



OPEN

Infection dynamics of COVID-19 virus under lockdown and reopening

Jakub Svoboda¹, Josef Tkadlec², Andreas Pavlogiannis³, Krishnendu Chatterjee¹ & Martin A. Nowak^{2,4}✉

Motivated by COVID-19, we develop and analyze a simple stochastic model for the spread of disease in human population. We track how the number of infected and critically ill people develops over time in order to estimate the demand that is imposed on the hospital system. To keep this demand under control, we consider a class of simple policies for slowing down and reopening society and we compare their efficiency in mitigating the spread of the virus from several different points of view. We find that in order to avoid overwhelming of the hospital system, a policy must impose a harsh lockdown or it must react swiftly (or both). While reacting swiftly is universally beneficial, being harsh pays off only when the country is patient about reopening and when the neighboring countries coordinate their mitigation efforts. Our work highlights the importance of acting decisively when closing down and the importance of patience and coordination between neighboring countries when reopening.

Severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) is the virus that causes the current pandemic of coronavirus disease 2019 (COVID-19). The infection was first identified in December 2019 in Wuhan (China) and has since spread globally. By November 2021 more than 250 million people have been tested positive for the virus and more than 5 million people have died from complications caused by the virus. The large majority of cases result in recovery after mild or no symptoms. The coronavirus pandemic has led to an unprecedented global response in quarantine measures, social distancing, travel restrictions and shutting down of economic activity.

The use of mathematics in the study of dynamics of infectious diseases has a long tradition^{1–13}. Mathematical epidemiology was immediately applied to the coronavirus pandemic^{14–16}, resulting in an immense amount of literature on the subject. As a quick sample, there is ample work on understanding how the disease progresses within an individual^{17–20}, how it spreads through the community^{21–27}, and on the effects of various non-pharmaceutical interventions under varying circumstances^{28–35}. Other work highlighted the difficulties of estimating the parameters and the future course of an ongoing epidemic^{36,37}. In contrast, in this work we focus on a simple class of policies for closing down and reopening the society throughout the course of a pandemic, thereby complementing the rich literature on covid modelling.

We consider infection dynamics of coronavirus in a population of size N . The population represents a community (city, state, or a country). Initially, all people are uninfected. Then, we add one (or several) infected individuals and follow the stochastic trajectories of viral spread. The process advances in discrete time steps that represent days. Individuals are in different states describing their status of being: susceptible (S), exposed (E), mildly ill/infectious (I), critically ill (C), and recovered/removed (R), see Fig. 1a. We assume that critical cases are hospitalized. The infection spreads whenever a susceptible person comes in contact with an infectious person. In this case, the infection is transmitted with probability p (see Fig. 1b). We denote the number of daily contacts per person by k_0 .

We assume that the community has a capacity c of hospital beds to treat the critical cases. When left unregulated, the disease would surge through the community and exceed that capacity c (see Fig. 1c). A country can mitigate the spread of the disease by introducing various non-pharmaceutical interventions, such as enforcing social distancing or shutting down non-essential businesses. We model such interventions by decreasing the number k_0 of daily contacts of an individual to a value $k < k_0$. We call the regime when the interventions are put in place a *lockdown*.

¹IST Austria, 3400 Klosterneuburg, Austria. ²Department of Mathematics, Harvard University, Cambridge, MA 02138, USA. ³Department of Computer Science, Aarhus University, 8200 Aarhus, Denmark. ⁴Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138, USA. ✉email: martin_nowak@harvard.edu

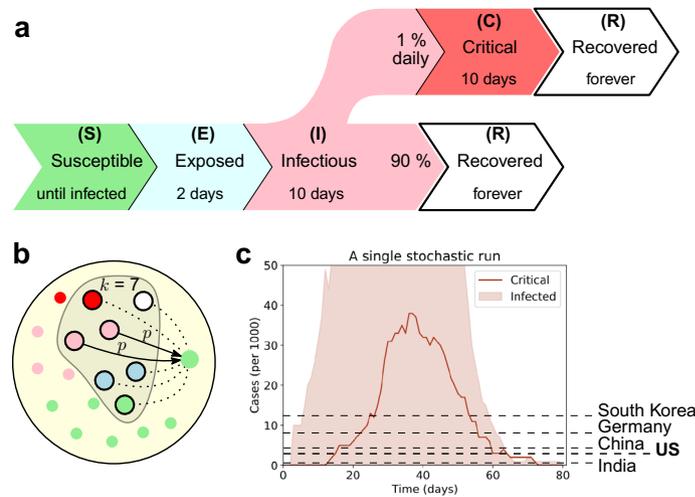


Figure 1. The disease spread without an intervention. **(a)** Two days after being exposed to the disease (E), the individual becomes infectious (I) as they develop mild condition. If a critical condition develops (C), the individual is hospitalized and isolated. We assume that all surviving individuals (R) acquire immunity. **(b)** A population of N individuals. Each day, an individual meets k other individuals. During a single meeting with an infectious person, a susceptible individual contracts a disease with a transmission probability p . **(c)** Without intervention, the disease surges through the community and the critical cases (curve) at its peak C_{\max} exceed the available hospital bed capacity c (dashed lines^{41,42}). Here $N = 1000$, $k = 15$, $p = 2\%$, thus the effective reproductive rate R_e is roughly $R_e \doteq 2.9$.

Here we study the key question of how quickly and how severely a community should lock down, and how patient it should be before reopening again. To that end, we consider a simple class of policies characterized by three parameters τ , k , d . The parameter τ is the number of critical cases that triggers the community to enter the lockdown. It models the cautiousness or agility of the policy. The parameter k is the number of daily contacts per person in a lockdown. It models the severity of the policy. Finally, the parameter d is the number of days the community needs to spend with critical cases below the trigger threshold τ before the lockdown is lifted. It models the patience of the policy. In other words, a policy $P(k, \tau, d)$ locks down to $k < k_0$ daily contacts per person once the number of critically ill individuals exceeds a given trigger threshold τ , and it reopens to k_0 daily contacts once the number of critically ill individuals remains under that threshold τ for d consecutive days. We evaluate the performance of such policies with respect to several measures. For instance, we consider the *peak load* C_{\max} , which is the expected number of critical cases at its maximum, and the *overflow probability* P_{fail} , which is the probability that the peak load C_{\max} exceeds the hospital bed capacity c available to treat the critically ill cases.

We find that the best performing policies are those that quickly transition to a severe lockdown and that are patient about reopening. However, when either quick or severe action is not feasible, as is often the case, a country can compensate by pressing more along the other dimension. This gives rise to a spectrum of possible policies for closing down and reopening the society. At one end of this spectrum, a *moderate low-trigger* policy (P_{ML}) imposes a gentle lockdown at the first sign of the onset of the disease. At the other end, a *severe high-trigger* policy (P_{SH}) remains open till the latest moment possible and then imposes a harsh lockdown. We find that, though comparable in some regards, those two policies are very different in terms of their long-term behavior and in terms of their sensitivity to policies employed by neighboring countries. Specifically, we argue that the moderate low-trigger policies are preferable assuming that the countries are able to coordinate and that an efficient vaccine is not distributed soon, whereas the severe high-trigger policies are preferable otherwise.

Model

In order to describe the spread of COVID-19, we consider a stochastic, discrete-time, individual-based SIR-like model.

Disease progression within an individual. We consider the following simple model. Initially, a typical individual is susceptible (S) and can contract the disease after contact with an infectious individual (see Fig. 1a). Immediately upon contracting the disease, an individual becomes exposed (E) and does not yet spread it. Later they become infectious (I) as they develop mild condition and then they either recover (R) or become critically ill (C). Critical individuals are hospitalized and isolated and they occupy part of the capacity c of the health system. Eventually, they either recover or die (R). We assume that recovered individuals acquire immunity. Upon contracting the disease, each transition occurs after a number of days that is given by a corresponding random variable $X_{E \rightarrow I}$, $X_{I \rightarrow R}$, $X_{I \rightarrow C}$, $X_{C \rightarrow R}$. For concreteness, we set the values based on the data on COVID-19:^{17–19,21,38–40} the pre-infectious period is $X_{E \rightarrow I} = 2$ days and the individual recovers from a mild condition upon $X_{I \rightarrow R} = 10$ days. During each of those 10 days, an individual might become critically ill with probability 1% (hence $X_{I \rightarrow C}$ is exponentially distributed with parameter 1% and roughly 10% of cases become critical). The critical cases

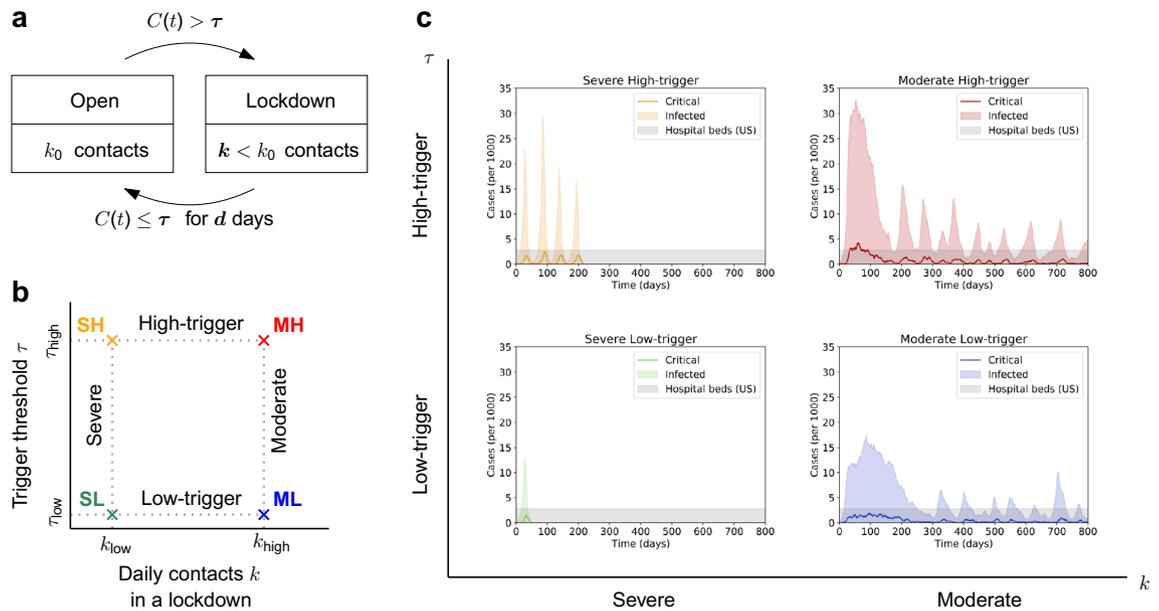


Figure 2. Basic policies. (a) Under a policy $P(\tau, k, d)$, the country locks down to $k < k_0$ daily contacts whenever the number $C(t)$ of critical cases exceeds a trigger threshold τ . It reopens (to k_0 daily contacts) once the number of critical cases stays below τ for d consecutive days. (b) We consider four different policies given by a combination of a trigger threshold (low trigger $\tau_{low} = 3$, high trigger $\tau_{high} = 12$) and a lockdown severity (severe $k_{low} = 1.25$, moderate $k_{high} = 6$), and common patience $d = 10$ days. (c) Representative runs under the four policies (for 800 days). While with the severe high-trigger policy P_{SH} (top left) all peaks are similar in shape, with the moderate low-trigger policy P_{ML} (bottom right) all subsequent peaks are much smaller than the first one. With the moderate high-trigger policy P_{MH} (top right) the capacity is exceeded and with the severe low-trigger policy P_{SL} (bottom left) the disease is quickly eradicated.

recover (or die) after $X_{C \rightarrow R} = 10$ more days. In the Supplementary Information, we also consider other choices of parameters and other distributions of variables, see Supplementary Figs. 2 and 3.

Disease spread through the population. We consider a population of N individuals. Each day, each individual comes in contact with k_0 other individuals (see Fig. 1b). When a susceptible individual (S) meets an infectious individual (I), he or she contracts the disease with transmission probability p and becomes exposed (E). Note that since we assume that recovered individuals acquire immunity, the disease eventually gets eradicated—at the latest, this happens once everyone is exposed and later recovers. We denote the number of critically ill individuals at day t by $C(t)$ and use analogous notation for other conditions. For concreteness^{41–43} we consider a population of $N = 20,000$ individuals, a health system capacity of 2.8 beds per 1000 individuals leading to $c = 56$ beds in total, $k_0 = 15$ daily interactions per person, and a transmission probability $p = 2\%$. All in all, this gives the epidemiological basic reproductive ratio R_0 equal to roughly $R_0 = k_0 \cdot p \cdot X_{I \rightarrow R} = 2.9$.

Policies. We consider a 3-parameter class of policies that a country can use to mitigate the spread of the disease. The policies toggle between two regimes—the default *open* regime and the temporary *lockdown* regime (see Fig. 2a). The policy $P(\tau, k, d)$ can be efficiently described using three parameters τ, k, d that describe how soon and severely the policy locks down, and how soon it reopens:

1. Once the number $C(t)$ of critically ill cases exceeds the trigger threshold τ , the policy locks down by reducing the number of daily contacts per person from k_0 to k .
2. Once the number $C(t)$ of critically ill cases remains under the trigger threshold τ for d consecutive days, the policy reopens by resetting the number of daily contacts per person back to k_0 .

The three parameters τ, k, d thus model three natural features of the policy: its “cautiousness”, that is, how easily it is triggered into a lockdown; the “severity” of its lockdowns; and its “patience” when reopening, respectively. We remark that we chose the number of critically ill people $C(t)$ rather than the number of infectious people $I(t)$ since the latter is not so easily accessible to policymakers.

Performance of the policy. In order to evaluate the performance of a policy, we study the following quantities. Denote by $M = \max\{C(t) \mid t \geq 0\}$ the random variable that corresponds to the number of critically ill people at their maximum, over the duration of the disease.

1. The (expected) *peak load* C_{\max} : that is, the expected number $C_{\max} = \mathbb{E}[M]$ of critically ill people at their maximum, over the duration of the disease. This represents the maximum demand on the health system (hospital beds).
2. The *overflow probability* p_{fail} : that is, the probability $p_{\text{fail}} = \Pr[M > c]$ of exceeding the available bed capacity c at some point throughout the course of the disease.
3. The *total load* C_{all} : that is, the total cumulative number of critical cases, over the duration of the disease. This can be used to estimate the total number of deaths.
4. The total (expected) *duration* D of the lockdowns: that is, the total expected number of days spent in the lockdown regime, until the disease is eventually eradicated within the community.

In the Supplementary Information, we also consider another measure, namely the total (expected) overflow of the bed capacity $\mathbb{E}[\sum_{t \geq 0} \max\{C(t) - c, 0\}]$, see Supplementary Fig. 1. In all cases, the lower the quantity the better the policy. Hence, we can think of all the quantities as costs of the policy. We note that the first three quantities can be viewed as costs related to health of the population. For a fixed lockdown severity k , the fourth quantity can be viewed as an economic cost of imposing a lockdown of that severity, since a lockdown reduces the economic activity of a community.

Results

We evaluate the performance of the above defined policies. First, we do this for a single country. Later, we do this for two neighboring countries. Recall that a policy $P(\tau, k, d)$ is given by three parameters: the threshold number of critical cases τ that triggers the policy to toggle to lockdown and back (“trigger value”); the number k of daily contacts per person during a lockdown (“severity”); and the number d of days required to remain below the trigger threshold τ before the society can reopen (“patience”), see Fig. 2a.

Four example policies. To illustrate the differences in performance of various policies, we first consider four specific policies that all share the patience parameter $d = 10$ days and that differ in the trigger value τ and in the severity k (see Fig. 2b). Specifically, in terms of the trigger value τ , we distinguish *low-trigger* policies ($\tau_{\text{low}} = 3$) from *high-trigger* policies ($\tau_{\text{high}} = 12$). Similarly, in terms of the number k of daily contacts in a lockdown, we distinguish *severe* policies ($k_{\text{low}} = 1.25$) from *moderate* policies ($k_{\text{high}} = 6$). All in all, this yields $2 \times 2 = 4$ combinations P_{SL} (severe, low-trigger), P_{SH} (severe, high-trigger), P_{ML} (moderate, low-trigger), and P_{MH} (moderate, high-trigger). We observe that the policies substantially alter how the number of infected and critical cases evolves in time (see Fig. 2c).

To explain the difference, it is instructive to think in terms of the effective reproductive number R_e that determines whether the number of infected individuals in a population is quickly surging ($R_e > 1$), disappearing ($R_e < 1$) or changing slowly ($R_e \approx 1$). Note that R_e is not constant in time—it crucially depends on the current number k of daily contacts (R_e decreases as k decreases) and also on the percentage x of immune individuals (R_e decreases as x increases). In the open society (k large) and with no immune individuals ($x = 0$) we have $R_0 > 1$, and hence the disease initially spreads quickly.

Severe policies. Under the two severe policies, the number of cases in time follows the familiar spikes: each lockdown is so harsh that as long as it is in place, we have $R_e < 1$ even when $x = 0$. Therefore, a few days upon imposing the lockdown (the pre-infectious period) the infected cases rapidly drop, then the critical cases drop too and the disease can possibly get eradicated in some communities. This happens over a short period of time and only a few people acquire immunity ($x \approx 0$). When the trigger $\tau = \tau_{\text{high}}$ is high, the patience of $d = 10$ days is insufficient to eradicate the disease completely, the lockdown is lifted too early, a subsequent spike of similar shape is likely, and the whole cycle repeats several times. When the trigger $\tau = \tau_{\text{low}}$ is low, waiting for $d = 10$ days will typically suffice to eradicate the disease completely within the community and no subsequent spikes occur.

Moderate policies. Under the two moderate policies, the typical stochastic trajectories are different. The lockdown is so gentle that when $x = 0$, we have $R_e \approx 1$. Hence, upon imposing a moderate lockdown, the number of ill individuals becomes roughly constant in time. But as time goes by and the individuals progress through the disease stages and acquire immunity ($x > 0$), the value of R_e decreases and the disease starts to die out ($R_e < 1$). Therefore, compared to the severe policies, the first peak is substantially broader. (However, the taller it is, the less apparent this distinction is.) Crucially, if the lockdown is lifted too soon and another outbreak occurs later, once the same moderate lockdown is imposed again, the immune subpopulation ($x > 0$) causes the disease to die out right away—in fact, it dies away faster and faster in every subsequent lockdown. Therefore, while the first outbreak might require a long lockdown phase, all subsequent outbreaks are dealt with promptly. In a sense, when the moderate lockdown is lifted for the first time, the population had already acquired herd immunity level for the interaction rate k of the moderate lockdown. This means that when the lockdown is put in place, the disease does not spread. However, once the lockdown is lifted and individuals start to interact more frequently, the disease could start spreading again. For our parameters, when the trigger $\tau = \tau_{\text{high}}$ is high, the lockdown is not strong enough and the bed capacity is (slightly) exceeded within the first peak. When the trigger $\tau = \tau_{\text{low}}$ is low, the critical cases stay safely below the bed capacity. In both cases, the patience $d = 10$ days is insufficient to completely eradicate the disease and subsequent peaks occur—all substantially smaller than the first one.

How parameters affect the performance. In order to understand the role of the three parameters (trigger value τ , severity k , patience d) on each of the four performance measures (peak size C_{\max} , overflow probability p_{fail} , total number of critical cases C_{all} , lockdown duration D), we run exhaustive computer experiments.

In Fig. 3, each row shows a different performance measure (y -axis) as a function of the number k of daily contacts in a lockdown (x -axis). The red dotted vertical line marks the number k^* of daily contacts that corresponds to $R_e = 1$ (when there are no immune individuals, $x = 0$). Within each row, the left panel shows policies that have low patience ($d = 7$ days) and the right panel shows policies that have high patience ($d = 70$ days). Within each panel, the blue curve shows the low-trigger policies ($\tau_{\text{low}} = 3$) and the green curve shows the high-trigger policies ($\tau_{\text{high}} = 12$).

Recall that for each performance measure, the lower the value, the better the performance. The effect of the trigger value can be seen by comparing the blue and the green curve: since the blue curve is typically lower, low-trigger policies are generally better than high-trigger policies. The effect of the lockdown severity can be seen by observing the performance curves as k decreases: since the curves decrease at lower k values, severe policies are generally better than moderate policies. The effect of the patience can be seen by comparing the left and the right panel: since the curves in the right panel are typically lower, patient policies are generally better than impatient policies.

There are two exceptions to those rules, both concerning the lockdown duration D : first, we observe that when k becomes too large, the duration D decreases. This is due to the fact that the lockdown becomes too weak and the disease quickly sweeps through the whole population. Second, we observe that when the patience is low and the lockdown is moderate, decreasing the trigger value τ actually leads to more time spent in the lockdown.

Key parameters for different regimes. Next, for each performance measure, we characterize which of the three parameters τ, k, d are key to substantially improving the performance and which of them are marginal.

Let k^* be the number of daily contacts that corresponds to $R_e = 1$ when there are no immune individuals ($x = 0$). For our parameters, we have $k^* \doteq 5.3$.

Peak size C_{\max} . We observe that (see Fig. 3a):

- τ : The low-trigger policies (blue curve) are consistently better than the high-trigger policies (green curve).
- k : For both trigger values, C_{\max} is roughly constant as long as $k < k^*$ but then it increases rapidly when $k > k^*$.
- d : The left and the right panel are comparable.

Hence, the important insight is to have the severity below a threshold ($k < k^*$) and to have the trigger τ low. The effects of the patience d and the severity k (given that $k < k^*$) are marginal.

The intuitive explanation is that when $k > k^*$, the lockdown is so weak that the disease still continues to spread, even when the lockdown is put in place. Hence, having $k < k^*$ is key. On the other hand, given that $k < k^*$, the actual value of k is not that important: when a lockdown is put in place, the peak size (and the moment when the peak occurs) have already been essentially determined, since most of the critical cases at the peak are due to individuals who have already been infected when the lockdown was put in place. Similarly, the patience d is not that important as it affects the number of peaks rather than their size. On the other hand, the trigger value τ is key: the nature of the exponential growth and the inherent delay due to pre-infectious period and non-critical infection translate the difference in trigger value τ to a difference in peak size C_{\max} .

Overflow probability p_{fail} . We observe that (see Fig. 3b):

- τ : The low-trigger policies (blue curve) are consistently better than the high-trigger policies (green curve).
- k : For low-trigger policies, p_{fail} exhibits a threshold behavior with respect to k . For high-trigger policies, p_{fail} increases when $k < k^*$ and increases rapidly when $k > k^*$.
- d : For high-trigger policies (green curve), increasing the patience d decreases p_{fail} (when $k < k^*$).

Hence, the important insight is, again, to have the severity below a threshold (k below k^* or just very slightly above) and to have the trigger τ low. When the trigger τ is high, both decreasing k and increasing d help, but even the combined effect is negligible compared to the effect of having τ low.

The intuitive explanation is that, in large populations, most stochastic trajectories are qualitatively similar. Hence, even though the peak size is a random variable, it is narrowly concentrated around its average value. Thus, whenever the average peak size C_{\max} slightly exceeds the available bed capacity c , the overflow probability is almost 1. And, vice versa, whenever the average peak size C_{\max} is slightly lower than the available bed capacity c , the overflow probability is almost 0. In other words, increasing the bed capacity c typically does not decrease p_{fail} —unless we exactly cross from the regime $c < p_{\text{fail}}$ to $c > p_{\text{fail}}$, in which case it helps dramatically. In a sense, the overflow probability exaggerates the difference between C_{\max} and c .

Total critical cases C_{all} . We observe that (see Fig. 3c):

- τ : When impatient (left panel), the low-trigger policies (blue curve) are consistently better than the high-trigger policies (green curve).

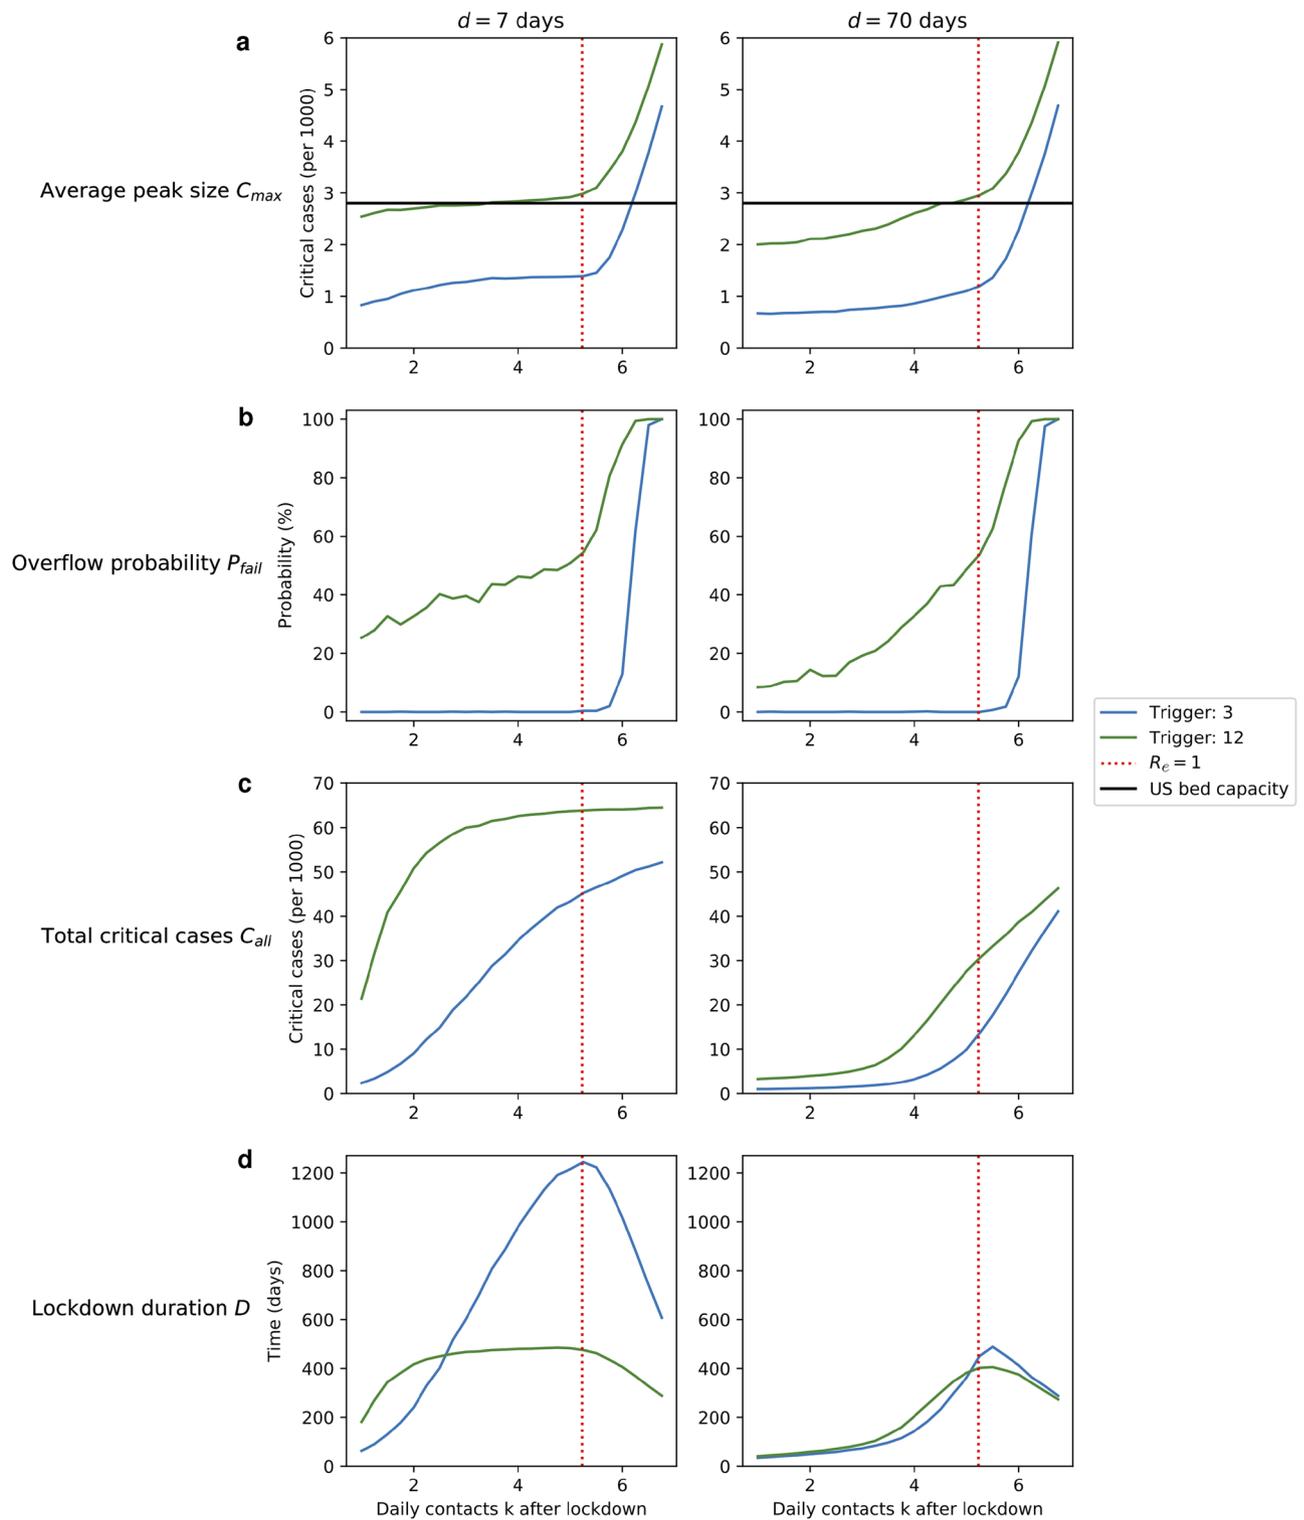


Figure 3. Performance of low-trigger and high-trigger policies. The four performance measures: **(a)** average peak size C_{max} , **(b)** overflow probability p_{fail} , **(c)** total critical cases C_{all} , and **(d)** lockdown duration D (10^5 runs). In each panel, we vary the number k of daily contacts (x-axis) and consider the performance (cost) of the low-trigger policies ($\tau_{low} = 3$, blue) and of the high-trigger policies ($\tau_{high} = 12$, green), when the patience parameter is low ($d = 7$ days, left column) and high ($d = 70$, right column). The dotted red line shows the number k^* of daily contacts that corresponds to the effective reproductive rate R_e equal to 1 (when no individuals have yet recovered). Generally speaking, it is beneficial to have the trigger value τ low (blue curves are below green ones), to impose severe rather than moderate lockdown (all curves are increasing functions of k for $k \leq k^*$), and to be patient (the curves in the right panels are lower). For C_{max} and p_{fail} , the key is to have the trigger value τ low. For C_{all} and D , the key is to have the patience d high.

- k : For impatient high-trigger policies (green curve in the left panel), C_{all} is roughly constant unless k is small.
- d : Increasing the patience d decreases C_{all} .

Hence, the important insight is to have the patience d high. In order to achieve comparable results with low patience, one has to have a low trigger τ and a very severe lockdown.

The intuitive explanation is that high patience d is key because it decreases the chance of a premature reopening and thereby reduces the number of times a lockdown has to be put in place. With a sufficiently high patience, the disease gets eradicated within the community upon completing the first peak. With low patience d (and, thus, many peaks), the only way to avoid many critical cases is to make sure each peak is small. This requires a low trigger value τ and severe lockdown k every time the trigger value is reached. As a final remark, note that when impatient and high-trigger, the lockdowns have to be extremely severe to help even a little: this is the only way to get at least some hope that the disease gets eradicated by the short time the lockdown will be lifted.

Total lockdown duration D . Since a severe lockdown is very different from a moderate lockdown, here we focus our comparison on only those lockdowns that have the same severity. We observe that (see Fig. 3d):

- τ : When impatient (left panel), low-trigger policies (blue curve) are better assuming k is small, otherwise high-trigger policies (green curve) are better. When patient (right panel), low- and high-trigger policies are comparable.
- d : Increasing the patience d decreases D .

Hence, for severe lockdowns ($k < k^*$), the important insight is, again, to have the patience d high. In those cases, the effect of the trigger value is marginal.

The intuitive explanation for severe lockdowns of fixed severity $k < k^*$ is that the key to minimizing the total lockdown duration is to minimize the probability q of a subsequent outbreak (that would occur if the lockdown were lifted too early). Since in a severe lockdown, the numbers of infected and exposed individuals decay exponentially ($k < k^*$), there are two ways to decrease q (and thereby D) by a constant factor: either to increase the patience by a constant number of days, or to decrease the trigger value τ by a constant factor. The former is less costly, so in this case having high patience is the single most important aspect.

Summary. Here we summarize three findings from the above paragraphs. First, any successful policy must impose lockdowns that restrict the number k of daily contacts under the threshold k^* . Second, in terms of minimizing either the peak size C_{max} or the overflow probability p_{fail} , it is crucial to employ a policy with low trigger threshold τ . Third, in terms of minimizing either the total number C_{all} of critical cases or the lockdown duration D , it is crucial to employ policies with high patience d . The precise optimal value of patience depends on the severity of the lockdown and on the trigger threshold τ .

Two countries. Some interaction among communities is inevitable. While the inter-community interaction can be limited by closing borders between countries and imposing quarantine upon entry, it can not be completely disregarded. We study how the policies perform in the environment where different communities might employ different policies.

To model this, we consider two communities Country 1 and Country 2 that experience the onset of the disease at the same time. Occasionally, individuals from different countries meet. Namely, we assume that for each individual, a small portion $q \in (0, 1)$ of their interactions are with individuals in the other country.

We consider two very different policies and study how their performance depends on the policy employed by the neighboring country and on the interaction rate q between the two countries (see Fig. 4).

Specifically, we consider a moderate low-trigger policy $P_{ML} = P(\tau = 3, k = 6, d = 54)$ and a severe high-trigger policy $P_{SH} = P(\tau = 12, k = 1.25, d = 27)$. The parameters τ and k are chosen such that both policies have the same probability 10% of exceeding the US hospital bed capacity within their first peak. The parameter d is chosen such that, on average, 90% of the critical cases occur within the first peak. In other words, the policy reopens when it expects that 90% of all the critical cases have already been hospitalized.

When both countries employ the same policy, the performance is essentially the same as for a larger country employing that policy. Also, in the limit $q \rightarrow 0$ the two countries do not interact and the performance of a country is independent of the policy of the neighbor (this remains true for $q < 10^{-5}$). However, when q is non-negligible, one country uses P_{ML} (“ P_{ML} -country”) and the other one uses P_{SH} (“ P_{SH} -country”), the two policies clash. Specifically, we make the following observations about the overflow probability p_{fail} and the average peak size C_{max} (see Fig. 4):

1. For q small, the green curve increases as q increases: note that the infectious subpopulation of the neighboring P_{ML} -country is non-negligible for an extended period of time (at least throughout the first broad peak). Thus, as q increases, the individuals in the P_{SH} -country get repeatedly infected due to interactions with the P_{ML} -country. Most such new infections cause a new spike for the P_{SH} -country. Each such spike might exceed the previously largest peak and/or the available capacity. (Moreover, each such spike leads to new critical cases and it has to be contained by another lockdown phase so it is costly in terms of C_{all} and D too.) This effect is visible for interaction rates as small as $q = 10^{-4}$.
2. For q large, the blue and green curve decrease as q increases: two countries employing different policies typically reach their peaks at a different point in time. Thus, when one country is peaking, the other country likely

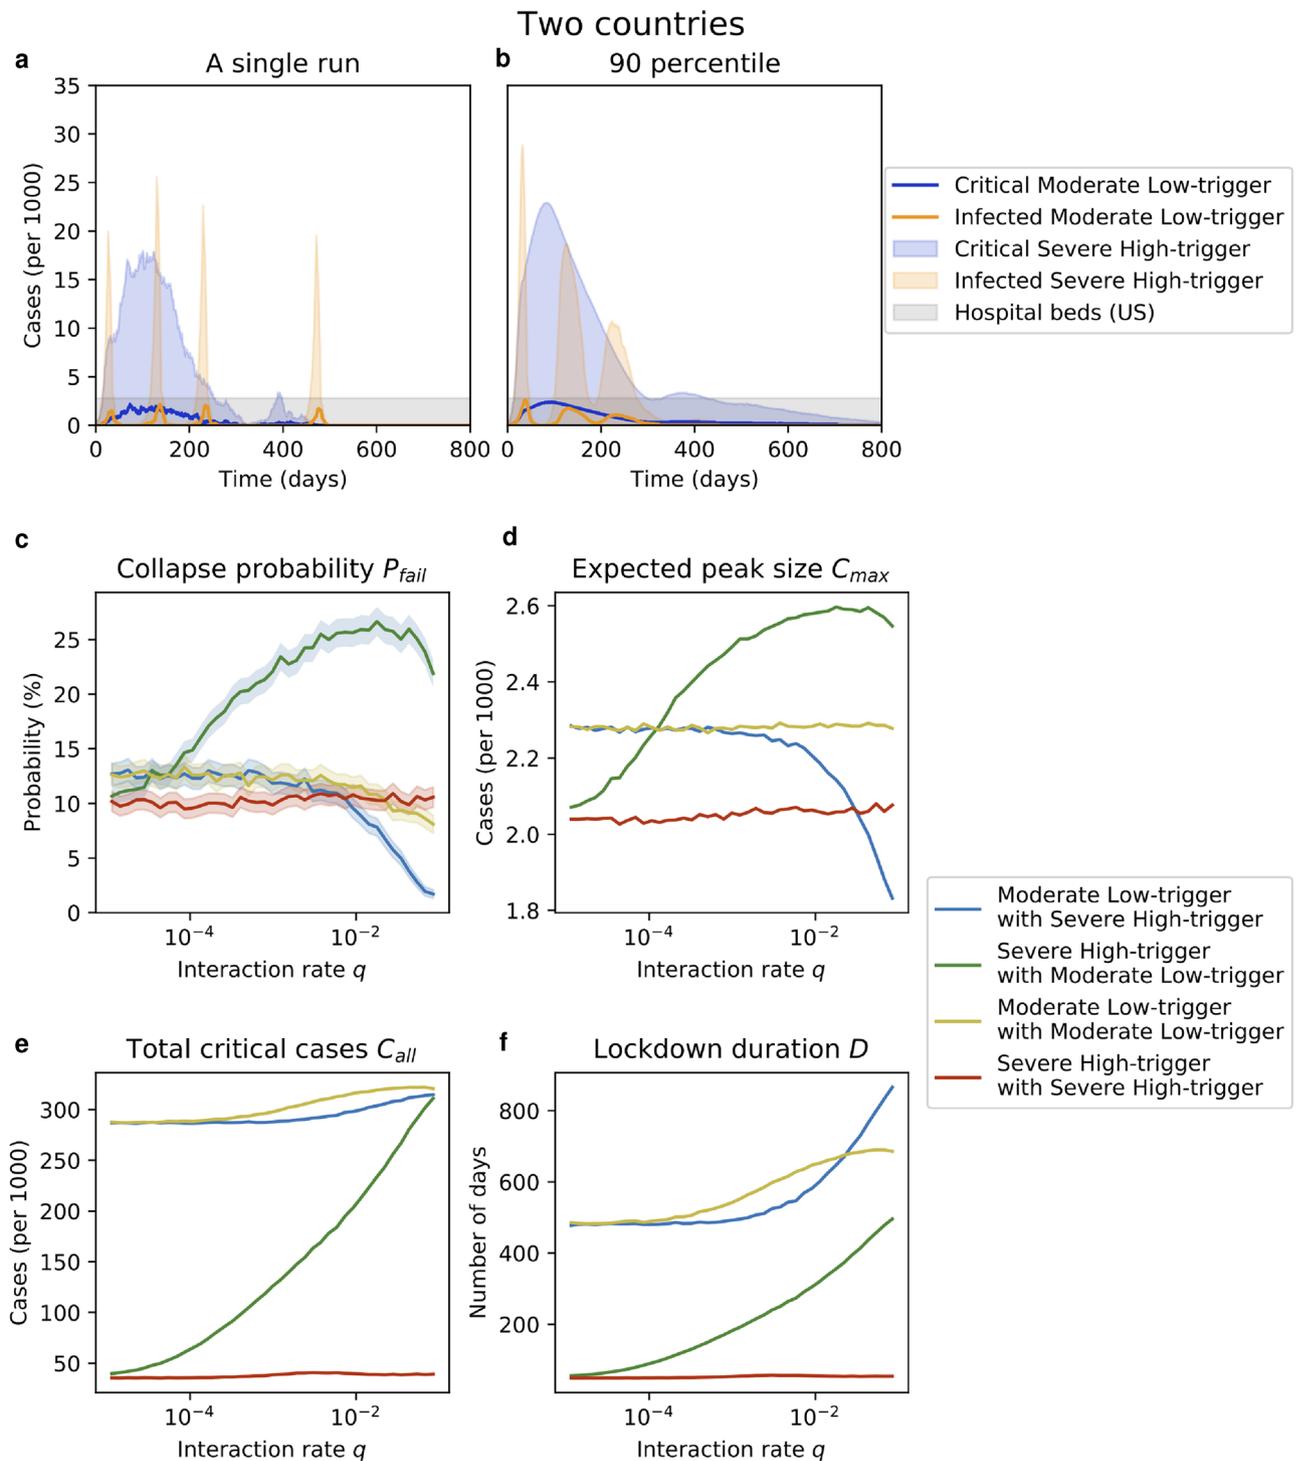


Figure 4. Two countries. (a) When two countries interact with a positive rate q they might reinfest each other. Here $q = 5 \cdot 10^{-4}$, and the countries employ different policies: P_{ML} (blue) and P_{SH} (yellow). (b) A 90% percentile: on any given day, 90% of the runs are below the respective curves (10^5 runs). (c–f) The performance of a P_{ML} policy against a P_{SH} policy (blue), P_{SH} vs. P_{ML} (green), P_{ML} vs. P_{ML} (yellow) and P_{SH} vs. P_{SH} (red), averaged over 10^4 runs. We vary the interaction rate q on a log-scale and measure: (c) the overflow probability P_{fail} (95% confidence intervals are shaded); (d) the expected peak size C_{max} ; (e) the total number C_{all} of critical cases; and (f) the lockdown duration D . A country employing P_{SH} does great when its neighbor employs P_{SH} (red) but bad when the neighbor employs P_{ML} (green). A country employing P_{ML} does comparably well, regardless of whether the neighbor employs P_{ML} (yellow) or P_{SH} (blue).

- has fewer infectious individuals and an interaction with that country will help alleviate the size of the peak in the first country. This, in turn, decreases the maximum peak load C_{\max} and the overflow probability p_{fail} . This effect is visible roughly for $q > 10^{-3}$ (blue curve), resp. $q > 10^{-1}$ (green curve).
- The yellow and the red curves are roughly constant, except that the yellow curve goes down in terms of the overflow probability when q is large: when two neighboring countries employ the same policy, the extra occasional mixing due to interacting individuals makes both countries behave in a slightly more average way. This does not change the expected size of the peak, and hence C_{\max} is constant, regardless of the interaction rate q . However, this process of “averaging out” does make the extreme behavior, such as overflowing of the available hospital capacity, somewhat less likely. This effect of diminished overflow probability is observed for two P_{ML} countries when $q > 10^{-2}$.

This leads to an interesting phenomenon resembling a social dilemma^{44,45}. Consider two countries with a non-negligible interaction rate $q > 10^{-3}$. First, if either of the two countries primarily cares about keeping the overflow probability low then that country would employ the policy P_{ML} rather than the policy P_{SH} , no matter which policy the other country is using (indeed blue is below red and yellow is below green). Second, given that one country uses the policy P_{ML} , the other country will use the policy P_{ML} too (yellow is below green). In terms of the overflow probability alone, this is an acceptable outcome (red and yellow are comparable) but in terms of the total number of critical cases C_{all} , this is undesirable: by employing the policy P_{ML} rather than P_{SH} , both countries increase the total number of their critical cases (and, possibly, deaths) by an order of magnitude.

Discussion

Motivated by the COVID-19 pandemic, we studied a simple stochastic model of a disease progression in a population of interacting individuals. We focused on a 3-parameter family of policies that can be used to mitigate the disease spread and evaluated the performance of those policies with respect to several measures, such as the number of critical cases at its maximum or the probability that this number exceeds the available hospital bed capacity. The three parameters describing the policies correspond to the agility when closing down, the severity of the lockdown, and the patience when reopening. We identified which parameters are important in which regime and explained why some policies are performing better than others. We note that understanding the dynamics of periodic lockdown is important in case the virus escapes vaccination or for future epidemics.

We highlight two different types of realistic policies, called *moderate low-trigger* (P_{ML}) and *severe high-trigger* (P_{SH}). With both policies, the probability p_{fail} of ever exceeding the available hospital beds is kept below a specified threshold (here arbitrarily set to 10%), but the two policies are very different: the P_{SH} policy is characterized by imposing a harsh, short lockdown (“severe”) at the last moment possible (“high-trigger”), whereas the P_{ML} policy is proactive and imposes gentle, longer lockdowns (“moderate”) at the first signal of an approaching outbreak (“low-trigger”).

Due to the above differences, both policies have their advantages and disadvantages. The P_{SH} policy minimizes the total number of critically ill cases and the total amount of time spent in lockdown. However, the lockdowns are severe and the society is more susceptible to any subsequent outbreaks. To avoid such recurring outbreaks, the authorities must be patient when reopening and any neighboring countries must coordinate when releasing their measures. The P_{ML} policy maintains its moderate lockdown for substantially longer but its performance is substantially more robust with respect to how soon the society reopens and with respect to what policies are employed by the neighboring countries. Moreover, upon completing the first lockdown phase, all subsequent lockdowns (if any) are shorter and involve substantially fewer critical cases than the first phase. Thus, the P_{ML} policy can be seen as minimizing the long-term risks under the pessimistic scenario that an ultimate long-term solution (such as a majority of the population being vaccinated) is not achieved any time soon. On the other hand, the P_{SH} policy is optimistic about the future and optimizes the short-term performance.

Our setup is intentionally simplified in many regards such as the disease progression within an individual, the disease spread throughout the population, and the class of policies we consider for closing down and reopening. We highlight four possible extensions worth pursuing in subsequent work. First, regarding the disease progression within an individual, one can distinguish more types of individuals (e.g. those who require only hospital beds and those who moreover require ventilation) and lift the assumption that the individuals acquire lifelong immunity. We note that our model already implicitly includes asymptomatic carriers, since those are equivalent to individuals who developed a mild condition and then recovered. Second, regarding the disease spread in a population, one can consider a population structure, e.g. described by a graph whose edge weights determine the daily pairwise transmission probabilities. This would allow one to investigate the effects of localized interventions such as contact tracing. Third, one can consider more complicated policies, e.g. policies that allow for a gradual reopening or policies that, when deciding whether and how much to reopen, take into account additional information, such as the situation in the neighboring countries and/or the outcomes of testing done earlier. Fourth, on top of considering the health viewpoint, one can incorporate the economic viewpoint by introducing an appropriate notion of economic cost of a lockdown of varying degree of severity.

Received: 23 March 2021; Accepted: 5 January 2022

Published online: 27 January 2022

References

- Bernoulli, D. & Blower S. An attempt at a new analysis of the mortality caused by smallpox and of the advantages of inoculation to prevent it. *Rev. Med. Virol.* **14**, 275–288 (2004).

2. Hamer, W. H. *Epidemic Disease in England: The Evidence of Variability and of Persistency of Type* (Bedford Press, 1906).
3. Kermack, W. O. & McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. Ser. A Contain. Pap. Math. Phys. Character* **115**, 700–721 (1927).
4. Bailey, N. T. *et al.* *The Mathematical Theory of Infectious Diseases and Its Applications* (Charles Griffin & Company Ltd, 1975).
5. Diekmann, O., Heesterbeek, J. A. P. & Metz, J. A. On the definition and the computation of the basic reproduction ratio r_0 in models for infectious diseases in heterogeneous populations. *J. Math. Biol.* **28**, 365–382 (1990).
6. Anderson, R. M. & May, R. M. *Infectious Diseases of Humans: Dynamics and Control* (Oxford University Press, 1992).
7. Dietz, K. The estimation of the basic reproduction number for infectious diseases. *Stat. Methods Med. Res.* **2**, 23–41 (1993).
8. Nowak, M. A. & May, R. M. Superinfection and the evolution of parasite virulence. *Proc. R. Soc. Lond. Ser. B Biol. Sci.* **255**, 81–89 (1994).
9. Nowak, M. A. & Bangham, C. R. Population dynamics of immune responses to persistent viruses. *Science* **272**, 74–79 (1996).
10. Diekmann, O. & Heesterbeek, J. *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation. Wiley Series in Mathematical and Computational Biology* (Wiley, 2000).
11. Nowak, M. & May, R. M. *Virus Dynamics: Mathematical Principles of Immunology and Virology* (Oxford University Press, 2000).
12. Hethcote, H. W. The mathematics of infectious diseases. *SIAM Rev.* **42**, 599–653 (2000).
13. Brauer, F., Castillo-Chavez, C. & Castillo-Chavez, C. *Mathematical Models in Population Biology and Epidemiology* Vol. 2 (Springer, 2012).
14. Anderson, R. M., Heesterbeek, H., Klinkenberg, D. & Hollingsworth, T. D. How will country-based mitigation measures influence the course of the COVID-19 epidemic?. *Lancet* **395**, 931–934 (2020).
15. Ferguson, N. *et al.* Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. Tech. Rep., Imperial College (2020).
16. Fauci, A. S., Lane, H. C. & Redfield, R. R. Covid-19-navigating the uncharted (2020).
17. Guan, W.-J. *et al.* Clinical characteristics of coronavirus disease 2019 in China. *N. Engl. J. Med.* **382**, 1708–1720 (2020).
18. Zhou, F. *et al.* Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: A retrospective cohort study. *Lancet* **395**, 1054–1062 (2020).
19. Grasselli, G. *et al.* Baseline characteristics and outcomes of 1591 patients infected with SARS-CoV-2 admitted to ICUs of the Lombardy region, Italy. *JAMA* **323**, 1574–1581 (2020).
20. Bar-On, Y. M., Flamholz, A., Phillips, R. & Milo, R. Science forum: SARS-CoV-2 (COVID-19) by the numbers. *eLife* **9**, e57309 (2020).
21. Li, Q. *et al.* Early transmission dynamics in Wuhan, China, of novel coronavirus-infected pneumonia. *N. Engl. J. Med.* **382**, 1199–1207 (2020).
22. Xu, B. *et al.* Epidemiological data from the COVID-19 outbreak, real-time case information. *Sci. Data* **7**, 1–6 (2020).
23. Wu, J. T. *et al.* Estimating clinical severity of COVID-19 from the transmission dynamics in Wuhan, China. *Nat. Med.* **26**, 506–510 (2020).
24. Kucharski, A. J. *et al.* Early dynamics of transmission and control of COVID-19: A mathematical modelling study. *Lancet Infect. Dis.* **20**, 553–558 (2020).
25. Leung, K., Wu, J. T., Liu, D. & Leung, G. M. First-wave COVID-19 transmissibility and severity in China outside Hubei after control measures, and second-wave scenario planning: A modelling impact assessment. *Lancet* **395**, 1382–1393 (2020).
26. Zhang, J. *et al.* Changes in contact patterns shape the dynamics of the COVID-19 outbreak in China. *Science* **368**, 1481–1486 (2020).
27. Rader, B. *et al.* Crowding and the shape of COVID-19 epidemics. *Nat. Med.* **26**, 1829–1834 (2020).
28. Kraemer, M. U. *et al.* The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science* **368**, 493–497 (2020).
29. Tian, H. *et al.* An investigation of transmission control measures during the first 50 days of the COVID-19 epidemic in China. *Science* **368**, 638–642 (2020).
30. Ferretti, L. *et al.* Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing. *Science* **368**, eabb6936 (2020).
31. Flaxman, S. *et al.* Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* **584**, 257–261 (2020).
32. Kissler, S. M., Tedijanto, C., Goldstein, E., Grad, Y. H. & Lipsitch, M. Projecting the transmission dynamics of SARS-CoV-2 through the postpandemic period. *Science* **368**, 860–868 (2020).
33. Walker, P. G. *et al.* The impact of COVID-19 and strategies for mitigation and suppression in low-and middle-income countries. *Science* **369**, 413–422 (2020).
34. Morris, D. H., Rossine, F. W., Plotkin, J. B. & Levin, S. A. Optimal, near-optimal, and robust epidemic control. arXiv preprint [arXiv:2004.02209](https://arxiv.org/abs/2004.02209) (2020).
35. Ashcroft, P., Lehtinen, S., Angst, D. C., Low, N. & Bonhoeffer, S. Quantifying the impact of quarantine duration on COVID-19 transmission. *eLife* **10**, e63704 (2021).
36. Gostic, K. M. *et al.* Practical considerations for measuring the effective reproductive number R_t . *PLOS Comput. Biol.* **16**, e1008409 (2020).
37. Castro, M., Ares, S., Cuesta, J. A. & Manrubia, S. The turning point and end of an expanding epidemic cannot be precisely forecast. *Proc. Natl. Acad. Sci.* **117**, 26190–26196 (2020).
38. Rajgor, D. D., Lee, M. H., Archuleta, S., Bagdasarian, N. & Quek, S. C. The many estimates of the COVID-19 case fatality rate. *Lancet Infect. Dis.* **20**, 776–777 (2020).
39. Verity, R. *et al.* Estimates of the severity of coronavirus disease 2019: A model-based analysis. *Lancet Infect. Dis.* **20**, 669–677 (2020).
40. Ruan, S. Likelihood of survival of coronavirus disease 2019. *Lancet Infect. Dis.* **20**, 630–631 (2020).
41. Rhodes, A. *et al.* The variability of critical care bed numbers in Europe. *Intensive Care Med.* **38**, 1647–1653 (2012).
42. Ma, X. & Vervoort, D. Critical care capacity during the COVID-19 pandemic: Global availability of intensive care beds. *J. Crit. Care* **58**, 96 (2020).
43. Phua, J. *et al.* Intensive care management of coronavirus disease 2019 (COVID-19): Challenges and recommendations. *Lancet Respir. Med.* **8**, 506–517 (2020).
44. Nowak, M. A. Five rules for the evolution of cooperation. *Science* **314**, 1560–1563 (2006).
45. Sigmund, K. *The Calculus of Selfishness* (Princeton University Press, 2016).

Acknowledgements

K.C. acknowledges support from ERC Consolidator Grant No. (863818: ForM-SMart). A.P. acknowledges support from FWF Grant No. J-4220. M.A.N. acknowledges support from Office of Naval Research grant N00014-16-1-2914 and from the John Templeton Foundation.

Author contributions

J.S., J.T., A.P., K.C. and M.A.N. designed the research, performed the analysis and wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-05333-5>.

Correspondence and requests for materials should be addressed to M.A.N.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022