FEDERAL RESERVE BANK OF MINNEAPOLIS | Research Division

STAFF REPORT
No. 618

Behavior and the Transmission of COVID-19

February 2021

Andrew G. Atkeson
University of California, Los Angeles, Federal Reserve Bank of Minneapolis, and NBER

Karen Kopecky
Federal Reserve Bank of Atlanta

Tao Zha
Federal Reserve Bank of Atlanta, Emory University, and NBER

DOI: https://doi.org/10.21034/sr.618
Keywords: COVID; Behavior; Epidemics
JEL classification: E1, E17, I10, I18

Behavior and the Transmission of COVID-19*

Andrew G. Atkeson† Karen Kopecky‡ Tao Zha§

December 31, 2020

Abstract

We show that a simple model of COVID-19 that incorporates feedback from disease prevalence to disease transmission through an endogenous response of human behavior does a remarkable job fitting the main features of the data on the growth rates of daily deaths observed across a large number countries and states of the United States from March to November of 2020. This finding, however, suggests a new empirical puzzle. Using an accounting procedure akin to that used for Business Cycle Accounting as in Chari et al. (2007), we show that when the parameters of the behavioral response of transmission to disease prevalence are estimated from the early phase of the epidemic, very large wedges that shift disease transmission rates holding disease prevalence fixed are required both across regions and within a region over time for the model to match the data on deaths from COVID-19 as an equilibrium outcome exactly. We show that these wedges correspond to large shifts in model forecasts for the long-run attack rate of COVID-19 both across locations and over time. Future research should focus on understanding the sources in these wedges in the relationship between disease prevalence and disease transmission.

*We are grateful to Hongyi Fu for superlative research assistance and Ben Moll for very useful discussions. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Banks of Atlanta and Minneapolis, the Federal Reserve System, or the National Bureau of Economic Research.

†Department of Economics, University of California, Los Angeles, NBER, and Federal Reserve Bank of Minneapolis, e-mail: andy@atkeson.net

‡Federal Reserve Bank of Atlanta, e-mail: karen.kopecky@atl.frb.org

§Federal Reserve Bank of Atlanta, Emory University, and NBER, e-mail: zmail@tzha.net

1 Introduction

Since the outbreak of the COVID-19 pandemic in early 2020, epidemiologists and economists have raced to develop models of the disease to be used for forecasting the progression of the epidemic, for understanding the interaction of the epidemic with the economy, and for evaluating the effectiveness of various interventions aimed at mitigating the spread of the disease.[1] The models developed by economists often differ from those developed by epidemiologists in that they include equations intended to capture the impact of endogenous changes in human behavior undertaken in response to the epidemic on the progression of the epidemic itself.[2] We refer to such models as *behavioral* SIR, or BSIR, models.

In this paper we develop and implement an accounting procedure to decompose the variation in transmission rates needed to account for the data on COVID deaths for a large number of countries and states of the United States into a portion that is accounted for by a simple behavioral SIR model[3] and a portion that must be accounted for by a wedge between the transmission rate implied by the BSIR model and that observed in the data.[4] To implement this procedure, we first estimate the parameters that govern the relationship between disease transmission and disease prevalence for each region using data from the initial phase of the pandemic. Then,

[1]The CDC published a Science Agenda for COVID-19 on November 12, 2020 and named improvements in mathematical modeling as an important scientific objective going forward. See Priority Area 1, Objective 4 in https://www.cdc.gov/coronavirus/2019-ncov/more/science-agenda-covid19.html.

[2]Tomas Phillipson pioneered the application of such models to the HIV epidemic. See, for example, Phillipson and Posner (1993). For examples of economic models applied to COVID-19, see Eksin et al. (2019), Keppo et al. (2020) (slides available at https://www.lonessmith.com/wp-content/uploads/2020/04/pandemic-slides.pdf), Farboodi et al. (2020), Eichenbaum et al. (2020), Guerrieri et al. (2020), Kaplan et al. (2020), Toxvaerd (2020), Acemoglu et al. (2020), Krueger et al. (2020), Gans (2020) and many others.

[3]The model we use is derived from one first presented by Gianluca Violante at the Spring NBER EFG meeting in slides available at http://conference.nber.org/confer/2020/EFGs20/Violante.pdf. Our model is also very similar to that presented by John Cochrane's at https://johnhcochrane.blogspot.com/2020/05/an-sir-model-with-behavior.html. See also Eksin et al. (2019).

[4]Our COVID-accounting procedure follows the methodology for Business Cycle Accounting laid out by Chari et al. (2007).

given this parameterization of our simple BSIR model, we back out the "wedges" to the equations used to model the relationship between disease incidence and disease transmission needed to have the model account for the observed data on deaths as an equilibrium outcome exactly in each location that we study. We argue that these wedges are of interest because they correspond to shifts in the transmission rate of COVID-19 holding behavior fixed. We also show that these wedges correspond in magnitude to shifts in the model-implied forecast for the long-run attack rate for COVID-19 both across locations and over time.

One central prediction of BSIR models is that the growth rate of infections and daily deaths from the disease should fall in the initial phase of the pandemic much more rapidly than would be predicted from standard SIR epidemiological models. Behavioral models predict that humans endogenously reduce their interactions with each other in response to rising disease prevalence, thus reducing the transmission rate of the disease well before the population approaches herd immunity.[5] A second central prediction of these BSIR models is that, after an initial peak of infections and daily deaths, the growth rate of infections and daily deaths from the disease should remain relatively close to zero. This is due to an equilibrium outcome determined by the endogenous response of human behavior to the prevalence of the disease. Shocks that raise infections and daily deaths should lead to an endogenous reduction in human activity, reducing transmission, and shocks that reduce infections and daily deaths should lead to an endogenous increase in human activity, increasing transmission. Thus, the endogenous behavioral response should lead, in equilibrium, to a linear, rather than exponential, rate of increase of cumulative cases and deaths from the disease.

We show that data on the growth rate of daily deaths from 69 countries and 34 states of the United States from March into November of 2020 conform to these predictions of our simple BSIR model remarkably well, even when the model does not include wedges to force the model to match the data. We show in the data that

[5]Eksin et al. (2019) apply a simple BSIR model to the study of data on prior epidemics and emphasize that standard SIR models routinely overpredict the severity of epidemics.

the growth rate of daily deaths in these locations started at high and highly dispersed levels in the early days of the epidemic and then fell toward zero fairly rapidly. The growth rates of daily deaths observed in these locations since this initial phase of the epidemic have remained in a relatively narrow range around zero into mid-November. This has remained true even with substantial second and third waves of daily deaths observed in many locations this Fall and Winter.

Despite this remarkable match between the predictions of a BSIR model for the growth rates of daily deaths and the evolution of deaths from the COVID-19 pandemic observed to date, other important features of the evolution of the pandemic in many locations remain unexplained. The level of daily deaths per capita from the disease has varied tremendously both across locations at a point in time and across time in a given location. Many locations in Europe, such as Italy for example, saw a big initial peak of daily deaths in the Spring, then experienced a substantial decline in daily deaths to very low levels in the Summer, and now in the Fall, have seen daily deaths return to their Spring peaks. Japan shows a similar pattern of multiple waves of daily deaths, although at substantially lower levels throughout than have been seen in Italy. Other locations, such as Arizona for example, showed a slow and steady growth in daily deaths for several months in the initial phases of their epidemics. When viewed through the lens of a behavioral SIR model, what forces are driving these different disease outcomes across locations?

We use our accounting procedure to show that many of the patterns in daily deaths observed in countries and states of the United States cannot be accounted for by our simple BSIR model without large dispersion in wedges to the relationship between disease prevalence and disease transmission across locations at a point in time and large variation in these wedges in a given location over time. Moreover, the wedges uncovered by our accounting procedure, interpreted as shifts in the transmission rate of COVID-19 holding disease incidence fixed, are, after the initial phase of the pandemic, very large in magnitude when interpreted in terms of the model-implied forecast for the long-run attack rate. The intuition for the result that large wedges are needed to account for the data is simple: because human behavior adjusts

endogenously in a BSIR model to disease incidence, the magnitude of wedge needed to match a given change in the growth rate of daily deaths is much larger than would be the case if there were no endogenous offsetting response of behavior.

We see this result as an empirical puzzle: what might account for the very large wedges required to have a BSIR model match the data? A great deal of research in economics on the early phase of the COVID epidemic points to spontaneous changes in human behavior as central to the changing dynamics of the epidemic relative to the direct impact of government policies.[6] We confirm this finding regarding the early phase of the pandemic with our accounting procedure: our BSIR model without wedges accounts for the data on deaths in the early phase of the epidemic quite well. In contrast, after the initial phase of the epidemic, we find that the clear negative relationships between disease incidence and the growth rate of daily deaths observed in the data for many locations disappear — the wedges play a much more important role in shaping outcomes.

Given the simplicity of our model, this cross-section and time-series variation in "wedges" impacting COVID transmission could represent a wide array of factors that might account for the discrepancies between the predictions of our simple BSIR model and the data. Our hope (and our aim in future research) is to promote further empirical study of what factors might account for these "wedges", both in terms of their sign and their magnitudes. Are government-mandated lockdowns or mask mandates key? Is seasonality in virus transmission key? Are these wedges the result of model mis-specification? Is adding behavioral or network heterogeneity to the model key for accounting for these "wedges"? Future research should aim at shedding light on the answers to these questions.

In this paper, we examine the implications of one potential factor that might account for the wedges in our model — *pandemic fatigue*. By this term, we mean a substantial reduction in the responsiveness of behavior to disease prevalence. We offer

[6]See, for example, Cronin and Evans (2020), Goolsbee and Syverson (2020), Fetzer et al. (2020), Gupta et al. (2020), Arnon et al. (2020), Correia et al. (2020), Sheridan et al. (2020), and Chapter 2 of the October 2020 edition of the IMF's *World Economic Outlook*.

this measurement of the wedges implied by an alternative specification of our model as an exercise to illustrate what might be learned from our procedure. We show that if we recalculate the wedges implied by our model under the assumption that the reduced-form response of disease transmission to disease prevalence falls substantially in each location after the initial phase of the pandemic, then the dispersion of wedges required to have this alternative model match the data exactly also falls substantially.

The flip side of this finding is that the median wedge required to have this alternative model match the data is now large and negative. In other words, for our model with pandemic fatigue to match the data on deaths, one must assume that disease transmission holding behavior fixed fell substantially across the large number of countries and states of the United States that we study. Absent future shifts in the wedges, such a model would imply that the forecast for the long-run attack rate for COVID-19 has fallen substantially in many locations across the world. If this alteration of the model-implied forecast for the long-run attack rate of COVID-19 were correct, then the model's predictions for the fraction of the world's population that needs to be vaccinated to control the further spread of the disease might be much lower than is currently estimated. We see this finding as illustrating what might be learned with further research into the source of the wedges uncovered by our accounting procedure.

The remainder of our paper is organized as follows.

In section 2, we present the epidemiological model that we use for measurement of the dynamics of the COVID-19 pandemic in the 69 countries and 34 states of the United States that we study. This model is a standard SIR model in which the transmission rate is allowed to vary over time in an arbitrary manner. We use this model to interpret data on cumulative deaths, daily deaths, and the growth rate of daily deaths in a region at a point in time in terms of the fractions of the population in a region that remain susceptible, are actively infected, or are recovered or dead at that point in time as well as the effective reproduction number and transmission rate of the disease in that region at that point in time. This measurement procedure

establishes the targets in the data we seek to account for with our BSIR model.

In section 3, we present the method we use to estimate the level of cumulative deaths, daily deaths, and the growth rate of daily deaths in each region over time. This estimation procedure is designed to overcome the problem that data on COVID deaths is reported with considerable noise and that we require consistent estimates of cumulative deaths, daily deaths, and the growth rate of daily deaths from these noisy data. In this estimation procedure, we fit a mixture of modified log-logistic densities to the reported data on daily deaths in each location. We use this mixture of modified log-logistic densities to compute the growth rates of daily deaths, the level of daily deaths, and the cumulative level of daily deaths for each location. We choose to use a mixture of modified log-logistic densities to fit the data on daily deaths in part because this distribution has the property that the implied growth rate of daily deaths remains bounded as is required by the structure of an SIR model with a time-varying transmission rate.

We then use this estimation procedure and our measurement model to establish four key facts about the evolution of the COVID pandemic seen in the countries and states of the United States that we study. We do this in subsection 3.1. These four key facts are as follows. First, the growth rate of daily deaths fell rapidly everywhere from high, and initially highly dispersed levels, down to much lower levels in the first 30 days after each region reached 25 cumulative deaths. Second, after this initial period, the growth rate of daily deaths in each location has hovered within a relatively narrow band that includes zero. Third, the dispersion in the cross section across regions of the growth rate of daily deaths fell rapidly in the first 10 days or so of the epidemic and has remained at a relatively low level since then. And finally, fourth, when interpreted through the lens of a variety of epidemiological models, these first three facts regarding the growth rates of daily deaths also apply to the estimated effective reproduction numbers and disease transmission rates implied by the deaths data and our simple SIR model for measurement.

In section 4, we then introduce our behavioral SIR model and analyze the dynamics

implied by this model. In our BSIR model, we assume two simple relationships that govern the transmission rate of disease. One relationship is that between the current severity of the epidemic measured as the level of daily deaths (or equivalently the current fraction of the population actively infected) and the level of human activity. The other relationship is between the level of human activity and the transmission rate of the epidemic. We introduce "wedges" in the model as shifts in these two relationships. We analyze the reduced form of the model that combines these two relationships into a single reduced-form relationship between the current severity of the epidemic measured as the level of daily deaths and the transmission rate of the disease and consider a single composite wedge to this relationship.[7]

In section 5, we first review some of the key implications of our simple BSIR model that can be derived analytically when the "wedges" are all equal to zero. We highlight that this model can account for a rapid decline in the growth rate of daily deaths (and the associated effective reproduction number) in the initial phase of a pandemic from high and highly dispersed levels. This model-implied decline in the growth rate of daily deaths occurs well before a substantial portion of the population becomes immune to the disease, and thus the BSIR model implies a much lower peak of infections and daily deaths than is predicted by an equivalent SIR model that does not incorporate an impact of changes in human behavior on disease transmission. We analyze the phase diagram that characterizes the model-implied dynamics of the disease to discuss how, absent "wedges", this BSIR model implies that, after an initial peak, the growth rate of daily deaths should settle in to a level a little below zero with a corresponding effective reproduction number a little below one.[8]

We then use this phase diagram to discuss some discrepancies between the model's implications for the growth rate of daily deaths and the data observed in many

[7]Note that we do not directly model optimizing behavior by rational agents as is done in many economics papers cited in footnote 2 above. While our modeling approach suffers in this respect that it is not fully structural, it does allow us to sidestep the issues connected with modeling agents' information and expectations in shaping individual behavior. We leave the measurement of wedges in such a model to future work.

[8]We thank Ben Moll for introducing us to this phase diagram.

countries and states of the United States if the model does not allow for "wedges". First, we use the phase diagram that characterizes the model implied dynamics of the disease to argue that by itself, the BSIR model cannot account for the rapid decline in daily deaths to very low levels found in many locations such as Italy, Japan, Spain, New York, Sweden, etc. that were first hit hard with the disease. To account for the patterns over time in these data on daily deaths, we must introduce "wedges" that reduce the disease transmission rate in the model below the levels that would be implied if the behavioral relationships assumed in the model were stable over time. Moreover, our BSIR model cannot account for steady positive growth rates of daily deaths observed in many locations, such as in Arizona, or in second or third waves occurring in many places after the summer of 2020. To account for the patterns over time in these data on daily deaths, we must introduce "wedges" that increase the disease transmission rate in the model above the levels that would be implied if the behavioral relationships assumed in the model were stable over time. We also show how these wedges correspond to shifts over time in the model-implied forecast for the model-implied long-run attack rate.

In section 6, we then turn to our accounting procedure. To set the parameters of the BSIR model that we use for our accounting exercise, we estimate the two behavioral parameters of our BSIR model for a large number of countries and states of the United States using data from the earliest phase of the pandemic. One of these parameters corresponds to the basic reproduction number of COVID-19 in each location at the start of the epidemic. The other corresponds to the semi-elasticity of the transmission rate with respect to the level of daily deaths implied by the reduced-form of the equilibrium of the model. We describe this estimation procedure in subsection 6.1.

We then use this estimated version of our model in each location to back out the "wedges" required to account for the data on deaths found using our measurement procedure above. in subsection 6.2, we summarize our main findings for the whole group of countries and US states that we study. In subsection 6.3, we look at results for Italy, Arizona, and Japan specifically to highlight variation in wedges over time

within a given region.

In subsection 6.4 we compute the wedges implied by an alternative specification of our model with pandemic fatigue.

In section 7, we conclude.

2 Interpreting deaths data with an SIR model

The SIR epidemiological model we use to interpret the data on deaths from COVID-19 is as follows. For notational simplicity, we suppress reference to the region i.

The population is set to N. At each moment of time, the population is divided into four categories (states) that sum to the total population. These states are susceptible S, infected I, resistant R , and dead D. Agents that are susceptible are at risk of getting the disease. Agents that are infected are contagious and may pass it on to others through some form of interaction with susceptible agents. Agents that are resistant are not at risk of getting the disease, either because they have immunity built up from a vaccine or from previous experience with this or similar diseases. Likewise, those who have died from the disease are no longer at risk of getting the disease. We normalize the total population $N = 1$, so all results regarding S, I, R and D should be interpreted as fractions of the relevant population.

We use $\mathcal{R}(t)$ to denote the *effective reproduction number* of the disease at date t. This effective reproduction number is the ratio of the rate at which infected agents infect susceptible agents to the recovery rate of infected agents from the disease at date t.

The equations of the model can be stated in terms of the effective reproduction number as[9]

$$dS(t)/dt = -\mathcal{R}(t)\gamma I(t) \tag{1}$$

[9] A discrete-time version of a SIR model is discussed in Fernandez-Villaverde and Jones (2020).

$$dI(t)/dt = (\mathcal{R}(t) - 1)\,\gamma I(t) \tag{2}$$

$$dR(t)/dt = (1 - \nu)\gamma I(t) \tag{3}$$

$$dD(t)/dt = \nu\gamma I(t) \tag{4}$$

The parameter γ governs the rate at which agents who are infected stop being infectious and hence stop transmitting the disease. We refer to this parameter as the *recovery rate*. This parameter is considered a fixed parameter determined by the biology of the disease. We denote the fatality rate from the disease by ν.

The parameter $\beta(t)$ is the rate at which infected agents spread the virus to others that they encounter at date t. We refer to this parameter as the *transmission rate*. We define the ratio $\beta(t)/\gamma$ to be the *normalized transmission rate*. It is standard to refer to the value of the normalized transmission rate at the start of the pandemic before any mitigation measures and use of prophylactics are undertaken as the *basic reproduction number* of the disease. We denote this basic reproduction number by $\mathcal{R}_0 \equiv \beta(0)/\gamma$.

We assume that infected agents interact randomly with other agents in a uniform manner so that the effective reproduction number of the disease is given by the product of the normalized transmission rate and the fraction of agents who remain susceptible to the disease:

$$\mathcal{R}(t) = \frac{\beta(t)}{\gamma}\frac{S(t)}{N - D(t)}. \tag{5}$$

We see from equation (5) that the effective reproduction number can fall either due to changes in the normalized transmission rate or changes in the fraction of the population remaining susceptible to the disease.

To invert this model to interpret data on deaths note that from (4), we have

$$I(t) = \frac{1}{\nu\gamma} dD(t)/dt. \tag{6}$$

Using (3) and (4) together and the assumption that $R(0) = D(0) = 0$, we have that

$$R(t) = \frac{1-\nu}{\nu} D(t). \tag{7}$$

Using that the states must sum to one, we have

$$S(t) = 1 - \frac{1}{\nu} D(t) - \frac{1}{\nu\gamma} \frac{dD(t)}{dt}. \tag{8}$$

To obtain the effective reproduction number implied by deaths data, note that from equations (2), (6), and the time derivative of this second equation, we have

$$\mathcal{R}(t) = 1 + \frac{1}{\gamma} \frac{\frac{d^2 D(t)}{dt^2}}{\frac{dD(t)}{dt}} \tag{9}$$

where the last term in this equation can be interpreted as the time derivative of the logarithm of daily deaths.

Note that this equation (9) implies that there is a linear relationship between the growth rate of daily deaths (here measured as the time derivative to the logarithm of daily deaths) and the model-implied effective reproduction number. This reproduction number is equal to one when the growth rate of daily deaths is equal to zero. The slope of this relationship is given by $1/\gamma$ corresponding to the number of days on average that an infected individual remains infectious. To compute estimates of the effective reproductive number that are consistent with our estimated paths for the growth rate of daily deaths, we set $\gamma = 0.2$. This value implies that if the growth rate of daily deaths is 30 percent per day initially, the basic reproduction number, the value of the effective number at date $t = 0$, is 2.5. These values are in line with

CDC estimates.[10].

Given these equations, one can obtain an estimate of the normalized transmission rate of the disease from equations (5) and (8). Thus, one can use this estimate to determine the extent to which the model-implied effective reproduction number has changed due to changes in the transmission rate versus a reduction in the fraction of the population remaining susceptible to the disease. For this exercise we also need to set a value for the infection fatality rate ν. We set $\nu = 0.005$. This level is equal to that recommended to modelers by the CDC for those aged 50-69. [11] The level of the fatality rate, if held constant, does not impact our estimates of the evolution of the effective reproduction number in equation 9. In the appendix we discuss the impact of changes in the fatality rate over time on the estimation of $\mathcal{R}(t)$.

3 Death Data and Estimation

We now discuss how we recover consistent estimates of the growth rate of daily deaths, the level of daily deaths, and the level of cumulative deaths from noisy reported data on daily deaths.

The data sources for daily deaths are New York Times for U.S. states and Johns Hopkins University for other countries. For each location i, our estimation period begins at the location specific date when cumulative deaths reached 25 in that location and ends on November 12, 2020. The 69 countries and countries that we study are Afghanistan, Algeria, Argentina, Armenia, Australia, Bangladesh, Belarus, Bel-

[10]See the range of estimates of the basic reproduction number for COVID-19 recommended by the CDC for use in models available in Table 1 at https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html. The range of estimates of $\mathcal{R}_0$ that the CDC recommends (as of September 2020) is from 2 to 4 with a most likely estimate of 2.5

[11]See Table 1 at https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html. See also O'Driscoll et al. (2020). We make the simplifying assumption that the infection fatality rate from COVID has remained constant over time. See https://www.nature.com/articles/d41586-020-03132-4 for a recent discussion of evidence that this fatality rate has declined somewhat over time.

gium, Bolivia, Bosnia, Brazil, Bulgaria, Canada, Chile, China, Colombia, Costa Rica, Czechia, Denmark, Dominican Republic, Ecuador, Egypt, El Salvador, Ethiopia, France, Germany, Guatemala, Honduras, Hungary, India, Indonesia, Iran, Ireland, Israel, Italy, Japan, Jordan, Kazakhstan, Kenya, Krygyzstan, Mexico, Moldova, Morocco, Nepal, Netherlands, Nigeria, North Macedonia, Oman, Pakistan, Panama, Paraguay, Peru, Philippines, Poland, Portugal, Romania, Russia, Saudi Arabia, South Africa, Spain, Sudan, Sweden, Switzerland, Tunisia, Turkey, the United Kingdom, and Ukraine. The 34 US states that we study are Alabama, Arizona, Arkansas, California, Colorado, Connecticut, Florida, Georgia, Illinois, Indiana, Iowa, Kansas, Kentucky, Louisiana, Maryland, Massachusetts, Michigan, Minnesota, Mississippi, Minnesota, Missouri, Nevada, New Jersey, New Mexico, New York, North Carolina, Ohio, Oklahoma, Pennsylvania, Rhode Island, South Carolina, Tennessee, Texas, Virginia, Washington and Wisconsin. The rest of the U.S. is counted as another region.

One problem that we face is that the data on daily deaths due to COVID-19 are noisy. In many regions, there are both substantial day-of-the-week effects and occasional large spikes in reported deaths due to, among other things, changes in the criteria used to classify deaths as being due to COVID-19. To estimate the *trend* growth of daily deaths from these noisy data, we assume that the trend path of daily deaths in each location is given by a mixture of modified log-logistic density functions which we estimate using a Bayesian procedure.

The modified log-logistic distribution that we use is given as follows. The jth distribution in the mixture of distributions for region i has cumulative distribution function $F_{i,j}(t) = 0$ for $t = 0$ and, for $t \geq 0$,

$$F_{i,j}(t) = \frac{H_{i,j}(t)}{1 + H_{i,j}(t)}$$

and density $f_{i,j}(t) = 0$ for $t < 0$ and, for $t \geq 0$,

$$f_{i,j}(t) = \frac{H'_{i,j}(t)}{(1 + H_{i,j}(t))^2} \tag{10}$$

with

$$H_{i,j}(t) = \left(\frac{t + q_{i,j}}{a_{i,j}}\right)^{b_{i,j}} - \left(\frac{q_{i,j}}{a_{i,j}}\right)^{b_{i,j}}$$

with parameters $a_{i,j}, b_{i,j}, q_{i,j} > 0$.

With this specification of the distribution of deaths, we have that the growth rate of deaths corresponding to the jth distribution in the mixture of distributions for region i is equal to zero for $t < 0$ and, for $t \geq 0$,

$$g_{i,j}(t) = \frac{f'_{i,j}(t)}{f_{i,j}(t)} = \frac{b_{i,j} - 1}{t + q_{i,j}} - 2\frac{\frac{b_{i,j}}{a_{i,j}}\left(\frac{t+q_{i,j}}{a_{i,j}}\right)^{b_{i,j}-1}}{1 + \left(\frac{t+q_{i,j}}{a_{i,j}}\right)^{b_{i,j}} - \left(\frac{q_{i,j}}{a_{i,j}}\right)^{b_{i,j}}} \tag{11}$$

We see here one feature of the modified log-logistic distribution is that, if $q_{i,j} > 0$, then the growth rate of the density remains bounded for all t as is required by any distribution of deaths produced as an outcome of an SIR model with bounded transmission rates.

We assume that observed daily deaths are the sum of a mixture of modified log-logistic density functions and a residual whose magnitude is regime-specific. The regime-switching residuals allow us to effectively deal with erratic noises in the data. The Bayesian procedure allows us to construct posterior probability bands around the estimate. It also allows us to derive smooth estimates of the first and higher derivatives of daily deaths that we need to recover estimates of the effective reproduction numbers and transmission rates of the disease from several structural SIR models.

Our estimation methodology begins by scaling the cumulative death data in location i so that $D_i(t)/(1 + d_i)D_i(T)$ lies between zero and one, where $D_i(T)$ is the

cumulative number of deaths in location i at the end of the estimation period and $d_i > 0$ is a scale parameter to be estimated as described in the appendix. Let $\Delta D_{i,t}^{\text{Data}}$ be the daily measured object for $dD_i(t)/dt$ and denote

$$\Delta \widetilde{D}_{i,t}^{\text{Data}} = \frac{\Delta D_{i,t}^{\text{Data}}}{(1+d_i) D_{i,T}^{\text{Data}}}.$$

Given the daily death data and the value of d_i, we run a non-linear regression with a mixture of modified log-logistic density functions and regime-switching heteroskedastic errors:

$$\Delta \widetilde{D}_{i,t}^{\text{Data}} = \sum_{j=1}^{\mathcal{J}} w_{i,j}\, f_{i,j}\left(t - t_{0,i} - c_{i,j}\right) + \sigma_{i,k_t} \varepsilon_{i,t},$$

where $\varepsilon_{i,t}$ is an iid standard normal random residual, weights $w_{i,j}$ are non-negative and sum to one across j, and $f_{i,j}$ is defined as in equation 10 with parameters $a_{i,j}, b_{i,j}, q_{i,j}$ when $t - t_{0,i} - c_{i,j} \geq 0$ and $t_{0,i}$ is the time when the cumulative death toll reached 25 for location i. The density $f_{i,j}$ is equal to zero when $t - t_{0,i} - c_{i,j} < 0$. The parameters $c_{i,j}$ control the date at which each density in the mixture starts relative to the date $t_{0,i}$ when cumulative deaths first reach 25 in region i.

The switching state $k_t \in \{1, \ldots, \mathcal{K}\}$ follows a Markov-switching process and can accommodate both a large surge in daily deaths and a low death volatility typically associated with a low number of deaths.[12]

Given the estimates of our death model parameters, we construct the estimates of death growth rates by taking the time derivative of our estimated mixture of modified log-logistic densities

$$\hat{g}_i(t) = \sum_{j=1}^{\mathcal{J}} w_{i,j}\, g_{i,j}(t) \frac{f_{i,j}\left(t - t_{0,i} - c_{i,j}\right)}{\sum_{j=1}^{\mathcal{J}} w_{i,j}\, f_{i,j}\left(t - t_{0,i} - c_{i,j}\right)}$$

where $g_{i,j}(t)$ is given in equation 11.

[12]For the selection of $\mathcal{K}$ and estimation details, see Appendix **??** and Atkeson et al. (2020).

Figure 1 shows the fit of our estimation to the noisy daily death data for three selected locations: Italy, Arizona, and Japan. We use these three regions as examples throughout the paper. The x-axis of these panels also show the starting date at which each of these regions first reached 25 cumulative deaths. The left panel shows the fit of the model to the data on daily deaths. The right panel shows the fit of the model to the data on cumulative deaths.

3.1 Findings

Our estimation results yield the following four stylized facts about the COVID-19 epidemic.

Fact 1. *The growth rate of daily deaths from COVID-19 fell rapidly everywhere within the first 30 days after each region reached 25 cumulative deaths.*

Fact 1 is shown in Figure 2. The solid black line shows that the median estimated growth rate of daily deaths fell rapidly from an initial level of about 12 percent to zero within the first 30 days of the estimation period. Notice that the 66% and 95% posterior probability intervals computed across all locations that we consider, the dashed lines in the figure, follow a similar pattern.[13]

Fact 2. *After this first period of rapid decline, the growth rate of daily deaths in all regions has hovered in an interval including zero.*

As Figure 2 shows, after the initial 30-40 day period, the median growth rate of daily deaths has been nearly flat and slightly below zero. The 66% and 95% posterior probability intervals have also been low relative to their initial range.

Fact 3. *The cross-regional standard deviation of growth rates of deaths fell rapidly in the first 20 days of the epidemic and has, subsequently, remained low relative to its initial level.*

[13]The posterior probability intervals in Figure 2 include both location uncertainty and sampling uncertainty. However, most of the cross-sectional dispersion in growth rates in the figure is driven by location uncertainty as sampling uncertainty within a location is small.

Fact 3 can also be seen in Figure 2. The shrinking of the two-thirds and 0.95% posterior probability intervals shows that the dispersion in death growth rates across locations fell sharply within the initial 20 days of the epidemic. Overall levels of dispersion remain small relative to their initial values since that initial phase of the pandemic, despite large second waves observed in the Fall in many locations.

Equation (5) illustrates that the effective reproduction number can fall due to both declines in the normalized transmission rate and declines in the fraction of the population remaining susceptible to the disease. We use our estimated paths for daily deaths and the equations of the SIR model to determine the relative contributions of each. We find that the rapid decline in daily death growth rates early on in the epidemic is primarily due to a rapid fall in the transmission rate of the disease as illustrated by the bottom panel of Figure 2. Disease transmission rates, like the effective reproduction numbers, fell early on in the epidemic from widely dispersed initial levels and and have since remained close to 1.

To explore the robustness of our results to model uncertainty, we consider three variants of the baseline SIR model: an SEIR model which extends the SIR model by assuming agents first become exposed to the disease before becoming infectious, an SIHR model which allows for a longer period between infection and death by adding a hospitalized state, and an SEIHR model which extends the SIR model by adding both the exposed state and the hospitalized state. Following a similar procedure as with the baseline SIR model, for each model extension we are able to express the effective reproduction number as a function of model parameters, as well as, daily deaths and its derivatives.

4 A Behavioral SIR model

Our behavioral SIR model is built on the SIR model in section 2 with dynamics given in equations 1 to 5. That model for measurement allows the transmission rate $\beta(t)$ to be an arbitrary function of time. The behavioral aspect of the BSIR model that

we study introduces a specific model of the evolution of the transmission rate $\beta(t)$ over time as an endogenous response to the prevalence of the disease. This model of the transmission rate is given by two relationships described as follows.

We assume that the transmission rate of the disease at time t in region i, denoted by $\beta_i(t)$, is a function of human activity in the region $Y_i(t)$ at that date

$$\beta_i(t) = \bar{\beta}_i Y_i(t)^{\alpha} \exp(\psi_{\beta,i}(t)) \tag{12}$$

The parameter $\bar{\beta}_i$ in equation 12 is a fixed coefficient that captures features of the population and environment of region i determined prior to the epidemic that might impact transmission. Factors considered in the literature include population density, modes of transportation (subway vs. car, etc.), household and demographic structure, cultural norms (bowing vs. shaking hands or kissing), temperature and humidity, etc.

The parameter α captures the elasticity of transmission with respect to activity. For simplicity, we assume that this parameter is common across regions. We set this parameters $\alpha = 2$ to capture the idea that the number of interactions between individuals goes up with the square of the activity level of all individuals.

The parameter $\psi_{\beta,i}(t)$ represents a potentially time-varying wedge shifting the region-specific relationship between activity and transmission. This wedge may represent the impact of policy and/or natural variation in the transmission of the virus over time.[14] When interpreting variation in $\psi_{\beta,i}(t)$ as representing the impact of policies, here we are thinking about policies such as mask-wearing, ventilation, physical distancing, redesign of workspaces, or other measures implemented after the start of the epidemic that reduce transmission given a fixed level of activity.

We normalize $\psi_{\beta,i}(0) = 0$ and we normalize the level of activity at the start of the

[14]Seasonality in virus transmission is one candidate for a natural force. See, for example https://www.sciencemag.org/news/2020/03/why-do-dozens-diseases-wax-and-wane-seasons-and-will-covid-19 for a discussion of seasonality in the transmission of COVID-19 and a variety of other viral diseases.

pandemic to $Y_i(0) = 1$. Given these normalizations, the parameter $\bar{\beta}_i$ determines the transmission rate of the virus in region i at the start of the epidemic, with the basic reproduction number of the virus in this region given by $\mathcal{R}_{0,i} = \bar{\beta}_i/\gamma$.

Next, we introduce the behavioral component of the model. We assume that individuals' decisions to engage in activity in region i at time t, $Y_i(t)$, are a declining function of the time derivative of cumulative deaths, $\dot{D}_i(t)$, which we refer to as the current level of *daily deaths*. We specify this function describing behavior as

$$Y_i(t) = \exp(-\kappa_i \dot{D}_i(t) + \psi_{y,i}(t)) \tag{13}$$

Here, $\kappa_i > 0$ represents the semi-elasticity of activity $Y_i(t)$ with respect to daily deaths. Note that we allow the semi-elasticity κ_i to vary by region. This semi-elasticity might vary depending both on individuals' opportunities to reduce activity by working from home, etc. and depending on their beliefs about the personal trade-offs involved in exposing themselves to virus transmission.

The variable $\psi_{y,i}(t)$ in equation 13 represents a time-varying shifter to the region-specific relationship between deaths and activity. We might interpret $\psi_{y,i}(t)$ as reflecting policies such as lockdowns or closures that would reduce activity below what agents might choose in a decentralized fashion. Note that if such policies are imposed conditional on the state of the disease, with restrictions on activity becoming more severe as the level of daily deaths rises (as in $\psi_{y,i}(t) = -\eta_i \dot{D}_i(t)$ with $\eta_i > 0$), then this dimension of policy heterogeneity across regions would also be a reason for why the semi-elasticity of activity with respect to daily deaths would vary across regions. Thus, we interpret differences in the semi-elasticity of activity with respect to daily deaths as due to either private behavior or systematic public policy.

Observe that by substituting equation 13 into equation 12, we get a reduced-form relationship between the current level of daily deaths and the transmission rate given by

$$\beta_i(t) = \bar{\beta}_i \exp(-\alpha \kappa_i \dot{D}_i(t) + \psi_i(t)) \tag{14}$$

where $\psi_i(t)$ is the composite wedge

$$\psi_i(t) \equiv \alpha\psi_{y,i}(t) + \psi_{\beta,i}(t) \tag{15}$$

In what follows, we focus on measuring this composite wedge $\psi_{i,t}$ without breaking it down into its components $\psi_{y,i}(t)$ and $\psi_{\beta,i}(t)$. We leave this further decomposition to future work.

Recall that since from equation 4, daily deaths $\dot{D}_i(t)$ are directly proportional to the fraction of the population that is currently actively infected $I_i(t)$, we can rewrite equation 14 as

$$\beta_i(t) = \bar{\beta}_i \exp(-\alpha\sigma_i I_i(t) + \psi_i(t)) \tag{16}$$

with $\sigma_i = \kappa_i\nu\gamma$.

To solve this model for a particular region i, we first specify initial conditions $S_i(0), I_i(0), R_i(0), D_i(0)$, parameters $\gamma, \nu, \bar{\beta}_i, \alpha, \sigma_i$, and time paths for the composite wedge $\psi_i(t)$. We then solve for the implied evolution of the disease using equations 1 - 5 the reduced form for transmission in equation 16. To solve for the path for activity $Y_i(t)$, one must separately specify paths for "wedges" $\psi_{y,i}(t)$ and $\psi_{\beta,i}(t)$.

5 BSIR model implied dynamics

We now study the epidemiological dynamics implied by our BSIR model. We are able to derive two analytical results by studying the phase diagram of the dynamics of the model with no wedges (so that $\psi_\beta(t) = \psi_y(t) = \psi(t) = 0$ for all t). We then demonstrate a third result with numerical examples. Finally we illustrate the role of time varying wedges, as modeled by changes over time in $\psi(t)$ on the model-implied dynamics of the epidemic, and demonstrate that the magnitude of the "wedges", as measured by the shift in transmission rates holding disease prevalence constant, are much larger than the corresponding equilibrium shifts in the effective reproduction number induced by these wedges.

5.1 Analytical results: single peak of and slow decline in daily deaths

The first of our two analytical results is that the BSIR model implies that the path of active infections $I(t)$ and daily deaths $\dot{D}(t)$ are both single-peaked. With no wedges, the model cannot generate multiple waves of infections and daily deaths. This is a property that our BSIR model shares with the standard SIR model.

The second analytical result is that, after the peak of infections and daily deaths, these outcomes cannot fall rapidly to a low level without substantial depletion of the pool of susceptible agents in the population $S(t)$. Thus, with no wedges, the model cannot account for the patterns of daily deaths seen in many of the regions that were initially hard hit followed by a steep decline in daily deaths to very low levels.

To establish these results, we develop the phase diagram representing the model dynamics implied by equations 1 to 5 and 16 as follows. Given that $N = S(t)+I(t)+R(t)+D(t)$ for all t and that the levels of $R(t)$ and $D(t)$ are directly proportional to each other when $R(0) = D(0) = 0$, we can summarize the state of the model and the associated model dynamics with the mapping between the state of the model as captured by the pair $(S(t), I(t))$ and the implied dynamics $(\dot{S}(t), \dot{I}(t))$. We do so in the phase diagram shown in Figure 3.

Figure 3 shows S/N on the x-axis and I/N on the y-axis. The arrows in the figure show the signs of $\dot{S}(t)$ and $\dot{I}(t)$ in each region of the state space $S/N \in (0, 1)$ and $I/N \geq 0$. From equation 1, we have that $\dot{S}(t) < 0$ for all values of (S, I) such that $S, I > 0$. Observe from equation 5 that the effective reproduction $\mathcal{R}(t) > 1$ when $\frac{\beta(t)}{\gamma}S(t) > 1$ and $\mathcal{R}(t) < 1$ when $\frac{\beta(t)}{\gamma}S(t) < 1$. Thus, from equation 2, we have that $\dot{I}(t) > 0$ when $\frac{\beta(t)}{\gamma}S(t) > 1$ and $\dot{I}(t) < 0$ when $\frac{\beta(t)}{\gamma}S(t) < 1$. From equation 16, we have that the transmission rate $\beta(t)$ is a strictly decreasing function of $I(t)$. Thus, equation 16 defines a locus of points $(\tilde{S}, \tilde{I})$ such that $\dot{I} = 0$. We show this locus of points as a black curve in Figure 3. We then have that $\dot{I} > 0$ when (S, I) lies below this locus of points and $\dot{I} < 0$ when (S, I) lies above this locus of points.

Note that this locus of points such that $\dot{I} = 0$ intersects the x-axis at a point $\bar{S}$ marked in the figure the level of which is determined by the basic reproduction number implied by our BSIR model of $\mathcal{R}_0 = \bar{\beta}/\gamma$. In fact, this point given by $\bar{S} = 1/\mathcal{R}_0$ also corresponds to the herd immunity threshold in a standard SIR model with a constant transmission rate $\bar{\beta}$. Thus, the long-run cumulative fraction of the population infected in our BSIR model is at least as large as the fraction corresponding to the herd immunity threshold in a standard SIR model. Hence, we interpret $1 - \bar{S}$ as our model's forecast for the long-run attack rate.

The red curve in Figure 3 shows the path of $(S(t), I(t))$ followed in the solution of the model. This pair of outcomes $(S(t), I(t))$ starts in the lower left corner of the figure with $S(0)$ very close to one and $I(0)$ positive but very close to zero. The fraction of the population infected rises rapidly initially and crosses over the black locus of points such that $\dot{I} = 0$ and then falls slowly, remaining above that locus until the fraction of agents remaining susceptible falls below the point $\bar{S}$ where the locus of points such that $\dot{I} = 0$ intersects the x-axis and $I(t)$ asymptotes to zero.

We obtain our two analytical results immediately from this phase diagram. First, the path of infections $I(t)$, and hence daily deaths $\dot{D}(t)$, is single-peaked. This follows from the observations that $I(t)$ rises until the red line crosses over the black locus of points such that $\dot{I} = 0$ and $I(t)$ falls after that point. Note that $I(t)$ cannot cross back over the black locus of points such that $\dot{I} = 0$ so it cannot start rising again.

Second, note that the fraction of the population infected $I(t)$, and the corresponding level of daily deaths $\dot{D}(t)$ from equation 4, cannot fall to a low level after reaching a first peak without a substantial decline the fraction $S(t)$ of the population susceptible to the disease. This result follows from the observation that $I(t)$ cannot cross back over the black locus of points such that $\dot{I} = 0$ from above. Thus, this locus of points defines a lower bound on the fraction on the population infected and the corresponding level of daily deaths that can be observed after a first peak of infections and deaths.

In Figure 4, we show a version of this phase diagram for our model with three

different specifications of the parameter σ governing the semi-elacticity of the transmission rate with respect to infections in equation 16. We see in this figure that the qualitative features of the solution of the model observed in Figure 3 hold even if this semi-elasticity varies.

5.2 Numerical Result: without wedges the BSIR model cannot generate relatively slow but steady growth in daily deaths.

We establish a third result with numerical examples. This third result is that the BSIR model cannot produce outcomes with a relatively slow but steady growth of daily deaths in a region (growth on the order of 2 to 5 percent per day over a month or so) unless the basic reproduction number in that region is relatively low — on the order of $\mathcal{R}_0 = 1.2$ or 1.3. Thus, the BSIR model cannot reproduce patterns of daily deaths seen in locations such as Arizona or Brazil unless COVID in such locations have quite low basic reproduction numbers relative to CDC estimates of this number.

In Figure 5, we show the paths of the basic reproduction number and daily deaths per million corresponding to the three solutions of our BSIR model shown in the phase diagram space in Figure 4. These three examples all set the basic reproduction number of the model to $\mathcal{R}_0 = 2.5$, in line with CDC preferred estimates of that number. We see in Figure 5 that while the paths of daily deaths per million implied by these three different specifications of our model differ substantially by the peak level of daily deaths, they all three imply similar paths for the effective reproduction number of the disease, starting at $\mathcal{R}_0 = 2.5$ and then falling rapidly below one. From equation 9, with our choice of $\gamma = 1/5$, these dynamics of the effective reproduction number correspond to a drop in the growth rate of daily deaths from 30 percent per day to a growth rate below zero over a period of about one month to six weeks.

Figure 5 illustrates our third result. Specifically, if our BSIR model has a basic reproduction number of 2.5, it cannot produce outcomes with a steady growth of daily

deaths in a region on the order of 2 to 5 percent per day (an effective reproduction number on the order of $\mathcal{R}(t) = 1.1$ to 1.25) over an extended period of time. Instead, as we see in the figure, the growth rate of daily deaths and the associated effective reproduction number falls rapidly over time from a high level to a level below zero.

In Figure 6 we show alternative simulations of the model with the basic reproduction number $\mathcal{R}_0 = 1.25$ rather than 2.5. In this case, we see that the model can produce a period of roughly 180 days of growth in daily deaths starting from a growth rate of daily deaths of 5 percent per day and falling to zero over the six month time period. But in this case, we assume that the basic reproduction number of the disease is much smaller than the preferred estimates of this number for COVID-19 in the literature.

5.3 BSIR model dynamics with wedges

We now consider the impact of variation over time in the wedge $\psi(t)$ in equation 16 linking the level of daily deaths and the transmission rate on the equilibrium dynamics implied by our model. Changes over time in $\psi(t)$ change our phase diagram shown in Figure 3 as follows. Changes in the wedge $\psi(t)$ generate shifts in the locus of points $(\tilde{S}, \tilde{I})$ such that $\dot{I} = 0$ shown in black in that figure to the left and right. In particular, the point $\bar{S}$ shifts over time with

$$\bar{S}(t) = \frac{\gamma}{\bar{\beta}} \exp(-\psi(t)) \tag{17}$$

Shifts over time in the wedge $\psi(t)$, our BSIR model can produce the a wide range of model outcomes. To illustrate that point, we consider the solution of our model with $\psi(t)$ set to follow a cosine wave over time with a frequency of one year. We illustrate the dynamics implied by the model over the course of three years in Figures 7 and 8. We set $\psi(t) = 0.35(\cos(2\pi t/365) - 1)$. The two black curves in Figure 7 represent the highest and lowest levels of the locus of points for which $\dot{I} = 0$. The red curve represents the solution of the model for $(S(t), I(t))$ over the course of three years.

The basic reproduction number in this simulation of the model oscillates between $\mathcal{R}_0 = 2.5$ and 1.25 (note that $\exp(-0.7) = 0.5$). We see in Figure 8 that the effective reproduction number in this simulation falls from its initially high value of 2.5 below one and then oscillates within a relatively narrow band around one after that, with daily deaths following a sequence of waves of slowly diminishing magnitude.

Note that at time $t = 0$, we have that $\bar{S}(0)$ is given by the inverse of the basic reproduction number of the disease $1/\mathcal{R}_0$. Thus, we interpret

$$\log\left(\bar{S}(t)\right) - \log\left(\bar{S}(0)\right) = \log\left(\mathcal{R}_0\right) - \log\left(\frac{\bar{\beta}}{\gamma}\exp(-\psi(t))\right) = \psi(t) \tag{18}$$

as a measure of the impact of the wedge ψ_t on the transmission rate of the virus holding fixed behavior. Note that the equilibrium response of virus transmission to a shift in the wedge $\psi(t)$ is given by $\log\left(\bar{\beta}\right) - \log\left(\beta(t)\right)$.

In the numerical example shown in Figures 7 and 8, the time variation in the wedge $\psi(t)$ results in a reduction in the transmission rate in half from peak to trough holding activity fixed. These fluctuation in the wedge $\psi(t)$ result in a shift in the fraction of the population left susceptible in the long run from $\bar{S}(t)$ from 0.4 when transmission is at its peak and 0.8 when transmission is at its trough. In terms of implied long-run outcomes, this difference in transmission rates is very large. If the wedge $\psi(t)$ were to remain at its peak level forever, then, in the long run, at least 60% of the population would contract the disease and, with a fatality rate of $\nu = 0.005$, then 0.3% of the population would be predicted to die from the disease, corresponding to roughly 990,000 deaths in the United States. In contrast, if the wedge $\psi(t)$ were to remain at its lowest level forever, then, in the long run, a bit more than 20% of the population would contract the disease and roughly 330,000 deaths would be expected in the United States. These calculations illustrate the point that the magnitude of the variation in transmission rates holding behavior fixed induced by our time varying wedge are much larger than the induced equilibrium variation in the effective reproduction number. These calculations also illustrate how the magnitude of the wedges in our model can be interpreted in terms of changes in the model's

forecast for the long-run attack rate of the disease.

6 Accounting for COVID dynamics with a BSIR model

We now present our procedure for using the BSIR model to account for the dynamics of the COVID epidemic for a large number of countries and states of the United States. After presenting our procedure, we present detailed results for three locations: Italy, Arizona, and Japan, to illustrate the performance of the model and the role of behavior and the wedge $\psi_i(t)$ in shaping the dynamics of the epidemic. We then summarize our findings for all regions considered.

The procedure we use to account for the dynamics of the COVID epidemic with our BSIR model has two steps. In the first step, we use data from the initial phase of the epidemic in each region i to estimate the model coefficients $\bar{\beta}_i$ and σ_i. We then back out the values of the "wedges" $\psi_i(t)$ needed so that our BSIR model matches exactly the data on COVID deaths in that region estimated using the procedure in section 3.

Specifically, once we have estimated parameters $\bar{\beta}_i$ and σ_i for each region, we then construct the wedges $\psi_i(t)$ as follows. Recall from section 2 that, given estimates of time paths for cumulative deaths $D_i(t)$, daily deaths $dD_i(t)/dt$, and the growth rate of daily deaths $(d^2D_i(t)/dt^2)/(dD_i(t)/dt)$ starting from time $t_{0,i}$ obtained as in section 3, we can construct SIR model-implied time paths for $S_i(t), I_i(t), R_i(t)$, and for the effective reproduction number $\mathcal{R}_i(t)$ from equations 6, 7, 8 and 9. The corresponding path for the transmission rate $\beta_i(t)$ is then obtained from from equations (5) and (8). Then, given this estimate of the actual time path of $\beta_i(t)$, we use equation 16 to recover the wedge $\psi_i(t)$ as

$$\psi_i(t) = \log\left(\frac{\beta_i(t)}{\bar{\beta}_i}\right) + \alpha\kappa_i \dot{D}_i^{Data}(t) \tag{19}$$

6.1 Estimation of the BSIR model

In each region i, we estimate σ_i and $\bar{\beta}_i$ via simulated methods of moments using the data on daily deaths in the region during the initial phase of the epidemic, period $t_{0,i}$ to $t_{1,i}$. Date $t_{0,i}$ is the date when cumulative deaths first reached 25 in the region. Absent any wedges, transmission rates in the BSIR model are a function of daily deaths as shown in Equation (14) above. We choose the end of the initial phase, $t_{1,i}$, in a region as the point where this functional relationship breaks down in the data. That is we compare $\dot{D}_i^{data}(t)$ and $\log \beta_i^{data}(t)$ which were estimated in Section 3 and choose $t_{1,i}$ to be the largest value of t such that for all $\dot{D}_i^{data}(t)$ with $t_{0,i} \leq t \leq t_{1,i}$ there is exactly one value of $\log \beta_i^{data}(t)$.

After setting the initial phase, for each region i, we choose σ_i to minimize the distance between daily deaths in the data as given by our Bayesian estimation procedure and daily deaths generated by the BSIR model with no wedges, or

$$\sum_{t=t_{0,i}}^{t_{1,i}} \left[\dot{D}_i^{data}(t) - \dot{D}_i(t)\right]^2$$

To solve the BSIR model, we set the initial fractions of susceptible, infected, recovered, and dead at date $t_{0,i}$ to their values based on estimated daily deaths in the data and the SIR model as described in Section 2. For each value of σ_i, $\bar{\beta}_i$ is set such that the transmission rate at date $t_{0,i}$ in the BSIR model corresponds to the rate implied by the SIR model and the death data. This is achieved by setting

$$\bar{\beta}_i = \beta_i^{data}(t_{0,i}) \exp(\alpha \kappa_i \dot{D}_i^{data}(t))$$

where $\kappa_i = \sigma_i / \nu\gamma$.

6.2 Summary of Findings

We now examine the fit of our BSIR model with and without wedges to the data on deaths for each location that we study. We summarize two main empirical findings in this subsection. In the next subsection, we look in greater detail at results for three specific regions. We then show the wedges implied by an alternative specification of our model with pandemic fatigue and briefly discuss an interpretation of these results.

The two main empirical findings that we summarize here are as follows.

First, the estimated BSIR model without wedges fits fairly well the main stylized facts regarding the distribution of growth rates of daily deaths across regions and over time shown in the data in Figure 2 and reviewed in section 3.1, both within the estimation period and outside the estimation period. We show this result in the top panel of Figure 9, which shows the distribution of growth rates of daily deaths for all the regions that we consider implied by our estimated BSIR model without wedges both within the estimation period and afterwards. Compare these predictions of our BSIR model without wedges shown in the top panel of 9 to the data on the distribution of growth rates of daily deaths shown in Figure 2. In both figures, we see that the growth rates of daily deaths was quite high and highly dispersed across regions in the initial phase of the pandemic. In both figures we see that the growth rates of daily deaths fell toward zero over the course of a month or less and that the dispersion of growth rates of daily deaths across regions has remained in a relatively narrow range that includes zero for the months that have followed this initial phase of the pandemic.

On closer inspection of these two figures, however, we see an important discrepancy between the predictions of our BSIR model without wedges for the growth rates of daily deaths from COVID and the data. As shown in Figure 9, the dispersion in the growth rates of daily deaths predicted by the model after the initial month or two of the pandemic is substantially smaller than the dispersion in the growth rates of

these deaths observed in the data in Figure 2. In the BSIR model without wedges, there are no realizations of growth rates substantially above or below zero, while in the data there are.

This observation leads us to our second finding regarding the magnitude of the wedges required to allow our BSIR model with wedges to match the data on daily deaths exactly. These wedges must be very large in magnitude. We show the distribution of those wedges $\psi_i(t)$ in the bottom panel of Figure 9. We see in that figure that the dispersion in wedges required to have the BSIR model fit the data grows substantially over time. To get a sense of the quantitative implications of these wedges, recall from equation 14, a wedge of $\psi_i(t)$ scales the transmission rate $\beta(t)$ holding fixed the level of daily deaths by $\exp(\psi_i(t))$. Thus, if $\psi_i(t) = 1$, the transmission rate holding fixed the level of daily deaths is shifted by a factor of more than 2.7. Likewise, from equation 18, this value of $\psi_i(t)$, if it were to persist forever, would shift the fraction of the population remaining susceptible at the end of the epidemic $\bar{S}_i(t)$ by a large amount. For example, if a region were to start with a basic reproduction number of $\mathcal{R}_0 = 2$ corresponding to a herd immunity threshold of $\bar{S}_i(0) = 0.5$, then a wedge of $\psi_i(t) = 1$ would shift that herd immunity threshold to $\bar{S}_i(t) = 0.18$ corresponding to a basic reproduction number over $\mathcal{R}_0 = 5$. The corresponding shift in the model's forecast in the long-run attack rate would be from 50% to 92%.

These figures highlight our two main empirical results: that a simple BSIR model matches the main features of the data on the growth rates of daily deaths from COVID from a large number of countries and states of the United States, but such a model requires large shifts in the model-implied transmission rate of COVID holding disease incidence fixed to match the data on daily deaths exactly. We explore these findings in greater detail for three specific regions next.

6.3 Results for Italy, Arizona, and Japan

We now look in detail at the results from our accounting procedure for three locations: Italy, Arizona, and Japan. These three regions experienced quite different patterns of daily deaths over time.

We see from the estimation results shown in Figure 2 that Italy experienced a large first peak of daily deaths and then saw a rapid decline in those deaths (relative to what would be predicted from a BSIR model without wedges) to very low levels in the summer. We see in the same figure that Arizona experienced relatively slow and steady growth in daily deaths in the first few months of the pandemic, reaching a first peak level of daily deaths only in midsummer. Finally, we see that Japan experienced a first peak in daily deaths relatively early on, but this peak level was much smaller than experienced in either Italy or Arizona. Japan then experienced a rapid decline in daily deaths to a very low level during the summer before experiencing a second wave of daily deaths into the early fall. In this subsection, we examine how our BSIR model with wedges accounts for these different patterns of the evolution of the epidemic.

We begin with Italy. In Figure 10, we show our estimation results for the BSIR model and a comparison of the predictions of our estimated model for the growth rate of daily deaths, the level of daily deaths, and cumulative deaths, with (in red) and without (in blue) wedges.

The two panels in the upper left portion of Figure 10 shows our estimation results for the equation 14. The blue line in these panels shows the predicted value of the log of the transmission rate given from equation 14 using the estimated parameters $\bar{\beta}_i$ and κ_i and data on daily deaths $\dot{D}_i^{data}(t)$ plotted against the level of daily deaths. The red curves in these panels shows the data on daily deaths and log transmission rates for each day $\dot{D}_i^{data}(t), \log\left(\beta_i^{data}(t)\right)$, with these variables estimated as in section 3. By comparing the blue and red curves in these panels, one can see the extent to which wedges $\psi_i(t)$ in equation 14 are needed to account for the variation in the

logarithm of the transmission rate relative to variation in the level of daily deaths. The panel on the left of the two panels in the upper left portion of Figure 10 shows the fit in the estimation period. The panel on the right shows the fit over the entire time period available. As one can see in these panels, the fit of the model without wedges is fairly good in the estimation period. After the estimation period, however, there is little apparent relationship between the data on the level of daily deaths and the logarithm of the transmission rate. This pattern suggestive of a change in regime shift in the relationship between daily deaths and the logarithm of the transmission rate and this pattern is quite common across all the countries and U.S. states that we examine.

The remaining panels in Figure 10 show the evolution over time of the BSIR model's predictions for the growth rate of daily deaths, the level of daily deaths, and cumulative deaths over time, from the estimated model without wedges in blue and the model with wedges in red. The time paths for these data are shown with a thin solid line in each panel. As one can see, the red line indicating the model solution with wedges lies on top of the data (by construction). The vertical black line in each of these panels indicates the end of the estimation period.

As is evident in Figure 10, the BSIR model without wedges (in blue) captures the main features of the dynamics of the growth rate of daily deaths over time. For Italy, this growth rate started at a high level and then fell rapidly over time to a level below zero. But small and persistent differences in growth rates translate into large differences in the levels of daily and cumulative deaths. In the data, the growth rate of daily deaths remained substantially below the level predicted by the BSIR model without wedges through the Summer, leading to a much more rapid decline in daily deaths after the initial peak and a much lower level of cumulative deaths by the end of the summer than would be predicted than would be predicted by the BSIR model without wedges.

In Figure 11, we show the simulated paths for $S_i(t)$ and $I_i(t)$ for Italy from the BSIR model without wedges in blue and with wedges in red in the phase plane. We

also show the locus of points for which $\dot{I} = 0$ for various dates implied by the time path of the wedge $\psi_i(t)$. The initial locus of points for which $\dot{I} = 0$ corresponding to $\psi_i(0) = 0$ is shown as a solid black line with no markers. For Italy, this line intersects the x-axis just below $S = 0.3$. This intersection corresponds to a basic reproduction number for Italy of $\mathcal{R}_0 = 3.44$. The blue curve in the figure shows the evolution of the epidemic implied by the BSIR model without wedges, with a first peak of infections smaller than what occurred in the data followed by a slow decline in the level of infections as the model outcome evolves slightly above this initial locus of points for which $\dot{I} = 0$.

In contrast with the curve in blue, we see from the curve in red that the BSIR model with wedges (and the data) had a higher initial peak of infections and then a rapid decline to a very low level. To have the BSIR model with wedges account for this pattern, we see first a rightward shift of the locus of points for which $\dot{I} = 0$ corresponding to negative values of $\psi_i(t)$ at seven and fourteen days after Italy first reached 25 cumulative deaths, and then an acceleration of transmission holding behavior fixed corresponding to positive values of $\psi_i(t)$ 21 and 28 days, followed to a dramatic decline in transmission holding behavior fixed to a level corresponding 84 days in to a value of $\bar{S}_i(t) = 0.97$ or a basic reproduction number barely above one.

The main conclusion we draw from Figure 11 is that the magnitude of the time variation in the wedges needed to account for the patterns of daily deaths observed in Italy are very large relative to variation over time in the growth rates of daily deaths and the corresponding effective reproduction numbers implied there. If one were to project long run outcomes for Italy based on the level of the wedge $\psi_i(t)$ at a point in time, one would forecast that anywhere between 3% and 80% of Italians would eventually contract the disease. Put another way, the transmission rate $\beta(t)$ holding behavior fixed implied by the BSIR model with his varying by a factor of five.

The results for Arizona are presented in Figures 12 and 13. Those for Japan are in Figures 14 and 15. The arrangement of the panels in each figure is the same as for

Italy. In the upper left panels of Figures 12 and 14, we see the fit of the estimated BSIR model reduced form relationship between daily deaths and the logarithm of the transmission rate is fairly good during the estimation period and then completely breaks down afterwards as was the case with Italy. Again, this pattern is pervasive across the locations that we study.

As shown in Figure 1, in Arizona, the rate of daily deaths built slowly to a late peak relative to Italy and Japan. As shown in Figure 12, the BSIR model without wedges (in blue) does not match this peak. Instead, wedges that raise the transmission rate and the growth rate of daily deaths after the estimation period are required to have the model match the data. This is the opposite of what we found for Italy, where the wedges initially reduced the transmission rate relative to what the model without wedges predicted after the estimation period.

In the phase diagram for Arizona in Figure 13, we again see how large the required wedges are when measured in terms of the corresponding fraction of the population left susceptible in the long run (indicated by where the shifting $\dot{I} = 0$ scheduled intersect the x-axis). At our estimated value of $\bar{\beta}_i$ for Arizona, the basic reproduction number at the start of the epidemic was on the order of 1.66, corresponding to a long-run attack rate of 40%. This level then shifts out to nearly 88% by day 105, corresponding to a basic reproduction number over 8, and it started shifting back substantially after that. So again, we see that the wedges needed to match the data correspond to very large shifts in the transmission rate holding disease prevalence constant.

As shown in Figure 1, qualitatively, Japan shows a pattern of multiple waves of daily deaths from COVID similar to that observed in Italy and other locations. In Figure 14, we see that this pattern is matched with wedges that first lower the transmission rate holding disease prevalence fixed and then raise it back up again.

But quantitatively, Japan is quite different from Italy in that it has not yet experienced a high level of deaths from COVID. We see why in the phase diagram in Figure 15. In that figure, we see that the transmission rate holding disease prevalence fixed,

as measured by the points at which the $\dot{I} = 0$ schedule intersects the x-axis, has never been high in Japan. This model implied long-run attack rate has been below 50% for the duration of the pandemic and it shifted to a very low level, less than 15%, around day 56. So again, the wedges required to match the data are large in magnitude.

6.4 Pandemic Fatigue

We now consider the wedges implied by an alternative specification of our model that includes pandemic fatigue. We model pandemic fatigue as a reduction in each region in the reduced-form semi-elasticity of the transmission rate to disease incidence when measured by daily deaths corresponding to the parameter $\alpha\kappa_i$ in equation 14. Note that the reduction in this parameter can be interpreted as a change in the sensitivity of the transmission rate with respect to activity (the parameter α in equation 12) or in the sensitivity of activity to disease prevalence (the parameter κ_i in equation 13) or some combination of the two. We consider specifically a reduction in $\alpha\kappa_i$ in each region by a factor of five relative to what we estimate in the initial phase of the epidemic. Thus we compute the wedges implied by our model for each region from equation 19 using the value of $\bar{\beta}_i$ estimated from the initial phase of the epidemic throughout and using the estimated value of $\alpha\kappa_i$ during the estimation period and replacing this parameter by one-fifth of its estimated value for all dates after the estimation period.

The evolution of the cross-section distribution of wedges implied by this alternative specification of our model is shown in Figure 16. The black solid line in both charts represents the median posterior estimate. The two red dash-dotted bands in this chart contain two thirds of the posterior probability at each point in time and the two blue dashed bands, 0.90 of the posterior probability. This figure should be compared to the bottom panel in Figure 9 which shows the same evolution of the distribution of wedges implied by our baseline model.

What we see from this comparison of the wedges implied by our baseline and alternative model with pandemic fatigue is that the dispersion of the wedges required to have this alternative model with pandemic fatigue match the data outside the estimation period is considerably reduced. That is, the distance between the two red dashed-dotted lines and between the two blue dashed lines in Figure 16 is much smaller than that between the corresponding pairs of lines in the bottom panel of Figure 9.

We also see in Figure 16, however, that the levels of the median and percentiles of distribution of wedges outside of the estimation period in this alternative specification of our model with pandemic fatigue are shifted substantially downward (toward negative wedges) relative to the level of those shown for our baseline model in Figure 9. The intuition for this finding can be seen in Figure 4 which shows the phase diagram for our model with different reduced form semi-elasticities of disease transmission with respect to disease prevalence. As shown in this figure, if this semi-elasticity is reduced, all else equal, infections (and hence deaths) should rise, and hence the measured transmission rate $\beta_i(t)$ should also rise. Since this rise in transmission rates is not seen in the data, negative wedges are required lowering the transmission rate holding behavior fixed to have the alternative model with pandemic fatigue match the data on deaths and transmission.

Note from Figure 16 that the magnitude of this median wedge is quite large — on the order of $\psi(t) = -0.5$. For a region that started the pandemic with a basic reproduction number of $\mathcal{R}_0 = 2.5$ and hence a model-implied long run attack rate of 60%, a wedge of this magnitude corresponds to a reduction of the model-implied long-run attack rate to 34%. If such a model-implied forecast were correct, it would correspond to a substantial reduction in the portion of the population needed to be vaccinated to bring the epidemic to a halt.

7 Conclusion

Qualitatively, the push by economists to introduce theories of behavior into epidemiological models of COVID-19 has been a big empirical success. As we demonstrate, even a simple BSIR model matches the main features of the dynamics of the growth rate of deaths observed in many locations around the world.

But, as our accounting procedure shows, much of the dynamics of the level of daily and cumulative deaths are left unexplained by a simple BSIR model. Such a model much be augmented with very large wedges to the transmission rate holding disease prevalence constant to match the data on deaths. What these wedges stand in for is an important topic for future research.

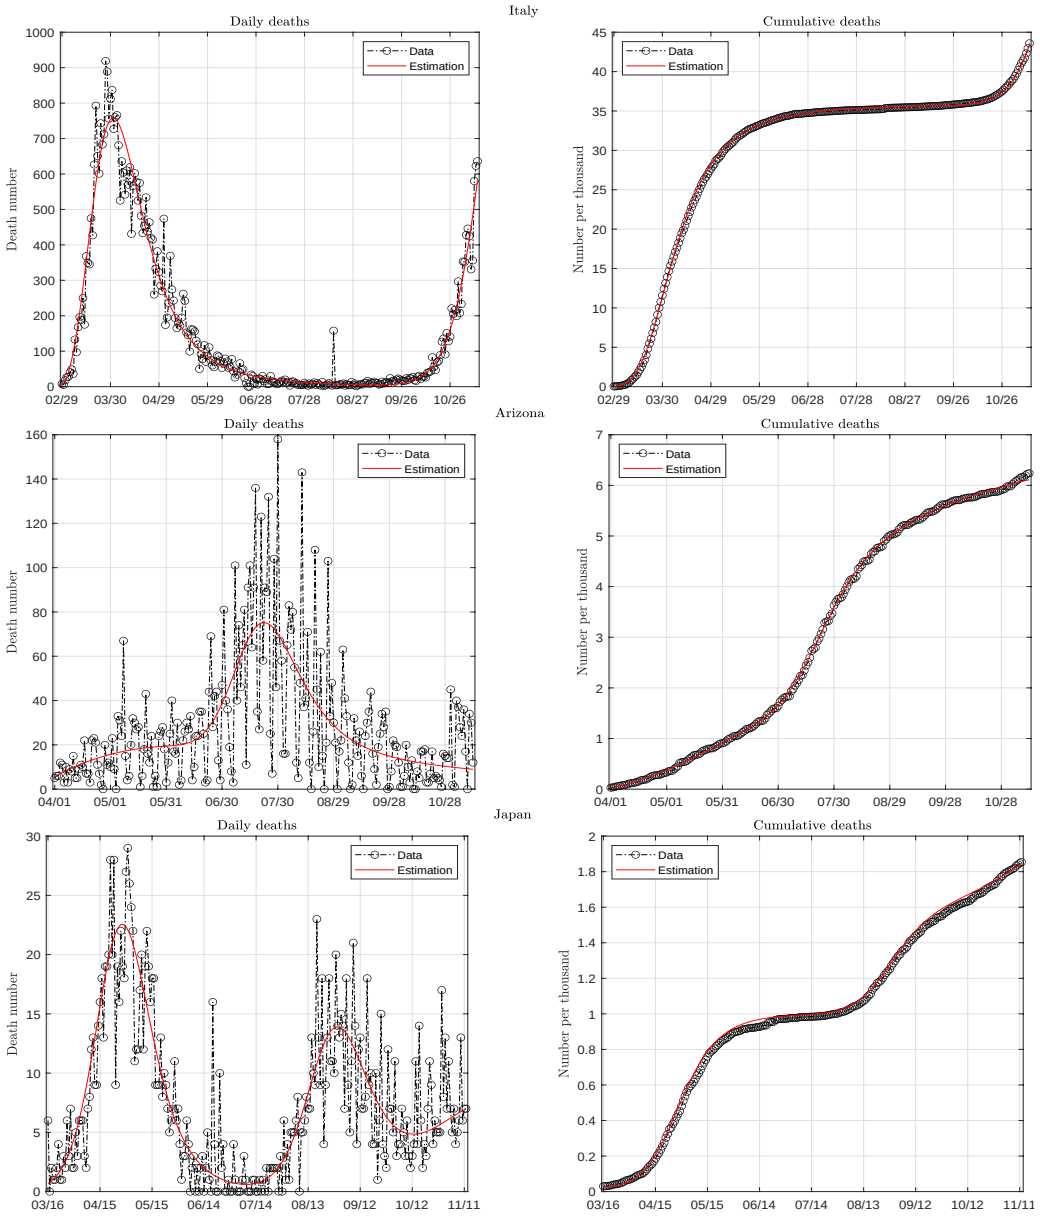

Figure 1: Estimates of Daily and Cumulative Deaths for Italy, Arizona, and Japan through November 12.

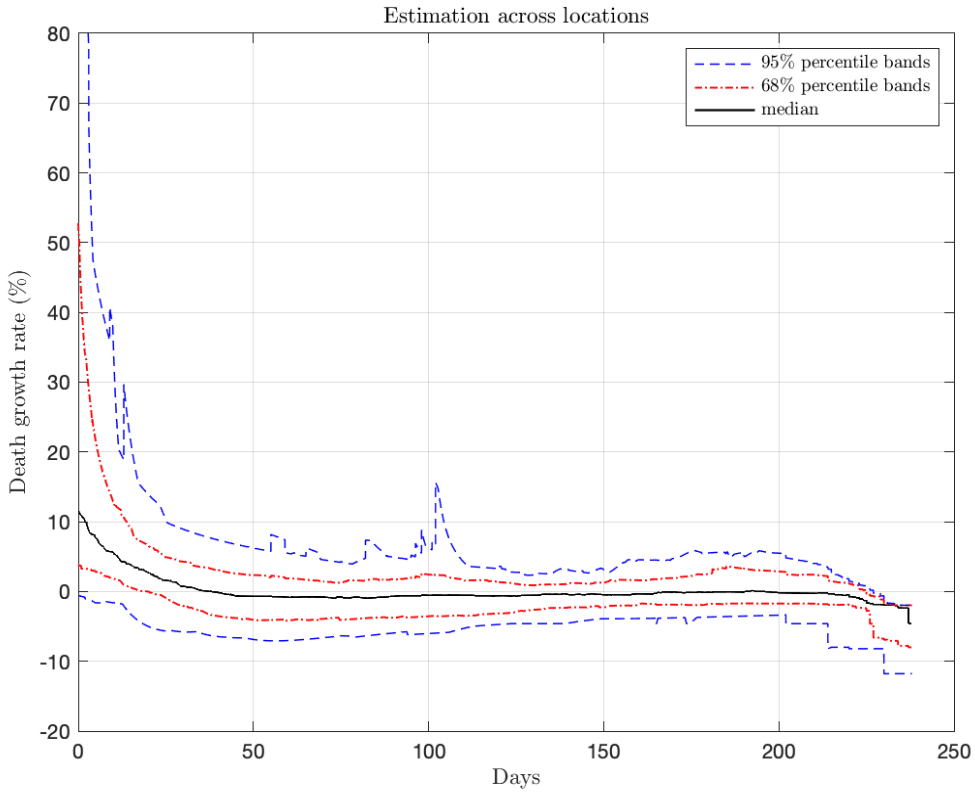

Figure 2: Location and sampling uncertainty. The black solid line represents the median posterior estimate. The two dash-dotted bands in both charts contain two thirds of the posterior probability at each point in time and the two dashed bands, 0.95% of the posterior probability. The growth rates of death is estimated according to the fitted mixture of modified log-logistic density functions. Day 0 is the earliest date when the cumulative death toll reached 25 in each location.

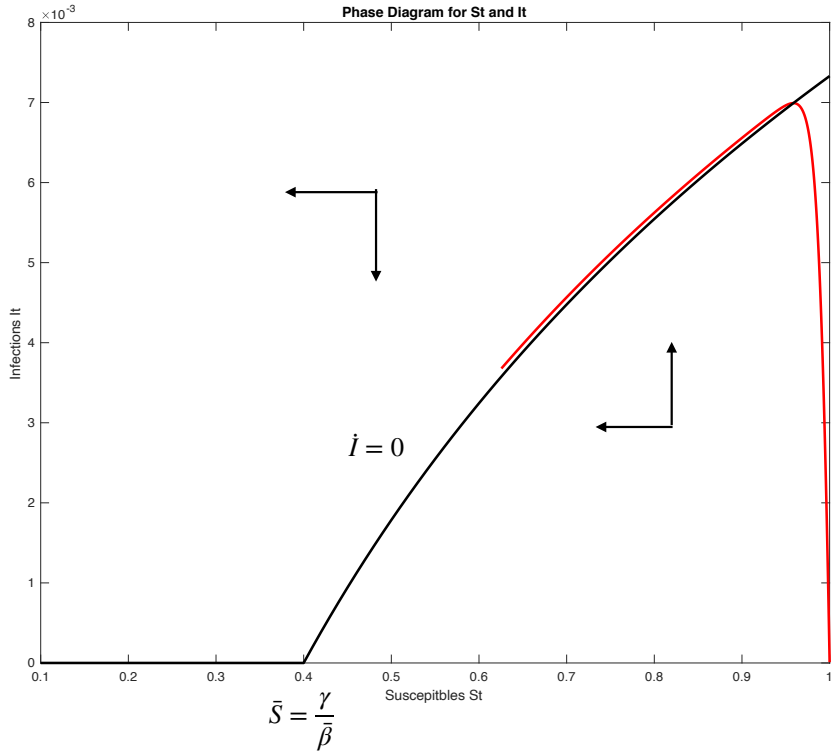

Figure 3: Phase diagram for our BSIR model. The fraction of the population susceptible S_t is on the x-axis, the fraction of the population currently infected is on the y-axis. The black curve is the locus of points (S, I) such that $\dot{I} = 0$. The red curve shows the model implied path of (S_t, I_t). When (S, I) lie below the black curve, $\dot{I} > 0$ and $\dot{S} < 0$. When (S, I) lie above the black curve, $\dot{I} < 0$ and $\dot{S} < 0$. The point at which the black curve intersects the x-axis is determined by the basic reproduction number $\bar{\beta}/\gamma$ with the formula given by $\bar{S} = \gamma/\bar{\beta}$. Model steady-states all have $I = 0$ and hence lie along the x-axis.

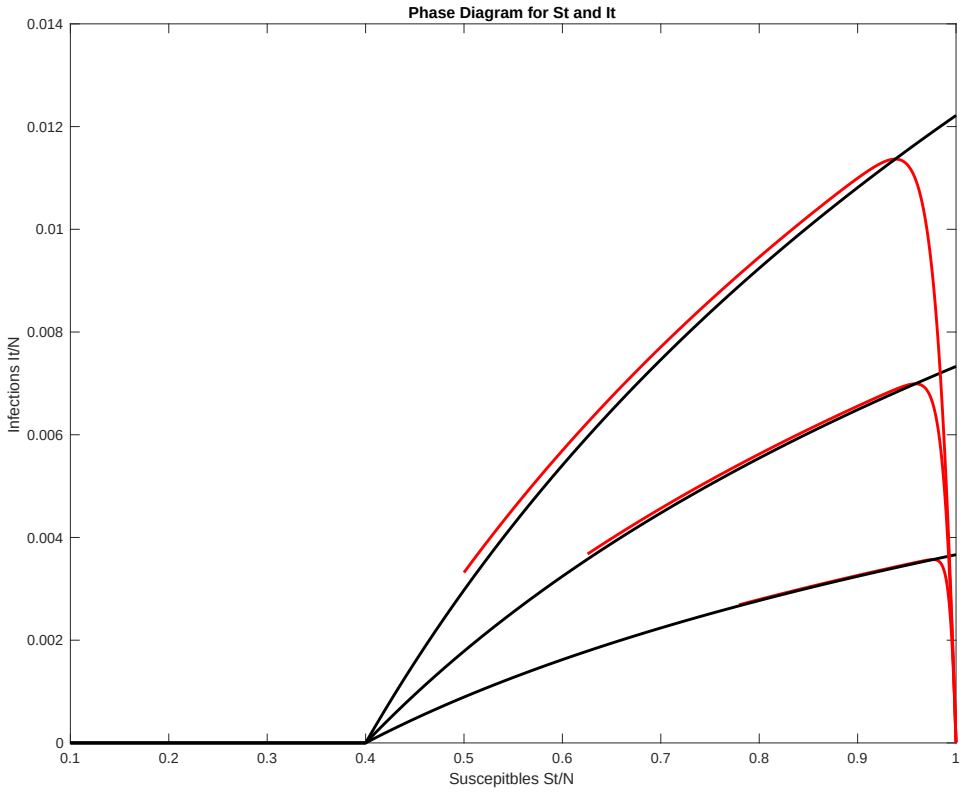

Figure 4: Phase diagram for our BSIR model for various values of the semi-elasticity σ of the transmission rate with respect to infections. The fraction of the population susceptible S_t is on the x-axis, the fraction of the population currently infected is on the y-axis. The black curve is the locus of points (S, I) such that $\dot{I} = 0$. The red curve shows the model implied paths of (S_t, I_t). When (S, I) lie below the black curve, $\dot{I} > 0$ and $\dot{S} < 0$. When (S, I) lie above the black curve, $\dot{I} < 0$ and $\dot{S} < 0$. The point at which the black curves intersect the x-axis is determined by the basic reproduction number $\bar{\beta}/\gamma$ with the formula given by $\bar{S} = \gamma/\bar{\beta}$. This basic reproduction number is held constant across all three model simulations. Model steady-states all have $I = 0$ and hence lie along the x-axis.

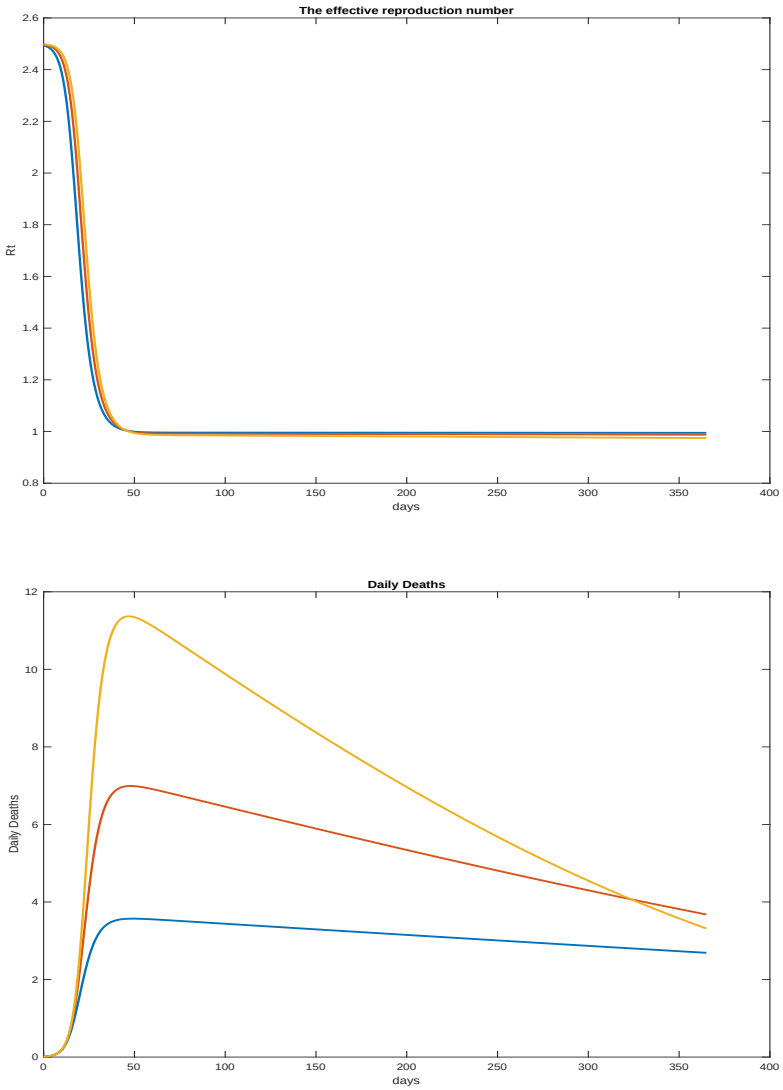

Figure 5: The path of the effective reproduction number and daily deaths per million of population for three simulations of our BSIR model with different values of the semi-elasticity of transmission with respect to daily deaths as shown in Figure 4.

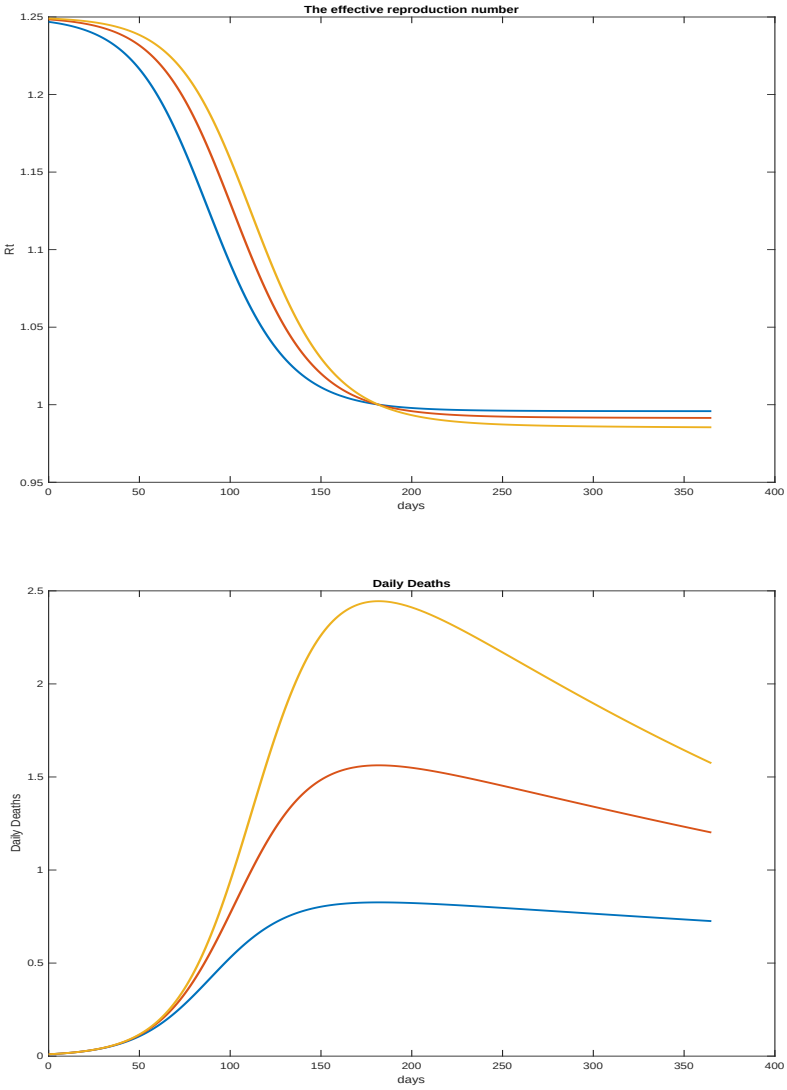

Figure 6: The path of the effective reproduction number and daily deaths per million of population for three simulations of our BSIR model with different values of the semi-elasticity of transmission with respect to daily deaths. These specifications of the model differ from those in Figure 5 in that these specifications have a basic reproduction number of $\mathcal{R}_0 = 1.25$.

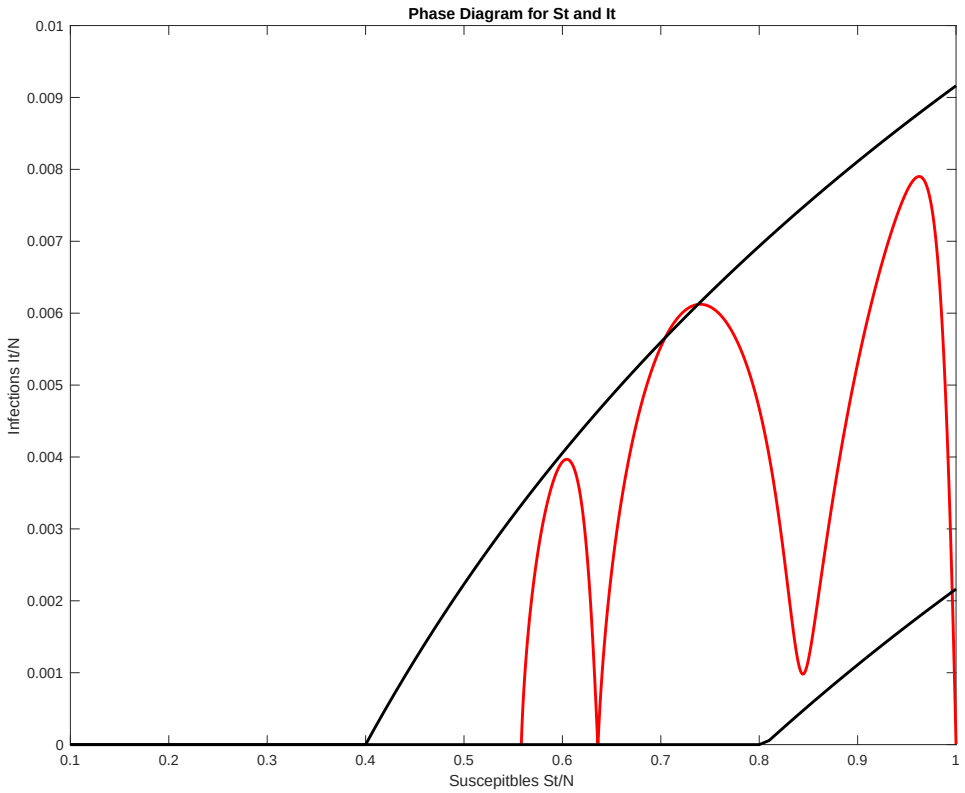

Figure 7: Phase diagram for our BSIR model with $\psi(t)$ following a cosine wave given by $\psi(t) = 0.35(\cos(2\pi t/365) - 1)$. The two black curves represent the highest and lowest levels of the locus of points for which $\dot{I} = 0$. The red curve represents the solution of the model for $(S(t), I(t))$ over the course of three years.

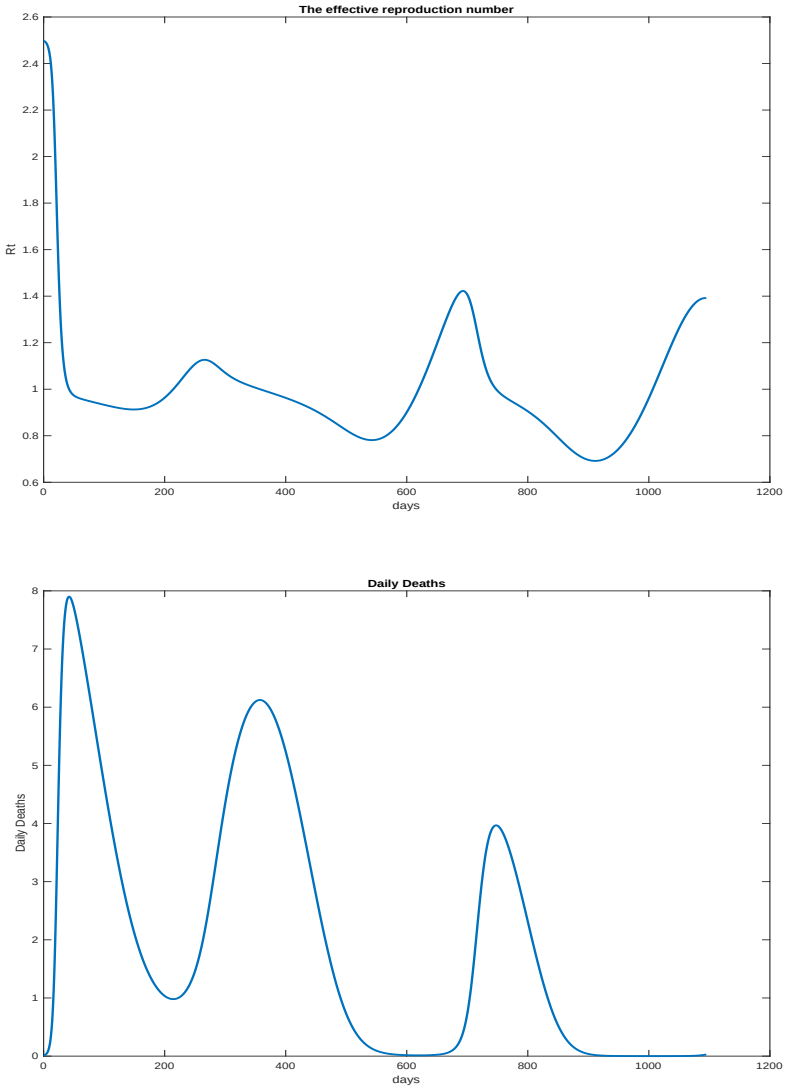

Figure 8: The path of the effective reproduction number and daily deaths per million of population for a three year long simulation our BSIR model with $\psi(t)$ following a cosine wave given by $\psi(t) = 0.35(\cos(2\pi t/365) - 1)$. The basic reproduction number of the model in these simulations oscillates between 2.5 and 1.25.

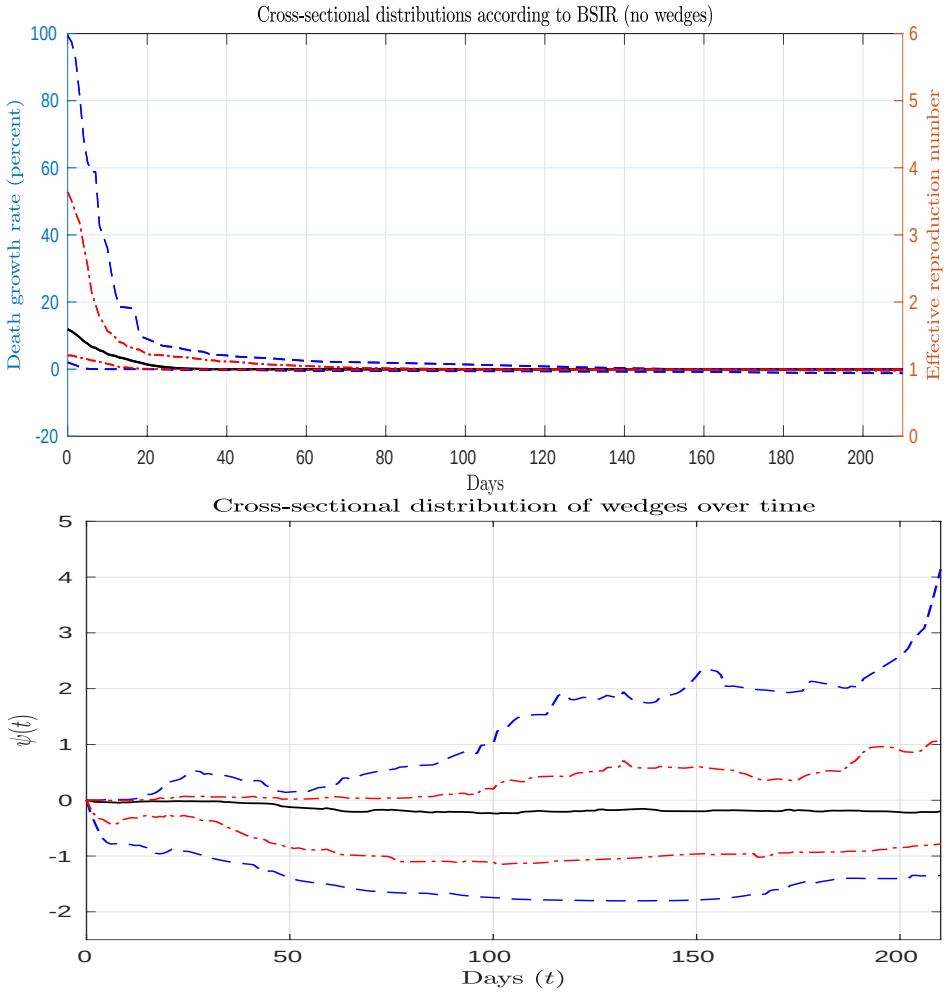

Figure 9: Top panel of this figure shows the distribution of growth rates of daily deaths (on the left axis) and the associated effective reproduction number (on the right axis) predicted by our estimated BSIR model without wedges. The bottom panel of this figure shows the distribution of wedges $\psi_i(t)$ required to make the BSIR model fit the observed data on deaths. The black solid line in both charts represents the median posterior estimate. The two red dash-dotted bands in both charts contain two thirds of the posterior probability at each point in time and the two blue dashed bands, 0.90 of the posterior probability.

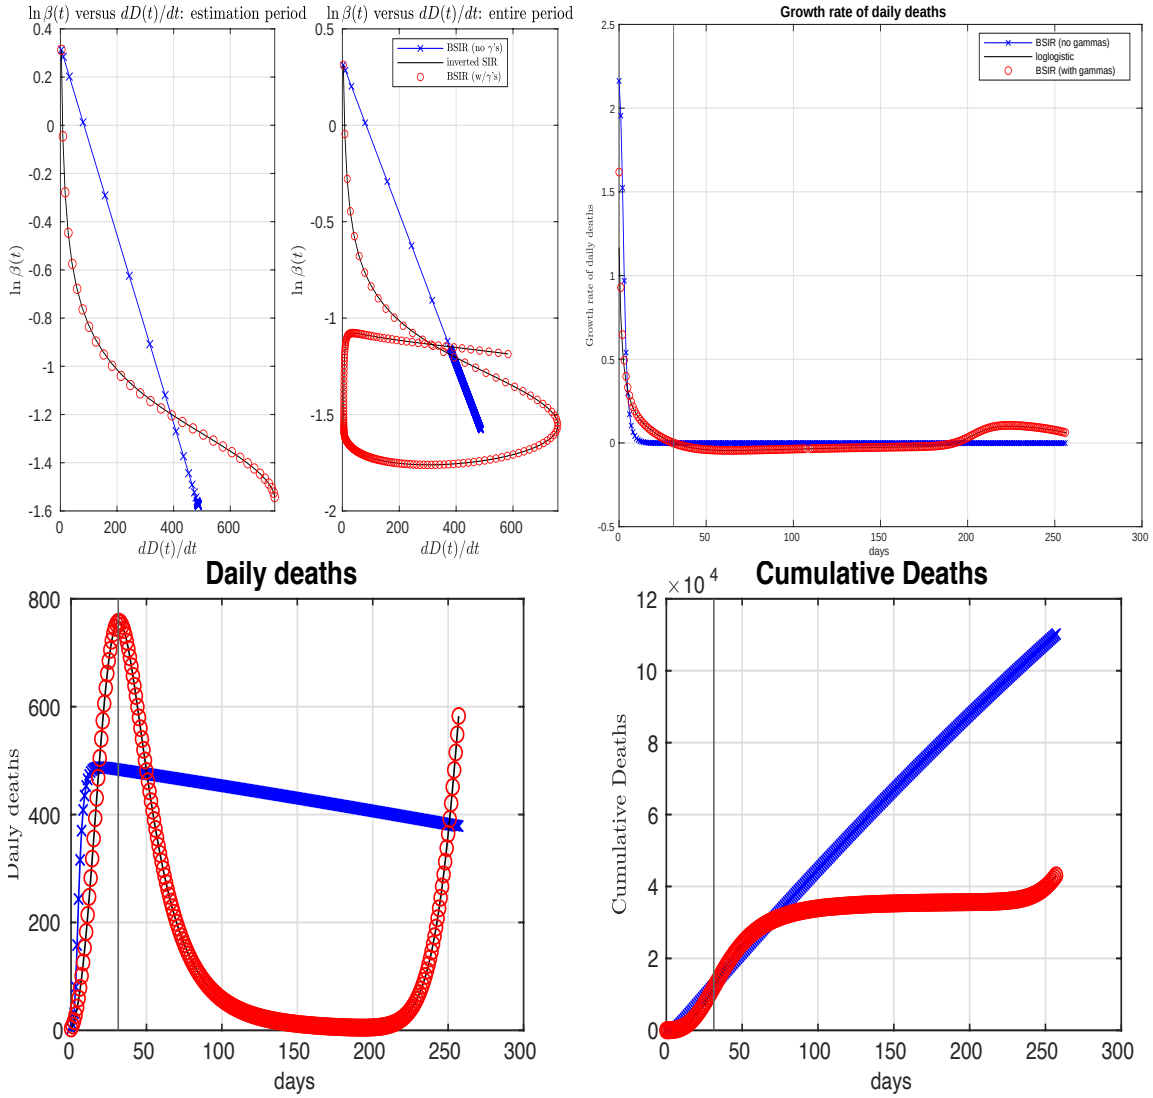

Figure 10: The BSIR model without (in blue) and with (in red) wedges for Italy. The upper left graphs show in blue the relationship between daily deaths and the logarithm of the transmission rate implied by equation 14 and our estimates of $\bar{\beta}_i$ and κ_i without wedges. The red curve shows the pairs $(\dot{D}_i(t), \log(\beta_i(t))$ from the data. The subpanel on the right shows the entire time period, the one on the left for the estimation period. The upper right graph shows the evolution of the growth rate of daily deaths implied by the BSIR model without wedges (in blue) and with wedges (in red), with the estimates from the data shown as a solid black line. (THIS PANEL NEEDS UPDATING). The lower graphs show the evolution of daily deaths and cumulative deaths implied by the BSIR model without wedges (in blue) and with wedges (in red), with the estimates from the data shown as a solid black line. The vertical black lines in these figures shows the end of the estimation period.

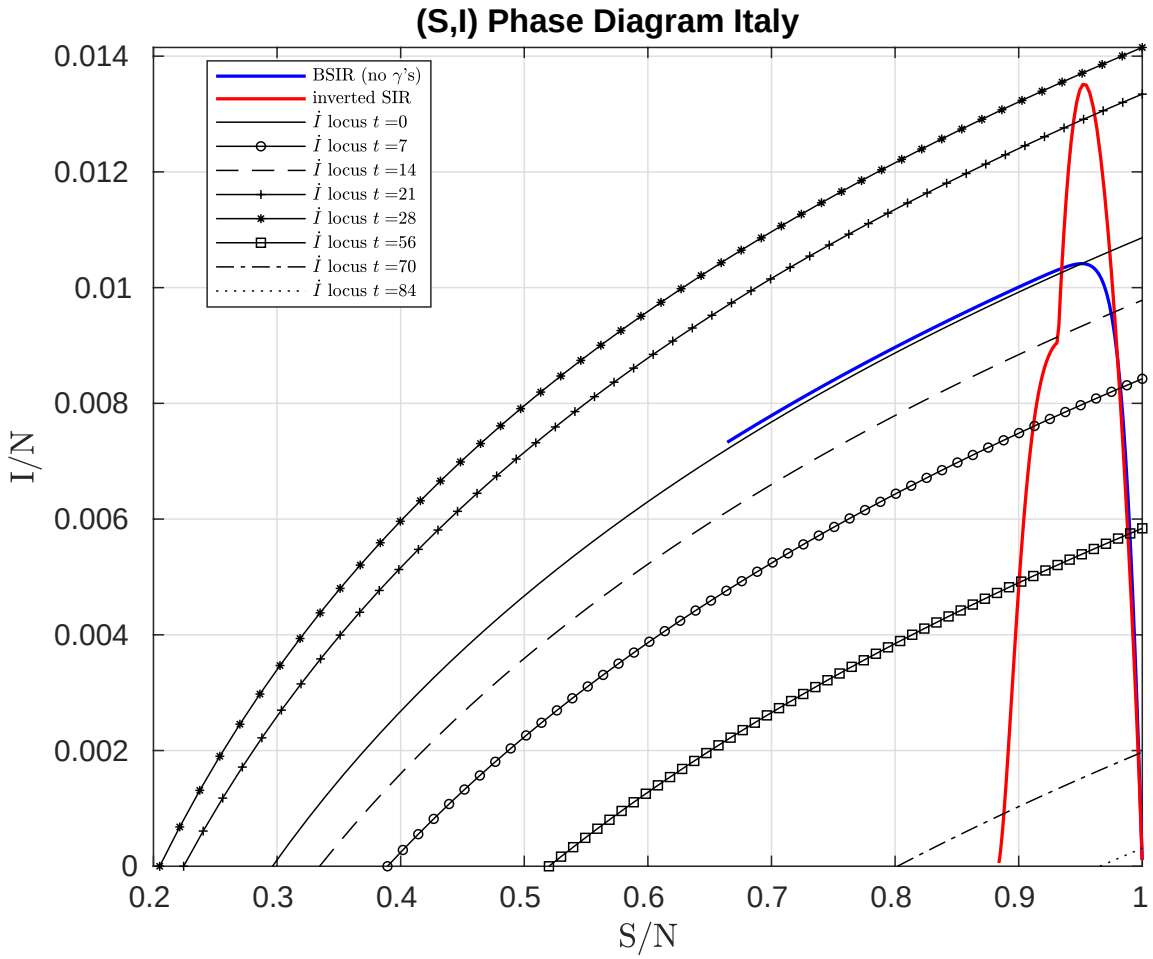

Figure 11: Phase diagram for our BSIR model with wedges for Italy. The fraction of the population susceptible S_t is on the x-axis, the fraction of the population currently infected is on the y-axis.

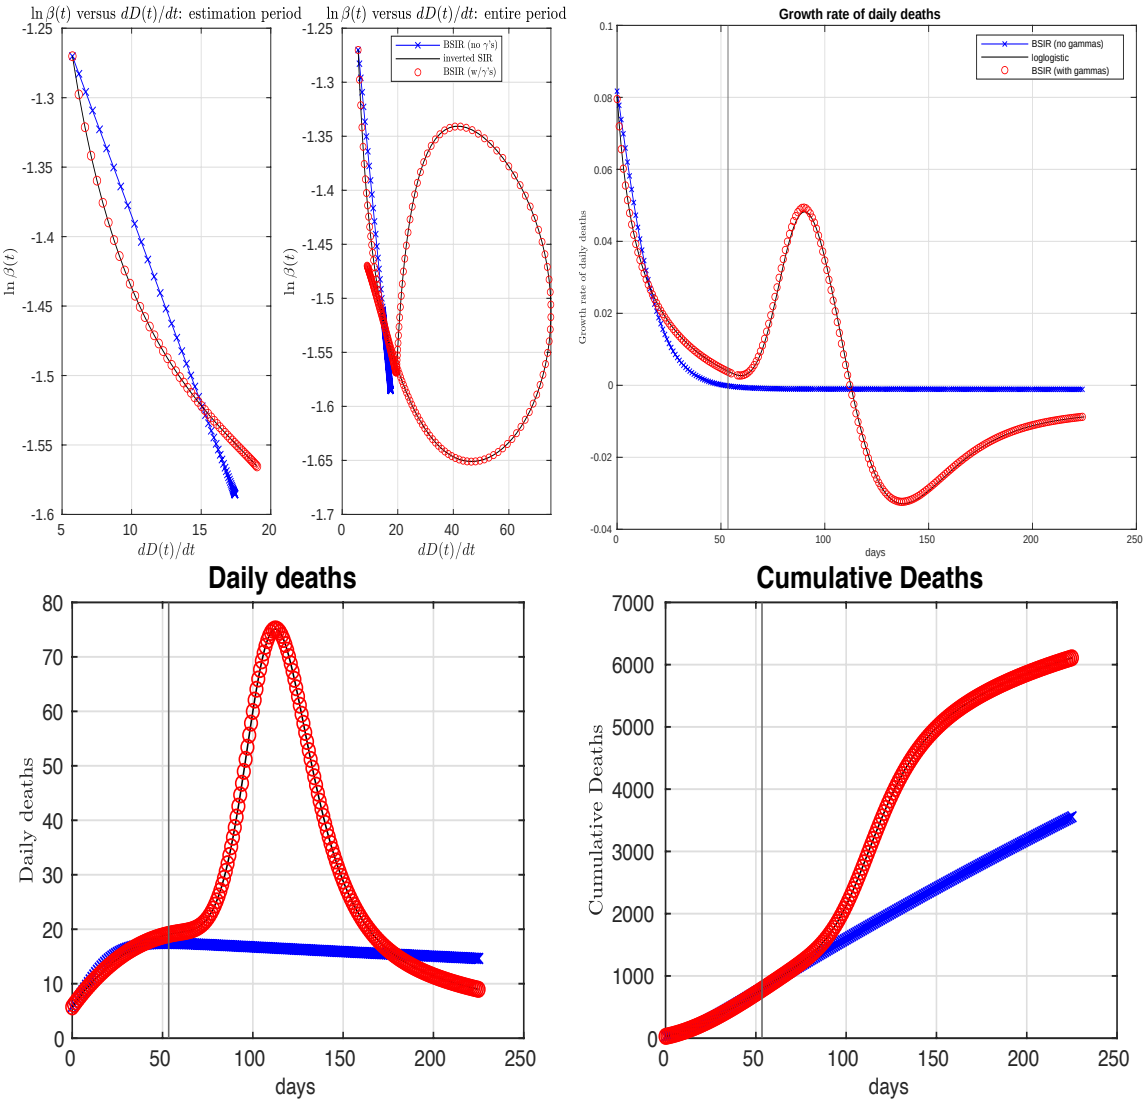

Figure 12: The BSIR model without (in blue) and with (in red) wedges for Arizona. The upper left graphs show in blue the relationship between daily deaths and the logarithm of the transmission rate implied by equation 14 and our estimates of $\bar{\beta}_i$ and κ_i without wedges. The red curve shows the pairs $(\dot{D}_i(t), \log(\beta_i(t))$ from the data. The subpanel on the right shows the entire time period, the one on the left for the estimation period. The upper right graph shows the evolution of the growth rate of daily deaths implied by the BSIR model without wedges (in blue) and with wedges (in red), with the estimates from the data shown as a solid black line. The lower graphs show the evolution of daily deaths and cumulative deaths implied by the BSIR model without wedges (in blue) and with wedges (in red), with the estimates from the data shown as a solid black line. The vertical black lines in these figures shows the end of the estimation period.

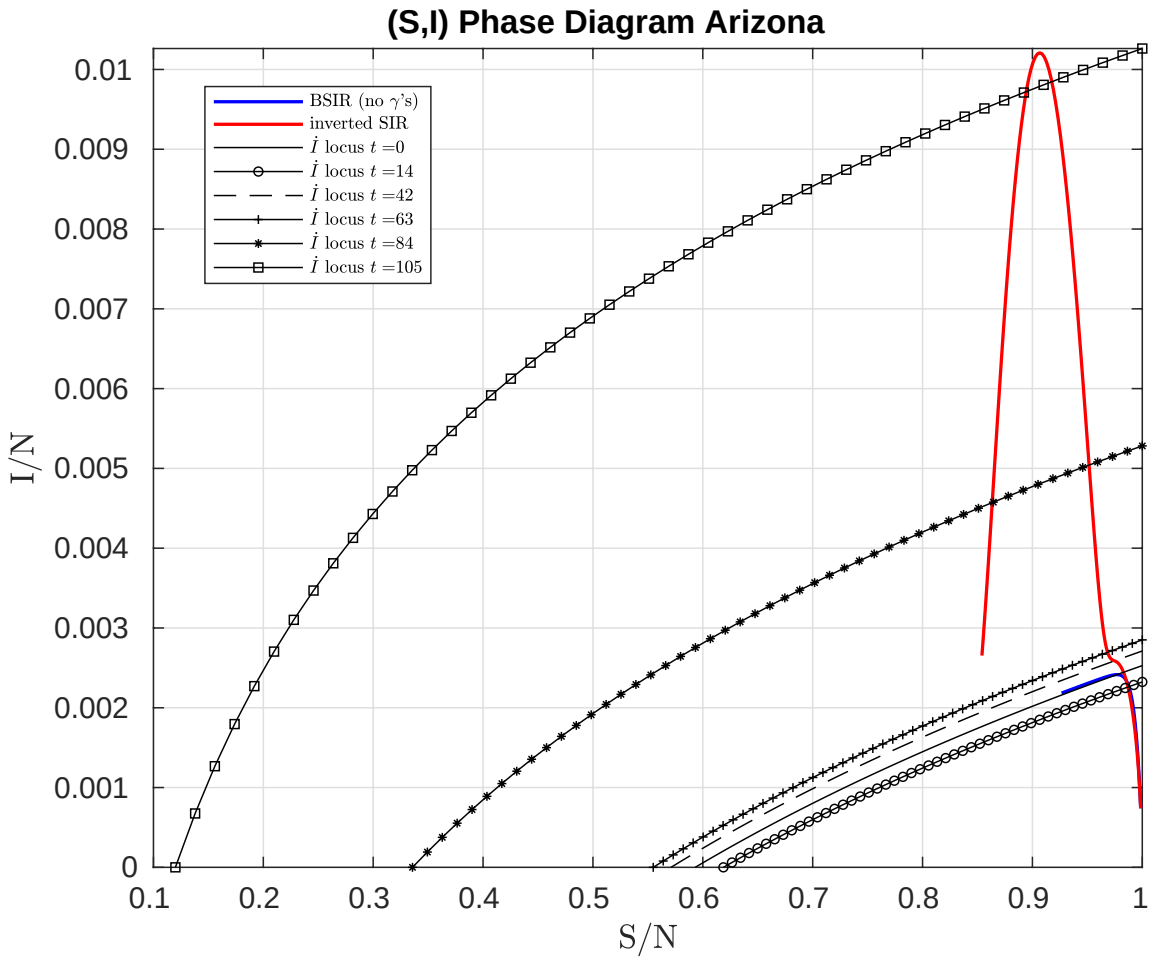

Figure 13: Phase diagram for our BSIR model with wedges for Arizona. The fraction of the population susceptible S_t is on the x-axis, the fraction of the population currently infected is on the y-axis.

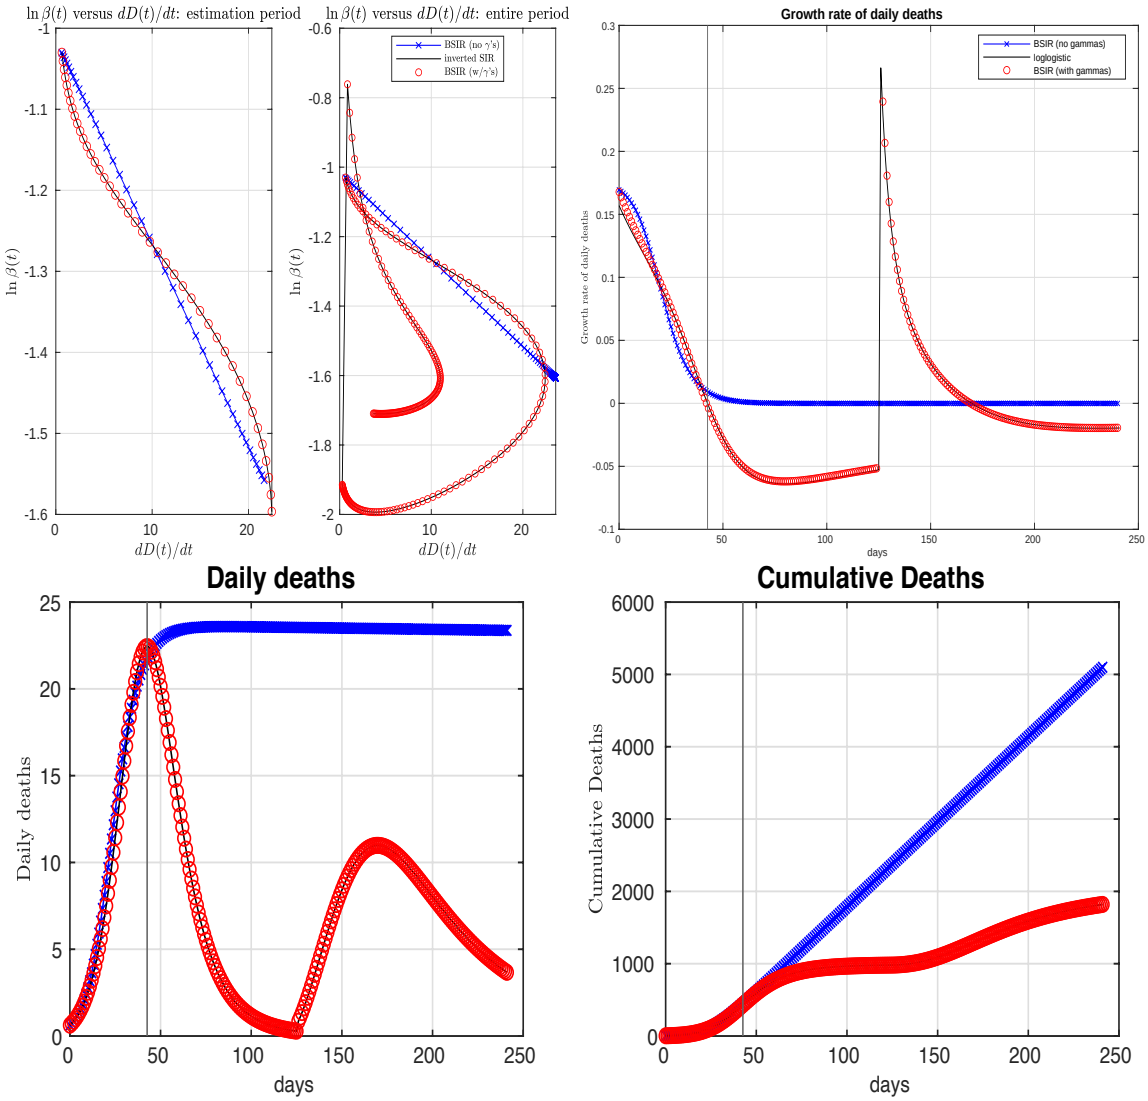

Figure 14: The BSIR model without (in blue) and with (in red) wedges for Japan. The upper left graphs show in blue the relationship between daily deaths and the logarithm of the transmission rate implied by equation 14 and our estimates of $\overline{\beta}_i$ and κ_i without wedges. The red curve shows the pairs $(\dot{D}_i(t), \log(\beta_i(t))$ from the data. The subpanel on the right shows the entire time period, the one on the left for the estimation period. The upper right graph shows the evolution of the growth rate of daily deaths implied by the BSIR model without wedges (in blue) and with wedges (in red), with the estimates from the data shown as a solid black line. The lower graphs show the evolution of daily deaths and cumulative deaths implied by the BSIR model without wedges (in blue) and with wedges (in red), with the estimates from the data shown as a solid black line. The vertical black lines in these figures shows the end of the estimation period.

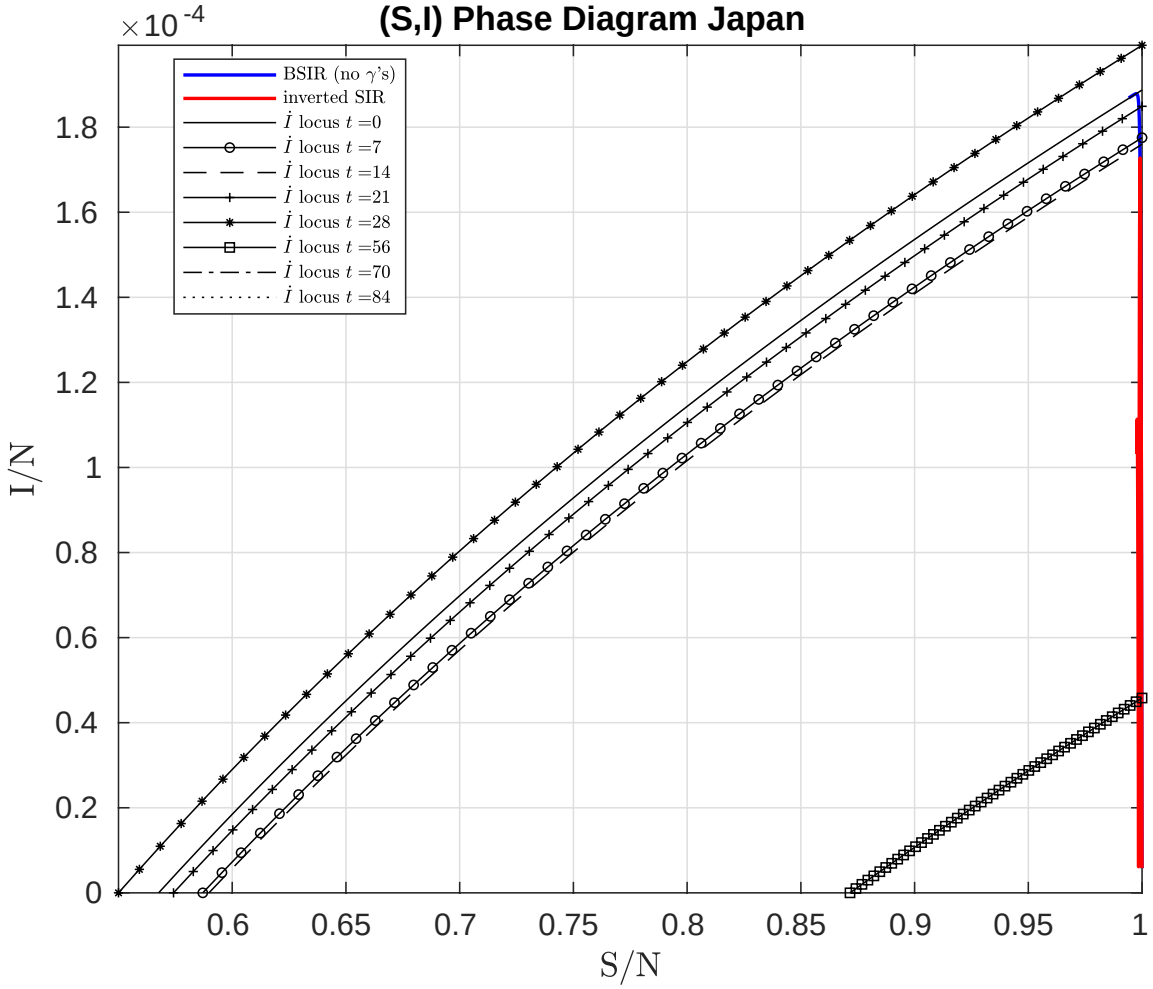

Figure 15: Phase diagram for our BSIR model with wedges for Japan. The fraction of the population susceptible S_t is on the x-axis, the fraction of the population currently infected is on the y-axis.

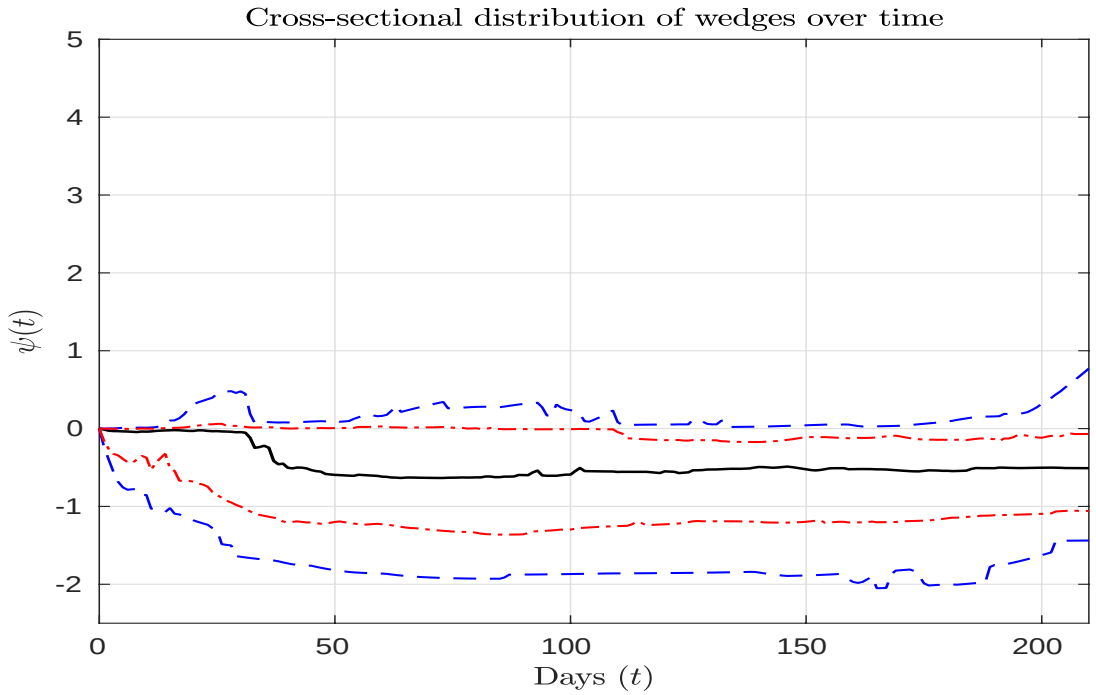

Figure 16: This figure shows the distribution of wedges $\psi_i(t)$ required to make the BSIR model with pandemic fatigue fit the observed data on deaths. The black solid line in both charts represents the median posterior estimate. The two red dash-dotted bands in this chart contain two thirds of the posterior probability at each point in time and the two blue dashed bands, 0.90 of the posterior probability. This figure should be compared to the bottom panel in Figure 9 which shows the same evolution of the distribution of wedges implied by our baseline model.

References

Daron Acemoglu, Victor Chernozhukov, Iván Werning, and Michael D Whinston. Optimal targeted lockdowns in a multi-group sir model. Working Paper 27102, National Bureau of Economic Research, May 2020.

Alex Arnon, John Ricco, and Kent Smetters. Epidemiological and economic effects of lockdown. *Brookings Papers on Economic Activity*, September 2020.

Andrew Atkeson, Karen Kopecky, and Tao Zha. Estimating and forecasting disease scenarios for covid-19 with an sir model. Working Paper 27335, National Bureau of Economic Research, June 2020. URL http://www.nber.org/papers/w27335.

V. V. Chari, Patrick J. Kehoe, and Ellen R McGrattan. Business cycle accounting. *Econometrica*, 75(3):781–836, May 2007.

Sergio Correia, Stephen Luck, and Emil Verner. Pandemics depress the economy, public health interventions do not: Evidence from the 1918 flu,. June 2020.

Christopher J Cronin and William N Evans. Private precaution and public restrictions: What drives social distancing and industry foot traffic in the covid-19 era? Working Paper 27531, National Bureau of Economic Research, July 2020. URL http://www.nber.org/papers/w27531.

Martin S Eichenbaum, Sergio Rebelo, and Mathias Trabandt. The macroeconomics of epidemics. Working Paper 26882, National Bureau of Economic Research, March 2020. URL http://www.nber.org/papers/w26882.

Ceyhun Eksin, Keith Paarporn, and Joshua S. Weitz. Systematic biases in disease forecasting – The role of behavior change. *Epidemics*, 27:96–105, June 2019.

Maryam Farboodi, Gregor Jarosch, and Robert Shimer. Internal and external effects of social distancing in a pandemic. Working Paper 27059, National Bureau of Economic Research, April 2020. URL http://www.nber.org/papers/w27059.

Jesus Fernandez-Villaverde and Charles I. Jones. Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities. Working Paper 27128, NBER, April 2020.

Thiemo R Fetzer, Marc Witte, Lukas Hensel, Jon Jachimowicz, Johannes Haushofer, Andriy Ivchenko, Stefano Caria, Elena Reutskaja, Christopher P Roth, Stefano Fiorin, Margarita Gómez, Gordon Kraft-Todd, Friedrich M Götz, and Erez Yoeli. Global behaviors and perceptions at the onset of the covid-19 pandemic. Working Paper 27082, National Bureau of Economic Research, May 2020. URL http://www.nber.org/papers/w27082.

Joshua S Gans. The economic consequences of r = 1: Towards a workable behavioural epidemiological model of pandemics. Working Paper 27632, National Bureau of Economic Research, July 2020.

Austan Goolsbee and Chad Syverson. Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. Working Paper 27432, National Bureau of Economic Research, June 2020. URL http://www.nber.org/papers/w27432.

Veronica Guerrieri, Guido Lorenzoni, Ludwig Straub, and Iván Werning. Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? Working Paper 26918, National Bureau of Economic Research, April 2020. URL http://www.nber.org/papers/w26918.

Sumedha Gupta, Thuy D Nguyen, Felipe Lozano Rojas, Shyam Raman, Byungkyu Lee, Ana Bento, Kosali I Simon, and Coady Wing. Tracking public and private responses to the covid-19 epidemic: Evidence from state and local government actions. Working Paper 27027, National Bureau of Economic Research, April 2020. URL http://www.nber.org/papers/w27027.

Greg Kaplan, Benjamin Moll, and Gianluca Violante. Pandemics according to hank. May 2020.

Jussi Keppo, Elena Quercioli, Marianna Kudlyak, Andrea Wilson, and Lones Smith. The behavioral sir model, with applications to the swine flu and covid-19 pandemics. April 2020.

Dirk Krueger, Harald Uhlig, and Taojun Xie. Macroeconomic dynamics and reallocation in an epidemic: Evaluating the "swedish solution". Working Paper 27047, National Bureau of Economic Research, April, 2020.

Megan O'Driscoll, Gabriel Ribeiro Dos Santos, Lin Wang, Derek A. T. Cummings, Andrew S. Azman, Juliette Paireau, Arnaud Fontanet, Simon Cauchemez, and Henrik Salje. Age-specific mortality and immunity patterns of sars-cov-2. *Nature*, 2020.

Tomas Phillipson and Richard A. Posner. *Private Choices and Public Health: the AIDS Epidemic in an Economic Perspective*. Harvard University Press, 1993.

Adam Sheridan, Asger Lau Andersen, Emil Toft Hansen, and Niels Johannesen. Social distancing laws cause only small losses of economic activity during the covid-19 pandemic in scandinavia. *Proceedings of the National Academy of Sciences*, 117 (34):20468–20473, 2020 2020.

Flavio Toxvaerd. Equilibrium social distancing. Working Paper 2020/8, Cambridge-INET, March 2020.