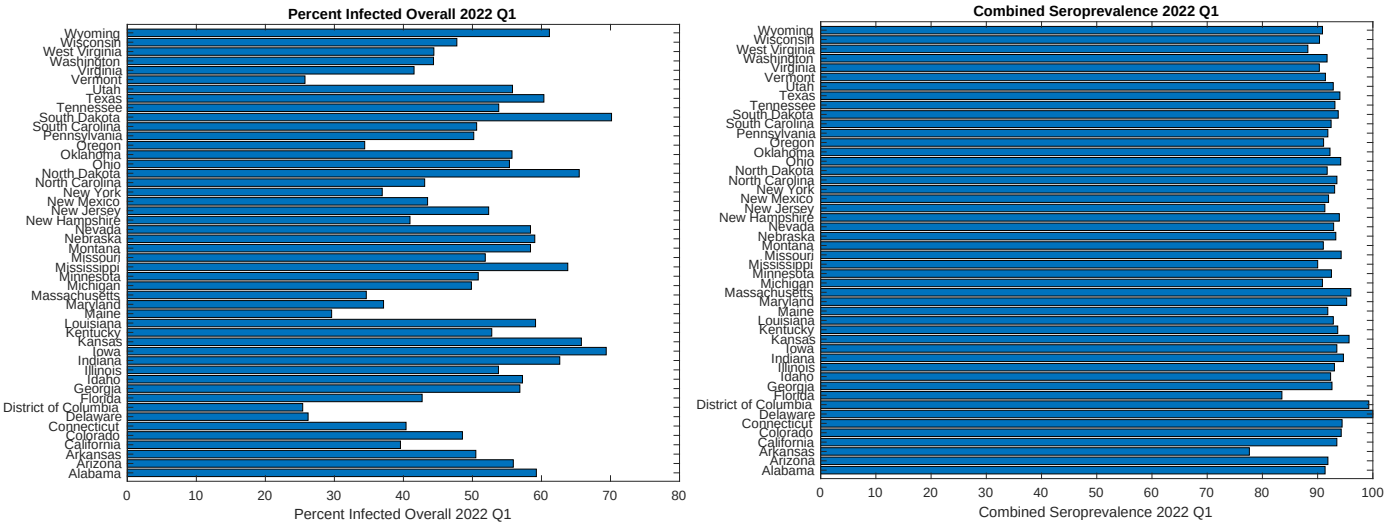

Figure B.11: Left Panel: State-level measures of percent infected in Blood Donor Serology Survey First Quarter 2022
Right Panel: State-level measures of Combined seroprevalence in Blood Donor Serology Survey First Quarter 2022

fatality rates can be both real and due to errors in measurement.) From these figures, we argue that the serology data provide a meaningful measure of the progress of the COVID epidemic at the state level.

After the first large Omicron wave in early 2022, the impact of further increases in measures of infections from serology data on further COVID deaths gets much weaker, consistent with the view that vaccination prior to a first infection with COVID protected substantially against COVID mortality from Omicron. In fact, after the first quarter of 2022, the differences in infection rates across states fell dramatically as Omicron managed to infect the vaccinated at a high rate.

In Figure B.13, we show a scatter plot of the state-level measures of the portion of the population previously infected from the 2022 Q1 Blood Donor Survey and the same measure at the state-level from the 2022 Q4 Blood Donor Survey. In that figure, we see that the percentage of the population infected at the state level in the first quarter of 2022 ranged from a low of roughly 25% to a high of roughly 70%. By the end of 2022, this dispersion in the percentage of the population infected shrank considerably, with the low end in particular now over 60%.

We thus conclude that this relationship between cumulative infections and deaths breaks down after the first big Omicron wave. Over the past two years, the incremental growth of cumulative deaths has been uniformly much smaller across U.S. states despite huge and differential growth in infections measured

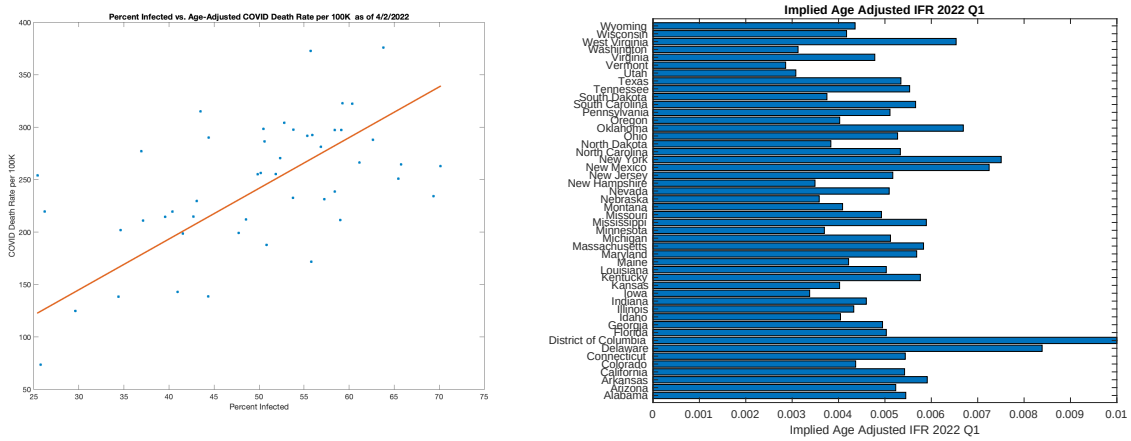

Figure B.12: Left Panel: Scatter plot of state-level cumulative COVID mortality as of April 2, 2022 vs. percent infected in Blood Donor Serology Survey First Quarter 2022 Right Panel: State level Infection Fatality Rates implied by these serology and deaths data.

in the serology data over the course of the remainder of 2022. We argue that these observations of relatively uniform outcomes over the past two years are best accounted for by the observation that by the end of the first quarter of 2022, nearly everyone had been vaccinated or infected or both so that the mortality impact of further infections has been much reduced and now depends largely on the biologically determined speed with which the protection against severe disease offered by vaccination or prior infection wanes or fails due to immune evasion by new variants.

B.2 Moderate Differences in Behavior and Transmissibility Can Account for These Cross-State Outcomes

We now use our model to consider the range of variation in the strength of behavior and of structural factors impacting transmission required to account for this dispersion in outcomes for COVID-19 mortality across states.

We first simulate our model with all baseline parameters except that we consider the strength of behavior as indexed by $\kappa(t)$ to be either twice its baseline value (strong behavior) or half its baseline value (weak behavior). Results are reported in the first two rows of Table 2, with the predicted dynamics of COVID deaths shown in Figure B.14. We see in this figure that in both cases, the model produces the linear growth of cumulative deaths seen in the state level data, just with different

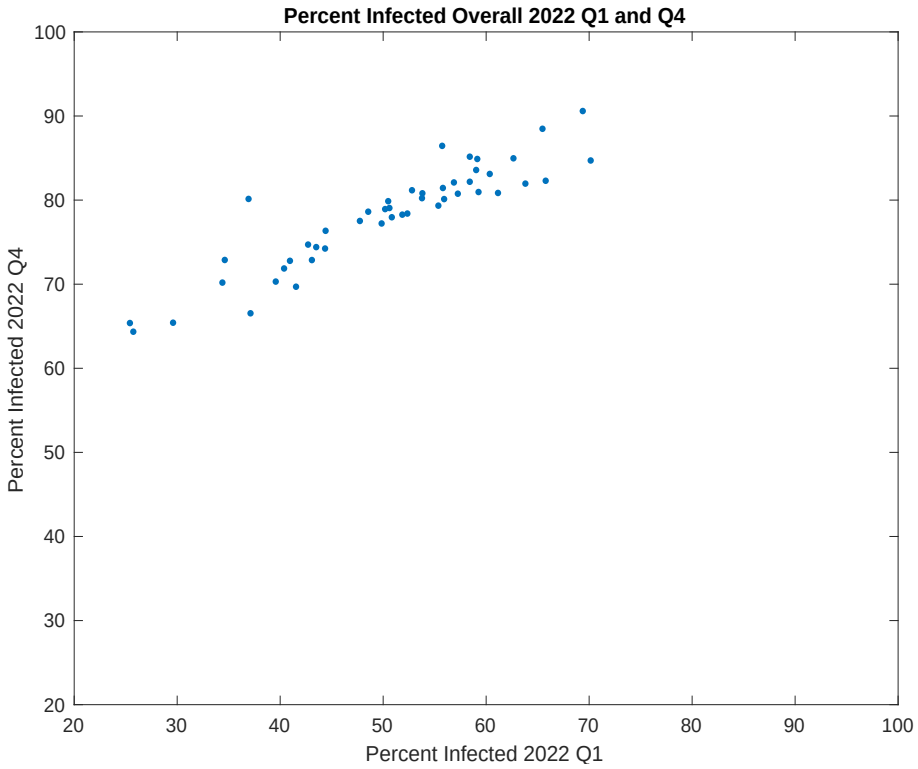

Figure B.13: Scatter plot of Blood Donor serology data on percent infected by state in the fourth quarter of 2022 vs. percent infected by state in the first quarter of 2022

slopes. This variation in behavior also produces a wide range of mortality outcomes when cumulated over time.

We then simulate our model with all baseline parameters except that we consider the transmission constants $\bar{\beta}_i$ to be either 1.5 times or 0.66 times their baseline values, corresponding to faster or slower transmission.[8] Results are reported in the first two rows of Table 2, with the predicted dynamics of COVID deaths shown in Figure B.15.

What is to be learned from this dispersion in infection and mortality outcomes across U.S. states? To what extent did specific actions or policies or patterns of private behavior account for these different outcomes across states? Or inherent differences in infection fatality rates? These questions are hard to answer given all of the confounding factors that also influenced infection and mortality outcomes

[8]See Ives and Bozzuto (2021) and Sy, White, and Nichols (2021) for estimates of the range of variation in COVID transmission rates across regions of the U.S.

Table 2: Model Implied Cumulative COVID Deaths:
Alternative Scenarios Capturing Cross-State Outcomes

Weaker behavior ($0.5 \times \kappa_t$), baseline transmission, with vaccines	1,581,000
Stronger behavior ($2 \times \kappa_t$), baseline transmission, with vaccines	863,000
Faster transmission ($1.5 \times \beta_i$), baseline behavior, with vaccines	1,558,000
Slower transmission ($0.66 \times \beta_i$), baseline behavior, with vaccines	764,000

across states. In particular, it is difficult to assess to what extent it was ex-ante differences in structural factors such as density or weather across states states that might have led to faster or slower transmission or differences in public and private behavior that caused the different growth rates of infections and deaths observed across states.

We do not attempt such an analysis in this paper. Moreover, we argue that the universally strong behavioral response in all 50 states is the most striking feature of the state level data.

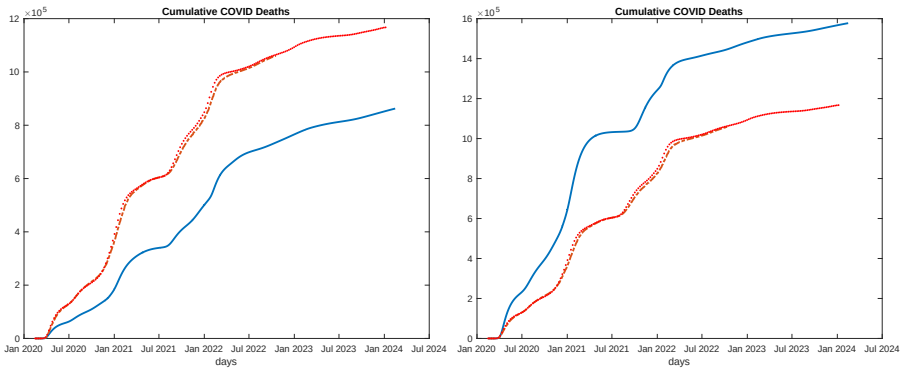

Figure B.14: Model with strong and weak behavioral response Left Panel: Strong behavior Cumulative COVID deaths in blue and data in red. Right Panel: Weak behavior Cumulative COVID deaths in blue and data in red

C Model Appendix

This appendix presents the model and parameters used in this paper. This model is an update of the model presented in the appendix to “Behavior and the Dynamics of Epidemics" by Andrew Atkeson for the Brookings Panel on Economic Activity Spring 2021 and in “The Impact of Behavior and Vaccines on U.S. Cumulative Deaths from COVID-19" by Andrew Atkeson available as NBER Working Paper

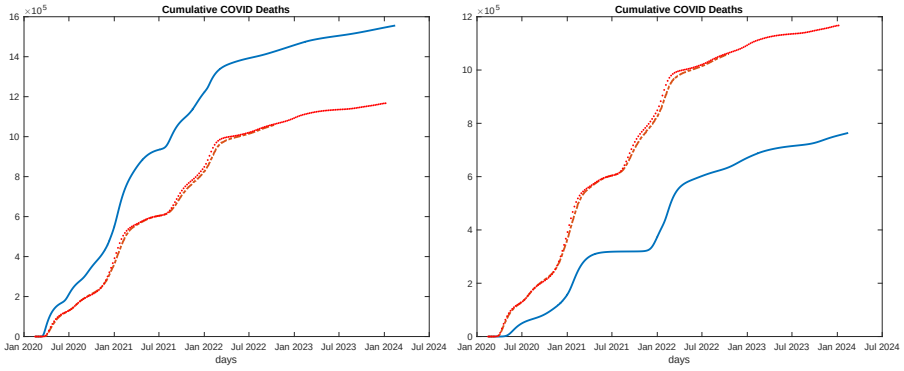

Figure B.15: Model with fast and slow transmission Left Panel: Fast Transmission Cumulative COVID deaths in blue and data in red. Right Panel: Slow Transmission Cumulative COVID deaths in blue and data in red

31525. This model is based closely on that presented in "A Parsimonious Behavioral SEIR Model of the 2020 COVID Epidemic in the United States and United Kingdom" which is available as NBER working paper 28434 and as Federal Reserve Bank of Minneapolis Staff Report 619. This appendix discusses the model extended to include vaccines and the potential for waning immunity, as well as the arrival of the Delta and Omicron variants. It is applied to the United States.

This appendix has three parts. In section C.1, we present the equations of the model. We also compare the structure of this model with that of a simpler behavioral SIRD model as analyzed in Atkeson, Kopecky, and Zha (2021)[9] and Droste and Stock (2021)[10] .

In section C.2, we discuss the values of the parameters and the rationale behind the choice of these parameters. The model is fit to US data on daily deaths from COVID as well as the serology estimates of the cumulative portion of the population infected and vaccinated prior to infection. Several parameters are set to match recommendations from the Center for Disease Control for modeling of COVID-19.

A full version of the MATLAB codes to run this model are posted with the final paper.

[9]Atkeson, Kopecky, and Zha "Behavior and the Transmission of COVID-19" forthcoming, *American Economic Review Papers and Proceedings* with the longer version available here `https://www.minneapolisfed.org/research/staff-reports/behavior-and-the-transmission-of-covid19`

[10]Droste and Stock "Adapting to the COVID-19 Pandemic" (2021) *American Economic Review Papers and Proceedings*

C.1 Model

The model is as follows.

The SEIHR model extends the SIR model by adding both the exposed state E and the hospitalized state H. In this version of the model the total population N is given by the sum of susceptible agents in state S, exposed in state E, infected in I, hospitalized in H, recovered in R, and dead in D. We do not consider population growth in the model.

The compartments E and I are further broken down by variant i, where i indexes the original variant, and the Alpha, Delta, and Omicron variants. The rate at which agents leave the E_i compartmenst for both the normal and more transmissible variants is σ and the rate at which agents leave the I_i compartments for all variants is γ. We also include compartments E_i and I_i corresponding to those experiencing breakthrough Omicron infections. These individuals are modeled as having immunity to previous variants but not to Omicron. The mean generation time for the model is then $1/\sigma + 1/\gamma$. As discussed below, the choice of these parameters is guided by CDC recommendations for these disease parameters.

As agents leave the I_i compartment, fraction η_i go into the hospitalized compartment H and $1 - \eta_i$ transition directly to the recovered compartment R_i. The rate at which agents leave the H compartment is ζ. We assume that all agents leaving the H compartment die. Thus, the overall infection fatality rate for variants is given by η_i and the mean time in the H compartment corresponding to illness and delays in reporting deaths is $1/\zeta$. Note that with these assumptions, it is not appropriate to compare the model's predictions for hospitalizations to data. Instead, this H compartment simply serves to introduce a delay between infection and death.

Those who recover from an infection with the original variant or the Alpha or Delta variants flow into compartment R and are immune from a second infection with these variants. This immunity wanes at rate ξ. They are also susceptible to breakthrough infections with Omicron as discussed below. There is a compartment V introduced to count those who have protection from vaccination prior to a first infection with COVID. Agents flow into this compartment from the compartment S as they are vaccinated and flow out with waning immunity (at rate ξ) and breakthrough infections from Omicron. There is a separate recovered compartment R_O for those recovered from an infection with Omicron. This is introduced to allow for faster waning immunity from an Omicron infection at rate ξ_o.

To introduce breakthrough infections for the Omicron variant, We assume that those infected with Omicron transmit their infection to those in the removed compartment R (those recovered from an infection with a prior variant) and the vaccinated prior to first infection compartment V at a fraction ν_O of the rate at which they transmit their infection to those in the susceptible compartment S. The infection fatality rate from Omicron for those with no prior infection

or vaccination is η_O while that for those infected with Omicron but with prior immunity (i.e. coming out of the R or V compartments) is η_{OR}. We include separate compartments for those exposed to and infected with Omicron depending on whether they came out of the S compartment or the R compartment to allow for these separate infection fatality rates. We assume that those infected with Omicron who do not die transit to a separate compartment removed compartment R_O indicating immunity from prior infection with Omicron. We assume that this protection against a second Omicron infection and serious disease wanes at the rate ξ_0.

The transmission rate of the original variant is denoted by $\beta(t)$. Those for the new variants are denoted by $\beta_i(t)$. New variants are introduced by setting $\bar{E}_i(t) = 1/population$ in the equations below for several days around a specified date t_i and equal to zero otherwise. This allows for a discrete jump from zero exposed to particular variant to a positive number. The window of days for this introduction needs to be chosen so that the differential equation solver (which samples on discrete dates) picks up the introduction of the variant.

The dynamics of the model are given by

$$\frac{dS(t)}{dt} = -\left(\beta(t)I(t) + \sum_{i=A,D}\beta_i(t)I_i(t) + \beta_O(I_O(t) + I_{OO}(t))\right)S(t) -$$

$$\lambda(t)S(t) + \xi(R(t) + V(t)) + \xi_O R_O(t)$$

Here the original variant is denoted without a subscript, $i = A, D$ refers to Alpha and Delta, $i = O$ refers to Omicron, and $I_{OO}(t)$ refers to those with a breakthrough Omicron infection. The parameter $\lambda(t)$ is the vaccination rate. The inflows $\xi(R(t) + V(t))$ are $\xi_O R_O(t)$ are due to waning immunity. Note that these individuals are also susceptible to severe disease and death in a manner equivalent to a completely naive individual.

The outflows from the S compartment are distributed as follows.

$$\frac{dE(t)}{dt} = \beta(t)I(t)S(t) - \sigma E(t)$$

For $i = A, D$

$$\frac{dE_i(t)}{dt} = \beta_i(t)I_i(t)S(t) - \sigma E_i(t) + \bar{E}_i(t)$$

For $i = O$

$$\frac{dE_O(t)}{dt} = \beta_O\nu_O(I_O(t) + I_{OO}(t))(R(t) + V(t)) - \sigma E_O(t) + \bar{E}_O(t)$$

Here the terms $\bar{E}_i(t)$ are used to introduce agents infected with new variants on particular dates. In the code, these terms are zero except for a window of several

days in which the new variant is introduced with this term equal to 1/332,000,000 in this two-day time window. For $i = OO$ representing breakthrough Omicron infections we have

$$\frac{dE_{OO}(t)}{dt} = \beta_O(t)(I_O(t) + I_{OO}(t))\nu_O(R(t) + V(t)) - \sigma E_{OR}(t)$$

Outflows from the exposed states move to corresponding infectious states as follows

$$\frac{dI(t)}{dt} = \sigma E(t) - \gamma I(t),$$

For $i = A, D, O, OO$

$$\frac{dI_i(t)}{dt} = \sigma E_i(t) - \gamma I_i(t)$$

Infection leads to hospitalization (and death) as follows. Infection fatality rates are denoted by η_i.

$$\frac{dH(t)}{dt} = \eta(t)\gamma(I(t) + \sum_{i=A,D} I_i(t)) + \eta_O\gamma I_O + \eta_{OO}\gamma I_{OO} - \zeta H(t)$$

The infection fatality rate for the original, Alpha, and Delta variants, denoted by $\eta(t)$ is allowed to vary by time to reflect the apparent decline in the infection fatality rate from 2020 to 2021 implied by the serology data. The infection fatality rates for an initial infection with Omicron η_O and a breakthrough infection with Omicron η_{OO} are assumed to be constant over time.

Deaths are recorded as agents flow out of the H compartment

$$\frac{dD(t)}{dt} = \zeta H(t)$$

Those who do not die from their infection with the original, Alpha, or Delta variants flow direclty into the recovered compartment

$$\frac{dR(t)}{dt} = (1 - \eta)\gamma(I(t) + \sum_{i=A,D} I_i(t)) - \beta_O\nu_O(I_O(t) + I_{OO}(t))R(t) - \xi R(t)$$

Note the two outflows from this compartment are due to breakthrough infections with Omicron and waning immunity.

$$\frac{dR_O(t)}{dt} = (1 - \eta_O)\gamma I_O + (1 - \eta_{OO})\gamma I_{OO} - \xi_O R_O(t)$$

All those who do not die of Omicron flow into the recovered from Omicron compartment. The evolution of those vaccinated prior to their first COVID infection is given by

$$\frac{dV(t)}{dt} = \lambda(t)S(t) - \beta_O(t)(I_O(t) + I_{OO}(t))\nu_O V(t) - \xi V(t)$$

To measure cumulative first infections (corresponding to the serology data) we introduce a compartment $CI(t)$ that evolves as follows

$$\frac{dCI(t)}{dt} = \left(\beta(t)I(t) + \sum_{i=A,D} \beta_i(t)I_i(t)\right) +$$

$$\beta_O\left(I_O(t) + I_{OO}(t)\right)S(t) + \beta_O(t)(I_O(t) + I_{OO}(t))\nu_O V(t) - \xi CI(t)$$

The reduced-form for the behavioral response of the transmission rate to the level of daily deaths is given by

$$\beta(t) = \bar{\beta}\exp(-\kappa(t)\frac{dD(t)}{dt} + \psi(t)) \tag{1}$$

and for variants

$$\beta_i(t) = \bar{\beta}_i\exp(-\kappa(t)\frac{dD(t)}{dt} + \psi(t))$$

where the parameters $\bar{\beta}$ and $\bar{\beta}_i$ control the baseline transmissibility of the original and subsequent variants of COVID, the parameter $\psi(t)$ is used to introduce seasonality in transmission, and $\kappa(t)$ is the semi-elasticity of transmission with respect to the level of daily deaths. Note that the relative transmissibility of each variant for any level of daily deaths and point in the seasonal cycle is given by $\bar{\beta}_i/\bar{\beta}$.

To model seasonality in the transmission of the virus, we set

$$\psi(t) = seasonalsize * (\cos((t + seasonalposition) * 2\pi/365) - 1)/2$$

where *seasonalsize* controls the magnitude of the seasonal fluctuations in transmissibility holding behavior fixed and *seasonalposition* controls the location of the seasonal peak in transmission. Note that t is indexed to $t = 0$ on February 15, 2020.

To model the change in $\kappa(t)$, I set

$$\kappa(t) = \bar{\kappa} * (1 - normcdf(t, fatiguemean, fatiguesig))+$$

$$fatiguesize * \bar{\kappa} * normcdf(t, fatiguemean, fatiguesig)$$

where $\bar{\kappa}$ sets the initial semi-elasticity of transmission with respect to daily deaths, $fatiguesize$ sets the percentage reduction in this semi-elasticity in the long run, *normcdf* is the normal CDF, $fatiguemean$ sets the date at which the transition in $\kappa(t)$ from its initial to new long run level is halfway complete, and $fatiguesig$ sets the speed with which that transition occurs.

Initial conditions for all simulations are $E(0) > 0$, $E_i(0) = I(0) = I_i(0) = R(0) = R_O(0) = V(0) = H(0) = D(0) = CI(0) = 0$, $S(0) = 1 - E(0)$. For the United States, $E(0) = 33$ on February 15, 2020 out of a population of 332 million. The model is simulated for four years after its starting date.

C.2 Parameters

In this section we discuss the choice of parameters of the model.

The four variants of COVID-19 considered in the model are the original variant, and the Alpha, Delta, and Omicron Variants. We assume that for all four variants, the rate at which agents flow from exposed (E) to infectious (I) compartments is given by $\sigma = 0.425$ and the rate at which they flow out of the infectious compartment is $\gamma = 0.4$. The parameter σ corresponds to an expected time before an exposed agent becomes infectious of 2.35 days and the parameter γ corresponds to an expected time for which an infected individual is infectious of 2.5 days. The generation time is defined as the average time between which one infected individual shows symptoms and a person infected by that individual shows symptoms. These two parameters together imply an average generation time of $1/\sigma + 1/\gamma = 4.85$ days.[11] As mentioned above, this generation time sets the time-scale of the epidemic implied by the model.

The rate at which agents flow out of the hospitalized compartment is $\zeta = 1/30$, corresponding to a mean time between infection and reported death of 30 days. The rates at which immunity wanes (so agents from from the R or RO compartments back to the S compartment) are set at $\xi = \xi_O = 1/(3 \times 365)$ corresponding to a mean time in these compartments (and thus with protection from severe disease) of three years.

The baseline transmission rates β_i for the four variants are given by $\bar{\beta} = 3\gamma$ for the original variant, $\bar{\beta}_A = 4\gamma$ for Alpha, $\bar{\beta}_D = 8\gamma$ for Delta, and $\bar{\beta}_O = 15.3\gamma$. The scalar in front of γ corresponds in the model to the basic reproduction number for each variant. Omicron is also assumed to breakthrough to infect those in the R compartment. The parameter ν_O governing the rate of these breakthrough infections is set to 1/10.

The infection fatality rate for the original variant is set to $\eta(t) = 0.01$ from February 15 through mid December 2020. After that time, the infection fatality rate for the original variant, Alpha, and Delta is $\eta(t) = 0.05$. This pattern of declining IFRs is required to simultaneously match the estimates of cumulative COVID deaths and COVID infections from the serology data. The infection fatality rate for someone infected with Omicron out of the S compartment (and thus either with no prior infection or vaccination or whose protection against severe disease has waned) is $\eta_O(t) = 0.0015$ corresponding to 30% of the IFR of the Delta variant that it displaced. This infection fatality rate declines gradually through 2022 and 2023. The infection fatality rate for breakthrough infections with Omicron is set at $\eta_{OO} = 2.25 \times 10^{-5}$ or 1.5% of that of first infections with

[11] See `https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html`. On that webpage, the CDC notes a mean time of approximately six days between symptom onset in one person to symptom onset in another person infected by that individual.

Omicron.

The Alpha variant is introduced on November 30, 2020. The Delta variant is introduced on April 25, 2021. The Omicron variant is introduced on November 13, 2021

Note that our parameter choices for variant transmissibility $\bar{\beta}_i$ and infection fatality rates are chosen to match the time patterns of growth rates of infections (from the serology data) and deaths. As is standard for models of this kind, a number of additional parameters impact the model implications for growth rates of infections and deaths, including the parameters σ and γ impacting the generation interval and the parameter ν_O governing the probability of breakthrough infections. Further work is required to reconcile these parameters and our model specification with other data measuring the timing and speed of the emergence of new variants measured with testing data. The model implications for the prevalence of each variant at each moment in time is given by the vector of $I_i(t)$.

The initial semi-elasticity of transmission with respect to daily deaths (measured as a fraction of the population) for the United States is $\bar{\kappa} = 250000$. To model the onset of pandemic fatigue in the United States, We set $fatiguesize = 0.375$ and $fatiguemean = 285$ and $fatiguesig = 15$. These parameters imply that $\kappa(t)$ falls from its original value of $\bar{\kappa}$ to 35% of that value in mid to late November of 2020. This behavioral parameter then remains constant for the remainder of the simulation.

To model seasonality of transmission in the United States, we set $seasonalsize = 0.35$ (line 74) and $seasonalposition = 20$ (line 77). This seasonal variation in the parameter $\psi(t)$ leads to variation over time in the basic reproduction number of the virus as discussed in the Spring 2021 version of this paper.

To model the impact of vaccines, we set $\lambda(t) = 0.0065$ starting on January 1, 2021 and zero before that date. Vaccines are administered at this rate for the first 185 days of 2021. The rate of vaccination then drops to $\lambda(t) = 0.0065/5$ until the end of 2022 and then $\lambda(t)$ is set to zero after that. In our model, agents in compartment $V(t)$ enjoy full protection from infection by the Alpha and Delta variants and substantial protection against death from Omicron. Thus, we regard the number of agents in this compartment as representing the population that is both vaccinated prior to a first COVID infection and that gained protection from that vaccination. We assume that this is 75% of the total vaccinated. Thus, when we compare the model implications for $V(t)$ to the measures from the serology data on those vaccinated but not infected, we plot $V(t)/0.75$ as a measure of the total population vaccinated.

References

Atkeson, Andrew. 2023. "The Impact of Vaccines and Behavior on U.S. Cumulative Deaths from COVID-19." Working Paper 31525, National Bureau of Economic Research.

Atkeson, Andrew G., Karen A. Kopecky, and Tao Zha. 2023. "Four Stylized Facts About Covid-19." *International Economic Review* .

Barro, Robert J. 2022. "Vaccination rates and COVID outcomes across U.S. states." *Economics and Human Biology* 47 (101201):1–10.

Bollyky, Thomas J, Emma Castro, Aleksandr Y Aravkin, Kayleigh Bhangdia, Jeremy Dalos, and Erin N Hulland et al. 2023. "Assessing COVID-19 pandemic policies and behaviours and their economic and educational trade-offs across US states from Jan 1, 2020, to July 31, 2022: an observational analysis." *The Lancet* 401:1341–1360.

Buss, LF, Prete CA Jr, Abrahim CMM, Mendrone A Jr, Salomon T, de Almeida-Neto C, França RFO, Belotti MC, Carvalho MPSS, and EC Sabino et al. 2021. "Three-quarters attack rate of SARS-CoV-2 in the Brazilian Amazon during a largely unmitigated epidemic." *Science* 371 (6526):288–292.

Chitwood, Melanie H., Marcus Russi, Kenneth Gunasekera, Joshua Havumaki, Fayette Klaassen, Virginia E. Pitzer, Joshua A. Salomon, Nicole A. Swartwood, Joshua L. Warren, Daniel M. Weinberger, Ted Cohen, and Nicolas A. Menzies. 2022. "Reconstructing the course of the COVID-19 epidemic over 2020 for US states and counties: Results of a Bayesian evidence synthesis model." *PLoS Computational Biology* .

Daskalakis, Demetre, Nikki Romanik, and Ashish K. Jha. 2024. "Lessons from the Mpox Response." *Journal of the American Medical Association* Published online (doi:10.1001/jama.2023.27868).

Eales, Oliver, David Haw, Haowei Wang, Christina Atchison, Deborah Ashby, and Graham S. Cooke et al. 2023. "Dynamics of SARS-CoV-2 infection hospitalisation and infection fatality ratios over 23 months in England." *PLOS Biology* 21 (5:e3002118):https://doi.org/10.1371/journal.pbio.3002118.

Ferguson, Neil M and et al. 2020. "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand." Report 9, Imperial College COVID-19 Response Team.

Ives, Anthony R. and Claudio Bozzuto. 2021. "Estimating and explaining the spread of COVID-19 at the county level in the USA." *Communications Biology* 4 (1). URL `http://dx.doi.org/10.1038/s42003-020-01609-6`.

Jia, Katherine M., William P. Hanage, Marc Lipsitch, Amelia G. Johnson, Avnika B. Amin, Akilah R. Ali, Heather M. Scobie, and David L. Swerdlow. 2023. "Estimated preventable COVID-19-associated deaths due to non-vaccination in the United States." *European Journal of Epidemiology* 38 (11):1125–1128.

Jones, Malia, Karim Khader, and Westyn Branch-Elliman. 2022. "Estimated Impact of the US COVID-19 Vaccination Campaign—Getting to 94 percent of Deaths Prevented." *JAMA Network Open* 5 (7):e2220391.

Kerpen, Phil, Stephen Moore, and Casey B. Mulligan. 2023. "A Final Report Card on the States' Response to COVID-19." *International Journal of the Economics of Business* 30 (2):139–158.

Klaassen, Fayette, Melanie H. Chitwood, Ted Cohen, Virginia E. Pitzer, Marcus Russi, Nicole A. Swartwood, Joshua A. Salomon, and Nicolas A. Menzies. 2023. "Population Immunity to Pre-Omicron and Omicron Severe Acute Respiratory Syndrome Coronavirus 2 Variants in US States and Counties Through 1 December 2021." *Clinical Infectious Diseases* 76:e350–e359.

Moghadas, Seyed M, Pratha Sah, Thomas N Vilches, and Alison P Galvani. 2021. "Can the USA return to pre-COVID-19 normal by July 4?" *The Lancet Infectious Diseases* 21 (8):P1073–1074,.

Stadlbauer, Daniel, Jessica Tan, and Kaijun Jiang et al. 2021. "Repeated cross-sectional sero-monitoring of SARS-CoV-2 in New York City." *Nature* 590:146–150.

Steele, Molly K., Alexia Couture, and Carrie Reed et al. 2022. "Estimated Number of COVID-19 Infections, Hospitalizations, and Deaths Prevented Among Vaccinated Persons in the US, December 2020 to September 2021." *JAMA Network Open* 5 (7):e2220385.

Sy, Karla Therese L., Laura F. White, and Brooke E. Nichols. 2021. "Population density and basic reproductive number of COVID-19 across United States counties." *PLOS ONE* 16 (4):e0249271.

Tokars, Jerome I., Melissa A. Rolfes, Ivo M. Foppa, and Carrie Reed. 2018. "An evaluation and update of methods for estimating the number of influenza cases averted by vaccination in the United States." *Vaccine* 36 (48):7331–7337.

Zhong, Ming, Tamara Glazer, Meghana Kshirsagar, Richard Johnston, Rahul Dodhia, Allen Kim, Divya Michael, Santiago Salcido, Sameer Nair-Desai, Thomas C. Tsai, Stefanie Friedhoff, William B Weeks, and Juan M. Lavista Ferres. 2023. "Estimating vaccine-preventable COVID-19 deaths under counterfactual vaccination scenarios in the United States." Preprint doi.org/10.1101/2022.05.19.22275310, Research Square.