OPTIMAL CONTROL OF AN EPIDEMIC THROUGH SOCIAL DISTANCING

By

Thomas Kruse and Philipp Strack

April 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2229

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Optimal Control of an Epidemic through Social Distancing*

Thomas Kruse† Philipp Strack‡

April 2020

Abstract

We analyze how to optimally engage in social distancing (SD) in order to minimize the spread of an infectious disease. We identify conditions under which the optimal policy is single-peaked, i.e., first engages in increasingly more social distancing and subsequently decreases its intensity. We show that the optimal policy might delay measures that decrease the transmission rate substantially to create "herd-immunity" and that engaging in social distancing sub-optimally early can increase the number of fatalities. Finally, we find that optimal social distancing can be an effective measure in substantially reducing the death rate of a disease.

Keywords— Social Distancing, SIR model, Time-Optimal Control of an Epidemic

1 Introduction

This paper analyzes how to optimally engage in measures to contain the spread of an infectious disease. We formalize this question in the context of a standard model from epidemiology, the Susceptible-Infected-Recovered (SIR) model (Kermack and McKendrick, 1927). This model divides the population into three groups susceptible, infected and recovered, and people transition from one group into another at given exogenously specified rates depending on the size of each sub-population. We extend this model by allowing an additional parameter controlled by the planer that affects the rate at which the disease is transmitted. We think of this parameter as capturing political measures such as social distancing, and the lockdown of businesses, schools, universities and other institutions. While such measures reduce the spread of the

*We want to thank Stefan Ankirchner and Aleh Tsyvinski for helpful discussions. We plan on updating this paper frequently, the current version of the paper can be found at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3553363

†Giessen University, thomas.kruse@math.uni-giessen.de

‡Yale University, philipp.strack@yale.edu

disease, they often come at a substantial economic and social cost. We model this trade-off by considering a planner who faces convex cost in the number of infected (capturing the number of people whose death is caused by the disease) and the reduction in transmission rate (capturing the cost of shutting down society).

Our analysis identifies several features of any optimal policy. First, whenever a constant fraction of those who are infected dies, the optimal policy is single peaked in the sense that first the measures to reduce the transmission rate are escalated until some point in time, and after this point in time these measures are reduced. Second, if the cost of reducing the transmission rate is linear, meaning that closing half of society for two days is equally costly as closing all of society for one day, only the most extreme policies are used. Either, the planner imposes the maximal possible lockdown or no restrictions at all. Intuitively, the planner can achieve a greater effect by imposing a more extreme policy for a shorter time and thus does not find it optimal to use intermediate policies. These results imply that for linear cost the optimal policy has a simple structure and consists of three phase: first it imposes no restrictions then it imposes as many restrictions as possible, and finally in the third phase imposes no restrictions at all. This result drastically simplifies the search for an optimal policy as the planner has to only optimize over the start and end time of the social distancing period.

We then calibrate our model to the current Covid-19 epidemic to illustrate some further insights. We first characterize the optimal timing of the social distancing period given that the planner has access to a certain budget of days of social distancing. We find that the optimal social distancing is often substantially delayed. For example, if the planer has a budget of 100 days of social distancing in the next 360 days after 0.1% of the population are infected it is optimal to delay social distancing by 50 days. This initial period of letting the disease spread uncontrolled is useful as it creates "herd immunity" and thereby reduces the overall severity of the epidemic. We show by an example that the benefit of herd immunity is so strong that sometimes more social distancing can *increase* the number of people that die from the epidemic. We show that in this example more people die when social distancing is imposed from day 0-100 compared to day 50-100. As this example suggests, benefit of optimally timing social distancing measures is often large and we illustrate this by comparing social distancing in the first t days after 0.1 percent are infected to t days of optimally timed social distancing. Finally, we quantify the optimal amount of social distancing. We find that for parameters commonly used to describe the spread of Covid-19 that when one assumes a value of a life of 10 million and that social distancing reduces the transmission rate by 60% that the optimal policy starts social distancing almost immediately and maintains it for around 300 days.

Related Literature Our theoretical results extend the literature on the optimal control of SIR epidemic models (for an overview see chapter 5 in Wickwire, 1977). Abakuks (1973) consider the question of how to optimally isolate infectious population if infectious population can be instantaneously isolated. Abakuks (1972, 1974) determine the optimal vaccination strategy in the same framework. Morton and Wickwire (1974); Wickwire (1975) extend the previous work on vaccination and isolation by considering flow controls. Behncke (2000) considers more general functional forms and Hansen and Day (2011) allows for hard bounds

on the control, while considering vaccination and isolation policies simultaneously. To the best of our knowledge the first paper that considers social distancing or lockdown policies is (Behncke, 2000, Chapter 4). This paper suggests an elegant reduction of the problem by establishing that the optimal policy depends only on the shadow price difference between infected and susceptible. This insight is also an important building block for our analysis (see Proposition 1). Based on it we are able to reduce the dimension of the dynamics of the problem and provide a much more explicit characterization of the optimal policy in the linear case (see Proposition 2 and 3).

Finally, our paper also relates to the recent literature that numerically studies optimal policies for the current epidemic of Covid-19 in the context of SIR models (Alvarez et al., 2020; Kissler et al., 2020; Toda, 2020). Alvarez et al. (2020) numerically characterize the optimal lockdown policy for the current covid pandemic in a similar SIR model. The paper Kantner (2020) considers an extended SEIR model and numerically analyzes social distancing policies that minimize disease-related deaths while establishing a desired degree of herd immunity at the same time. Kissler et al. (2020) numerically compare various lockdown policies and allow for seasonality effects. Toda (2020) estimates the transmission rate in the context of an SIR model with fixed transmission rate for various countries, compares various SD policies numerically, and considers asset prices during an epidemic. While it is not a goal of this paper to make any recommendations for the current Covid-19 epidemic we hope that the formal analysis and insights this paper contributes will be useful in the rapidly evolving discussion of how to optimally react to the Covid-19 epidemic (Atkeson, 2020; Barro et al., 2020; Dewatripont et al., 2020; Piguillem et al., 2020; Stock, 2020; de Walque et al., 2020).

2 The Evolution of an Epidemic

The SIR Model To model the spread of an infectious disease we rely on a basic model from epidemiology, the *Susceptible Infected Recovered* (SIR) model introduced in Kermack and McKendrick (1927). We divide society into three groups: susceptible s, infected i, and the rest which is either immune to the disease as they recovered from it or died. We denote by $s(t)$ the fraction of the population that is healthy, but susceptible to disease at time t, and by $i(t)$ the fraction of the population that is infected. The SIR model assumes the number of people that gets infected, by a single infected person is deterministic proportional to the fraction of society $s(t)$ that is still susceptible to the disease. Intuitively, if only a small fraction of society is susceptible to the disease it is unlikely that an infected person meets a susceptible person. The mass of healthy people that become infected during dt thus equals

$$\beta(t)i(t)s(t)\,,$$

where the transmission rate $\beta(t)$ captures both how infectious the disease, as well as measures society has taken to influence the speed at which the disease spreads (like social distancing). Infected become non-

infected, by either recovering from the disease, or dying of it at rate $\gamma > 0$, such that during a short time span dt, the fraction of infected is reduced by $\gamma i(t)$.The susceptible and infected populations $(s(t), i(t))_t$ thus for every $t \in [0, \infty)$ evolve according to the following dynamics

$$\begin{aligned} s'(t) &= -\beta(t) i(t) s(t), & s(0) &= s_0, \\ i'(t) &= \beta(t) i(t) s(t) - \gamma i(t), & i(0) &= i_0, \end{aligned} \tag{1}$$

where $s_0, i_0 \in (0,1)$ are given initial values satisfying $s_0 + i_0 \leq 1$.

Control of the Transmission Rate The time-dependent transmission rate $\beta \colon [0,\infty) \to B$ takes values in an compact interval $B = [\underline{b}, \overline{b}] \subset (0,\infty)$. We denote by $\overline{b}$ the maximal transmission rate and by $\underline{b}$ the minimal transmission rate that can be achieved through some policy measures. The set of admissible controls $\mathscr{B}$ consists of all measurable functions $\beta \colon [0,\infty) \to B$.

We introduce two cost functions $v \colon [0,1] \to [0,\infty)$ and $c \colon B \to [0,\infty)$ and for fixed time horizon $T \in [0,\infty)$ we consider the cost functional

$$J(\beta) = \int_0^T v(i(t)) + c(\beta(t))\, dt. \tag{2}$$

The cost $v(i)$ measures the number of people that die per unit of time if a share i of the population is infected. We suppose that v is convex, continuously differentiable and strictly increasing. Convexity of v captures the fact that the probability of dying from the disease might be higher if a large share of the population is infected and the hospital system is overwhelmed. We note that v can not only capture the people who die of the disease directly, but also those who die because other medical conditions remain untreated as an indirect consequence of the disease.

The cost function c captures the economic and social cost of measures taken to reduce the transmission rate. For example if social distancing measures are imposed which require the closure of most businesses this comes at a substantial economic cost. We only make minimal assumption on c and assume that it is convex and continuous, and without loss normalize the cost associated with the highest transmission rate to zero, $c(\overline{b}) = 0 > c(b)$ for all $b \in [\underline{b}, \overline{b})$.

The planner trades-off the number of people who die as a direct (or indirect) consequence of the disease with the economic and social cost of reducing the transmission rate. A policy β^* is *optimal* if it minimizes J over $\mathscr{B}$

$$\beta^* \in \operatorname*{arg\,min}_{\beta \in \mathscr{B}} J(\beta)\,. \tag{3}$$

3 The Optimal Policy

The next result shows existence of an optimal policy and provides necessary conditions that any solution of the optimal control problem (3) must satisfy.

Proposition 1. *An optimal policy exists. Let $\beta^* \in \mathscr{B}$ be such an optimal strategy and denote by $s^*, i^* \colon [0,T] \to [0,1]$ the associated state processes satisfying (1). Then there exists a function $\eta^* : [0,T] \to \mathbb{R}$ with $\eta^*(T) = 0$ such that for almost all $t \in [0,T]$ it holds*

$$\begin{aligned}(\eta^*)'(t) &= \eta^*(t)\beta^*(t)i^*(t) + \frac{v(i^*(t)) + c(\beta^*(t)) - v(i^*(T))}{i^*(t)} - v'(i^*(t)),\\ (\beta^*)(t) &\in \arg\min_{b \in B} \left[\eta^*(t)i^*(t)s^*(t)b + c(b)\right].\end{aligned} \tag{4}$$

Moreover, we have $\eta^*(t) > 0$ for all $t \in [0,T)$.

The proofs of all results of the paper are presented in the Appendix. The proof of Proposition 1 relies on a sequence of auxilliary results we establish in the appendix using standard arguments from control theory that can, e.g., be found in Clarke (2013). The existence of an optimal policy follows as the convexity of c and B ensures compactness of the policy space which leads to the existence of an optimal policy. Pontryagin's optimality principle then yields that for every optimal policy there exist two Lagrange multipliers λ_1^*, λ_2^* such that the optimal control is only a function of these multipliers. An argument similar to Behncke (2000) implies that these two Lagrange multipliers can be summarized into a single variable $\eta^* = \lambda_2^* - \lambda_1^* > 0$ that completely determines the optimal policy according to (4). This Lagrange multiplier $\eta^*(t)$ has a clear interpretation as marginal increase in the cost from infecting susceptible population. The fact that $\eta^*(t) > 0$ reflects the fact that the planner always benefits from having fewer infected. Finally, Proposition 1 provides a novel characterization of η^* as the solution to an ODE, which allows us to explicitly derive features of the optimal policy in Section 3.1.

3.1 Linear Costs

In this section we impose additional linearity assumptions on the cost to provide further insight into the structure of the optimal policy. Again we suppose that $\beta^* \in \mathscr{B}$ is an optimal control and denote by $s^*, i^* \colon [0,T] \to [0,1]$ the associated state processes.

Our first type of result assumes that v is linear, which means that the fraction of infected that die from the disease is independent of the total fraction of the population that is infected at any point in time. This assumption rules out capacity effects that arise from the overload of the medical system. It is thus a reasonable assumption if the number of infected is kept within levels that do not overburden the health system.

Proposition 2. *Suppose that v is linear.*[1] *Then β^* is quasi-convex, i.e., first decreasing and then increasing.*

Proposition 2 establishes that any optimal policy is single peaked, in the sense that the measures to decrease the transmission rate are first escalated until some point in time and then reduced over time. Any policy where a reduction in measures is followed by an increase is suboptimal.

[1] There exists $\alpha > 0$ such that $v(i) = \alpha i$

Our next result establishes that if both costs v and c are linear then the optimal policy involves only the two most extreme controls. The assumption that the cost c of measures that reduce the transmission rate is linear has a simple interpretation in the context of social distancing: Shutting down half of the economy for two days is equally costly as shutting down the whole economy for a single day.[2] While we think that there is no normative reason for this assumption we think of it as a natural baseline for the analysis.

Proposition 3. *Suppose that v and c are linear.*[3] *Then for any optimal control β^* there exists $0 \leq t_1^* \leq t_2^* \leq T$ such that for a.e. $t \in [0,1]$*

$$\beta^*(t) = \begin{cases} \overline{b} & \text{for } t \in [0, t_1^*) \\ \underline{b} & \text{for } t \in [t_1^*, t_2^*] \\ \overline{b} & \text{for } t \in (t_2^*, T] \end{cases}.$$

Proposition 3 drastically simplifies the search for an optimal policy as it implies that any optimal policy is characterized by the two points in time (t_1^*, t_2^*). Note that the proposition does not rule out that any of the intervals is empty. In particular, reducing the transmission rate by the maximal amount at every point in time as well as taking no measures at all to reduce the transmission rate can be optimal. The main insight of the proposition is that under plausible assumption it is never optimal to use intermediate measure for a longer time (i.e. closing only parts of the economy) as doing so is dominated by implementing maximal measures for a shorter time.

4 An Illustration

We next illustrate how our results can be used to derive policy advice for fighting an epidemic. In this illustration we aim to choose parameters in line with the COVID-19 pandemic. An important disclaimer is that at the current point in time there is substantial uncertainty about the true parameters governing the spread of COVID-19 which substantially influence the optimal policies.

Parametric Assumptions We assume that the average length of an infection equals 18 days ($\gamma = 1/18$) and that the social planner has access to two policies $0 < \underline{b} < \overline{b}$ corresponding to *social distancing* (SD) $\underline{b}$ and *no social distancing* (NSD) $\overline{b}$. We set $\overline{b}$ to 0.16 in line with a reproduction rate of R_0 of $\overline{b}/\gamma = 2.88$. We assume that enacting social distancing reduces the number of contacts by 60% and set $\underline{b} = 0.4\overline{b}$ consistent with $R_0 = \underline{b}/\gamma = 1.152$. The fraction of infected that dies equals 0.8%, below 20 times[4] the critical care

[2] This implicitly assumes that the transmission rate β depends linearly on the shut down of the economy.

[3] There exists $\alpha, \delta > 0$ such that $v(i) = \alpha i$ and $c(\beta) = \delta(\overline{b} - \beta)$.

[4] This is we implicitly assume that 5% of infections are sufficiently severe that they need hospitalization and access to critical care.

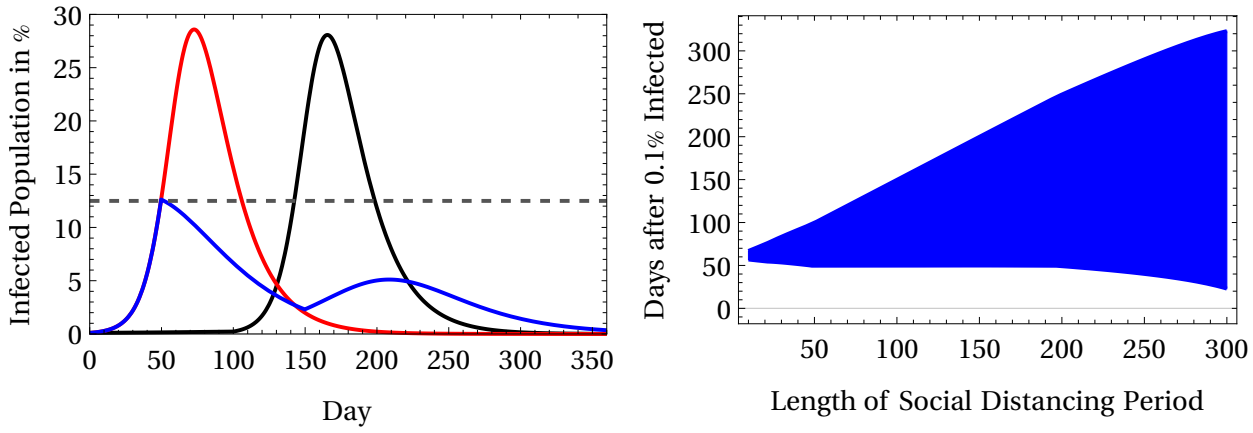

Figure 1: On the left: The number of infected over time without SD (in red), with SD from day 0-100 (in black), and for the optimal SD period from day 48 to day 148 in blue. On the right: the optimal period of SD as a function of the length of SD. The dashed line marks the point beyond which the fraction of infected who die increases due to an overload of the medical system.

bed capacity $\kappa = 0.000347$[5] and then grows linearly such that if 20% of the population is simultaneously infected 5% of infected die[6]

$$v(i) = (\gamma i) \times \left[0.008 + \frac{0.042(\gamma i - 20\kappa)}{\gamma 0.2 - 20\kappa} \mathbf{1}_{\gamma i \geq 20\kappa}\right].$$

Throughout our simulations we assume that at day zero, 0.1% of the population is infected. Finally, we assume that a cure and a vaccine for the disease arrive in one year (360 days) and no one dies of the disease afterwards.

The Optimal Timing of Social Distancing We begin by analysing the optimal timing of social distancing. In order to do so we first suppose that the planner has a fixed budget of days of social distancing and answer the question during which time period he optimally engages in social distancing. We only consider policies that consist of three subsequent periods, first NSD, follows by a period of SD, and a period of NSD. For example, consider the case where the planer has a budget of 100 days of SD. In this case the optimal policy is to start social distancing on day 48 and end it on day 148. As one can see in the left graph of Figure 1 this leads to a substantially flatter curve of infected over time than social distancing in the first 100

[5]The number of cricitcal care beds per population equals $\kappa = 0.000347$ for the US, $\kappa = 0.000292$ for Germany, and $\kappa = 0.000125$ for Italy. See https://www.sccm.org/getattachment/Blog/March-2020/United-States-Resource-Availability-for-COVID-19/United-States-Resource-Availability-for-COVID-19.pdf?lang=en-US.

[6]Note, that $v(i)$ aims not only at capturing the people who die directly as a consequence of the disease, but also those who die as they do not have access to critical care as a consequence of the overloaded medical system. Our assumption implies that if 50% of the population is simultaneously infected the death rate increases to 11.5%. We note that these are extremely pessimistic assumptions if the number of infected is substantially underestimated.

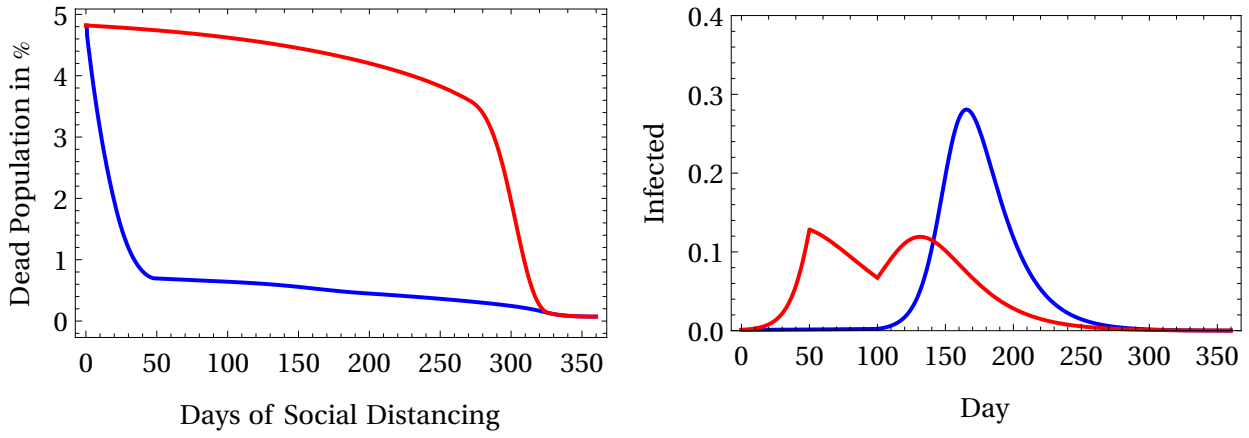

Figure 2: On the Left: Fraction of the population that dies within 360 days of 0.1% infected as a function of the length of social distancing for the optimal timing of social distancing (in blue) and social distancing starting at day 0 (in red). On the right: infected over time, if SD is exercises from day 50-100 (in blue) and from day 0-100 (in red).

days (in black) or no social distancing (in red). Interestingly, the effect of suboptimal social distancing is marginal in the sense that while it initially reduces the number of infected substantially, it essentially only delays the peak of infected, but does not substantially flatten it. This leads to a substantial reduction in the implied death rate within a year: 0.6% under optimal social distancing, 4.6% with social distancing in the first hundred days, and 4.8% without social distancing.

We next analyse how the optimal timing of social distancing depends on the length of social distancing. As one can see in the right graph of Figure 1 it is optimal to delay social distancing beyond the date where 0.1% of the population is infected. For example even if it is optimal for the planner to engage in 300 days of social distancing within the next year it is only optimal to start social distancing after 25 days. This observation might be surprising as it implies that if it is not optimal to maintain permanent social distancing (until the arrival of a vaccine/cure), then it is optimal to delay the period of SD.

The Value of Social Distancing Whether or not the planner wants to engage in SD is an orthogonal question to the optimal timing of SD. To study this question we plot in Figure 2 the death rate withing a year as a function of the number of days of social distancing. As one can see in the figure social distancing can be an effective measure to prevent the death of population. For example, 50 days of optimally timed social distancing (from day 50 to day 100) reduce the death rate by roughly 4%. The figure however shows that without optimal timing SD is much less effective and to achieve an equal reduction in the example one needs more than 300 days of SD.

We asses the value of optimal social distancing in prevented dead where we plot the percentage of dead population prevented per day of SD as a function of the number of days the planner engages in SD. The initial efficacy of social distancing is around 0.2% per day of SD and then decreases to 0.05% around 100

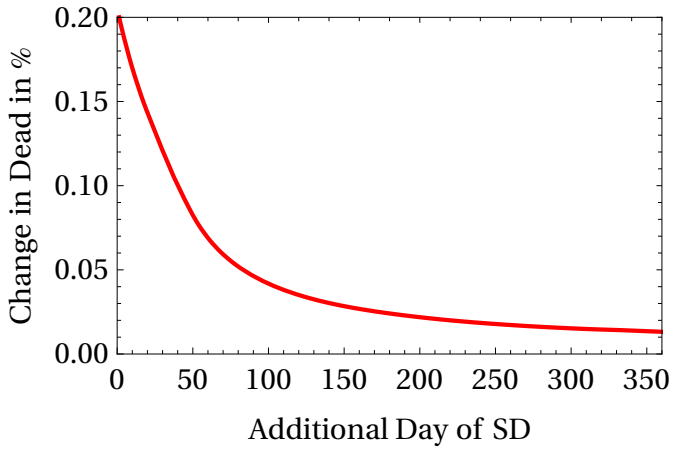

Figure 3: Prevented dead population per day of optimal social distancing.

days. If one assigns a value of around 10 million dollar to a statistical life as it is typically estimated in the literature[7] then one life corresponds to around 148 US per capita GDPs[8], which implies that the planner should be willing to endure a day of SD to save 0.0019% of the population. Thus, for the commonly assumed value of a statistical life the planner should engage in constant SD until a vaccine or cure is found. Figure 3 shows that this conclusion is robust and stays valid even if one assigns just a tenth of the commonly assumed value to a life, i.e. 1 million $.

Social Distancing can Lead to more Dead We next illustrate how suboptimally timed social distancing can actually *increase* the number of fatalities as a consequence of the epidemic. A particular example of this is shown in the right graph of Figure 2 which shows that social distancing from day 50-100 can lead to a substantially *flatter* curve than SD from day 0-100. The reason for this perhaps surprising phenomenon is that by not engaging in SD early, many more people are infected before day 100 which are then later on immune. The "herd-immunity" that is created this way leads to a substantially lower peak in infections and fewer dead (0.7% of the population vs 4.6%).

5 Extensions

5.1 Random Arrival of a Vaccine and Cure

In this section we introduce a variant of the model of Section 2 where the time until a vaccine or cure is available is random. More formally, we let $\tau\colon \Omega \to [0,T]$ be a bounded random variable on a probability space $(\Omega, \mathscr{F}, \mathbb{P})$. We assume that τ has a continuous density function $p\colon [0,T] \to [0,\infty)$ and we denote by

[7] See for example Viscusi and Aldy (2003).

[8] See https://en.wikipedia.org/wiki/List_of_countries_by_GDP_(PPP)_per_capita.

$F(t) = \mathbb{P}[\tau \leq t] = \int_0^t p(s)\,ds$ its distribution function[9]. The expected costs of a strategy $\beta \in \mathscr{B}$ are given by

$$J(\beta) = \mathbb{E}\left[\int_0^{\tau} v(i(t)) + c(\beta(t))\,dt\right]. \tag{5}$$

These can be transformed to

$$J(\beta) = \mathbb{E}\left[\int_0^{T} \mathbf{1}_{[0,\tau)}(s)(v(i(t)) + c(\beta(t)))\,dt\right] = \int_0^T (1-F(t))(v(i(t)) + c(\beta(t)))\,dt. \tag{6}$$

We obtain the following variant of Proposition 1[10].

Proposition 4. *An optimal policy exists. Let $\beta^* \in \mathscr{B}$ be such an optimal strategy and denote by $s^*, i^* \colon [0,T] \to [0,1]$ the associated state processes satisfying (1). Then there exists a function $\eta^* : [0,T] \to \mathbb{R}$ with $\eta^*(T) = 0$ such that for almost all $t \in [0,T]$ it holds*

$$(\eta^*)'(t) = \frac{(1-F(t))(v(i^*(t)) + c(\beta^*(t)) - v'(i^*(t))i^*(t)) - \int_t^T (v(i^*(s)) + c(\beta^*(s)))p(s)\,ds}{i^*(t)} + \eta^*(t)\beta^*(t)i^*(t) \tag{7}$$

and

$$\beta^*(t) \in \operatorname*{arg\,min}_{b \in B} \left[\eta^*(t)i^*(t)s^*(t)b + (1-F(t))c(b)\right]. \tag{8}$$

Moreover, we have $\eta^*(t) > 0$ for all $t \in [0,T)$.

5.2 Vaccine, but no Cure

Throughout, we made the simplifying assumption that a vaccine and a cure arrive simultaneously. While this assumption simplified the analysis, it is easy to extend the model such that there is no cure at the time a vaccine arrives. As after the comprehensive vaccination of the population no new infected would be added, the share of infected would evolve according to $i'(t) = -\gamma i(t)$ after time T and thus be given by $i(t) = i(T)e^{-\gamma(t-T)}$. The share of the population that would die after the arrival of the vaccine in this case would thus be given by

$$\bar{v}(i(T)) = \int_T^{\infty} v\left(i(T)e^{-\gamma(t-T)}\right)dt = \int_0^{i(T)} \frac{v(z)}{z\gamma}dz. \tag{9}$$

For example in the linear case where $v(i) = \alpha i$ this would simplify to $\bar{v}(i) = \frac{\alpha}{\gamma}i$. The objective function would thus be adjusted by a terminal cost

$$J(\beta) = \int_0^T v(i(t)) + c(\beta(t))dt + \bar{v}(i(T)).$$

[9] Without loss of generality we assume that T is the smallest upper bound of τ, i.e., $F(t) < 1$ for all $t < T$.

[10] See (Clarke, 2013, Corollary 22.6) for a statement of the maximum principle for time-dependent payoffs.

For the problem of minimizing J we obtain the following variant of Proposition 1.

Proposition 5. *An optimal policy exists. Let $\beta^* \in \mathscr{B}$ be such an optimal strategy and denote by $s^*, i^* \colon [0,T] \to [0,1]$ the associated state processes satisfying (1). Then there exists a function $\eta^* : [0,T] \to \mathbb{R}$ with $\eta^*(T) = \frac{v(i(T))}{\gamma i(T)}$ such that for almost all $t \in [0,T]$ it holds*

$$\begin{aligned}(\eta^*)'(t) &= \eta^*(t)\beta^*(t)i^*(t) + \frac{v(i^*(t)) + c(\beta^*(t)) - \min_{b\in B}\left(\frac{v(i^*(T))s^*(T)b}{\gamma} + c(b)\right)}{i^*(t)} - v'(i^*(t)),\\ (\beta^*)(t) &\in \arg\min_{b\in B}\left[\eta^*(t)i^*(t)s^*(t)b + c(b)\right].\end{aligned} \tag{10}$$

Moreover, we have $\eta^*(t) > 0$ for all $t \in [0,T)$.

6 Conclusion

We derived the optimal policy for social distancing during an epidemic. Our analysis revealed several features of the optimal policy. For cost linear in the number of infected, the optimal policy consists of two phases, a first phase where the measures taken to decrease the transmission rate are escalated and then a second phase where these measures are reduced. Furthermore, if the cost of reducing the transmission rate is linear, the optimal policy is always extreme. At any point in time either social distancing is carried out to the maximal extend possible or not at all. The intuitive reason for this result is that more extreme measures over a shorter time horizon are more effective than less extreme measures over a longer horizon. We illustrated through an example that the effectiveness of social distancing depends crucially on its optimal timing. Withing the context of this example optimal social distancing is often substantially delayed in order to generate herd immunity. Engaging in more, but too early social distancing can increase the peak number of infected and thereby the fatalities from the disease.

A Appendix

We split the appendix in two sections. Section A.1 proves a sequence of auxiliary results which together imply Proposition 1. Section A.2 provides a proof of Proposition 2 and 3.

A.1 Proof of Proposition 1

Lemma 6. *An optimal policy β^* exists that solves (3). Let $s^*, i^* \colon [0,T] \to [0,1]$ be the state processes associated with an optimal control satisfying (1). Then there exist absolutely continuous functions $\lambda_1^*, \lambda_2^* \colon [0,T] \to$*

$\mathbb{R}$ *which satisfy for almost all* $t \in [0,T]$ *the dynamics*

$$\begin{aligned} (\lambda_1^*)'(t) &= (\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)i^*(t), & \lambda_1^*(T) &= 0, \\ (\lambda_2^*)'(t) &= (\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)s^*(t) + \gamma\lambda_2^*(t) - v'(i^*(t)), & \lambda_2^*(T) &= 0\,, \end{aligned} \tag{11}$$

and the optimality condition

$$\beta^*(t) \in \operatorname*{arg\,min}_{b \in B} \left[(\lambda_2^*(t) - \lambda_1^*(t))i^*(t)s^*(t)b + c(b) \right]. \tag{12}$$

Moreover, there exists a constant $\overline{h} \in \mathbb{R}$ *such that for almost all* $t \in [0,T]$ *we have*

$$[\lambda_2^*(t) - \lambda_1^*(t)]\beta^*(t)i^*(t)s^*(t) - \gamma\lambda_2^*(t)i^*(t) + v(i^*(t)) + c(\beta^*(t)) = \overline{h}\,. \tag{13}$$

Proof of Lemma 6. Suppose an optimal policy β^* exists. The existence of λ_1^*, λ_2^* that satisfy (11), (12) and (13) follows from the Pontryagin principle (see, e.g., Clarke, 2013, Theorem 22.2 and Corollary 22.3). We show the existence of an optimal policy by verifying the conditions of Theorem 23.11 in Clarke (2013).

(a) $g(t,(s,i)) = \begin{pmatrix} -is \\ +is \end{pmatrix}$ which implies that $|g(t,(s,i))| \leq 2|is| < 2$.

(b) $B = [\underline{b}, \overline{b}]$ is closed and convex by definition.

(c) The sets $E = \{(s_0, i_0)\} \times \mathbb{R}_+$ and $Q = [0,T] \times [0,1]^2$ are closed and $\ell = 0$ is lower semicontinuous.

(d) The running cost $\beta \mapsto v(i) + c(\beta)$ is convex as c is convex. Furthermore, $v(i) + c(\beta) \geq 0$.

(e) The projection set is given by $\{(s_0, i_0)\}$ and thus bounded.

(f) As $\beta \in B$ it follows that $|\beta| \leq \overline{b}$. This verifies (f) (ii).

We have hence verified that there exists an optimal policy. □

Throughout we suppose that $\beta^* \in \mathscr{B}$ is an optimal control and denote by $s^*, i^* \colon [0,T] \to [0,1]$ the associated state processes satisfying (1). Moreover, we denote by $\lambda_1^*, \lambda_2^* \colon [0,T] \to \mathbb{R}$ the Lagrange variables from Lemma 6. Note that compactness of B and continuity of c ensure that for all $t \in [0,T]$ the function $b \mapsto [\lambda_2^*(t) - \lambda_1^*(t)]i^*(t)s^*(t)b + c(b)$ attains its minimum on B. By (12) this minimum is attained by $\beta^*(t)$ for almost all $t \in [0,T]$. By potentially changing β^* on a set of measure zero we suppose in the sequel that $\beta^*(t)$ attains the minimum for all $t \in [0,T]$ (i.e., (12) holds for all $t \in [0,T]$). Note that this change does not affect the trajectories of s^*, i^*, λ_1^* and λ_2^*.

We introduce the new Lagrange variable

$$\eta^*(t) = \lambda_2^*(t) - \lambda_1^*(t)\,. \tag{14}$$

The variable $\eta^*(t)$ has a clear interpretation: it measures the marginal change in the cost with respect to infecting susceptible population. Intuitively speaking, $\eta^*(t)$ measures the additional cost if one additional person is infected at time t given the optimal policy is used. Note that by (12) at each time t the optimal

control $\beta^*(t)$ depends on $\lambda_1^*(t)$ and $\lambda_2^*(t)$ only through their difference $\eta^*(t) = \lambda_2^*(t) - \lambda_1^*(t)$.

The next result shows that towards the end of the time horizon it becomes optimal to use the control $\overline{b}$. Moreover, it identifies the constant $\overline{h}$ from Lemma 6.

Lemma 7. *Let $\beta^* \in \mathscr{B}$ be an optimal control and suppose that the optimality condition (12) holds for all $t \in [0,T]$. Suppose that $t_0 \in [0,T]$ satisfies $\eta^*(t_0) \leq 0$. Then it holds that $\lim_{t\to t_0} \beta^*(t) = \beta^*(t_0) = \overline{b}$. In particular, we have that $\lim_{t\to T} \beta^*(t) = \overline{b}$ and that $\overline{h} = v(i^*(T))$ for the constant $\overline{h}$ from Lemma 6.*

Proof of Lemma 7. First note that the assumption $\eta^*(t_0) \leq 0$ ensures that the function $b \mapsto \eta^*(t) i^*(t) s^*(t) b + c(b)$ attains its global minimum on B at $\overline{b}$. Hence (12) implies that $\beta^*(t_0) = \overline{b}$. Next let (t_n) be a sequence such that $t_n \to t_0$ as $n \to \infty$. Suppose by contradiction that there exists a subsequence such that $\lim_{n\to\infty} \beta^*(t_n) =: b_0 < \overline{b}$. Next note that (12) ensures for all $n \in \mathbb{N}$ that (recall that $c(\overline{b}) = 0$)

$$\eta^*(t_n) i^*(t_n) s^*(t_n) \beta^*(t_n) + c(\beta^*(t_n)) \leq \eta^*(t_n) i^*(t_n) s^*(t_n) \overline{b}. \tag{15}$$

This implies that

$$\eta_2^*(t_n) i^*(t_n) s^*(t_n) \geq \frac{c(\beta^*(t_n))}{\overline{b} - \beta^*(t_n)}. \tag{16}$$

Taking the limit $n \to \infty$ yields the contradiction

$$0 = \lim_{n\to\infty} \eta^*(t_n) i^*(t_n) s^*(t_n) \geq \lim_{n\to\infty} \frac{c(\beta^*(t_n))}{\overline{b} - \beta^*(t_n)} = \frac{c(b_0)}{\overline{b} - b_0} > 0. \tag{17}$$

Therefore, we have $\lim_{t\to t_0} \beta^*(t) = \overline{b} = \beta^*(t_0)$. Taking the limit $t \to T$ in (13) and using that $\eta^*(T) = 0$ implies $\overline{h} = v(i^*(T))$. □

The next result shows that the cost of additional infected η^* is characterized by an ordinary differential equation (ODE) that does not depend on λ_1^* and λ_2^*. Moreover, we show that both λ_1^* and λ_2^* can be recovered from η^*.

Lemma 8. *The variable η^* solves*

$$(\eta^*)'(t) = \eta^*(t) \beta^*(t) i^*(t) + \frac{v(i^*(t)) + c(\beta^*(t)) - v(i^*(T))}{i^*(t)} - v'(i^*(t)), \quad \eta^*(T) = 0. \tag{18}$$

Conversely, suppose that $i, s, \beta, \eta \colon [0,T] \to \mathbb{R}$ satisfy

$$\begin{aligned}
s'(t) &= -\beta(t) i(t) s(t), \quad s(0) = s_0, \\
i'(t) &= \beta(t) i(t) s(t) - \gamma i(t), \quad i(0) = i_0, \\
\eta'(t) &= \eta(t) \beta(t) i(t) + \frac{v(i(t)) + c(\beta(t)) - v(i(T))}{i(t)} - v'(i(t)), \quad \eta(T) = 0, \\
\beta(t) &\in \operatorname*{arg\,min}_{b \in B} \left[\eta(t) i(t) s(t) b + c(b) \right],
\end{aligned} \tag{19}$$

then

$$\begin{aligned}\lambda_1(t) &= \frac{1}{\gamma}\left[\eta(t)\beta(t)s(t) + \frac{v(i(t)) + c(\beta(t)) - v(i(T))}{i(t)}\right] - \eta(t),\\ \lambda_2(t) &= \frac{1}{\gamma}\left[\eta(t)\beta(t)s(t) + \frac{v(i(t)) + c(\beta(t)) - v(i(T))}{i(t)}\right]\end{aligned} \tag{20}$$

solves (11).

Proof of Lemma 8. First note that it follows from (13) and Lemma 7 that

$$\eta^*(t)\beta^*(t)i^*(t)s^*(t) - \gamma\lambda_2^*(t)i^*(t) = v(i^*(T)) - v(i^*(t)) - c(\beta^*(t)). \tag{21}$$

Then (11) implies that

$$\begin{aligned}(\eta^*)'(t) &= (\lambda_2^*)'(t) - (\lambda_1^*)'(t) = -\eta^*(t)\beta^*(t)s^*(t) + \gamma\lambda_2^*(t) - v'(i^*(t)) + \eta^*(t)\beta^*(t)i^*(t)\\ &= \eta^*(t)\beta^*(t)i^*(t) - v'(i^*(t)) - \frac{v(i^*(T)) - v(i^*(t)) - c(\beta^*(t))}{i^*(t)}.\end{aligned} \tag{22}$$

Next suppose that s, i and η solve (19) and that λ_1 and λ_2 are given by (20). Observe that it follows similarly as in Lemma 7 that $\beta(T) = \overline{b}$ and consequently that $\lambda_1(T) = \lambda_2(T) = 0$. Next note that the envelope theorem ensures that

$$\frac{\partial}{\partial t}\left[\eta(t)i(t)s(t)\beta(t) + c(\beta(t))\right] = \frac{\partial}{\partial t}\min_{b\in B}\left[\eta(t)i(t)s(t)b + c(b)\right] = \beta(t)\frac{\partial}{\partial t}\left[\eta(t)i(t)s(t)\right]. \tag{23}$$

Then it holds that

$$\begin{aligned}\lambda_2'(t) &= \frac{\partial}{\partial t}\left[\frac{\eta(t)\beta(t)i(t)s(t) + v(i(t)) + c(\beta(t)) - v(i(T))}{\gamma i(t)}\right]\\ &= \frac{\beta(t)\frac{\partial}{\partial t}\left[\eta(t)i(t)s(t)\right] + v'(i(t))i'(t)}{\gamma i(t)}\\ &\quad - \frac{(\eta(t)\beta(t)i(t)s(t) + v(i(t)) + c(\beta(t)) - v(i(T)))i'(t)}{\gamma(i(t))^2}\\ &= \frac{1}{\gamma}(\beta(t)\eta'(t)s(t) + \beta(t)\eta(t)s'(t))\\ &\quad + \frac{i'(t)}{\gamma i(t)}\left(v'(i(t)) - \frac{v(i(t)) + c(\beta(t)) - v(i(T))}{i(t)}\right)\\ &= \frac{1}{\gamma}(\beta(t)\eta'(t)s(t) + \beta(t)\eta(t)s'(t)) + \frac{i'(t)}{\gamma i(t)}\left(\eta(t)\beta(t)i(t) - \eta'(t)\right)\\ &= \frac{\beta(t)\eta(t)}{\gamma}(s'(t) + i'(t)) + \frac{\eta'(t)}{\gamma}\left(\beta(t)s(t) - \frac{i'(t)}{i(t)}\right)\\ &= -\beta(t)\eta(t)i(t) + \eta'(t).\end{aligned}$$

Therefore we obtain that

$$\begin{aligned}\lambda_2'(t) &= \frac{v(i(t))+c(\beta(t))}{i(t)} - v'(i(t)) = \gamma\lambda_2(t) - \eta(t)\beta(t)s(t) - v'(i(t)) \\ &= (\lambda_1(t)-\lambda_2(t))\,\beta(t)s(t) + \gamma\lambda_2(t) - v'(i(t))\,.\end{aligned}$$

Similarly, λ_1 satisfies

$$\lambda_1'(t) = \lambda_2'(t) - \eta'(t) = [\eta'(t) - \eta(t)\beta(t)i(t)] - \eta'(t) = (\lambda_1(t)-\lambda_2(t))\beta(t)i(t)\,. \qquad \square$$

Lemma 9 (More Infected are Costly)**.** *The function η^* satisfies $\eta^*(t) > 0$ for all $t \in [0,T)$.*

Proof of Lemma 9. Suppose that there exists $t \in [0,T)$ such that $\eta^*(t) \le 0$. Then Lemma 7 shows that $\beta^*(t) = \bar{b}$. Moreover, Lemma 7 ensures that β^* is continuous at t and hence η^* is differentiable at t. Then (18) shows (recall that $c(\bar{b}) = 0$)

$$(\eta^*)'(t) = \eta^*(t)\beta^*(t)i^*(t) + \frac{v(i^*(t)) - v(i^*(T)) - v'(i^*(t))i^*(t)}{i^*(t)}. \tag{24}$$

Since $v' \ge 0$ we thus obtain that

$$(\eta^*)'(t) \le \eta^*(t)\beta^*(t)i^*(t) + \frac{v(i^*(t)) - v(i^*(T)) - v'(i^*(t))(i^*(t) - i^*(T))}{i^*(t)}. \tag{25}$$

Convexity of v thus implies that

$$(\eta^*)'(t) \le \eta^*(t)\beta^*(t)i^*(t) \le 0. \tag{26}$$

We conclude from the terminal condition $\eta^*(T) = 0$ that $\eta^*(t) \ge 0$ for all $t \in [0,T]$. If $v' > 0$ then (25) is a strict inequality and hence we have $(\eta^*)'(t) < 0$ for all $t \in [0,T)$ with $\eta^*(t) \le 0$. Again we conclude from the terminal condition $\eta^*(T) = 0$ that η^* is strictly positive on $[0,T)$. $\square$

Since $(\lambda_1^*)'(t) = -\eta^*(t)\beta^*(t)i^*(t)$ we obtain from Lemma 9 that λ_1^* is decreasing in time and, in particular, that λ_1^* is non-negative. This means that a marginal increase of the susceptible population (while keeping the infected population constant marginally increases the costs. This marginal effect decreases over time and vanishes at time T.

A.2 Proof of Proposition 2 and 3

Proof of Proposition 2. We introduce the function $g(t) = \eta^*(t)i^*(t)s^*(t)$, $t \in [0,T]$ and compute $g'(t)$

$$\begin{aligned}
g'(t) &= (\eta^*)'(t)i^*(t)s^*(t) + (\eta^*(t))(i^*)'(t)s^*(t) + \eta^*(t)i^*(t)(s^*)'(t) \\
&= [(\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)s^*(t) + \lambda_2^*(t)\gamma - v'(i^*(t)) - (\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)i^*(t)]i^*(t)s^*(t) \\
&\quad + (\lambda_2^*(t) - \lambda_1^*(t))s^*(t)[\beta^*(t)i^*(t)s^*(t) - \gamma i^*(t)] \\
&\quad - (\lambda_2^*(t) - \lambda_1^*(t))i^*(t)\beta^*(t)i^*(t)s^*(t) \qquad (27) \\
&= (\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)[-(s^*(t))^2 i^*(t) + (i^*(t))^2 s^*(t) + (s^*(t))^2 i^*(t) - (i^*(t))^2 s^*(t)] \\
&\quad + [\gamma\lambda_2^*(t) - v'(i^*(t))]i^*(t)s^*(t) - \gamma(\lambda_2^*(t) - \lambda_1^*(t))s^*(t)i^*(t) \\
&= [\gamma\lambda_1^*(t) - v'(i^*(t))]i^*(t)s^*(t).
\end{aligned}$$

Since $v'(i) = \alpha$ we have by Lemma 9 that

$$\frac{\partial}{\partial t}[\gamma\lambda_1^*(t) - v'(i^*(t))] = -\gamma\eta^*(t)\beta^*(t)i^*(t) < 0\,. \qquad (28)$$

This together with (27) shows that g' changes its sign at most once and that this change (if existent) is from positive to negative. It follows from (28) that g' can not be equal to zero on any interval. This implies that g is first strictly increasing and then strictly decreasing, i.e., strictly quasi-concave. Next take two points in time $t, t' \in [0,T]$. Then (12) shows that

$$g(t)\beta^*(t) - c(\beta^*(t)) \le g(t)\beta^*(t') - c(\beta^*(t')) \text{ and } g(t')\beta^*(t') - c(\beta^*(t')) \le g(t')\beta^*(t) - c(\beta^*(t)). \qquad (29)$$

Adding these two inequalities yields that

$$(g(t) - g(t'))(\beta^*(t) - \beta^*(t')) \le 0 \qquad (30)$$

Since g is strictly quasi-concave we obtain that β^* is quasi-convex. □

Proof of Proposition 3. Let $v(i) = \alpha i$ and $c(\beta) = \delta(\overline{b} - \beta)$. It follows from (12) that any optimal control satisfies

$$\beta^*(t) = \begin{cases} \overline{b} & \text{if } g(t) < \delta \\ \underline{b} & \text{if } g(t) > \delta\,. \end{cases}$$

As argued in the proof of Proposition 2, g' is strictly quasi-convex which implies the result with $t_1^* = \inf\{t \ge 0\colon g(t) \ge \delta\}$ and $t_2^* = \sup\{t \le T\colon g(t) \ge \delta\}$. □

References

A. Abakuks. *Some optimal isolation and immunisation policies for epidemics.* PhD thesis, University of Sussex, 1972.

A. Abakuks. An optimal isolation policy for an epidemic. *Journal of Applied Probability*, 10(2):247–262, 1973.

A. Abakuks. Optimal immunisation policies for epidemics. *Advances in Applied Probability*, 6(3):494–511, 1974.

F. E. Alvarez, D. Argente, and F. Lippi. A simple planning problem for covid-19 lockdown. Technical report, National Bureau of Economic Research, 2020.

A. Atkeson. What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios. Technical report, National Bureau of Economic Research, 2020.

R. J. Barro, J. F. Ursua, and J. Weng. The coronavirus and the great influenza epidemic-lessons from the. 2020.

H. Behncke. Optimal control of deterministic epidemics. *Optimal control applications and methods*, 21(6): 269–285, 2000.

F. Clarke. *Functional analysis, calculus of variations and optimal control*, volume 264. Springer Science & Business Media, 2013.

D. de Walque, J. Friedman, R. Gatti, and A. Mattoo. How two tests can help contain covid-19 and revive the economy, 2020.

M. Dewatripont, M. Goldman, E. Muraille, and J.-P. Platteau. Rapidly identifying workers who are immune to covid-19 and virus-free is a priority for restarting the economy. *VOX CEPR Policy Portal*, 23, 2020.

E. Hansen and T. Day. Optimal control of epidemics with limited resources. *Journal of mathematical biology*, 62(3):423–451, 2011.

M. Kantner. Beyond just "flattening the curve": Optimal control of epidemics with purely non-pharmaceutical interventions. *arXiv Preprint 2004.09471*, 2020.

W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.

S. M. Kissler, C. Tedijanto, M. Lipsitch, and Y. Grad. Social distancing strategies for curbing the covid-19 epidemic. *medRxiv*, 2020.

R. Morton and K. H. Wickwire. On the optimal control of a deterministic epidemic. *Advances in Applied Probability*, 6(4):622–635, 1974.

F. Piguillem, L. Shi, et al. The optimal covid-19 quarantine and testing policies. Technical report, Einaudi Institute for Economics and Finance (EIEF), 2020.

J. H. Stock. Data gaps and the policy response to the novel coronavirus. Technical report, National Bureau of Economic Research, 2020.

A. A. Toda. Susceptible-infected-recovered (sir) dynamics of covid-19 and economic impact. *arXiv preprint arXiv:2003.11221*, 2020.

W. K. Viscusi and J. E. Aldy. The value of a statistical life: a critical review of market estimates throughout the world. *Journal of risk and uncertainty*, 27(1):5–76, 2003.

K. Wickwire. A note on the optimal control of carrier-borne epidemics. *Journal of Applied probability*, 12 (3):565–568, 1975.

K. Wickwire. Mathematical models for the control of pests and infectious diseases: a survey. *Theoretical population biology*, 11(2):182–238, 1977.