



Working Paper Series

Social Distancing, Quarantine, Contact Tracing, and Testing: Implications of an Augmented SEIR Model

WP 20-04

Andreas Hornstein
Federal Reserve Bank of Richmond

This paper can be downloaded without charge
from: <http://www.richmondfed.org/publications/>



Richmond • Baltimore • Charlotte

Social Distancing, Quarantine, Contact Tracing, and Testing: Implications of an Augmented SEIR Model *

Andreas Hornstein
Federal Reserve Bank of Richmond
andreas.hornstein@rich.frb.org

First Version March 25, 2020
This Version May 8, 2020

WORK IN PROGRESS
COMMENTS WELCOME

Abstract

I modify the basic SEIR model to incorporate demand for health care. The model is used to study the relative effectiveness of policy interventions that include social distancing, quarantine, contact tracing, and random testing. A version of the model that is calibrated to the Ferguson et al. (2020) model suggests that permanent, high-intensity social distancing reduces mortality rates and peak ICU demand substantially, but that a policy that relaxes high-intensity social distancing over time in the context of a permanent efficient quarantine regime is even more effective and also likely to be less disruptive for the economy. Adding contact tracing and random testing to this policy further improves outcomes. However, given the uncertainty surrounding the disease parameters, especially the transmission rate of the disease and the effectiveness of policies, the uncertainty for health outcomes is very large.

***Important qualification: I am an economist and not an epidemiologist, so take anything stated here with a very large grain of salt.** This revision corrects an algebra mistake in previous versions that significantly affected results. I would like to thank Alex Wolman and Zhilan Feng for helpful comments and Elaine Wissuchek for research assistance. Any opinions expressed are mine and do not reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

Contents

1	Introduction	3
1.1	Related work in epidemiology	5
1.2	Related recent work by economists using SIR-type models	5
2	The basic SEIR model	6
3	An extended SEIR model with hospitalizations and death	7
4	Calibration	10
5	Experiments	12
5.1	Effectiveness of social distancing	13
5.2	Effectiveness of quarantine	14
5.3	Effectiveness of combined policies	15
5.4	Implications for employment	17
6	Caveats	18
6.1	Higher basic reproduction rate	18
6.2	Uncertainty	20
7	Conclusion	21
A	Appendix	24
A.1	Reproduction rates	24
A.1.1	Basic reproduction rate \mathcal{R}_0 for SIR model	24
A.1.2	Basic reproduction rate in \mathcal{R}_0 for SEIR model	24
A.1.3	New infections with quarantine	25
A.1.4	Probability of recovery without developing symptoms	26
A.2	Social distancing	26
A.3	Seeding the initial condition	27
A.4	Representing parameter uncertainty	27

1 Introduction

So far the primary response to the coronavirus pandemic, high-intensity social distancing, has been extremely disruptive for any economy where it has been applied. The question becomes whether the response can be maintained for an extended time without large negative effects for social, economic, and health outcomes. If high-intensity social distancing cannot be a permanent response to limit the spread of the coronavirus, then it is likely that the fallout of the pandemic might be dampened now but ultimately only delayed. Or are there alternative policy options that would be less disruptive for the economy but still contain the spread of the disease?

In trying to come up with an answer, I have to acknowledge that I am not an epidemiologist and very likely do not have a full appreciation of the literature. So I may be reinventing the wheel.

Ferguson et al. (2020) study the possible containment of the virus in a large-scale pandemic model emphasizing social distancing. Shen, Taleb and Bar Yam (2020) argue that this approach omits effective methods, such as testing for the virus and tracing contacts of known infected individuals. Modeling these methods could reduce the number of predicted deaths. To evaluate this criticism, I modify a simple susceptible-exposed-infected-recovered (SEIR) model to provide a stylized version that abstracts from all the demographic detail of the model of Ferguson et al. (2020). The model includes asymptomatic and symptomatic individuals who spread the disease and hospitalized individuals who require more or less intensive medical care. Symptomatic individuals are assumed to be known and can be quarantined. Furthermore, previously infected contacts of newly symptomatic individuals can be traced, and some can be quarantined too. Finally, random tests can be performed on the general population to find asymptomatic but infectious individuals. As in the standard SEIR model, health-state changes follow Poisson processes. The model is calibrated based on information in Ferguson et al. (2020).

With a baseline infection fatality rate of about 1 percent, the consequences from no intervention are dire: about 1 percent of the population is at risk of dying. For the UK that means about 600 thousand deaths, and for the US it means about 3.25 million deaths. I consider various interventions that involve social distancing, quarantine, contact tracing, and random testing to ameliorate this outcome. For the calibrated stylized model, I find that

- high-intensity social distancing (SD) is effective in the sense that it lowers cumulative deaths to less than 0.1 percent of the population, but it is only effective if it is permanent;

- permanent efficient quarantine is less effective than SD, it lowers cumulative deaths to 0.25 percent of the population, but when augmented with an efficient tracing process for previous contacts of newly symptomatic individuals, it is about as effective as permanent high-intensity SD;
- combining permanent high-intensity quarantine with a gradual relaxation of high-intensity SD is noticeably more effective than a policy of permanent high-intensity SD. At the same time this combination is presumably less disruptive for the overall economy and likely to reduce employment by less;
- adding contact tracing or random testing to the combination of permanent quarantine and gradual relaxation of SD further improves outcomes, but more for tracing than for testing.

To summarize, for a simple SEIR model that is calibrated to the Ferguson et al. (2020) study, there are alternative policies to permanent SD that provide health outcomes that are at least as good and potentially less disruptive. All of these policies attempt to reduce the rate at which the disease spreads, a summary statistic of which is the basic reproduction rate. Independent of whether the simple SEIR model is appropriate, there is a large degree of uncertainty associated with the effectiveness of any of these policies in the model. Most of this uncertainty is related to what we do (not) know about the parameters that characterize the spread of the disease. In a robustness analysis, I find that

- the model cannot match the sharp increase in cumulative deaths observed for the US and UK from late March to mid-April 2020 if it is parameterized to widely used estimates of the basic reproduction rate;
- the model can match the sharp increase in cumulative deaths if more recent estimates of higher reproduction rates are used, but for this case all policies become correspondingly less effective;
- more generally, given the large uncertainty surrounding parameter estimates for the disease process, the uncertainty about health outcomes predicted by the model is equally large. In the model the main driver of this outcome uncertainty is the uncertainty surrounding the basic reproduction rate.

One can have well-founded reservations on the use of the kind of model described here for policy analysis, and Jewell, Lewnard and Jewell (2020) provide an extensive list of these reservations. On the other hand, short of running actual ‘experiments’ on an economy,

models like the one described here provide some guidance on possible outcomes for these policy interventions. Nevertheless, predictions on the relative efficiency of policy measures should be interpreted in the context of other work and past experience.

1.1 Related work in epidemiology

We work with an augmented version of the standard SEIR model of disease diffusion with Poisson arrival rates for health-state changes and implied exponential distributions for stage duration. While analytically convenient, the assumption of constant hazard rates for transitioning between disease stages in a SEIR model leads to outcomes that do not match the actual spread patterns for many infectious diseases. For example, Wearing, Rohani and Keeling (2005) and Feng, Xu and Zhao (2007) argue that relative to the observed diffusion of infectious diseases, standard SEIR-type models for which health-state transitions follow Poisson processes understate peak infection periods and overstate the duration of the process. They suggest that SEIR-models with gamma distributions for the stage distributions provide a better match of actual disease diffusion. But Feng (2007) also argues that in the presence of policy interventions, like quarantine, this simple ranking of the disease process for exponential and gamma distributions may no longer hold. These qualifications should be kept in mind when interpreting the numerical results from our SEIR model.

Most epidemiological work on quarantine and contact tracing models these interventions as setting aside a fraction of newly infected individuals and gradually moving them to a quarantine state, similar to the transition between health states. The effectiveness of these interventions is then determined by the share and speed parameters, see for example Wearing et al. (2005) or Feng (2007). Lipsitch et al. (2003) use a similar approach to study the issue of contact tracing in the context of the SARS epidemic.

Compared to this epidemiological work, the approach taken here to model quarantine and tracing is more reduced form: a share of infected individuals is identified, and they are immediately quarantined, but only a fraction of quarantined individuals can be excluded from the infectious pool.

1.2 Related recent work by economists using SIR-type models

Eichenbaum, Rebelo and Trabandt (2020) study the impact of SIR-type dynamics on employment and output in a simple macro model with some endogenous response of meeting rates to the disease. Atkeson (2020) studies the impact of SD on deaths in a simple SIR-model. Alvarez, Argente and Lippi (2020) and Farboodi, Jarosch and Shimer (2020) study the optimal application of social distancing measures in a SIR model without and with an en-

ogenous response of individuals to the emergence of the disease. Fernandez-Villaverde and Jones (2020) estimate time-varying transmission rates in a SIR-model by matching observed time paths of cumulative deaths in different localities.

Piguillem and Shi (2020) and Berger, Herkenhoff and Mongey (2020) are closest to this paper. They study optimal quarantine and testing in a SEIR-type model but do not include contact tracing. Berger et al. (2020) use a time-delayed quarantine model similar to the standard epidemiological literature, whereas the quarantine model in Piguillem and Shi (2020) is similar to the one we are using. The calibration in neither paper is tied as closely to Ferguson et al. (2020) as this paper is. Stock (2020) discusses the limitations of random testing of the general population to obtain better estimates of the asymptomatic share in the population.

New papers on the implications of the coronavirus for the economy are appearing daily, so this survey is already outdated.

2 The basic SEIR model

Define the stock of susceptible population S , infected and infectious population I , and recovered population R . Total population is

$$N = S + I + R.$$

Individuals transition sequentially between the states determined by Poisson processes with given arrival rates. Assume that the disease transmission rate for a given encounter is α , that the recovery rate from the disease is γ , and that recovered individuals are immune to the disease. See Figure 1 (a) for a graphic representation.

Total disease transmission, M , following from meetings between the susceptible and infected population is then,

$$M = \alpha \frac{IS}{N}.$$

The dynamics of $x = (S, I, R)$ are described by the differential equations

$$\begin{aligned} \dot{S} &= -\alpha \frac{IS}{N} \\ \dot{I} &= \alpha \frac{IS}{N} - \gamma I \\ \dot{R} &= \gamma I. \end{aligned}$$

The growth rate of the infectious group is

$$\hat{I} = \left(\frac{\alpha S}{\gamma N} - 1 \right) \gamma.$$

Assume that the initial value for the population share of susceptible individuals when the process starts is essentially one, $S(0) \approx N$. Therefore the number of infected people is initially increasing if

$$\mathcal{R}_0 = \frac{\alpha}{\gamma} > 1.$$

The ratio \mathcal{R}_0 is called the basic reproduction number because it is approximately the average number of new infections before recovery from an infected individual at time zero,

$$\int_0^\infty \left[\alpha \frac{S(\tau)}{N} \tau \right] \gamma e^{-\gamma\tau} d\tau \approx \frac{\alpha}{\gamma} = \mathcal{R}_0,$$

where the first term in the integral is the average number of infections over a time interval τ and the second term is the probability of staying infectious for that time.

A standard extension of the SIR model places an exposed state that is not infectious, E , between the susceptible and the infectious group. This is called the SEIR model. Introducing the exposed state changes the dynamics of the model, e.g., it tends to change peak infection rates, but it usually does not affect terminal outcomes much. Let ϕ denote the rate at which exposed individuals become infectious, normalize the population at one, $N = 1$, and interpret the variables $x = (S, E, I, R)$ as population shares. Then the modified SEIR system is

$$\begin{aligned} \dot{S} &= -\alpha ES \\ \dot{E} &= \alpha ES - \phi E \\ \dot{I} &= \phi E - \gamma I \\ \dot{R} &= \gamma I. \end{aligned}$$

The system of differential equations is straightforward to solve, e.g., using MATLAB's ode45 routine starting with an initial condition $x_0 = (S_0, E_0, I_0, R_0)$.

3 An extended SEIR model with hospitalizations and death

I now extend the basic SEIR model to provide a stylized representation of the pandemic model in Ferguson et al. (2020). The pandemic model of Ferguson et al. (2020) contains a detailed description of the demographics of the population, its age distribution, locations, etc. Our stylized model will not contain any of that detail. What the model takes from Ferguson et al. (2020) is the basic mechanics of how the disease spreads from exposure to asymptomatic infection to symptomatic infection, hospitalization, and finally recovery or death. This abstraction makes it easy to explore the relative merits of various policy

measures, such as social distancing, quarantine, contact tracing, and random testing in a unified framework.

We start with the SEIR model. Susceptible individuals are exposed to the infection but are not immediately infectious. Exposed individuals become infectious, but they initially do not show any symptoms. After some time, asymptomatic infected individuals do show symptoms of the disease and are triaged depending on their condition. Most do not require hospitalization, but some do, in severe cases in ICUs. All infected individuals either recover over time and become immune, or they die.

Figure 1 (b) provides a graphic representation of this process. The stock of exposed individuals is E , the inflow of newly exposed individuals is M , and the rate at which exposed individuals become infectious without symptoms is ϕ . Asymptomatic individuals recover at rates γ , and they become symptomatic at rate β . For a fraction ω of newly symptomatic individuals, the condition is serious enough to be hospitalized. In addition, a fraction η of the hospitalized individuals require ICU treatment. Hospitalized individuals recover at rates γ respectively γ_{ICU} , and they die at rates δ respectively δ_{ICU} . Asymptomatic and symptomatic individuals who are not hospitalized also recover or die at rates γ respectively δ .¹

The following system of differential equations provides the formal representation of the process dynamics.

$$\begin{aligned}
\dot{S} &= -M \\
\dot{E} &= M - \phi E - q_{TE} - q_{FE} \\
\dot{I}_A &= \phi E - (\beta + \gamma) I_A - q_{TA} - q_{FA} \\
\dot{E}_T &= q_{TE} + q_{FE} - \phi E_T \\
\dot{I}_{AT} &= q_{TA} + q_{FA} + \phi E_T - (\beta + \gamma) I_{AT} \\
\dot{I}_S &= (1 - \omega)\beta(I_A + I_{AT}) - (\gamma + \delta)I_S \\
\dot{H}_B &= (1 - \eta)\omega\beta(I_A + I_{AT}) - (\gamma + \delta)H_B \\
\dot{H}_I &= \eta\omega\beta(I_A + I_{AT}) - (\gamma_{ICU} + \delta_{ICU})H_I \\
\dot{R} &= \gamma(I_A + I_{AT} + I_S + H_B) + \gamma_{ICU}H_I \\
\dot{D} &= \delta(I_S + H_B) + \delta_{ICU}H_I
\end{aligned}$$

¹Total deaths are small enough such that the implicit assumption of a constant population is not too distorting.

The flow terms q_{TE} , q_{TA} , q_{FE} , and q_{FA} , and the stocks E_T and I_{AT} refer to the identification of exposed and asymptomatic individuals through tracing and/or random testing discussed below.

Policy interventions, such as social distancing and quarantining known infected individuals, are modeled through their impact on the flow of new infections. As in the basic SIR model, the flow of new infections is proportional to the product of susceptible individuals and infectious individuals, but quarantine can reduce the number of infected individuals who can meet the susceptible population. We assume that symptomatic individuals are always known and that tracing and random testing can identify some of the exposed and asymptomatic individuals, E_T and I_{AT} . Let ε_i denote the effectiveness of quarantine for the known infected population groups, $i \in \{S, B, ICU, AT\}$, and also assume that symptomatic infected are more infectious than asymptomatic infected at the rate σ , then the effective pool of infectious individuals that meets the susceptible population and the inflow of newly infected individuals are²

$$\begin{aligned} I^* &= I_A + (1 - \varepsilon_{AT})I_{AT} + \sigma \left[(1 - \varepsilon_S)I_S + \sum_{i=B,ICU} (1 - \varepsilon_i)H_i \right], \\ M &= \alpha SI^*. \end{aligned}$$

Social distancing is assumed to directly reduce the rate at which individuals, infectious and susceptible, contact each other. Let ψ denote the relative contact rate for an individual, that is, $\psi \leq 1$ and $\psi = 1$, in the absence of SD. Then the transmission flow is

$$M = \alpha_0(\psi S)(\psi I^*) = \alpha_0 \psi^2 SI^*,$$

where α_0 is the disease transmission rate without any SD measures. In the following we will use $\alpha = \alpha_0 \psi^2$ as the effective transmission rate.

Social distancing is thus potentially a very effective way to contain the spread of the disease since a reduction of contact rates applies to all individuals, infectious and non-infectious. Therefore a reduction of contact rates implies a squared reduction of transmission rates. Social distancing is also ‘easy’ to implement since all individuals are supposed to reduce their contact rates, that is, no particular information is required. This indiscriminate reduction of contact rates also makes SD very disruptive for the economy.

Quarantine methods on the other hand target individuals who are infectious, that is, they require information on an individual’s health status. As long as the health status is

²This is a simplified version of the quarantine model used in the epidemiological literature in the sense that identified people are added instantaneously to the quarantine pool, but some infections seep out of that pool. The epidemiological literature I am aware of assumes that infected individuals join the quarantine pool gradually following a Poisson process, but then quarantine is perfect. For example, Feng (2007).

observable, that is, for symptomatic individuals, it is relatively straightforward to implement, though not costless. The problem with a disease like COVID-19 is that a large share of infectious individuals, current estimates are around 50 percent, may never show symptoms. Thus even if one were able to quarantine all symptomatic individuals, one would only be able to reduce the pool of infectious individuals by 50 percent. On the other hand, quarantine is somewhat more efficient than that since symptomatic individuals are presumably more infectious than asymptomatic individuals. Contact tracing and random testing are attempts to reduce the pool of infectious individuals even more.

Tracing of asymptomatic infected individuals is modeled as follows. The average number of people an asymptomatic individual has infected and who are still in the exposed resp. asymptomatic state when he or she becomes symptomatic is \mathcal{R}_{ATE} resp. \mathcal{R}_{ATA} , derived in the Appendix. If ε_T is the efficiency of tracing, then the inflow of newly identified exposed and asymptomatic individuals through tracing is

$$q_{TE} = \varepsilon_T \mathcal{R}_{ATE} \beta I_S \text{ and } q_{TA} = \varepsilon_T \mathcal{R}_{ATA} \beta I_S.$$

We essentially assume that tracing does not require time, but is instantaneous.³

Testing is modeled as follows. Let f be the flow rate at which not yet identified asymptomatic people are randomly tested. Assume that asymptomatic infected can be identified through tests, but not merely exposed individuals. Also assume that recovered individuals are not tested. Then the share of identified asymptomatic in a random test is⁴

$$p_F = \frac{I_A}{S + E + I_A}.$$

The inflow of newly identified exposed and asymptomatic individuals through random testing is

$$q_{FA} = p_F f (1 + \varepsilon_T \mathcal{R}_{ATA}) \text{ and } q_{FE} = p_F f \varepsilon_T \mathcal{R}_{ATE},$$

where we allow for the possibility that previous contacts of newly identified asymptomatic individuals are then also traced.

4 Calibration

I parameterize the model following Ferguson et al. (2020) as much as possible, that is, unless otherwise noted all listed statistics are from Ferguson et al. (2020). The unit time interval is a year.

³It is straightforward to introduce a time delay for the recovery of tracked individuals. Again, we model the efficiency of tracing not through the rate at which potentially traceable individuals enter the quarantine pool, but through the size of the captured pool, see footnote 2.

⁴This potentially overstates the effectiveness of random testing with incomplete quarantine to the extent that the infectious pool also contains symptomatic individuals.

- The basic reproduction rate is $\mathcal{R}_0 = 2.4$. This estimate is consistent with the assessment of Fauci, Lane and Redfield (2020).
- The incubation period is 5.1 days, $\phi = 1/(5/365)$.
- Symptomatic infections are 50% more infectious than asymptomatic infections, $\sigma = 1.5$
- 4.4 percent of newly symptomatic infected are hospitalized, $\omega = 0.044$
- 30 percent of hospitalized infected require ICU, $\eta = 0.3$
- The mean duration of a hospital stay is 10.4 days
 - Non-ICU for 8 days, $\gamma_B = 1/(8/365)$
 - ICU for 16 days, of which 10 days are on ICU. We set $\gamma_{ICU} = 1/(16/365)$, which overstates the time ICU requirement by about 50 percent.
 - We set the recovery rates of non-hospitalized infected to the same as the one of non-ICU hospitalized, $\gamma = \gamma_B$
- 50 percent of infected in ICU die, $p_{D,ICU} = 0.5$. In the appendix we derive the probability for death in ICU, $P_{D,ICU}(\delta_{ICU}, \gamma_{ICU})$. We can solve $p_{D,ICU} = P_{D,ICU}(\delta_{ICU}, \gamma_{ICU})$ for δ_{ICU} .
- 40 percent to 50 percent of infected are never identified, mainly because they are asymptomatic, $p_{AR} = 0.5$. In the Appendix we derive the probability that an asymptomatic infected recovers before showing symptoms as a function of the rate of becoming symptomatic, and the recovery and death rates, $P_{AR}(\beta, \gamma, \delta)$. We can solve $p_{AR} = P_{AR}(\beta, \delta, \gamma)$ for β .
- The unconditional infection fatality ratio (IFR) is 0.9 percent, $p_I = 0.009$. We adjust the death rate for non-ICU infected, δ , such that the overall terminal fatality rate without intervention is close to p_I .
- Two-thirds of I_S self-isolate after one day, with a mean delay of five days. Since our quarantine does not involve any time delay, we assume that the baseline quarantine rate for non-hospitalized I_S is $\varepsilon_S = 1/3$.
- Quarantine: Baseline effectiveness for policy intervention is $\varepsilon_S = 0.5$, which is an average of the two options listed

- Case isolation at home (CI): I_S stay home for seven days, reduce contacts with non-household members by 75%. Compliance is 75%. $\varepsilon = 0.75 \times 0.75 = 0.6$
- Voluntary quarantine at home (VQ): All household members stay home for 14 days. Infection rate within households doubles, community contacts reduced by 75%. Compliance is 50%.
- Social distancing (SD) is assumed to reduce contact rates for workplace interactions by 25 percent and for social interactions by 75 percent. I use the 2018 American Time Use survey together with data on US employment rates to calculate the implications of these assumed reductions in contact rates for the average contact rate in the economy, Appendix A.2. The average contact rate ψ declines by about 60 percent, depending on what assumptions we make on the relative intensity of social and workplace interactions. This means that SD can reduce the transmission rate α and the reproduction rate \mathcal{R}_0 by about 80 percent.
- Finally, I made up the quarantine rate for hospitalized infected, $\varepsilon_i = 0.95$ for $i \in \{B, ICU\}$. These quarantine rates should be high, but medical staff gets infected.

5 Experiments

I consider various time-varying interventions affecting the basic reproduction rate, \mathcal{R}_0 , that is, infection rate α , the quarantine efficiency for non-hospitalized symptomatic infected, ε_S , and the tracing efficiency, ε_T . For SD and quarantine policies, we consider a permanent intervention, that is, a permanent change in the policy parameter, and a temporary intervention that returns the policy parameter to its initial value after some time. I then consider joint policies of SD and quarantine, augmented by tracing and testing.

We seed the initial condition following Ferguson et al. (2020) and assume that the first infection occurs January 1, 2020, and that infections double every five days. Taking the case fatality rate of 0.9%, we then match the number of deaths at the starting date of the simulation. For the UK and the USA, we take the starting date to be March 24, when the UK imposed a national lockdown.⁵ Up to that day, 335 deaths and 5,654 infections were reported in the UK. According to the seeding method, reported infections represented 9

⁵In the US, 21 states had issued stay-at-home orders by March 24, including California and the north-eastern states. An additional 19 states issued these orders by April 1. These orders cover most of the US population. Source: <https://www.kff.org/coronavirus-policy-watch/stay-at-home-orders-to-fight-covid19/>

percent of imputed infections in the UK.⁶ We also assume that initially there are one and a half times as many exposed individuals as there are imputed infected individuals.

The baseline outcome from the spread of the disease without any policy intervention is about 1 percent of the population dead since the assumed case fatality rate is about 1 percent. That means 600 thousand deaths in the UK and 3.25 million deaths in the USA. By how much can the various policy interventions reduce the total number of deaths?

The model specification assumes that ICU units are available for any infected individuals requiring intensive care. Fatality rates will be higher if demand for ICU units exceeds the number of available ICU units. So the impact of policies on the number of infected requiring ICU units is also important. There are about 4 thousand ICU units in the UK, about 0.006 percent of UK population, and 63 thousand ICU units in the US, about 0.095 percent of US population.⁷

In the following section, we consider the impact of variations in social distancing and the effectiveness of quarantine, tracing, and random testing measures to reduce cumulative deaths and peak ICU demand. These experiments are performed for the UK seeding, but the seeding does not make a big difference. We report the outcomes for population shares and occasionally compare the absolute numbers with Ferguson et al. (2020).

5.1 Effectiveness of social distancing

High-intensity SD, that is, large permanent reductions in the basic reproduction rate, has a large impact on fatalities and peak ICU usage. But even high-intensity SD interventions have to be permanent to be effective.

- We consider permanent SD interventions and SD interventions that are limited to six months, after which the reproduction rate returns to its base value. The results are displayed in Table 1 and Figure 2.⁸
- A permanent reduction of the reproduction rate by 75 percent reduces total deaths by a factor of 150, from 1 percent to 0.006 percent of the population, top panel of Table 1, column 5. In addition it cuts the peak demand for ICU units by a factor of more

⁶We could also seed the model with US data. On March 24, there were 471 cumulative deaths and 42,164 reported infections in the USA. Reported infections represent 43 percent of imputed infections in the USA. Peak infection rates and terminal conditions do not depend on the two initial conditions.

⁷For the UK, *Daily Telegraph*, March 25, 2020, <https://www.telegraph.co.uk/global-health/science-and-disease/hospitals-could-need-75-times-number-critical-care-beds-treat/>. For the US, medical intensive care and other ICUs for adults from <https://www.aha.org/statistics/fast-facts-us-hospitals> for the US.

⁸Recall that the percentage reduction of the reproduction rate is the squared percentage reduction of the contact rate.

than 50 to 0.001 percent of the population, Table 1, column 4. This peak ICU demand is below ICU capacity for either the UK or the US.

- SD interventions need not necessarily have to bring the basic reproduction rate below one to be effective. For example, a 50 percent reduction of the reproduction rate still leaves it above one, but it reduces total deaths by a factor of twenty.
- Temporary reductions of the basic reproduction rate have a minor impact on total deaths and peak ICU demand, they mostly delay them, see bottom panel of Table 1, columns 4 and 5, and Figure 2. Essentially, most people are still susceptible to the virus at the time SD is lifted, and the spread of the disease starts anew, Table 1, column 6.⁹
- It is not obvious how much of a reduction in the reproduction rate can be attained through SD. Using the assumptions of Ferguson et al. (2020), the reproduction rate can be reduced by about 80 percent, depending on the assumptions on the relative intensity of social and workplace interactions, Appendix A.2. But even a 75 percent reduction of the reproduction rate reduces total deaths to about 4 thousand in the UK and brings peak ICU demand below capacity. These numbers for deaths and ICU demand in the UK are substantially smaller than the numbers in Ferguson et al. (2020), who report cumulative deaths of 80 thousand to 100 thousand for policies that emphasize SD. Since we are interested in the impact of policy alternatives to SD for a calibration that starts with an SD policy whose implications are comparable to the ones discussed in Ferguson et al. (2020), from now on we assume that the impact of SD is more limited. In particular, we assume that SD reduces the reproduction rate only by 45 percent, resulting in cumulative deaths of about 80 thousand in the UK.

5.2 Effectiveness of quarantine

Efficient permanent quarantine on its own reduces fatalities and peak ICU demand substantially. When quarantine is combined with contact tracing, it yields results comparable to SD.

- We allow for the possibility of quarantining a fraction, ε_S , of the known symptomatic non-hospitalized individuals, and possibly trace previous contacts of newly symptomatic individuals. We only display results for a permanent quarantine regime, since

⁹In Piguillem and Shi (2020), a temporary SD policy is effective because they assume that a critical mass of infected individuals is needed for the disease to spread.

transitory quarantine policies are as ineffective as are transitory SD policies. The results are displayed in the top panel of Table 2 and Figure 3.

- Permanent strict quarantine that removes up to 90 percent of the known symptomatic infected individuals from the infectious pool reduces total deaths by 75 percent and brings peak ICU demand below capacity in the UK and US, top panel of Table 2, columns 4 and 5.
- Combining efficient quarantine with perfect contact tracing reduces the infectious pool by another factor of three, column 1 of Table 2. Quarantining traced asymptomatic individuals then cuts peak ICU demand and total deaths by another factor of four, Table 2, columns 4 and 5.

5.3 Effectiveness of combined policies

We now consider the impact on total deaths and peak ICU demand of four policy interventions that to various degrees combine elements of SD, quarantine, tracing, and testing, Table 3. As a reference point, we list the outcomes from no intervention in the first row of Table 3. The baseline policy is one of permanent high-intensity SD and temporary medium efficient quarantine based on Ferguson et al. (2020). We then consider alternative policies that combine a relaxation of SD over time with more efficient permanent quarantine regimes, augmented with efficient tracing and/or random testing. We find that in our calibrated stylized model, the alternative policies that combine efficient quarantine with tracing do equally well as SD in terms of reducing peak ICU demand and imply significantly lower total deaths than the baseline SD policy.

For our stylized version of the policy studied in Ferguson et al. (2020), we interpret the baseline policy as a permanent 45 percent reduction of the transmission rate α , combined with a temporary three-month increase of quarantine efficiency to $\varepsilon_S = 0.5$.¹⁰ Relative to no intervention, this policy reduces total deaths by a factor of ten and peak ICU demand by a factor of 50, Table 3, Policies 0 and 1. In absolute numbers, for the UK this means about 50 thousand deaths and 800 peak ICU demand. Recall that UK ICU capacity is estimated to be about 5 thousand. These projected numbers are lower than those projected in the Ferguson et al. (2020) study.¹¹

¹⁰See sections 4, 5.1, and Ferguson et al. (2020), Table 4, for the cases with general quarantine and SD. Ferguson et al. (2020) propose SD for at least five months, with subsequent relaxation and tightening contingent on ICU demand triggers. Effectively SD is in place for 80 percent of the time.

¹¹Ferguson et al. (2020), Table 4, for the cases with general quarantine and SD predicts total deaths of 100 thousand and peak ICU demand of 10 thousand. These numbers are predicted to be lower if additional policies targeting particular demographic groups are implemented.

We now consider alternative policies that relax SD over time, in the context of a permanent and efficient quarantine policy, backed up by efficient contact tracing and/or random testing. For this policy, we start with a two-month, 45 percent reduction of the transmission rate α through SD, followed by another three months with a 25 percent reduction of the transmission rate, and finally a permanent 5 percent reduction. All reductions are relative to the base level. Quarantine efficiency is permanently increased to 90 percent.

The first alternative policy combines a gradual relaxation of SD with an efficient quarantine regime, Table 3, Policy 2. For this policy, we assume that 90 percent of newly symptomatic individuals are known and quarantined. This policy reduces total deaths relative to the baseline SD policy by a factor of seven and yields similar peak ICU demand. As we now show, contact tracing and random testing yield only marginal improvements over this policy.

The second alternative policy backs up the efficient quarantine policy with an efficient tracing regime, Table 3, Policy 3. For this policy, we assume that 90 percent of previous contacts that a newly symptomatic individual has infected are traced and quarantined. This policy reduces total deaths relative to the baseline SD policy by a factor of eight and yields similar peak ICU demand.

We have not discussed how tracing is actually implemented. The contact-tracing process for a newly confirmed symptomatic patient consists of a detailed interview with the patient to find out where they have been and then reaching out to those people or the heads of organizations responsible for places, such as airlines, hotels, or religious organizations, that may have been affected. High-risk/close contacts are monitored by public health authorities and low-risk contacts are asked to self-monitor for symptoms in the process laid out by the CDC.¹² As far as we can tell, even among traced individuals only the ones showing symptoms are tested.

No matter how contact tracing is implemented, our assumptions that tracing is efficient and that individuals who have been identified through tracing can be quarantined the same way as symptomatic individuals are highly optimistic. Furthermore, contact tracing has been mainly used for less prevalent diseases and not for large-scale pandemics.

Consider now the alternative of backing up quarantine through random testing of asymptomatic individuals at a rate that would test the complete population within a year. For comparison, the US has been able to increase its testing rate from 50 thousand a day to 100 thousand a day from the middle of March to the middle of April. At that rate the US can test 10 percent of its population in a year. So our assumption on the testing rate would require another ten-fold increase. Table 3, Policy 4, displays the impact of high-intensity

¹²Landman (2020), Armbruster and Brandeau (2007)

random testing. In our stylized model, adding random testing to quarantine, at least for the rate considered here, is somewhat less effective than contact tracing, but total deaths are reduced by a similar magnitude as with tracing, and peak ICU demand is reduced as much as with tracing. Finally, adding random testing to tracing with quarantine has a negligible impact, Table 3, Policy 5.

The main reason why random testing is not very effective is that with an efficient quarantine policy in the background, the share of infectious asymptomatic individuals in the general population is not very large. The peak value of that share is less than 0.1 percent, Table 3, column 1, and the probability of finding an asymptomatic infectious individual through a random test is less than 0.01 percent. Testing every newly symptomatic individual alone would require testing less than 0.5 percent of the population in a year, well within the current capacity constraints for testing.

To summarize, the stylized model predicts that a policy with gradual relaxation of SD, combined with permanent high-efficiency quarantine and possibly tracing of infectious individuals reduces total deaths more and has the same impact on peak ICU demand as a policy of high-intensity permanent SD. A by-product of the successful reduction of new infections by all of these policies is that after more than a year almost all of the population remains susceptible to the virus, Table 3, last column. Thus, in the absence of a vaccine or effective treatment, these policies need to remain permanently in place.

5.4 Implications for employment

The purpose of this paper is to study the impact of policy alternatives to a high-intensity SD policy that are less disruptive for the workings of the overall economy. If we view current policy in the UK or US as representing high-intensity SD as described in the preceding exercises, that is, a reduction of individual contact rates by 25 percent with a corresponding reduction of the transmission rate by 45 percent, then this policy has been disruptive. Employment has declined by about 12 percent, and current estimates are for a total decline of 25 percent in the second quarter of 2020, see Appendix A.2.

In Figure 4, we plot ‘guesses’ of the impact of the policy alternatives on employment in the economy. The solid lines represent the population available for work in the economy, relative to normal at one. The dashed lines represent employment consistent with the available workforce and the extent of SD.

The available workforce consists of those who are healthy and not quarantined.¹³ For none of the policies we consider, the pure health effect on workforce availability is noticeable, and the pure health effect on employment is dwarfed by the disruptions of high-intensity SD.

The dashed lines in Figure 4 represent the joint impact of SD and other policies on employment. We take as given that a 25 percent reduction of contact rates reduces employment relative to available workforce by 25 percent. We then assume that smaller reductions of the transmission rates through SD reduce employment proportionally to the corresponding reduction in the contact rate. More or less by assumption (or interpolation), the alternative policies result in substantially better employment outcomes than the permanent high-intensity SD policy.

6 Caveats

I have used a stylized model to evaluate the relative efficiency of four policy interventions to contain the spread of the coronavirus: SD, quarantine, contact tracing, and random testing. The qualitative features of the relative efficiency of these policies are intuitive enough to expect that they would hold in more general models. How much one should trust the quantitative implications is a different issue.

The first thing to note is that the model was intentionally parameterized to replicate the Ferguson et al. (2020) model. To the extent that there is uncertainty about the ‘stylized facts’ in Ferguson et al. (2020), we will do a robustness exercise below. Second, and possibly more important, the disease does not spread as fast in the model as we observe in the data.

6.1 Higher basic reproduction rate

We have seeded the model to the 335 cumulative deaths in the UK on March 24. Three weeks later on April 14, cumulative deaths in the UK were 11,329. The model predicts, however, that after three more weeks, cumulative deaths without an intervention should have been about 4,600, and under a high-intensity SD policy they should have been about 3,300. The corresponding numbers for the US are actual cumulative deaths of 673 on March 24 and 21,972 on April 14. Seeding the model to the March 24 deaths, the model predicts 8,700 deaths for April 14 with no intervention and 5,800 deaths with a high-intensity SD policy. For both countries, the predicted increase of cumulative deaths is substantially below the actual increase of reported deaths.¹⁴

¹³We essentially assume a representative worker or that employed and non-employed are equally affected by the spread of the disease.

¹⁴The data are from the WHO website <https://covid19.who.int/region/euro/country/gb> and [../usa](https://covid19.who.int/region/usa/country/usa), April 22, 2020.

One way to account for the large increase of cumulative deaths from March 24 to April 14 is to work with a larger basic reproduction rate. Sanche, Lin, Xu, Romero-Severson, Hengartner and Ke (2020), for example, reconsider the emergence of COVID-19 in Wuhan and argue that it is twice as infectious as previous estimates suggested. They estimate the basic reproduction rate to be 5.7 and that infections double within 2.7 days. Similarly, Fernandez-Villaverde and Jones (2020) estimate a time-varying effective transmission rate α by matching cumulative deaths to the predictions of a SIR model. They find reproduction rates in excess of 4 for some US cities and European countries.¹⁵

We now replicate the comparison of alternative policies when we seed the model to the higher basic reproduction rate estimated by Sanche et al. (2020), keeping all other parameters unchanged. Again, we match the cumulative deaths on March 24. For the UK, the model now predicts cumulative deaths on April 14 of about 26,000 with no intervention and about 11,400 with the high-intensity SD policy. The corresponding cumulative deaths for the US on April 14 are now about 56,000 with no intervention and 20,200 with the high-intensity SD policy. Recall that we chose March 24 as a starting date because the UK adopted a national lockdown policy on that day, and a substantial share of US population was already subject to stay-at-home policies by March 24. The predicted increase in cumulative deaths associated with the high-intensity SD policy is then remarkably close to actual outcomes for both the UK and the US.

Table 4 displays the outcomes for the same policies we considered previously when the reproduction rate is twice as high as in the baseline analysis. If there is no intervention, peak infections and ICU demand triple, and deaths increase by 30 percent relative to the lower reproduction rate; Table 4, Policy 0. The main result for all policy interventions is that their ability to reduce the spread of the disease is greatly diminished. Permanent high-intensity SD now reduces cumulative deaths by only 10 percent, rather than a factor of ten as before. The alternative policies still improve on the high-intensity SD policy but by less. For example, they reduce cumulative deaths by an additional 10 percent, rather than a factor of seven. Finally, peak ICU demand now exceeds capacity for the UK, but it remains below capacity for the US.

With a higher reproduction rate, policies not only cannot reduce cumulative deaths that much, they also cannot slow down the rate at which deaths accumulate. The substantial

¹⁵Another reason why the stylized model might understate the increase in cumulative deaths could be related to the assumption that disease state changes follow a Poisson process. As mentioned in Section 1.1, a number of authors in the epidemiological literature argue that SEIR-type models with duration-dependent transition rates provide a better match for the dynamics of diseases like SARS, delivering a bigger peak and shorter duration, for example, Wearing et al. (2005) and Feng et al. (2007). But then Feng (2007) also argues that this simple ranking of models with duration (in)dependent transition rates may depend on the particular way policy interventions like quarantine are modeled.

run-up in cumulative deaths that the model generates for late March is only the precursor of more future deaths to come in the near future. Given the high rate at which the disease spreads, cumulative deaths attain their terminal value within 15 to 25 weeks, depending on the policy, Figure 5. This seems inconsistent with European countries and US states being able to flatten the path for cumulative deaths substantially. One way to account for this observation in the model might be further adjustments to the social distancing parameter.

6.2 Uncertainty

As we just saw, estimates of the basic reproduction rate are being revised upward, but estimates of other parameters, such as the incident fatality rate, also vary substantially. We do not really know what the share of exposed or asymptomatic individuals in the population is, etc. On the policy side, we do not really know what the implemented SD policies mean for the transmission rate of the disease. For example, if we target a 50 percent reduction of the transmission rate through such a policy, how do we know that that's what we get? To address some of these questions, we perform the following simulation study.

We classify parameters from our baseline calibration as being subject to low, medium, or high uncertainty. This means that percentage deviations of a parameter from its baseline value have a 5 percent, 10 percent, or 15 percent coefficient of variation. The classification is subjective but informed by the literature as summarized by the Robert Koch Institut, see Appendix A.4.¹⁶ For example, we consider the uncertainty surrounding the basic reproduction rate and the effectiveness of SD as large, but the uncertainty surrounding the mean recovery periods as small. That being said, the alternative basic reproduction rate we just discussed is very unlikely, even for the high uncertainty case. We then generate one million joint random draws on the parameters from gamma distributions, keep 5,000 of them, and calculate the implied time paths. As an illustration, in Figure A.4 we plot for the above-discussed high-intensity SD policy the time path of cumulative deaths for the fixed parameter values and the mean, median, and the symmetric ranges containing 33 percent and 66 percent of all realizations. We do this for four cases. The first case displays the joint uncertainty surrounding disease and policy parameters. The second case considers only uncertainty related to policy parameters, that is, we take all parameters but ε as fixed. The third case considers only uncertainty related to disease parameters, that is, we take the policy parameters ε as fixed. Finally, the fourth case illustrates the main source of outcome uncertainty, the basic transmission rate α_0 .

¹⁶https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Steckbrief.html, as of April 30, 2020.

Figure A.4 shows that for the stylized model and the particular SD policy the uncertainty surrounding outcomes for total deaths is large, and almost all of it can be attributed to the uncertainty surrounding the disease parameters, in particular, the basic transmission rate α_0 . Panel (a) of Figure A.4 shows that the outcome uncertainty associated with uncertainty in all parameters is large, the 66 percentage coverage area for total deaths after a year ranges from 0.02 percent to 0.9 percent. The latter is the no-intervention outcome for the baseline parameters. Even though the median outcome is close to but below the fixed-parameter path, the mean outcome is substantially larger than the fixed parameter path. In other words, the risks associated with uncertainty are weighted to the upside. Comparing panels (b) and (c) of Figure A.4, we see that almost all of the outcome uncertainty is associated with the disease parameter uncertainty rather than the uncertainty about policy parameters.¹⁷ Finally, comparing panels (c) and (d) of Figure A.4 shows that uncertainty in the basic transmission rate is the main driver of outcome uncertainty.

7 Conclusion

I have studied the effectiveness of alternative policies to contain the spread of a pandemic in a stylized model of the SEIR variety that is calibrated to the Ferguson et al. (2020) study. I find that a policy that combines a gradual relaxation of social distancing with an efficient quarantine, possibly augmented by contact tracing, improves noticeably on a policy of permanent high-intensity SD.

We should qualify the stylized model’s ability to make quantitative predictions on the spread of the disease. First, cumulative deaths in the model do not increase as fast as we observe for the UK and the US from late March to mid-April 2020. The model better matches this increase in cumulative deaths for a higher basic reproduction rate, consistent with recently revised estimates. But if COVID-19 is much more infectious than what we have assumed until now, then the effectiveness of all policies will be greatly reduced. More generally, the uncertainty surrounding all parameter estimates used to calibrate the model is large, and so is the implied uncertainty for policy outcomes. The most important contributor to outcome uncertainty, at least as it relates to cumulative deaths, appears to be the uncertainty about the disease transmission rate.

¹⁷Note the different scales for panels (b) and (c).

References

- Alvarez, Fernando, David Argente, and Francesco Lippi, “A Simple Planning Problem for COVID-19 Lockdown,” *Working Paper*, April 2020.
- Armbruster, B., and ML Brandeau, “Contact Tracing to Control Infectious Disease: When Enough Is Enough,” *Health Care Manag Sci.*, 2007, 10 (4), 341–355.
- Atkeson, Andrew, “What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios,” Working Paper 26867, National Bureau of Economic Research March 2020.
- Berger, David W., Kyle F. Herkenhoff, and Simon Mongey, “An SEIR Infectious Disease Model with Testing and Conditional Quarantine,” Working Paper 26901, National Bureau of Economic Research March 2020.
- Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt, “The Macroeconomics of Epidemics,” Working Paper 26882, National Bureau of Economic Research March 2020.
- Farboodi, Maryam, Gregor Jarosch, and Robert Shimer, “Internal and External Effects of Social Distancing in a Pandemic,” *Working Paper*, April 2020.
- Fauci, Anthony S., H. Clifford Lane, and Robert R. Redfield, “Covid-19 — Navigating the Uncharted,” *New England Journal of Medicine*, 2020.
- Feng, Zhilan, “Final and Peak Epidemic Sizes for SEIR Models with Quarantine and Isolation,” *Mathematical Biosciences and Engineering*, 2007, 4, 675–686.
- , Dashun Xu, and Haiyun Zhao, “Epidemiological Models with Non-Exponentially Distributed Disease Stages and Applications to Disease Control,” *Bulletin of Mathematical Biology*, 2007, 69 (5), 1511–1536.
- Ferguson, Neil M., et al. “Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand.,” Technical Report 2020.
- Fernandez-Villaverde, Jesus, and Charles Jones, “Estimating and Simulating a SIRD Model of COVID-19,” *Working Paper*, 2020.
- Jewell, Nicholas P., Joseph A. Lewnard, and Britta L. Jewell, “Predictive Mathematical Models of the COVID-19 Pandemic: Underlying Principles and Value of Projections,” *JAMA*, April 2020.

- Landman, Keren, “How The Painstaking Work of Contact Tracing Can Slow the Spread of an Outbreak,” *National Public Radio*, 2020.
- Lipsitch, Marc, et al. “Transmission Dynamics and Control of Severe Acute Respiratory Syndrome,” *Science*, 2003, *300* (5627), 1966–1970.
- Piguillem, Facundo, and Liyan Shi, “Optimal COVID-19 Quarantine and Testing Policies,” EIEF Working Papers Series 2004, Einaudi Institute for Economics and Finance (EIEF) 2020.
- Sanche, Steven, Y.T. Lin, C. Xu, E. Romero-Severson, N. Hengartner, and R. Ke, “High Contagiousness and Rapid Spread of Severe Acute Respiratory Syndrome Coronavirus 2,” *Emerg Infect Dis.*, July 2020.
- Shen, Chen, Nassim N. Taleb, and Yaneer Bar Yam, “Review of Ferguson et al ‘Impact of non-pharmaceutical interventions...’,” Technical Report 2020.
- Stock, James H., “Data Gaps and the Policy Response to the Novel Coronavirus,” Working Paper 26902, National Bureau of Economic Research March 2020.
- Wearing, Helen J., Pejman Rohani, and Matt J. Keeling, “Appropriate Models for the Management of Infectious Diseases,” *PLOS Medicine*, July 2005, *2* (7).

A Appendix

A.1 Reproduction rates

We now calculate the average new infections caused by a newly infectious agent. We start with the basic reproduction rate in the SIR model, then the basic reproduction rate in the SEIR model, and then calculate average new infections from an asymptomatic individual until he becomes symptomatic and is quarantined.

A.1.1 Basic reproduction rate \mathcal{R}_0 for SIR model

The individual is infectious at rate $S(\tau)\alpha$ until recovery or death ($\tilde{\gamma} = \gamma + \delta$).

$$\begin{aligned}\mathcal{R}_0 &= \int_0^\infty [S(\tau)\alpha\tau] [(\gamma e^{-\gamma\tau})e^{-\delta\tau} + (\delta e^{-\delta\tau})e^{-\tilde{\gamma}\tau}] d\tau \\ &\approx S(0) \int_0^\infty (\alpha\tau) (\tilde{\gamma}e^{-\tilde{\gamma}\tau}) d\tau \\ &\approx \alpha\tilde{\gamma} \int_0^\infty \tau e^{-\tilde{\gamma}\tau} d\tau\end{aligned}$$

For the first approximation, we assume that changes in the measure of susceptible individuals S are small over the time of an individual infection. For the second approximation, we assume that initially the share of susceptible individuals is close to one.

Note that

$$\int_0^t \tau e^{\alpha\tau} d\tau = \frac{1}{\alpha^2} [1 + e^{\alpha t} (\alpha t - 1)] \quad \text{and} \quad \lim_{t \rightarrow \infty} \int_0^t \tau e^{-\gamma\tau} d\tau = \frac{1}{\gamma^2}$$

Therefore

$$\mathcal{R}_0 = \frac{\alpha}{\tilde{\gamma}}$$

A.1.2 Basic reproduction rate in \mathcal{R}_0 for SEIR model

We consider the progression from an asymptomatic infectious individual to a symptomatic infectious one, working backwards.

The average number of new infections caused by a symptomatic individual, ignoring hospitalization, is

$$\begin{aligned}\mathcal{R}_{0S} &= S(t) \int_0^\infty [\sigma\alpha\tau] (\tilde{\gamma}_S e^{-\tilde{\gamma}_S\tau}) d\tau \\ &= S(t)\alpha \frac{\sigma}{\tilde{\gamma}_S}\end{aligned}$$

with $\tilde{\gamma}_S = \gamma_S + \delta$

The average number of new infections caused by an asymptomatic individual is

$$\begin{aligned}\mathcal{R}_{0A} &= S(t) \int_0^\infty [\alpha\tau + \mathcal{R}_{0S}] (\beta e^{-\beta\tau}) (e^{-\gamma_A\tau}) d\tau + S(t) \int_0^\infty (\alpha\tau) (\gamma_A e^{-\gamma_A\tau}) (e^{-\beta\tau}) d\tau \\ &= S(t)\alpha \frac{1}{(\beta + \gamma_A)} \left[1 + \frac{\sigma\beta}{\gamma_S(\beta + \gamma_A)} \right]\end{aligned}$$

A.1.3 New infections with quarantine

We consider an asymptomatic infectious individual, $(\alpha, \beta, \gamma_A)$, who is quarantined once he becomes symptomatic. For this case, we calculate the average number of exposed and infectious asymptomatic individuals that this individual has created.

By the time an asymptomatic individual becomes symptomatic, the average number that individual has infected is

$$\begin{aligned}\mathcal{R}_{AQ} &= S(t) \left[\int_0^\infty (\alpha\tau) (\beta e^{-\beta\tau}) (e^{-\gamma_A\tau}) d\tau \right] \\ &= S(t)\alpha \frac{\beta}{(\beta + \gamma_A)^2}\end{aligned}$$

The average number of individuals that the infectious agent has infected and who are not yet infectious at the time the agent becomes symptomatic is

$$\mathcal{R}_{ATE} = S(t) \int_0^\infty \left[\alpha \int_0^\tau e^{-\phi s} ds \right] [(\beta e^{-\beta\tau}) (e^{-\gamma_A\tau})] d\tau$$

The term in the first square bracket denotes the total who have been infected by the infectious individual at τ and who have not yet become infectious at that time. This can be rewritten as

$$\mathcal{R}_{ATE} = S(t)\alpha \frac{\beta}{(\beta + \gamma_A)(\beta + \gamma_A + \phi)}.$$

The average number of individuals that an infectious agent has infected and who are infectious but asymptomatic at the time the agent becomes symptomatic is

$$\mathcal{R}_{ATA} = S(t) \int_0^\infty \left[\alpha \int_0^\tau \left[\int_0^s \phi e^{-\phi v} e^{-\gamma_A(s-v)} dv \right] ds \right] [\beta e^{-(\beta + \gamma_A)\tau}] d\tau$$

The innermost integral is the probability that an individual who has been infected time s ago has become infectious in the meantime but also has not yet recovered at the time the original infectious individual becomes symptomatic. This can be rewritten as

$$\mathcal{R}_{ATA} = \alpha \frac{\beta\phi}{2(\beta + \gamma_A)^2(\beta + \gamma_A + \phi)}$$

A.1.4 Probability of recovery without developing symptoms

The probability of recovering while asymptomatic before becoming symptomatic

$$p_{AR} = \int_0^{\infty} (\gamma_A e^{-\gamma_A \tau}) e^{-\beta \tau} d\tau = \frac{\gamma_A}{\gamma_A + \beta}$$

A.2 Social distancing

- According to the American Time Use Survey for 2018, an employed person spends on average 6.3 hours working and 5.13 hours on social activities (purchasing, helping non-household members, education, participating in organizations, and leisure and sports). A non-employed person spends on average 0.12 hours on work related activities and 9.36 hours on social activities.
- Social interactions may be more or less intense than workplace interactions. Given the reports on super spreader events related to soccer games in Italy and churches in South Korea, social interactions may well be more intense than workplace interactions, suppose 50 percent more. This is the opposite of Eichenbaum et al. (2020) for which workplace infections dominate infections related to consumption or unspecified social interactions.
- Assume that 60 percent of the population are working. This corresponds to US employment rates.
- In the last two weeks of March and the first week of April, new unemployment insurance claims increased by about 18 million. On a payroll employment base of 151 million, this means that employment probably decreased by about 12 percent, and the employment rate declined to about 53 percent. Current estimates are for additional employment declines with a total employment decline of 25 percent. Taking this into account reduces social contacts per person by about 63 percent, an additional 2 percentage points.
- The following table lists the implied average contact rates and social reproduction factors for various assumptions on the relative intensity of social interactions, with and without taking into account changes in the employment rates. Contact rates may decline by about 60 percent, and implied reproduction rates may decline by about 80 percent.

S/W	Individual Contact Rate ψ Percent relative to normal		Reproduction Rate Factor α_S Fraction relative to normal	
	ω fixed	ω declines	ω fixed	ω declines
0.75	46.4	42.1	0.22	0.18
1.00	43.0	39.6	0.18	0.16
1.50	38.6	36.5	0.15	0.13

A.3 Seeding the initial condition

We start with initial cumulative deaths, $D(0)$. Assuming a seeding rate σ , such that infections are doubling every five days, and an unconditional case fatality rate δ , consistent with an unconditional case fatality probability $p_D = 0.009$, total cumulative deaths starting from $-\Delta$ are

$$D(0) = I(0) \delta \left(\frac{1 - e^{-\sigma\Delta}}{\sigma} \right)$$

We assume that infections start two and half months before the initial date, $\Delta = 2.5/12$.

A.4 Representing parameter uncertainty

Consider a parameter p and assume that the uncertainty about the parameter is represented by the following form

$$\begin{aligned} \ln p &= \ln \bar{p} + \ln X - E[\ln X] \\ \ln X &\sim \Gamma(k, \theta) \end{aligned}$$

where Γ denotes the Gamma distribution. Then

$$\begin{aligned} E[\ln p] &= \ln \bar{p} \\ Var(\ln p) &= Var(\ln X) \end{aligned}$$

The mean and variance of the gamma distribution are

$$\begin{aligned} \mu &= E[\ln X] = k\theta \\ \sigma^2 &= Var(\ln X) = k\theta^2 \end{aligned}$$

and the median ν is bounded by

$$\mu - 1/3 < \nu < \mu$$

So to get a symmetric distribution we need μ to be large. Let

$$S = k\theta$$

Suppose we fix the coefficient of variation for the observed variable

$$CoV = \frac{Std(\ln p)}{E[\ln p]} = \frac{Std(\ln X)}{\ln \bar{p}} = \frac{\sqrt{k}\theta}{\ln \bar{p}} = \frac{\sqrt{k}S/k}{\ln \bar{p}} = \frac{S}{\sqrt{k} \ln \bar{p}}$$

So the parameters of the gamma distribution are

$$\begin{aligned} k &= \left(\frac{S}{CoV \ln \bar{p}} \right)^2 \\ \theta &= \frac{S}{k} = S \left(\frac{CoV \ln \bar{p}}{S} \right)^2 = \frac{(CoV \ln \bar{p})^2}{S} \end{aligned}$$

The MATLAB usage of the gamma function is

$$\Gamma(a, b) = \Gamma(k, \theta)$$

We represent uncertainty through the CoV . The Robert Koch Institut (RKI) summarizes the available evidence on various characteristics of the coronavirus.¹⁸ For example, estimates of the basic reproduction rate \mathcal{R}_0 range from 2.4 to 3.3. If we interpret the range as representing a 2 standard deviation band around a mean of 2.8, then the CoV for percentage deviation is 13%. We interpret this CoV as representing the uncertainty surrounding the basic transmission rate α_0 , but we should note that \mathcal{R}_0 not only depends on the transmission rate, but also on the incubation time, recovery time, and relative infectiousness of symptomatic individuals. Since the RKI excludes studies with significantly higher values than 3.3 from its summary of the evidence, assuming a CoV of 15% for the basic transmission rate α_0 may not overstate its uncertainty by much. We classify uncertainty as high, $CoV = 15\%$, medium, $CoV = 10\%$, and low, $CoV = 5\%$ for the parameters

$$\begin{aligned} \text{High:} & \quad \alpha, \alpha_S, \phi, \beta, \sigma, \varepsilon_i \text{ for } i \in \{AT, S, B, ICU\}, \varepsilon_T \\ \text{Medium:} & \quad \delta, \delta_{ICU}, \omega, \eta, \\ \text{Low:} & \quad \gamma, \gamma_{ICU} \end{aligned}$$

¹⁸https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Steckbrief.html, April 30, 2020.

Table 1: Effectiveness of Social Distancing

Model	(1) Max I_A	(2) Max I_{AT}	(3) Max I_S	(4) Max H_{ICU}	(5) Term D	(6) Term S
Permanent Change						
$\mathcal{R}_0=2.40$	4.294	0.000	3.871	0.054	0.912	23.665
$\mathcal{R}_0=1.80$	1.556	0.000	1.444	0.020	0.627	47.476
$\mathcal{R}_0=1.20$	0.101	0.000	0.070	0.001	0.040	96.684
$\mathcal{R}_0=0.60$	0.100	0.000	0.056	0.001	0.006	99.521
Transitory Change						
$\mathcal{R}_0=2.40$	4.294	0.000	3.871	0.054	0.912	23.665
$\mathcal{R}_0=1.80$	1.656	0.000	1.529	0.021	0.802	32.813
$\mathcal{R}_0=1.20$	3.932	0.000	3.556	0.050	0.901	24.594
$\mathcal{R}_0=0.60$	4.179	0.000	3.768	0.053	0.907	23.977

Note. The rows list the replication rate \mathcal{R}_0 implied by reduction of contact rates ψ through SD. The first four columns are the peak shares of (1) asymptomatic infected, (2) known asymptomatic infected, (3) symptomatic at home, and (4) ICU units required. The last two columns are the terminal values after one and a half years for (5) cumulative deaths and (6) susceptible population. All variables are percent of total population. A temporary intervention reduces the basic reproduction rate for a six month period and then returns it to its baseline value of 2.4.

Table 2: Effectiveness of Quarantine ε_S and Tracing ε_T

Model	(1) Max I_A	(2) Max I_{AT}	(3) Max I_S	(4) Max H_{ICU}	(5) Term D	(6) Term S
No Contact Tracing $\varepsilon_T = 0$						
$\varepsilon_S=0.33$	4.294	0.000	3.871	0.054	0.912	23.665
$\varepsilon_S=0.50$	3.085	0.000	2.817	0.040	0.804	32.718
$\varepsilon_S=0.70$	1.571	0.000	1.456	0.020	0.599	49.830
$\varepsilon_S=0.80$	0.858	0.000	0.801	0.011	0.448	62.507
$\varepsilon_S=0.90$	0.293	0.000	0.274	0.004	0.247	79.293
Perfect Contact Tracing $\varepsilon_T = 1.0$						
$\varepsilon_S=0.33$	3.674	0.402	3.679	0.052	0.899	24.698
$\varepsilon_S=0.50$	2.421	0.297	2.490	0.035	0.772	35.326
$\varepsilon_S=0.70$	0.910	0.134	0.972	0.014	0.509	57.402
$\varepsilon_S=0.80$	0.303	0.050	0.331	0.005	0.290	75.677
$\varepsilon_S=0.90$	0.101	0.016	0.092	0.001	0.058	95.211

Note. See Notes for Table 1.

Table 3: Effectiveness of Alternative Policies

Model	(1) Max I_A	(2) Max I_{AT}	(3) Max I_S	(4) Max H_{ICU}	(5) Term D	(6) Term S
Policy 0	4.301	0.000	3.879	0.054	0.912	23.645
Policy 1	0.116	0.000	0.083	0.001	0.076	93.464
Policy 2	0.115	0.000	0.074	0.001	0.013	98.976
Policy 3	0.112	0.007	0.074	0.001	0.009	99.261
Policy 4	0.114	0.001	0.074	0.001	0.012	99.008
Policy 5	0.111	0.008	0.073	0.001	0.009	99.277

Note. The policies are defined on the intervals covering the first two months, the third through fifth month, and the remaining time. Policy 0 is the no-intervention case. The parameters for policy interventions are as follows

Policy 1: $\alpha = 0.55$, $\varepsilon_S = (0.5, 1/3, 1/3)$, $\varepsilon_{AT} = \varepsilon_T = f = 0$

Policy 2: $\alpha = (0.55, 0.75, 0.95)$, $\varepsilon_S = \varepsilon_{AT} = 0.9$, $\varepsilon_T = 0$, $f = 0$

Policy 3: $\alpha = (0.55, 0.75, 0.95)$, $\varepsilon_S = \varepsilon_{AT} = 0.9$, $\varepsilon_T = 0.9$, $f = 0$

Policy 4: $\alpha = (0.55, 0.75, 0.95)$, $\varepsilon_S = \varepsilon_{AT} = 0.9$, $\varepsilon_T = 0$, $f = 1.0$


Policy 5: $\alpha = (0.55, 0.75, 0.95)$, $\varepsilon_S = \varepsilon_{AT} = 0.9$, $\varepsilon_T = 0.9$, $f = 1.0$

Table 4: Effectiveness of Alternative Policies for High \mathcal{R}_0


Model	(1) Max I_A	(2) Max I_{AT}	(3) Max I_S	(4) Max H_{ICU}	(5) Term D	(6) Term S
Policy 0	15.089	0.000	11.711	0.164	1.180	1.193
Policy 1	6.679	0.000	5.857	0.082	1.050	12.107
Policy 2	5.115	0.000	4.522	0.063	0.932	21.980
Policy 3	3.114	0.551	3.305	0.046	0.877	26.600
Policy 4	4.940	0.097	4.443	0.062	0.921	22.870
Policy 5	2.941	0.592	3.194	0.045	0.865	27.541

Note. The basic reproduction rate is $\mathcal{R}_0 = 5.7$. All policies are defined as in Table 3.

Figure 1: The SEIR model




(a) SIR




(b) GSEIR

Figure 2: Effectiveness of Social Distancing \mathcal{R}_0




(a) ICU Hospital Beds




(b) Deaths

Note: See notes for Table 1. Solid lines represent permanent policies and dashed lines represent temporary policies. The shaded area denotes the first six months for which a temporary policy is in place.

Figure 3: Effectiveness of Quarantine ε_{IS}



(a) ICU Hospital Beds



(b) Deaths


Note: See notes for Table 1. Solid lines represent permanent policies and dashed lines represent temporary policies. The shaded area denotes the first six months for which a temporary policy is in place.

Figure 4: Employment Impact



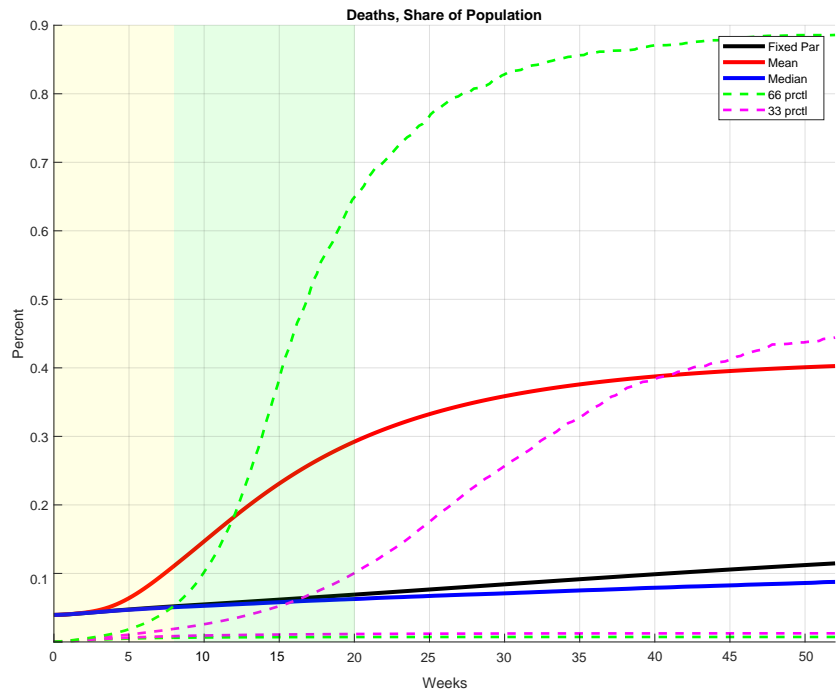
Note: The solid lines denote the population available for work, that is, not hospitalized and not quarantined. In terms of health outcomes all policies are about equally effective. The dashed lines denote the additional employment reduction associated with SD. For the policies see notes for Table 3. SD1Q1 is Policy 1, SD2Q2T2 is Policy 3, SD2Q2F2 is Policy 4, and SD2Q2T2F2 is Policy 5.

Figure 5: Deaths with Large \mathcal{R}_0




Note: The reproduction rate is $\mathcal{R}_0 = 5.7$. The policies correspond to the policies in Table 3: SD1Q1 is Policy 1, SD2Q2 is Policy 2, SD2Q2T2 is Policy 3, SD2Q2F2 is Policy 4, SD2Q2T2F2 is Policy 5. Some of the policies vary over time, and the shaded areas cover the first two months, and the third through fifth month for which the policies change.

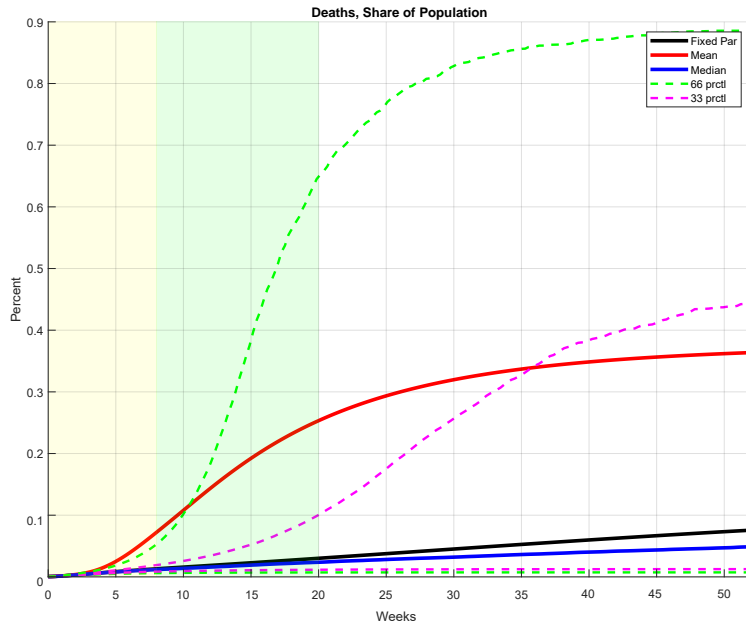
Figure 6: Impact of Parameter Uncertainty on Projected Deaths



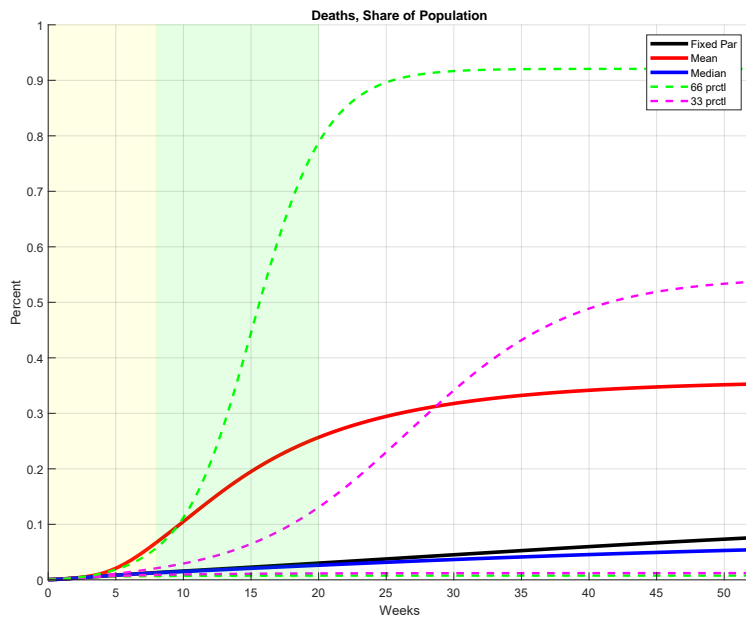
(a) All Parameters



(b) Policy Parameters



(c) Disease Parameters



(d) Contagiousness α_0

Note: Baseline policy is permanent high-intensity SD, combined with temporary medium-intensity quarantine. Solid black line is the outcome for the calibrated parameter values. Solid red and blue lines are the mean and median from the Monte Carlo simulations. The area between the dashed purple and green lines reflect the symmetric ranges that contain 33 percent, respectively 66 percent, of the realizations from the Monte Carlo simulations. Panel (a) allows for uncertainty in policy and disease parameters, panel (b) keeps the disease parameters fixed, panel (c) keeps the policy parameters ε fixed, and panel (d) keeps all parameters fixed except the disease transmission rate α_0 .