Online Appendix for "Behavior and the Dynamics of Epidemics" for the Spring 2021 BPEA *

Andrew G. Atkeson†

This appendix presents the model and parameters used in "Behavior and the Dynamics of Epidemics" by Andrew Atkeson for the Brookings Panel on Economic Activity Spring 2021. This model is based closely on that presented in "A Parsimonious Behavioral SEIR Model of the 2020 COVID Epidemic in the United States and United Kingdom" which is available as NBER working paper 28434 and as Federal Reserve Bank of Minneapolis Staff Report 619. This appendix discusses the model extended to include vaccines and the potential for waning immunity. It is applied to the United States.

This model is a an SEIR model (with compartments for agents who are susceptible, S, exposed, E, infectious, I, and recovered and hence removed R) modified to include a compartment for those infected agents who end up with serious disease. I refer to this compartment as H, for hospitalized. Agents who die from COVID are assumed to transition from infection I to death, D, through this compartment H. The expected time that agents spend in this compartment is set to 30 days to capture the delay between serious illness, death, and the reporting of that death.

Behavior in this model is assumed to respond to daily death rates as they are reported. It is assumed that behavior does not respond immediately to new infections

*All errors are mine. The views expressed here are entirely my own and not official statements of the Federal Reserve Bank of Minneapolis or the Federal Reserve.

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

as these are not directly observed. As discussed by John Cochrane[1] and Weitz et. al. 2020[2] the delay between infection and death introduced by this compartment H implies that this simple behavioral model has oscillatory endogenous dynamics that are helpful in allowing the model to reproduce the data with only a few shocks.

In the United States, we have seen public authorities tighten and loosen COVID mitigation policies in tandem with changes in the prevalence of this disease. I interpret this correlation between public policies and disease prevalence as arising from a public behavioral response to shifting political calculations as disease prevalence rises and falls, that is, as a social choice behavioral response. I interpret private behavioral responses as arising from rising concern over personal infection risk with rising disease prevalence. As described below, I thus interpret the reduced form behavioral response of transmission rates to disease prevalence as resulting from a combination of private and public reactions to disease prevalence. I do not attempt to distinguish the relative importance of these two responses.

The four shocks considered in this paper are as follows. First, I add a standard seasonal variation in the baseline transmission rate of the virus from a winter peak to a low in midsummer. Second, I introduce a one-time change in behavior modeled as a reduction in the semi-elasticity of the transmission rate with respect to the daily death rate from an initial level to a new, permanently lower level. I refer to this second shock as the onset of pandemic fatigue. When I fit the model to the United States, I assume that pandemic fatigue sets in late in 2020. Third, I introduce a more contagious variant of COVID to the United States through a single individual on December 1, 2020. The transmissibility of this variant is set to be 50% higher than the original variant. The model implies that this new variant becomes the dominant variant circulating in the United States by summer of 2021. The fourth shock is the introduction of vaccines starting on January 1, 2021. I assume an aggressive

[1]See https://johnhcochrane.blogspot.com/2020/05/an-sir-model-with-behavior.html

[2]Joshua Weitz, Sang Woo Park, Ceyhun Eksin, and Jonathan Dushoff, "Awareness-driven behavior changes can shift the shape of epidemics away from peaks and toward plateaus, shoulders, and oscillations" , *Proceedings of the National Academy of Science*, vol. 117, no. 51, December 22, 2020

vaccination program that proceeds at a rate to vaccinate just over 50% of the total population by July 1, 2021 and that continues indefinitely.

This appendix has three parts. In section 1, I present the equations of the model. I discuss the interpretation of the reduced form behavioral response of the transmission rate to disease prevalence measured by daily deaths as a combination of endogenous private and public responses to disease prevalence. I also compare the structure of this model with that of a simpler behavioral SIRD model as analyzed in Atkeson, Kopecky, and Zha (2021)[3] and Droste and Stock (2021)[4] .

In section 2, I discuss the values of the parameters. The model is fit to US data on daily deaths from COVID. Several parameters are set to match recommendations from the Center for Disease Control for modeling of COVID-19. Other parameters, in particular, the basic reproduction number of the original variant of the virus, the original semi-elasticity of the transmission rate to the level of daily deaths, and the size of the seasonal fluctuation in transmission, are chosen through a process of trial and error. To illustrate this fitting process, I show the sensitivity of the model predictions to variation in these three parameters around their baseline values.

In section 3, I discuss the shocks to the model. I illustrate the role that these shocks play in allowing the model to match the data on daily deaths from COVID by solving the model with each of these shocks turned off.

1 Model

The model is as follows.

The SEIHR model extends the SIR model by adding both the exposed state E

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

[4]Droste and Stock "Adapting to the Pandemic" forthcoming, *American Economic Review Papers and Proceedings*

and the hospitalized state H. In this version of the model the total population N is given by the sum of susceptible agents in state S, exposed in state E, infected in I, hospitalized in H, recovered in R, and dead in D.

The rate at which agents leave the E compartment for both the normal and more transmissible variants is σ and the rate at which agents leave the I compartment for both variants is γ. The mean generation time for the model is then $1/\sigma + 1/\gamma$. As discussed below, the choice of these parameters is guided by CDC recommendations for these disease parameters.

As agents leave the I compartment, fraction η go into the hospitalized compartment H and $1 - \eta$ transition directly to the recovered compartment R. The rate at which agents leave the H compartment is ζ. Of those leaving the H compartment fraction ν die and $1 - \nu$ survive. Thus, the overall infection fatality rate is given by $\eta\nu$ and the mean time in the H compartment corresponding to illness and delays in reporting deaths is $1/\zeta$.

To model the introduction of a new variant, I add separate compartments E_v and I_v for those exposed to and infectious with the new variant. The transmission rate of the original variant is denoted by $\beta(t)$. That for the new variant is denoted by $\beta_v(t)$. I will assume that the new variant is always 50% more transmissible than the original variant, That is, I set $\beta_v(t) = 1.5\beta(t)$ for all t. The new variant is introduced by setting $\bar{E}_v(t) = 1/population$ in the equations below for one day on a specified date t_v and equal to zero otherwise. This corresponds to the introduction from abroad of a single individual carrying this more transmissible variant. Note that this quantity is subtracted off of the change in the R compartment simply to keep the population constant. Since this shift is only one person for one day, it does not impact the quantitative implications of the model for large populations.

I model the impact of vaccines as moving agents from the susceptible compartment S directly to the removed compartment R at a rate λ per day. With this assumption, I impose that the vaccine blocks both transmission by the vaccinated and disease in the vaccinated. These assumptions are clearly optimistic. I model waning immunity

as a movement of agents from the removed compartment R back to the susceptible compartment at a rate ξ per day. I do not consider population growth in the model.

The dynamics of the model are given by

$$\frac{dS(t)}{dt} = -(\beta(t)I(t) + \beta_v(t)I_v(t))S(t) - \lambda(t)S(t) + \xi R(t)$$

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

$$\frac{dE_v(t)}{dt} = \beta_v(t)I_v(t)S(t) - \sigma E_v(t) + \bar{E}_v(t)$$

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

$$\frac{dI_v(t)}{dt} = \sigma E_v(t) - \gamma I_v(t)$$

$$\frac{dH(t)}{dt} = \eta\gamma(I(t) + I_v(t)) - \zeta H(t)$$

$$\frac{dR(t)}{dt} = (1-\nu)\zeta H(t) + (1-\eta)\gamma(I(t) + I_v(t)) - \bar{E}_v(t) + \lambda(t)S(t) - \xi R(t)$$

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

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

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

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

where the parameters $\bar{\beta}$ and $\bar{\beta}_v$ control the baseline transmissibility of the normal and variant of COVID, the parameter $\psi(t)$ is used to introduce seasonality in transmission, and $\kappa(t)$ is the semi-elasticity of transmission with respect to the level of daily deaths.

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

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

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

To model pandemic fatigue, I set

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

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

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

For all model forecasts, I leave the behavioral parameter κ which determined the semi-elasticity of the transmission rate with respect to daily deaths fixed at its final value at the end of the model estimation period of one year. Thus, I assume that there are no further changes in behavioral parameters going forward.

Initial conditions for all simulations are $E(0) > 0$, $E_v(0) = I(0) = I_v(0) = R(0) = H(0) = D(0) = 0$, $S(0) = 1 - E(0)$. For the United States, $E(0) = 33$ on February 15 out of a population of 330 million. The new more contagious variant is introduced with $\bar{E}_v(t) = 1/330,000,000$ for one day on December 1, 2020.

The MATLAB code to run this model and create all of the figures in this document and the main text is comprised of the following files. *BPEAAtkesonSpring2021.mlx* is a Matlab live script that is the main file. This file calls three Matlab scripts *bpeaodefile.m* is a function with all the differential equations described above, *bpearunthe-*

model.m contains the code to solve these differential equations starting from the initial conditions described above with various parameter configurations, and *us-data.m* contains the CDC data on daily COVID deaths in the United States used in the plots.

1.1 Behavior as a combination of private and public responses to disease prevalence

The reduced form response of the transmission rates of the original and new variants of the virus to daily deaths assumed in the model can be understood as arising from the combined responses of private and public actors to disease prevalence as follows.

This reduced form response of transmission to daily deaths can be obtained as a result of a two-equation system in which the transmission rate is given as a function of activity $Y(t)$ with

$$\beta(t) = \bar{\beta} Y(t)^{\alpha} \exp(v(t) + \psi(t))$$

and activity is given as a declining function of daily deaths

$$Y(t) = \exp(-\frac{\kappa_p(t)}{\alpha} \frac{dD(t)}{dt} + u(t))$$

In the first of these two equations, the parameters $\bar{\beta}$ and $\bar{\beta}_v$ are fixed coefficients that capture features of the virus and the population determined prior to the epidemic that might impact transmission of the original and more transmissible variants. 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.

The parameter $v(t)$ represents a potentially time-varying shock to the region-specific relationship between activity and transmission that may represent the impact

of policy over time. I normalize $v(0) = 0$. When interpreting variation in $v(t)$ as representing the impact of policies, here I am 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.

I normalize the level of activity at the start of the pandemic to $Y(0) = 1$. Given this normalization, the parameter $\bar{\beta}$ sets the transmission rate of the virus at the start of the epidemic. Specifically, $\bar{\beta}/\gamma$ and $\bar{\beta}_v/\gamma$ correspond to the *basic reproduction number* of the original and more transmissible variants at the peak of the seasonal cycle of transmissibility.

In the second of these two equations, I model individuals' decisions to engage in activity at time t, $Y(t)$, as a declining function of the level of daily deaths, $\dot{D}(t)$. The parameter $\frac{\kappa_p(t)}{\alpha}$ is the semi-elasticity of private behavior with respect to daily deaths. The variable $u(t)$ in this second equation represents a time-varying shock to the region-specific relationship between deaths and activity. I interpret $u(t)$ as reflecting public policies such as lockdowns or closures that would reduce activity below what agents might choose in a decentralized fashion.

I consider government policies that are "behavioral" in that they are responsive to the prevalence of the disease as measured by daily deaths as well as policies that depend only on time. To model the behavioral component of government policies, I assume that public policies impacting the relationship between activity and transmission $v(t)$ and between daily deaths and activity $u(t)$ are responsive to the level of daily deaths. Specifically, I assume that

$$v(t) = -\eta_v \frac{dD(t)}{dt}$$

$$u(t) = -\frac{\eta_u}{\alpha} \frac{dD(t)}{dt}$$

I assume that all government mandated mitigation that is dependent only on time enter the model through $\psi(t)$.

With these assumptions, we have that the behavioral decline in activity and disease transmission with daily deaths can be interpreted as arising either from a change in private behavior or public mandates that are conditioned on the prevalence of the disease. Specifically, with these assumptions, we obtain the reduced-form relationship between daily deaths and disease transmission in equation (1) with the overall semi-elasticity of transmission with respect to daily deaths given as a compound parameter capturing both private and public responses to daily deaths

$$\kappa = \kappa_p + \eta_v + \eta_u$$

The parameter $\psi(t)$ corresponds to changes in disease transmission that are unrelated to the level of daily deaths. In the baseline model, this variation is pure seasonal variation in disease transmission. When I consider additional government efforts at mitigation of transmission such as through testing and tracing and isolation of the infected, I model these interventions as a shift in $\psi(t)$ as well. In this case, this component of public policy should be interpreted as an intervention that is not directly responding to disease prevalence.

1.2 Comparison with behavioral SIRD models

Note here that the addition of the compartments E and H to this model together with the assumption that behavior responds to the level of daily deaths distinguishes the behavioral component of this model from those in the SIRD models in Atkeson, Kopecky, and Zha (2021) and Droste and Stock (2021). These additional compartments substantially alter the endogenous dynamics implied by this model relative to these models. The difference between the implied dynamics can be understood as follows.

In an SIRD model, the level of daily deaths is directly proportional to the current level of the infected population $I(t)$, so, an assumption that behavior reacts to disease prevalence measured by the level of daily deaths is equivalent to an assumption that

behavior reacts to the current level of active infections in the population. In turn, since the growth of the infected population is directly governed by the transmission rate $\beta(t)$, then the assumption, as above, that behavior reacts to daily deaths is equivalent to assuming a tight relationship between the level and growth rate of active infections. One way to interpret the findings in Atkeson, Kopecky, and Zha (2021) and Droste and Stock (2021) that large and frequent shocks are needed to fit a behavioral SIRD model to the data on daily deaths from the United States is as a finding that this simple behavioral model is "over-controlled". That is, for the simple behavioral SIRD model, the the model's endogenous dynamics do not permit significant fluctuations in the growth rate of daily deaths after the initial phase of the pandemic and hence such a model requires large and frequent shocks to match the observed fluctuations in the growth rates of daily deaths.

In contrast, in the model presented here, the level of daily deaths in not directly proportional to the infected population $I(t)$, but instead it is directly proportional to the level of hospitalizations $H(t)$, and this level of hospitalizations is, in turn, equivalent to a distributed lag of past levels of infections $I(t)$, with ζ controlling the length of these lags. This lag between infections and the behavioral response gives the model oscillatory dynamics. I show below how changes in this parameter ζ change the endogenous dynamics of the model in such a way to substantially reduce the need for shocks to the model.

That the addition of the E and H compartments allow for a good fit to the deaths data for the United States and United Kingdom is serendipitous. The COVID epidemic has displayed different dynamics in different locations. I have not had similar success in fitting this model to some other locations that I have tried. I leave it future research to determine whether further refinement of the structure of the model or additional shocks are needed to fit the wide range of COVID outcomes observed around the world.

2 Parameters

In this section I discuss the choice of parameters and shocks to the model. To illustrate how parameter values and shocks impact the implications of the model, I use the simulation of the model shown in Figure 4 of the main text as a baseline for comparison. This simulation of the model has the baseline parameter values described below and the shocks of seasonality in transmission, pandemic fatigue, and the new variant, and the introduction of vaccines. To ease comparison, I reproduce Figure 4 from the main text here in Figure A1.

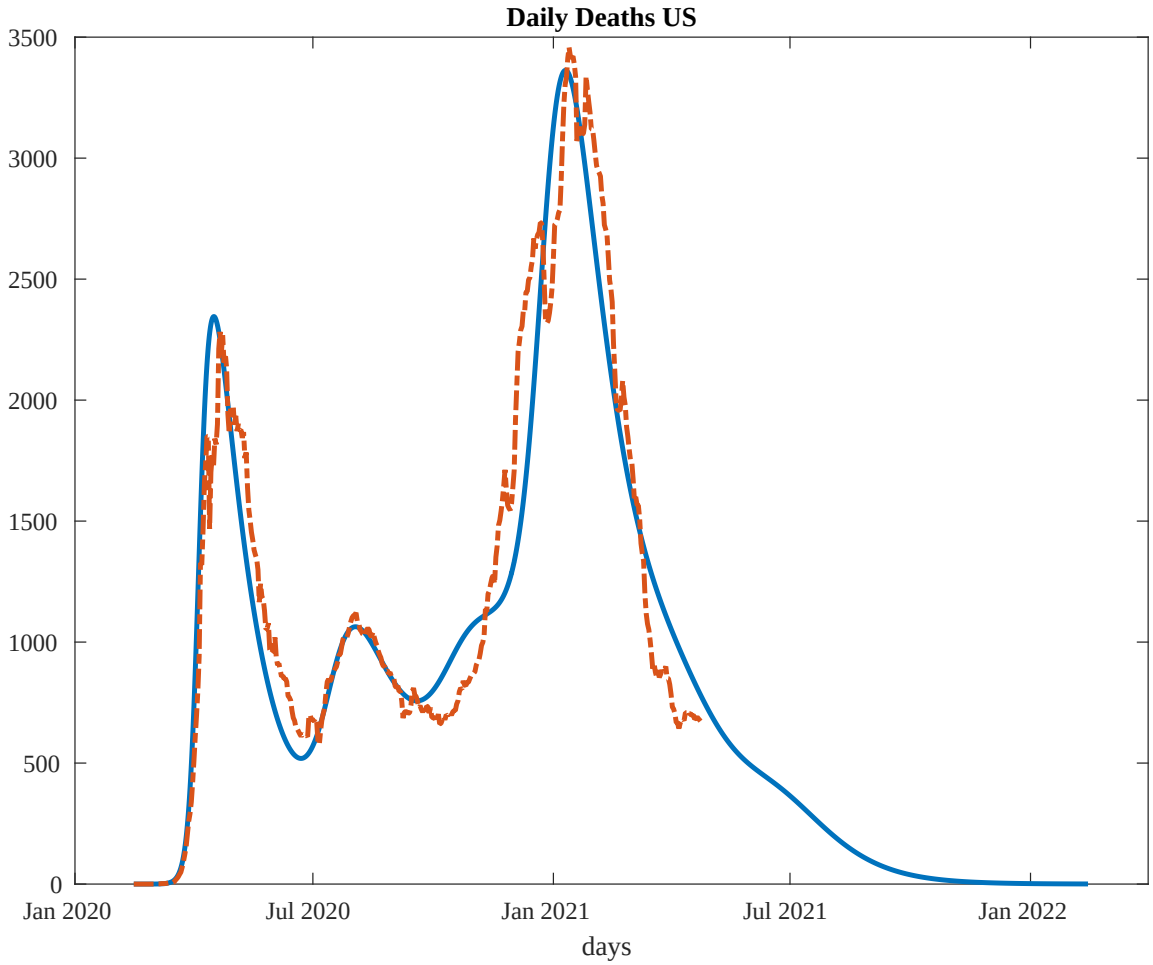

Figure A1: Baseline model implications corresponding to Figure 4 in the main text. This simulation has the baseline parameter valeus, the shocks due to seasonality, pandemic fatigue, and the introduction of a new, more transmissible variant, as well as the introduction of vaccines.

I group the parameters of the model into three sets.

The first set are those set from CDC model recommendations for the generation time of infections and the infection fatality rate. As is standard in an SIR model, the basic reproduction number, which is a ratio of two rates and is thus unitless, determines the overall shape of the epidemic. The generation time then determines the time-scale over which the epidemic plays out.

This first set of parameters is set as follows: $\gamma = 0.4$, $\sigma = 0.425$, $\eta = 0.025$, $\nu = 0.2$. The parameter σ corresponds to an expected time before an exposed agent becomes infectious of 2.35 days and the parameter γ corresponds to an expected time for which an infected individual is infectious of 2.5 days. The generation time is defined as the average time between which one infected individual shows symptoms and a person infected by that individual shows symptoms. These two parameters together imply an average generation time of $1/\sigma + 1/\gamma = 4.85$ days.[5] As mentioned above, this generation time sets the time-scale of the epidemic implied by the model.

The infection fatality rate implied by these parameters is $\eta\nu = 0.005$. This parameter is key for translating the model's implications for infections into implications for deaths. While this model is simplified in that it does not consider the age distribution of infections, it does do fairly well in matching the overall estimates of the burden of COVID provided by the CDC.[6] The decomposition of this infection fatality rate has fraction $\nu = 0.2$ of those hospitalized dying (again see the CDC model scenario page cited above for data on the fatality rate conditional on hospitalization) and thus fraction $\eta = 0.025$ resulting in hospitalizations. The decomposition of the infection fatality rate into a rate of hospitalization and a fatality rate conditional on hospitalization does not impact the prediction of the model for daily deaths, but is relevant if one were to attempt to use the model to match data on hospitalizations.

The second set of parameters is comprised of the rate ζ at which those hospitalized flow either to death or recovery. This rate is chosen to have an average stay in compartment H of 30 days, which corresponds to an average stay in the hospital of two weeks for those with serious illness and a reporting delay of deaths of two weeks.[7] This assumption combined with the average generation time assumed above

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

[6]The CDC estimates that 83 million Americans had been infected by the end of December 2020. Total COVID deaths reached 445,000 30 days later, giving an average estimated infection fatality rate including the delay from infection to death of slightly over 0.005. See https://www.cdc.gov/coronavirus/2019-ncov/cases-updates/burden.html

[7]See https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html. On that web-

implies that roughly six generations of infections pass on average before reported daily deaths rise provoking a behavioral response.

The delay between infection and reported death introduced by this compartment H implies that this simple behavioral model has oscillatory endogenous dynamics that are helpful in allowing the model to reproduce the data with only a few shocks. To illustrate the role of this delay in generating endogenous model dynamics, in Figure A2, I show the implications of the model with all the parameters set as in Figure A1 except that I set $\zeta = 10$ so that the average stay in the H compartment is only a few hours. In comparing Figure A2 below and Figure A1, we see that if we have only a short stay in the H compartment, the model does not match the waves in the first six months of the epidemic while the baseline model shown in Figure A1 does.

page, the CDC notes a median time from symptom onset to death of approximately two weeks and a median time from death to reporting just under three weeks.

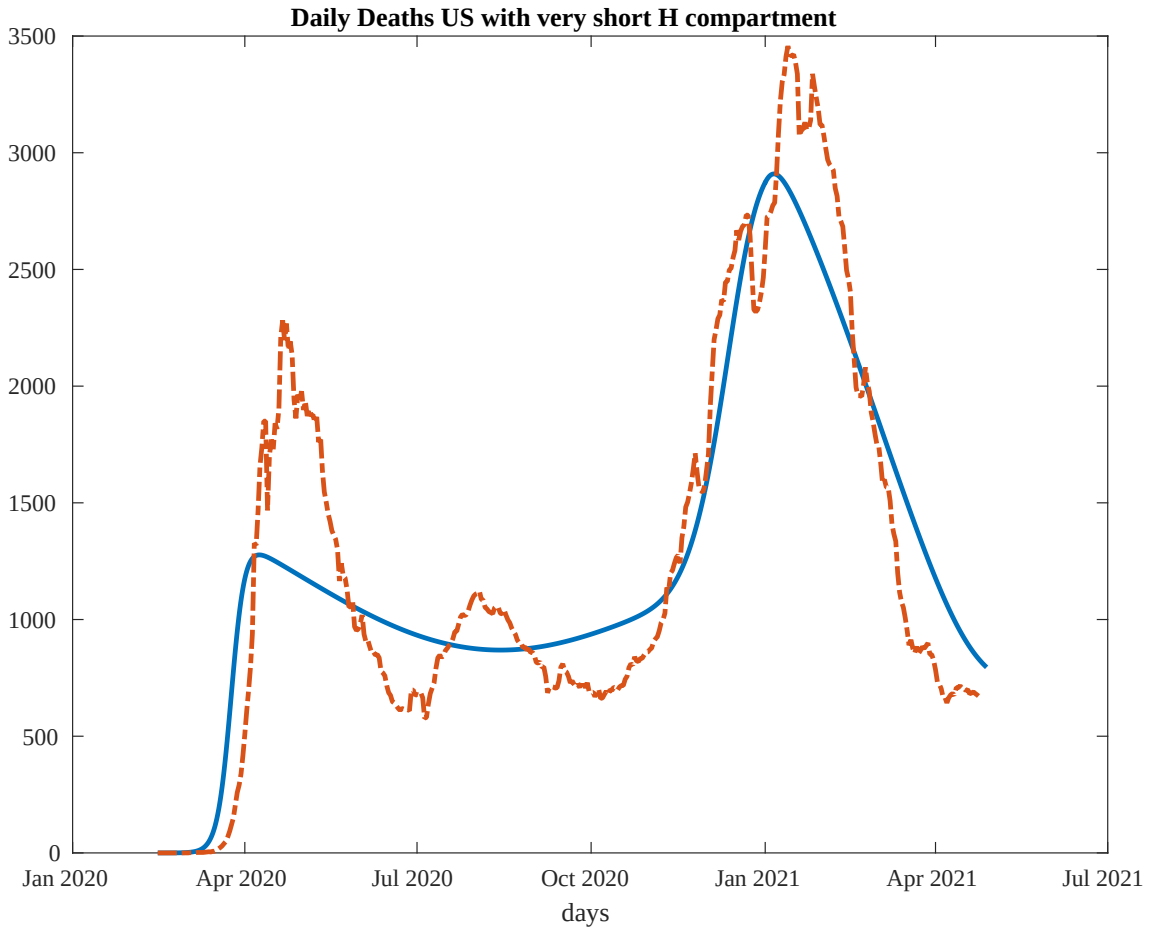

Figure A2: Model implications with a very short duration of the H compartment

The third set of parameters are those that I choose through a process of trial and error to fit the model to data from the first six months of the epidemic in the United States. These are the basic reproduction number, the original semi-elasticity of transmission with respect to daily deaths, and the size of the seasonal fluctuation in transmission rates. Note from Figure A1 that the model fits the path of daily deaths quite well over the first six months of the epidemic with the only shock being the seasonality in transmission. The assumed shocks for pandemic fatigue, the new variant, and vaccines, do not impact the model's predictions over these first six months.

The basic reproduction number of the virus at peak transmissibility is $\mathcal{R}_0(t) = \bar{\beta}/\gamma$ for the original virus and $\mathcal{R}_0(t) = \bar{\beta}_v/\gamma$ for the new variant. For the United States, I set $\bar{\beta} = 3\gamma$ giving a peak basic reproduction number at pre-pandemic levels of behavior in Winter of 3. This number is well within the range of estimates of this parameter from the early phase of the pandemic.[8] As discussed above, this parameter was chosen by trial and error. In Figure A3 I show how the model implications for the path of daily deaths varies if this parameter is increased or decreased by 20% (from 3.6 to 3 to 2.4). As can be seen in this figure, increasing this basic reproduction number increases the speed and height of the first wave of deaths.

[8]See https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html.

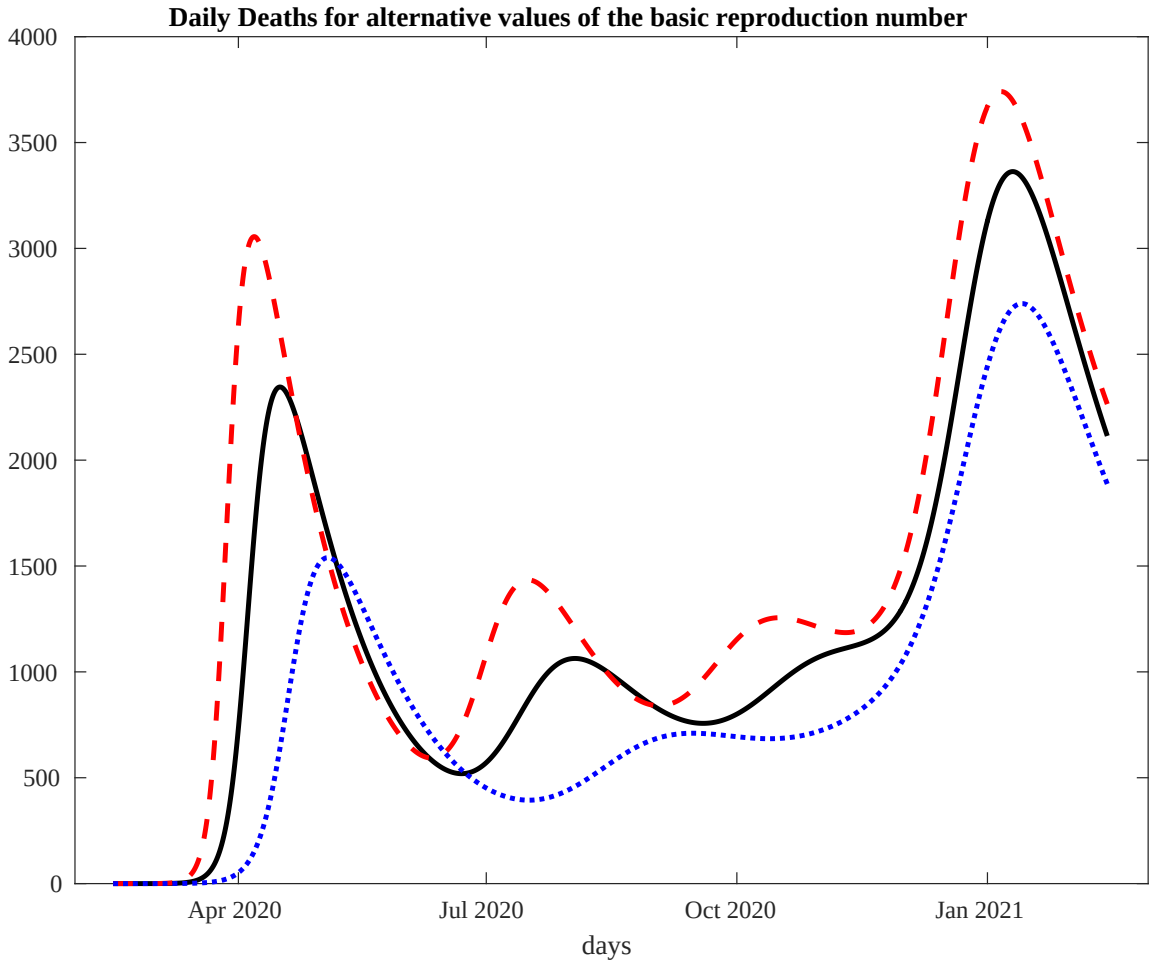

Figure A3: Implied path for daily deaths with the basic reproduction number $\bar{\beta}/\gamma = 3$ and with this parameter set 20% above and below this baseline value.

The initial semi-elasticity of transmission with respect to daily deaths (measured as a fraction of the population) for the United States is $\bar{\kappa} = 250000$. This parameter value was also chosen through trial and error. In Figure A4, I show the implications of the model with this initial semi-elasticity set to this baseline value and to values 20% above and below this baseline value. As can be seen in this figure, this semi-elasticity primarily impacts the level of the daily deaths in a roughly parallel manner.

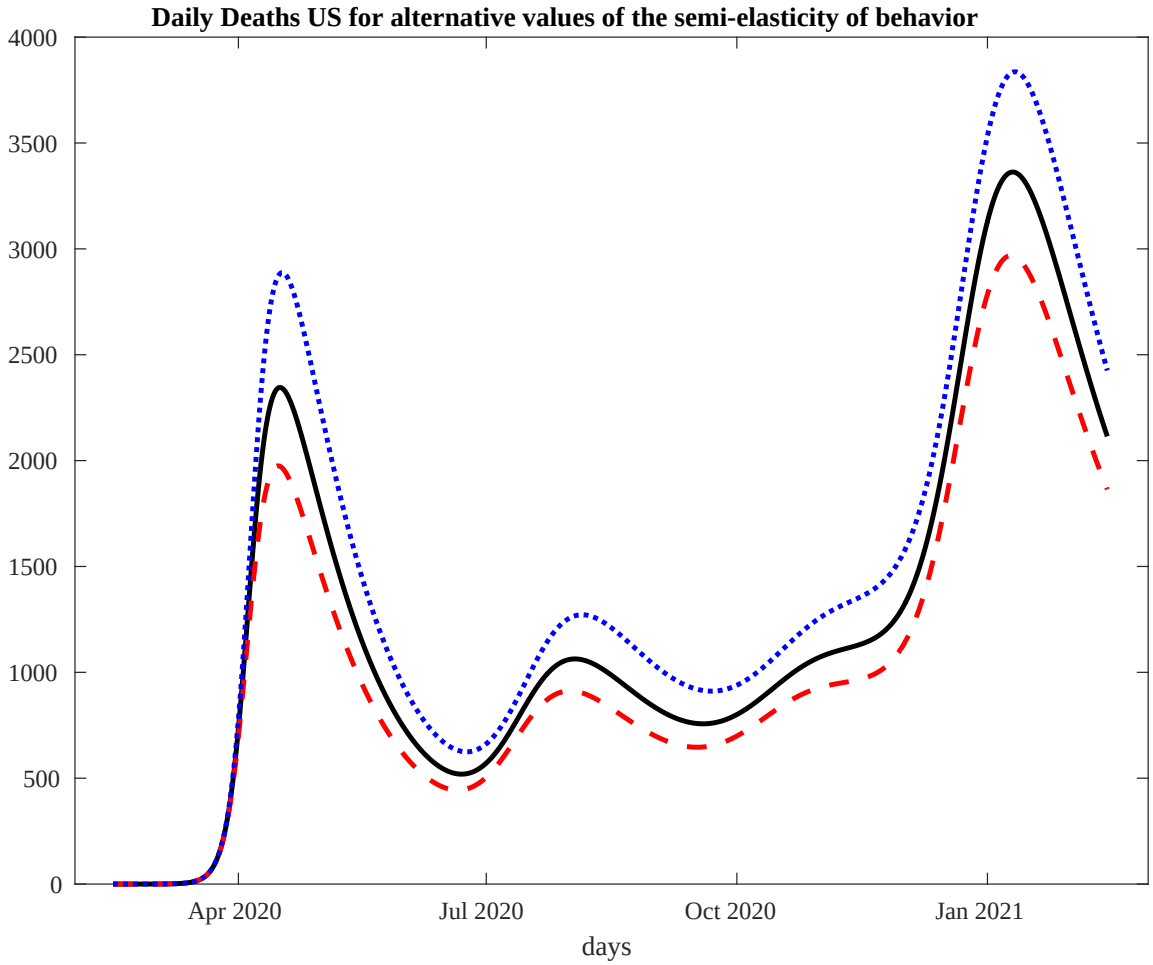

Figure A4: Implied path for daily deaths with the semi-elasticity $\bar{\kappa} = 250000$ and with this parameter set 20% above and below this baseline value.

To model seasonality of transmission in the United States, I set $seasonalsize = 0.35$ and $seasonalposition = 20$. This seasonal variation in the parameter $\psi(t)$ leads to variation over time in the basic reproduction number of the virus. Figure A5 shows the basic reproduction number for the original variant corresponding to no reduction in transmission due to a behavioral response (the transmissibility of the virus with behavior at pre-pandemic patterns of behavior). We see that the assumed pattern for seasonality introduces a reduction in transmissibility of the virus by a factor of $\exp(-0.35)$ holding behavior fixed from the winter peak in late January and

the summer low in late July.

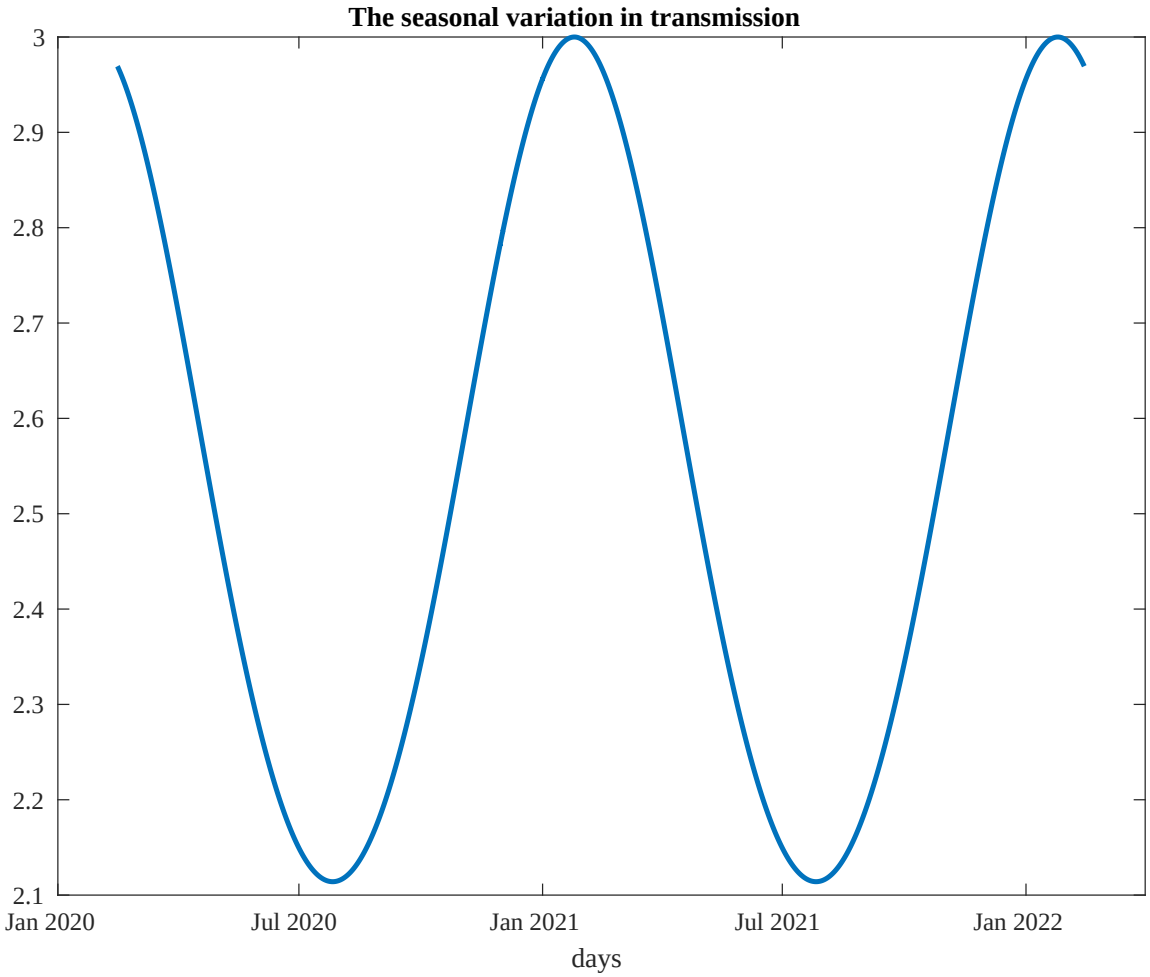

Figure A5: Assumed seasonality in the basic reproduction number.

In Figure A6 I show the implications of the baseline model from Figure A1 but with no seasonality in transmission. By comparing this Figure A6 to Figure A1, we see that seasonal variation in transmission is important for matching the drop in daily deaths observed over the summer.

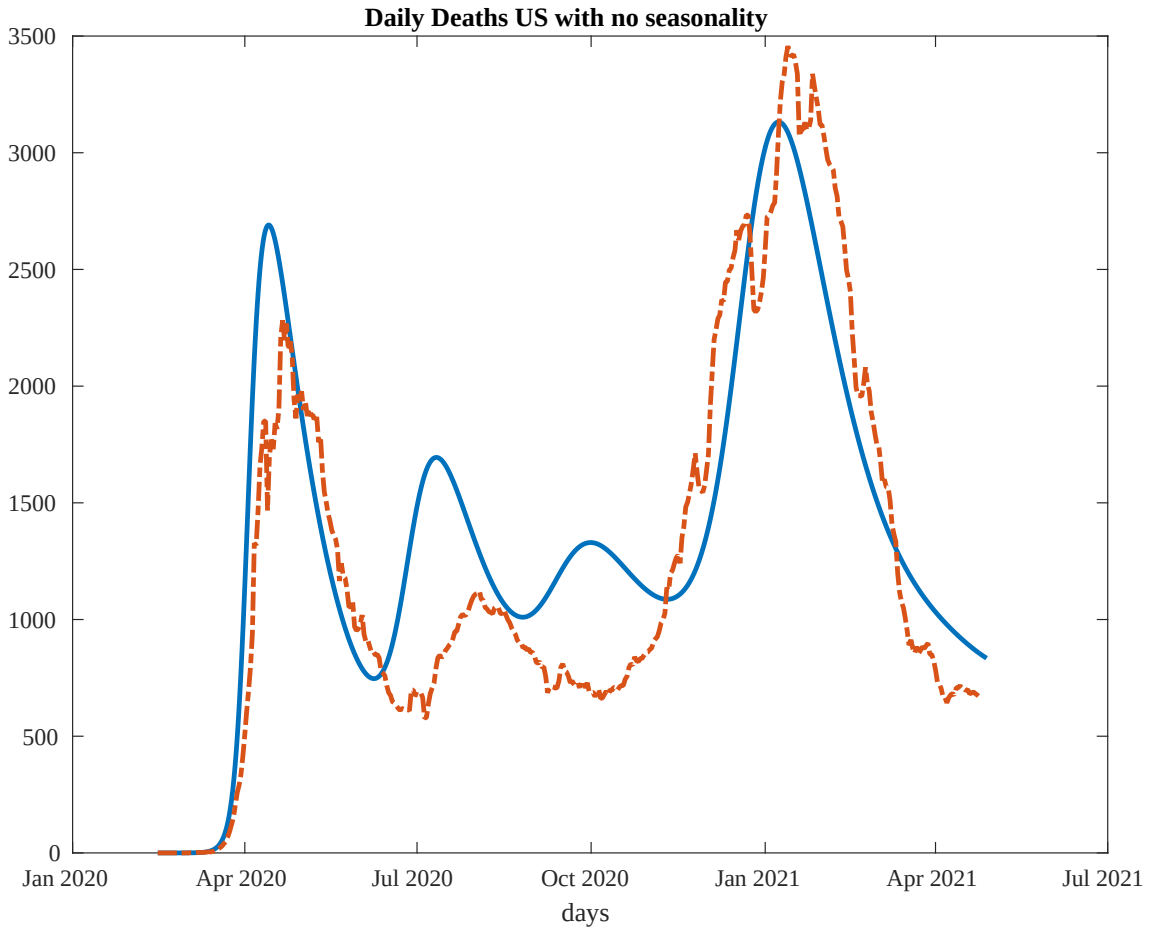

Figure A6: Implications of the model with vaccines, the new variant, and pandemic fatigue, but no seasonality in transmission.

The magnitude of the seasonal fluctuation in transmission was chosen by trial and error. While the seasonal variation in transmission is important for matching the pattern of daily deaths in the data, the model's implications are fairly insensitive to the precise magnitude of the seasonal. In Figure A7, I show the implications of the model for daily deaths with the size of the seasonal variation in transmission set to its baseline value and 20% above and below this value.

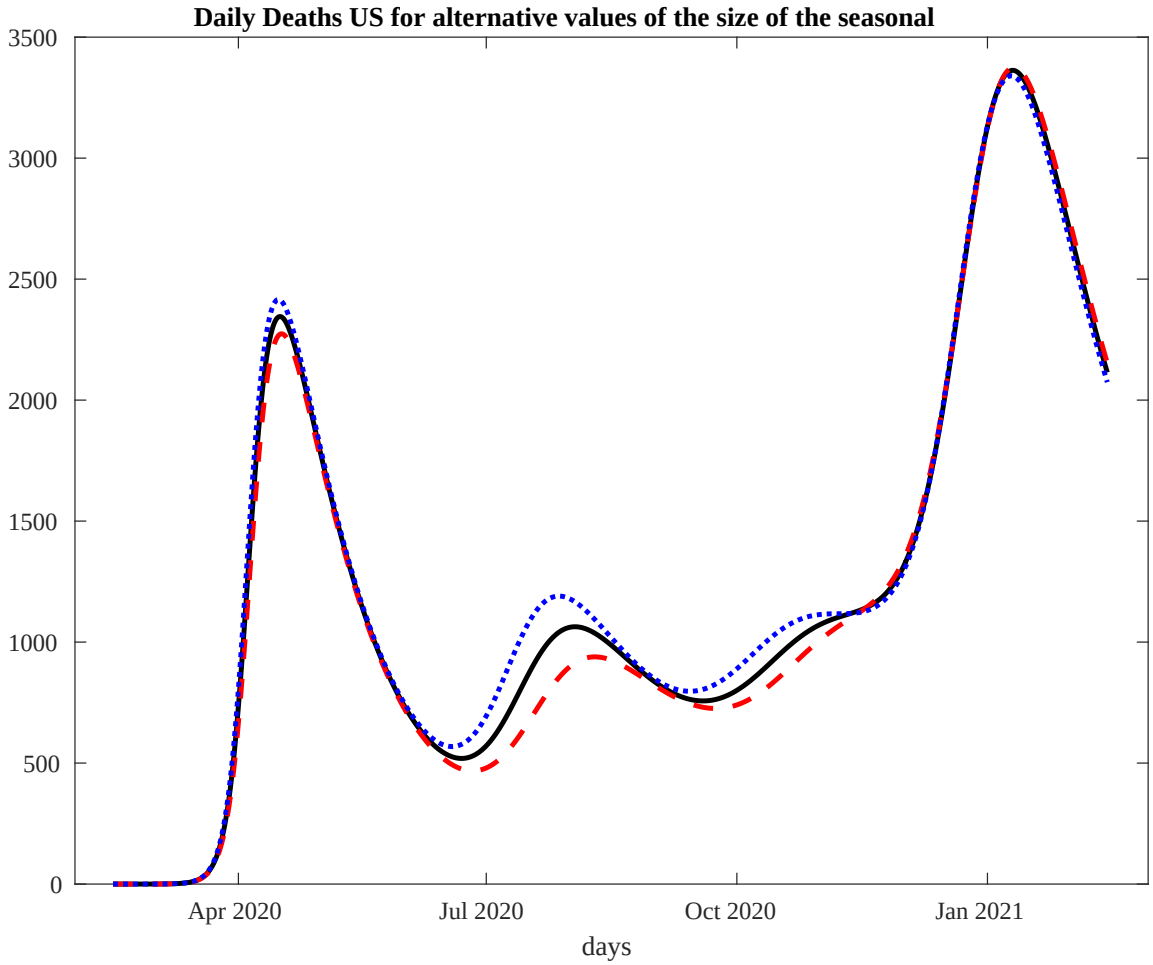

Figure A7: Implications of the model with the size of the seasonal variation in transmission 20% above and below its baseline value.

3 Shocks

The model has several shocks in addition to the seasonal fluctuation in transmission rates. These include pandemic fatigue, the introduction of a new variant on December 1, 2020, the introduction of vaccines, and the counterfactuals of extra mitigation measures that are not dependent on disease prevalence and of waning immunity. I

describe the specification and impact of these shocks and counterfactual experiments here.

The initial semi-elasticity of transmission with respect to daily deaths (measured as a fraction of the population) for the United States is $\bar{\kappa} = 250000$. To model the onset of pandemic fatigue in the United States, I set $fatiguesize = 0.375$, $fatiguemean = 285$ and $fatiguesig = 15$. Figure A8 shows the the semi-elasticity of the transmission rate with respect to the level of daily deaths. We see in that figure that this semi-elasticity is assumed to fall to 37.5% of its original level late in the year.

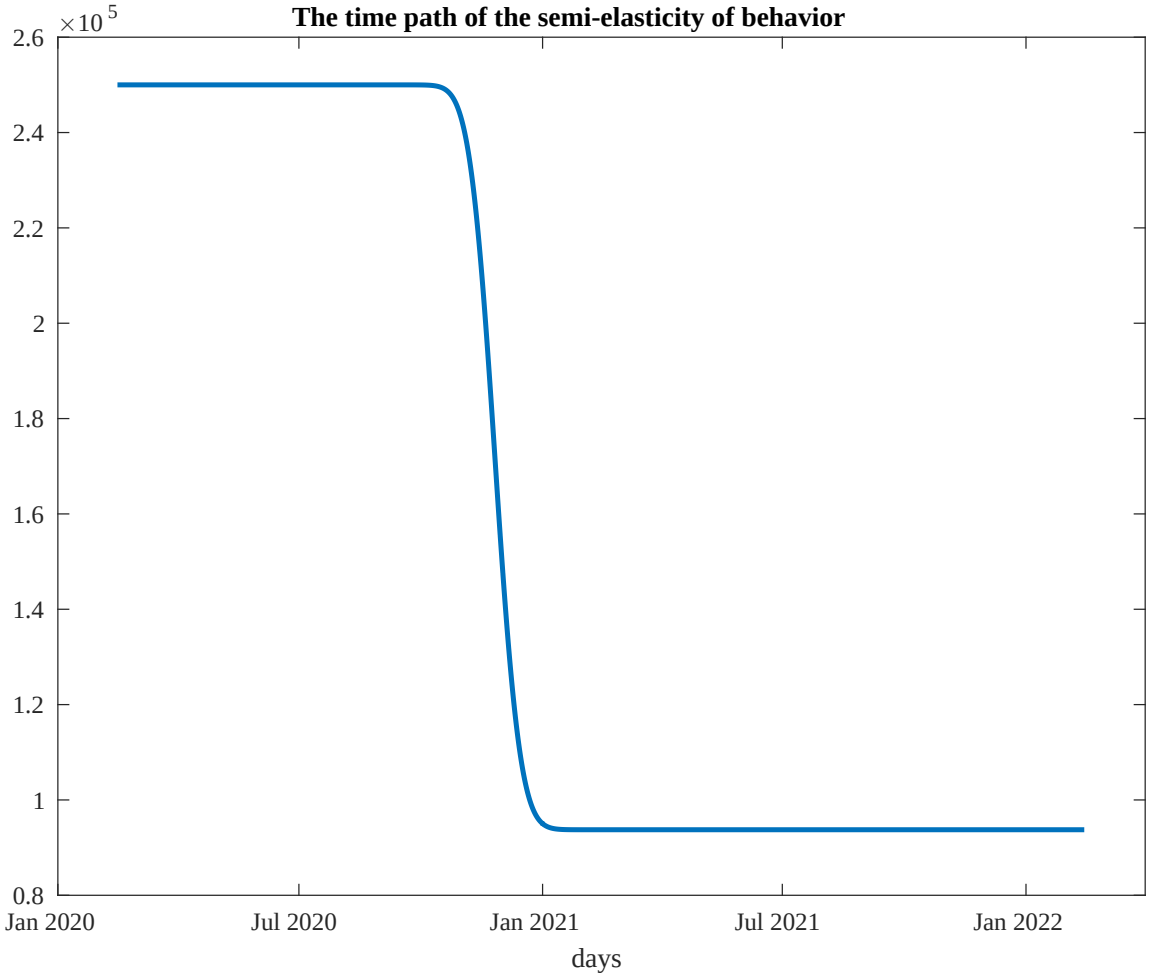

Figure A8: Assumed pandemic fatigue. The blue line shows the evolution of the semi-elasticity of transmission with respect to daily deaths relative to its initial level.

In Figure A9 I show the implications of the baseline model with no pandemic fatigue. By comparing this figure with Figure A1, we see that pandemic fatigue is quite important in generating the large wave of deaths in the winter of 2020. Essentially, the initial behavioral response in the model is too strong to allow for a significant winter wave even with the seasonal rise in transmission.

Note from Figure A8 that I leave the semi-elasticity of transmission with respect to behavior at its new, lower level, throughout the remainder of the Spring of 2021.

The match between the model in this time period shown in Figure A1 suggests that the behavioral reaction of transmission to the reduction in daily deaths induced by the introduction of vaccines has remained as it was at the end of 2020 through the first three months of 2021.

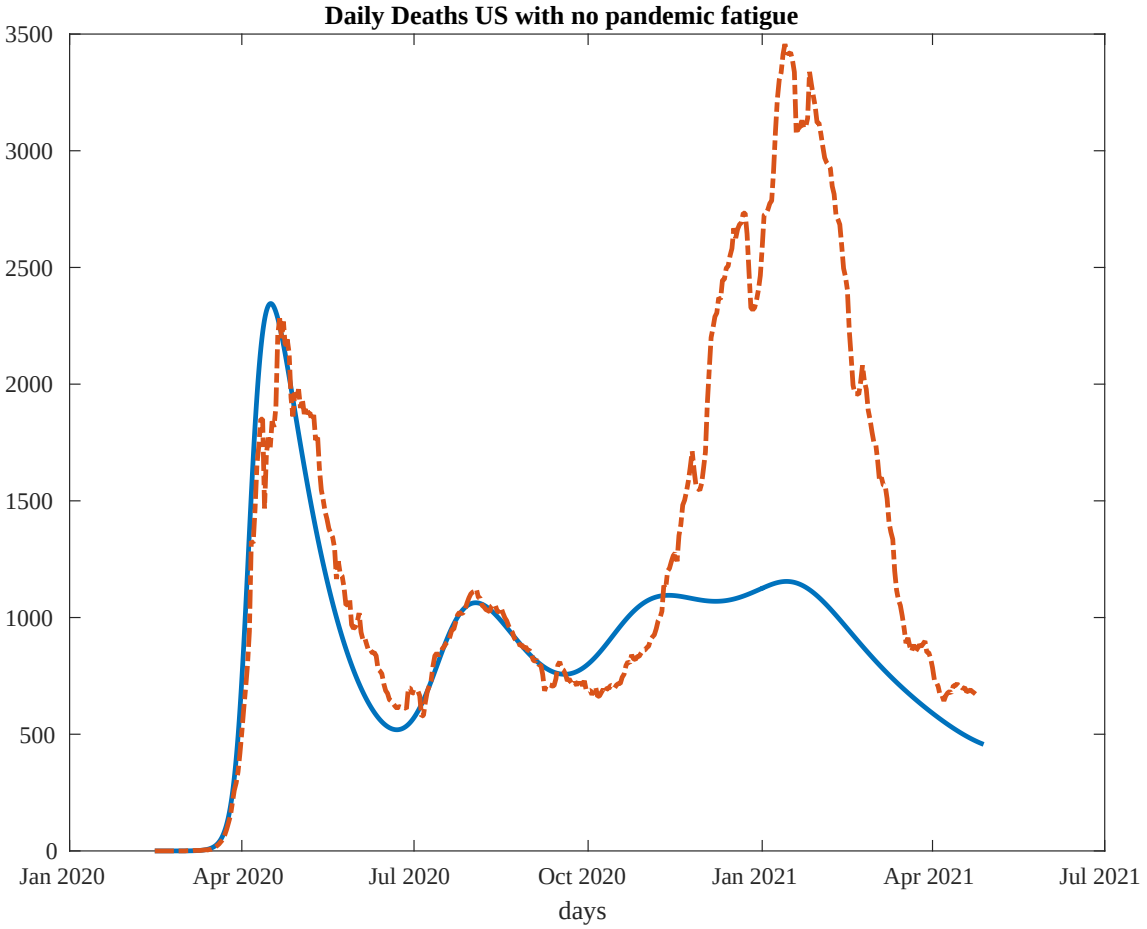

Figure A9: Implications of the model with no pandemic fatigue

To model the transmissibility of the new variant of the virus, I set $\bar{\beta}_v = 4.5\gamma$ giving a basic reproduction in Winter of 4.5. This implies that the new variant is 50% more transmissible than the original variant at any given level of activity. Note that the other epidemiological parameters associated with this new variant are assumed to stay the same, including the infection fatality rate.

The assumption that the new variant of the virus is 50% more transmissible than the original variant of the virus controls how fast it takes over as the dominant variant among those currently infected. In Figure A10 I show the model implications for the fraction of current infections that are due to the new variant of the virus. As can be seen in this figure, under our baseline simulation of the model with vaccines, the new variant becomes dominant quickly in the spring of 2021 after starting from a single infection on December 1, 2020.[9]

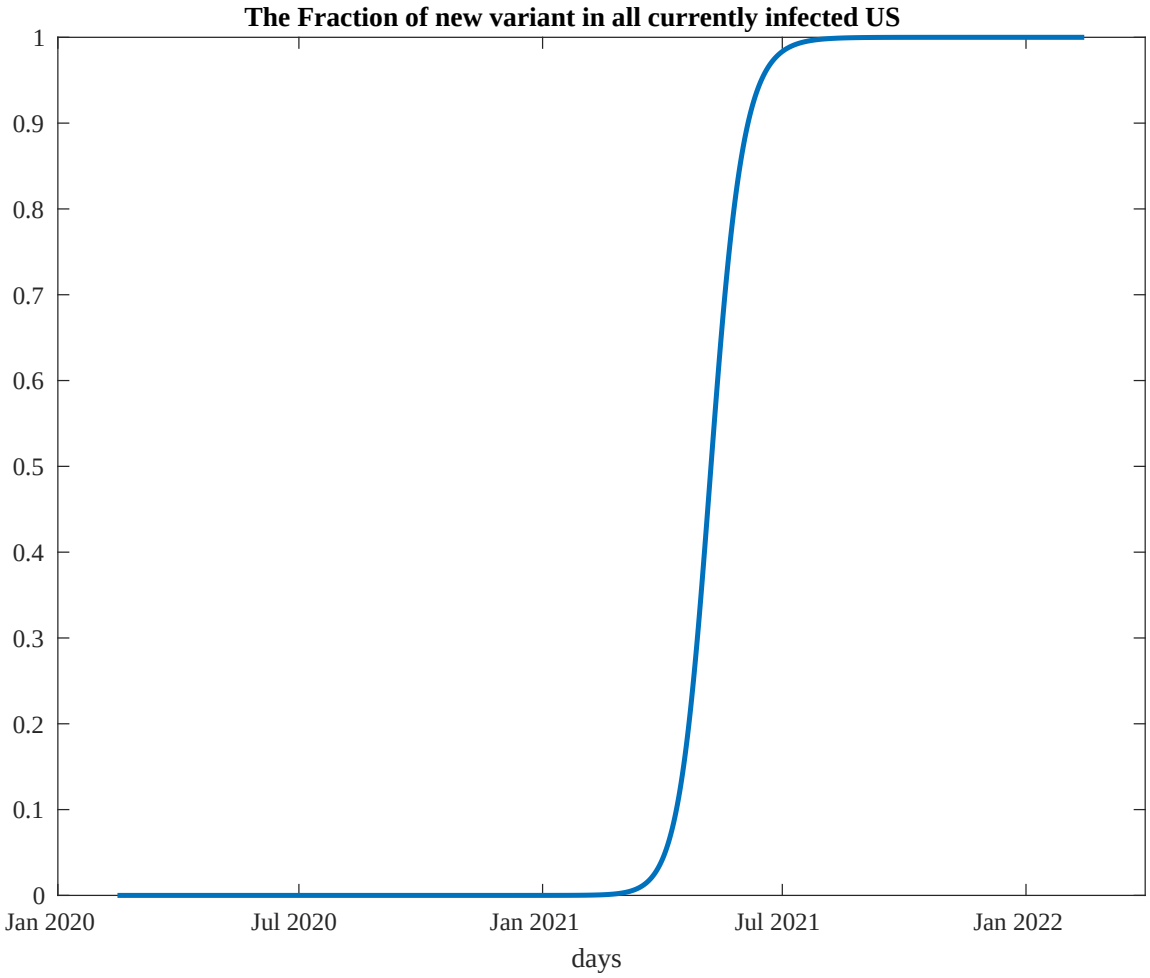

Figure A10: Current infections with the new variant as a fraction of total current infections

[9]CDC estimates of the prevalence of variants this Spring is available here https://covid.cdc.gov/covid-data-tracker under the page labelled "Genomic Suvelliance: variant proportions".

In simulations in which I include a vaccine, I set $\lambda(t) = 0.004$ starting on January 1, 2021. I assume that vaccinations are offered to the general population. This implies that in a population of 330 million, the daily number of vaccines administered is close to 1.3 million. In comparing this number to data on vaccinations, one must take into account that most of the vaccines administered require two doses for full effect. This assumption implies that roughly 51% of the population is fully vaccinated by July 1, 2021. Note that Figure 1 in the main text shows the implications of the model with the baseline parameters but no vaccines. A comparison of Figures 1 and 5 in the main text suggest that the introduction of vaccines has had an important impact on the death toll from COVID in March and April of 2021.

In simulations in which I assume that the Federal government takes steps to mitigate the transmission of the virus through a program such as testing and tracing that is not contingent on the level of daily deaths but instead is contingent on fixed dates, I subtract 0.5 from the seasonal pattern in $\psi(t)$ for those dates. The implications of the model with this additional mitigation are shown in the main text in Figures 3 in a simulation with no vaccines. The implications of the model with additional mitigation and the introduction of vaccines are shown in Figure A11. As is evident in this figure, the extra mitigation substantially reduces the death toll in the first year of the epidemic. As discussed in the main text, in this simulation, the predicted cumulative death toll over a five year period is 302,000.

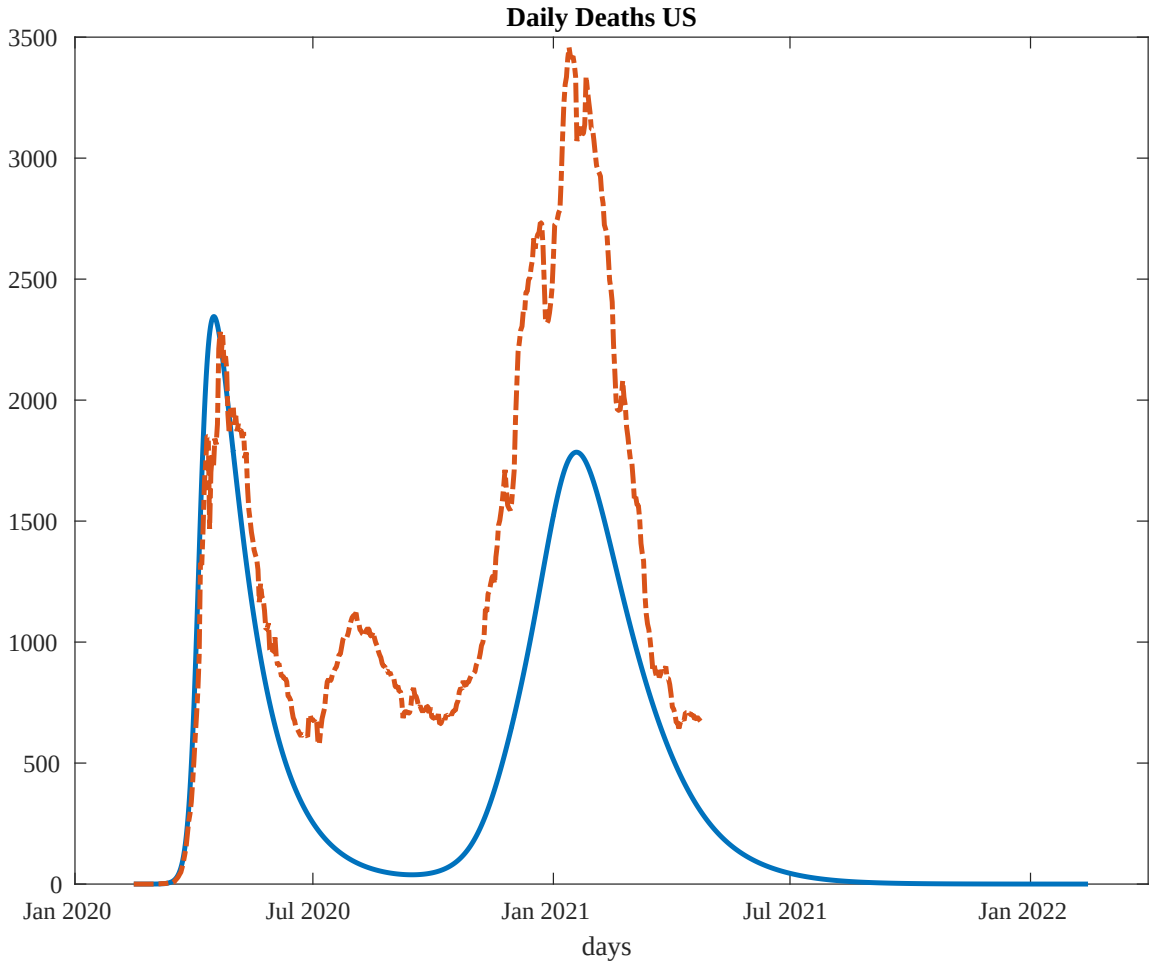

Figure A11: Implications of the model with extra mitigation and vaccines

In simulations in which I assume waning immunity, I set $\xi = 1/547.5$ which corresponds to an expected time to loss of immunity of 18 months. In these simulations, I also assume a higher rate of transmission for the new variant. In this simulation shown in Figure 5 in the main text, I set $\bar{\beta}_v/\gamma = 5$ rather than 4.5.