Cost/benefit analysis of Covid-19 pandemic suppression using an SEIR model

Jan Maciejowski, Robert Rowthorn, Scott Sheffield, David Vines and Annie Williamson

6th April 2021

Cost/benefit analysis of Covid-19 pandemic suppression using an SEIR model

Jan Maciejowski[1], Robert Rowthorn[2], Scott Sheffield[3], David Vines[4], and Annie Williamson[5]

[1]Dept of Engineering and Pembroke College, University of Cambridge, `jmm1@cam.ac.uk`
[2]Faculty of Economics and King's College, University of Cambridge, `rer3@cam.ac.uk`
[3]Dept of Mathematics, MIT, `sheffield@math.mit.edu`
[4]Dept of Economics, Balliol College, and INET, University of Oxford, `david.vines@economics.ox.ac.uk`
[5]Barts and the London School of Medicine and Dentistry, Queen Mary University of London, `anniewilliamson6@gmail.com`

April 6, 2021

Abstract

In the United Kingdom, as elsewhere, there has been a debate on how to escape from lockdown without provoking a resurgence of the Covid-19 disease. This paper presents a simple cost-benefit analysis inspired by optimal control theory, incorporating an SEIR model of disease propagation. Our calibration accords with UK experience and simulations start from the beginning of January 2021. The optimal path for government intervention is computed under different assumptions about the value of life. We examine how the test & trace system and vaccination affect this optimal path and show that these policies are complements. Test & trace has most effect at the beginning when prevalence is high, whereas vaccination only affects infections gradually. Comparing optimal paths, we show that economic cost is much lower under vaccination alone than under test & trace alone. Deaths, in contrast, are somewhat higher under vaccination. In general, the greater is the value of life, the longer is the optimal lockdown. Under certain conditions, this relationship is discontinuous: a small increase in the value of life leads to a jump in the optimal length of lockdown.

Keywords: Covid-19, cost-benefit, lockdown, SEIR, pandemic.

1   Introduction

How quickly should the Covid lockdown be ended in the United Kingdom? How many people will die as a result? What does the trade-off between deaths and the costs of lockdown look like? Is there a stark choice between a policy of living with Covid and one of eliminating Covid? Many people are asking these questions.

There is a need for an effective synthesis of epidemiological and economic perspectives to answer these questions. A report by the Institute for Government (IoG) "Science advice in a crisis" has called for such work. The authors claim [21, page 21] that "integration [between Treasury ministers and officials and the Scientific Advisory Group for Emergencies (SAGE)] was not helped by the fact that economic and epidemiological modelling is largely done separately.... [S]everal academics have argued that given time, modelling could be brought together if the government funded economists and epidemiologists were to collaborate".

We were amongst the first to answer the above questions by bringing together an SIR epidemiological model and the kind of cost benefit analysis normally carried out by economists [19, 20, 23, 24]. There are now many other papers doing this, see for example [3, 9]. This paper is a further step towards breaking down the silos which the IoG report identified.

Since we wrote our earlier papers, the Covid situation in the UK has changed in four important ways. First, new variants of Covid have arrived which spread more rapidly. Second, it has become apparent that those who are infected take some time to become infectious. Third, the UK's test-and-trace system seems to be working better, and numbers are available about the extent of this improvement. Finally, vaccination has arrived and is being rolled out quickly. We show how a test-and-trace system, and a vaccination process, can radically change the answers to the above questions.

Some of the papers written by economists about policy responses to Covid contain sophisticated "micro-foundations" and justify government intervention on the grounds that utility maximising individuals may under-protect themselves against the risk of infection: they may ignore the fact that by getting infected themselves they increase the risk that others will become infected. Individuals may also reduce their expenditure during the pandemic, thereby damaging the economy. A good example of such a micro-founded approach is [2]. By contrast most authors, whilst recognising the importance of micro-foundations, do not seek to derive individual behaviour from utility maximisation. Instead, they build and calibrate simple models of disease dynamics and then suppose that government interventions will influence these dynamics in a specified fashion. This is the approach adopted in [20] and which we adopt here.

Our first step in [20] was to explore the trade-off between deaths and economic costs, using an SIR model, where economic costs include not only the medical costs and output loss due to illness but also the costs of the restrictive policies adopted by government to combat the disease . Aspects of this trade-off had already been identified by the Imperial College group in their now-famous paper [8]. This trade-off has also been explored in several other papers [2, 6, 7].

The central feature of [20] was to make use of this trade-off in a dynamic cost-benefit framework that enabled us to determine the optimal length of the initial lockdown and the optimal severity of subsequent policies. We showed that what is optimal

depends crucially on the value attached to human life. Our results, produced early in the epidemic, provided striking evidence that the UK seemed to be attaching a surprisingly high value to life. We also quantified the benefits that would be gained by imposing lockdown earlier rather than later - even by just one week - and we explored the benefits of a test-and-trace system.

In this paper we attempt to quantify how, in the new circumstances described above, the optimal length of initial lockdown, and the optimal severity of subsequent policies, depend on the value attached to human life. This is a complex piece of analysis, for two reasons. First, the epidemiological model which we use – an SEIR model - is inter-temporal, in a highly non-linear way. Second, finding the best point on a trade-off between economic costs and deaths requires a choice of policy intervention across many different time periods. The combination of non-linearity and a wide spectrum of policy choices requires optimal control techniques if one is to find an optimal outcome.

Perhaps our most surprising result — albeit confirming what we already found in [20] — is that a small change in the value of life can lead to a radically different optimal length of lockdown, and optimal severity of subsequent policy. These also depend on the efficiency of the test-and-trace programme and on the pace of vaccination. Clearly both matter. But, as we show, test-and-trace is more effective at reducing deaths, whereas vaccination is more effective at reducing economic costs; we explain why this is the case. Our results therefore show that test-and-trace and vaccination are complementary policies rather than substitutes. When beginning with a high number of infections, it is necessary to use both policies, as is now happening in the UK. We deliberately carry out our analysis looking forward from the present, rather than looking backwards to what was happening a year ago. To be precise, we start our simulations from the beginning of January 2021 when the post-Christmas spike in infections was already in place, and when vaccination had just begun.

Some authors assume that only two options are available: lockdown or mitigation, where mitigation involves involves shielding the elderly and vulnerable, leaving most of the rest of the population to do as they please. An example is the paper by Scherbina [22], which assumes that government policy is a period of lockdown followed by mitigation thereafter. The objective thus becomes a relatively simple one: to find the optimal length of the lockdown.

In this paper, we assume that the government has a continuum of options available ranging from lockdown to no intervention at all. It can vary the time profile of its interventions in any way it chooses. Optimisation is over a two-year period without discounting.

As already noted, a thorny issue is the value of life for use in the cost-benefit calculations. In particular, what weight should be attached to the fact that most Covid-19 fatalities are old or in poor health? The conventional VSL (value of social life) approach ignores age [16] and studies that use this approach may assign a value of £2 million or more to each Covid fatality. At the other end of the spectrum are studies which rely on criteria such as quality adjusted years of life (QALY) or life-time production. The value of life is estimated for each age group and a weighted average is then computed using as weights the shares of these age-groups in Covid fatalities. This typically results in a value of life in the range £200,000 to £300,000. We use the UK Treasury figure of £2 million for most of our examples. Because the valuation of life is

integral to cost-benefit analyses in a pandemic, we refrain from treating any one value as definitive. Rather, we vary across a spectrum of life values, finding that optimal policy is sensitive to the number chosen. And, of course, for any number chosen, the optimal policy will depend on the presence of vaccination and the operation of any test-and-trace regime.

The paper is laid out as follows. In Section 2 we set out our model, explaining why we use an SEIR model rather than an SIR model. Sections 3 and 4 define the "efficiency frontier", which summarises the inescapable trade-offs between deaths and economic costs, and describe how we identify the optimal path that determines where on this frontier society will lie. Sections 5 and 6 describe how we analyse test-and-trace and vaccination. In Section 7 we analyse our findings in detail. Section 8 concludes the paper. Appendix A describes in detail how we model test-and-trace. Appendix B briefly explains how we solve the optimal control problem that we formulate. Appendix C sets out the parameter values that we use and notes the relevant sources.

2 The Model

The analysis in this paper uses a version of the standard SEIR model of disease propagation. Ignoring births and non-Covid deaths, the initial population at time $t = 0$ will divide in the future into four groups of people: Susceptible, Exposed, Infected and Removed, denoted by $S(t), E(t), I(t)$ and $R(t)$. The removed category includes individuals who have been infected and recovered, together with those who have died. The latter are denoted by $D(t)$. We use this setup to capture how the epidemic causes people to progress from susceptible to exposed to infected to removed. The initial population is normalised to 1, so these various quantities can be interpreted as shares. The members of $E(t)$ are infected but not infectious, members of $I(t)$ are infectious and remain infectious until they recover or die. Infected individuals who recover acquire complete and permanent immunity, so the journey from $S(t)$ via $E(t)$ and $I(t)$ to $R(t)$ is in one direction only.

The dynamics of the disease are determined by the following equations:

$$\frac{dS(t)}{dt} = -\beta(t)S(t)I(t) \tag{1}$$

$$\frac{dE(t)}{dt} = \beta(t)S(t)I(t) - \sigma E(t) \tag{2}$$

$$\frac{dI(t)}{dt} = \sigma E(t) - \gamma I(t) \tag{3}$$

$$\frac{dR(t)}{dt} = \gamma I(t) \tag{4}$$

$$\frac{dD(t)}{dt} = \delta(t)\gamma I(t) \tag{5}$$

$$S(0) = S_0 \geq 0 \tag{6}$$

$$E(0) = E_0 \geq 0 \tag{7}$$

$$I(0) = I_0 \geq 0 \tag{8}$$

$$R(0) = R_0 \geq 0 \tag{9}$$

$$D(0) = D_0 \geq 0 \tag{10}$$

$$S(t) + E(t) + I(t) + R(t) = 1 \tag{11}$$

where $\sigma, \gamma > 0$ are constant. These constants indicate, respectively, the rate at which exposed individuals become infectious and the rate at which infectious individuals cease to be infectious. The proportion of Removed individuals that die is

$$\delta(t) = \delta_0 e^{-\eta_1 t} \qquad (\eta_1 \geq 0) \tag{12}$$

This formulation allows for medical improvements in the course of time, and will allow for preferential access to vaccination for vulnerable individuals. Note that there are only three genuinely independent state variables in this model. For example, if the trajectories of $S(t), E(t)$ and $I(t)$ are known, the trajectories of $R(t)$ and $D(t)$ are uniquely determined by equations (4) and (5). We adopt an SEIR model rather than an SIR model to allow for the fact that those who are infected take some time to become infectious.

Equation (1) indicates how the pool of susceptibles is depleted by the outflow of newly infected individuals. Assuming that social encounters are random, the probability that a susceptible individual will be infected in a given unit of time is proportional to the prevalence of infection in the population. Equations (2) and (3) indicate how the pools of exposed and infectious individuals vary in response to inflows and outflows.

The coefficient $\beta(t)$ in equations (1) and (2) is a variable which depends on the current intensity of social interaction. The latter depends, in turn, on the measures that the government puts in place to inhibit the spread of the disease. Specifically, we assume that

$$\beta(t) = [1 - q(t)]\beta_0 \tag{13}$$

where $q(t) \in [0, q_{max}]$ is an index of policy severity and $q_{max} < 1$ is an upper limit beyond which it is not feasible to increase $q(t)$.

There are many different economic costs associated with the Covid disease. Two that play a central role in our analysis are as follows. (1) Costs which depend only on the current prevalence of the disease; they include the cost of hospital treatment and output lost from illness. Such costs are equal to $\pi_I I(t)$ where $\pi_I > 0$ is constant. (2) Costs which depend only on the current scale of government intervention in the form of travel restrictions, business closures and the like. These costs are specified by a function $C(q(t))$ which is increasing in $q(t)$ and is zero for $q(t) = 0$.

Throughout this paper we shall assume that the cost-of-intervention function is of the form

$$C(q) = C_{max} \left(\frac{q}{q_{max}} \right)^{1+\phi}$$

where C_{max} is the cost when there is "full" lockdown, q_{max} is the value of q when there is full lockdown, and ϕ is an index of convexity. Figure 1 plots the cost function when $q_{max} = 0.8$, $C_{max} = 0.01$ and $\phi = 2$. These are the parameter values that we use in the simulations reported below. This function has the property that the marginal cost of government intervention is initially very low, but rises steeply as full lockdown is approached.
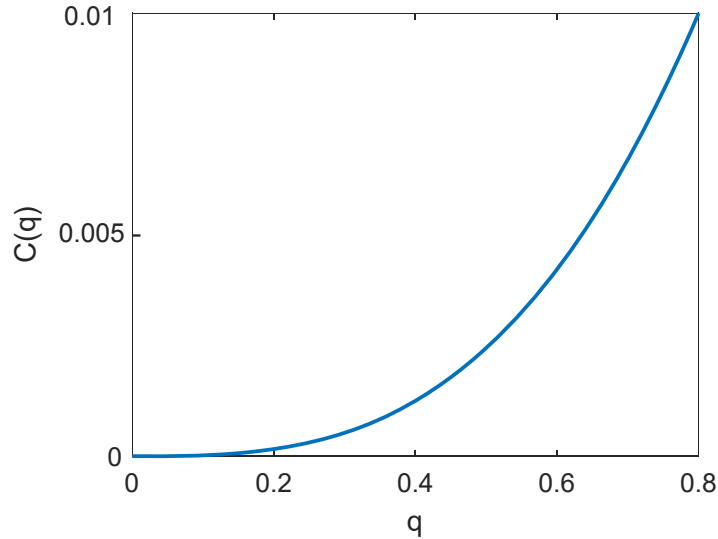

Figure 1: Graph of the cost function $C(q)$ with $q_{max} = 0.8$, $C_{max} = 0.01$ and $\phi = 2$.

Government policy is specified by a complete trajectory of $q(t)$ over the range $[0, T]$ where T is the fixed planning horizon. It is assumed that the government has perfect foresight. Given such a trajectory, the evolution of the system is determined by equations (1) to (5) and the initial conditions (6) to (10). The economic cost incurred

over the course of this trajectory is:

$$\mathscr{E} = \int_0^T [C(q(t)) + \pi_I I(t)] \, dt \tag{14}$$

Deaths arising from the Covid infection are included separately under the heading "value of life". We ignore collateral damage such as deaths caused by the diversion of medical resources to Covid patients. The total number of deaths is

$$D_T \triangleq D(T) = \int_0^T \delta(t) \gamma I(t) \, dt$$

Our study assumes a two-year scenario ($T = 730$ days). We have chosen parameter values to try to represent the UK at the start of 2021. The specific values that we have assumed are listed in Appendix C.

3   The efficiency frontier

A combination $(\mathscr{E}, D_T)$ is *efficient* if there is no achievable combination $(\hat{\mathscr{E}}, \hat{D})$ such that $\hat{\mathscr{E}} < \mathscr{E}, \hat{D} \leq D_T$ [1]. The trajectory of $q(t)$ that gives rise to an efficient combination is a thus a least cost trajectory. There is no allowable trajectory that provides a cheaper way of preventing total deaths from exceeding D_T.

In $(\mathscr{E}, D_T)$ space, the efficient combinations generate a downward sloping curve with functional form $\mathscr{E} = f(D_T)$. This is the *efficiency frontier*. Figure 2 provides an illustration based mostly on our model, with $\phi = 2$. The efficiency frontier is indicated by the solid line in this diagram.

Several points should be noted. The curve is truncated in the top-left because there is a maximum possible number of deaths achieved when $q(t) = 0$ throughout. The curve is truncated at the opposite end because there is a minimum number of deaths achieved when $q(t) = q_{max}$ throughout. A third, less obvious feature is a dent in the mid-range. Despite the convexity of the integrand $C(q(t)) + \pi_I I(t)$, the efficiency frontier is not universally convex. The significance of the dent is discussed below.

4   Policy Choice

The efficiency frontier spells out the inescapable trade-offs between economic costs and deaths. The government can, in principle, choose any point on this frontier. The point that it actually chooses depends on its preferences. Suppose the government seeks to minimise the following total cost function

$$J(\mathscr{E}, D_T) = \mathscr{E} + \pi_D D_T \tag{15}$$

For a given π_D the above equation generates a family of downward sloping indifference curves:

$$D_T = -\frac{1}{\pi_D} \mathscr{E} + \text{constant}$$

---

[1] This condition is equivalent in the present context to $\hat{\mathscr{E}} \leq \mathscr{E}, \hat{D} < D_T$.
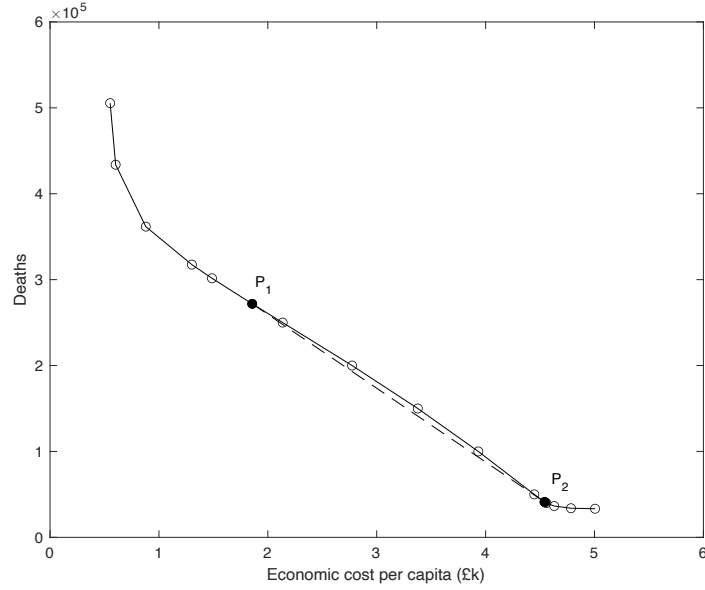
Figure 2: The efficiency frontier. The circles indicate computed points. The broken line spans the 'dent' discussed in the text.

At all points on a given indifference curve the total cost is the same. The optimum is found by locating the lowest indifference curve that intersects the efficiency frontier. Any point at which this intersection occurs is an optimum. For a given π_D, the intersection point will be unique if the efficiency frontier $\mathscr{E} = f(D_T)$ is strictly convex. In this case, by varying π_D we obtain a sequence of intersection points that trace out the whole of the efficiency frontier. At each point the relevant indifference curve (with slope $-1/\pi_D$) is tangent to the frontier.

Suppose the efficiency frontier is strictly convex apart from a single dent as in Figure 2. This case is illustrated more schematically in Figure 3. Starting in the top-left corner above the dent, as we increase π_D, the tangent rotates in an anti-clockwise direction. At first the tangent intersects the efficiency frontier at a single point such as Q_1. This point is optimal given π_D. All other points on the efficiency frontier lie above the tangent (indifference curve) at Q_1. However, as rotation continues, the point will come where the tangent bridges the gap left by the dent and intersects the efficiency frontier twice at P_1 and P_2. These points lie on the same tangent and therefore have the same cost. Both of them are optimal. As π_D increases still further, each tangent once again intersects the efficiency frontier in a single point uch as Q_2. This point is optimal for the relevant π_D.

The significance of the above discussion is as follows. With a linear cost function it is never optimal to choose a point within the dent. In the vicinity of the dent a small change in the value of life π_D may lead to a huge change in the optimal path and hence
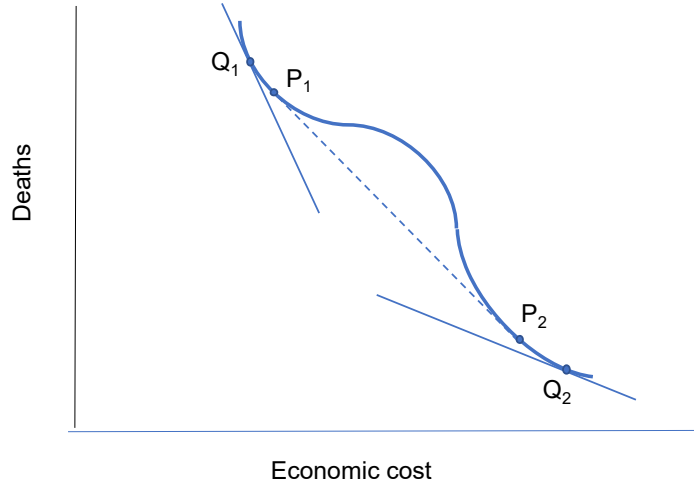
8

Figure 3: An exaggerated version of the efficiency frontier.

in the optimal combination of $\mathscr{E}$ and D_T. This is what Rowthorn and Maciejowski [20] found using an SIR model. The same is true here using an SEIR model. In Figure 2 the points P_1 and P_2 correspond to $\pi_D \approx 860$. Figures 4 and 5 show the radically different optimal paths when $\pi_D = 850$ and $\pi_D = 870$. The number of deaths is 272 thousand and 41 thousand, respectively. The corresponding economic costs are £1857 and £4541 (*per capita*). The difference is explained by the existence of a prolonged initial lockdown (approaching 100 days) at maximum intensity when $\pi_D = 870$ and no full intensity lockdown at all when $\pi_D = 850$.

Our optimisation approach assumes a given value of life π_D. When there is a dent in the efficiency frontier, the resulting optimum path may be highly sensitive to the choice of π_D. In the above example, a small increase in this parameter from 850 to 870 leads to a huge reduction in "optimal" deaths, albeit with a large increase in economic costs. This suggests a lack of robustness in the method. The "value of life" is a hazy concept at the best of times and it is desirable to have a choice criterion that does not rely on this parameter. The following is one possibility. Suppose the government decides on a socially acceptable number of deaths D^* and then chooses the trajectory of $q(t)$ that minimises the economic cost of achieving this outcome. This will take the system to a point $(\mathscr{E}, D_T) = (f(D^*), D^*)$ on the efficiency frontier. Depending on the size of D^* this point may lie within the dent. There is no discontinuity in the optimal solution, and in this sense the new criterion is more robust than the linear criterion. Note that for sufficiently small or large values of D^* the two criteria are equivalent, since both yield a solution that lies on the convex part of the efficiency frontier, away from the dent.
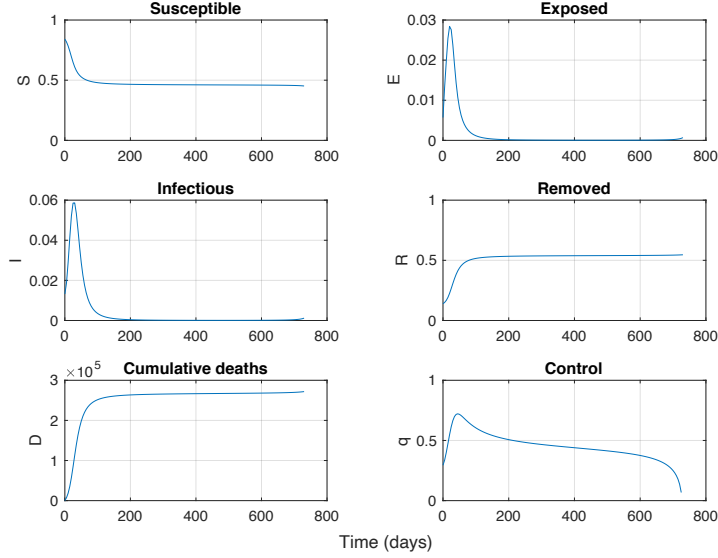
Figure 4: Trajectories of the optimal solution with $\pi_D = 850$.

5 Test and Trace

Although slow off the mark, the UK now has an extensive test-and-trace system. This system is expensive, costing £37 billion (£550 *per capita*) over two years, and its effectiveness is uncertain [5, page 3]. To model the operation of this system, we modify the dynamic equations (2), (3) and (4) as follows:

$$\frac{dE(t)}{dt} = \beta(t)S(t)I(t) - \sigma E(t) - r_E E(t) \qquad (16)$$

$$\frac{dI(t)}{dt} = \sigma E(t) - \gamma I(t) - r_I I(t) \qquad (17)$$

$$\frac{dR(t)}{dt} = \gamma I(t) + r_I I(t) + r_E E(t) \qquad (18)$$

Thus, constant fractions of $E(t)$ and $I(t)$ are isolated each day and are no longer able to infect the rest of the population. Using the limited information that is available, we assume that $r_E = 0.01$, $r_I = 0.025$. These fractions are weighted averages that reflect the degree to which infected individuals self-isolate. They assume that the rate of testing remains constant throughout the entire trajectory. Details are given in Appendix A.

Figure 6 is derived from Figure 2. It illustrates, amongst other things, how test-and-trace improves the choices facing the government. The proportion of the infected population that is isolated each day may appear small, but as can be seen from Figure 6, the cumulative effect of these removals is substantial. Despite its high cost and its
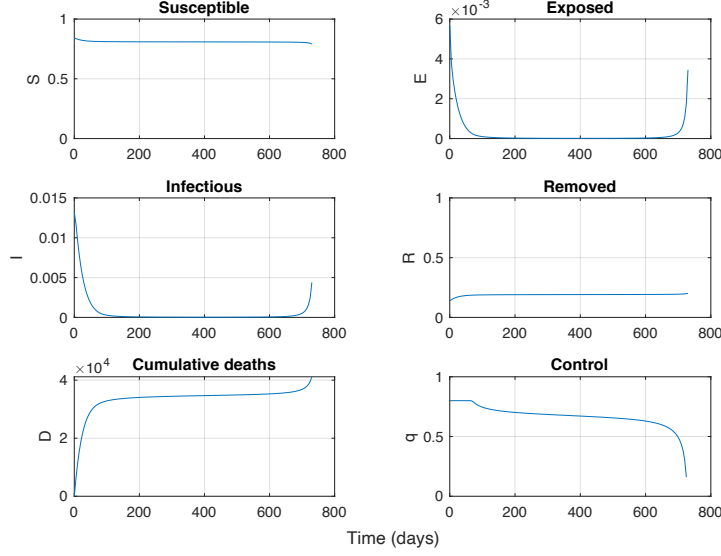
Figure 5: Trajectories of the optimal solution with $\pi_D = 870$.

imperfect implementation (see Appendix A), Test and Trace repays the expenditure on it several times over.

6 Vaccination

In this section we consider how vaccination affects the results. To model vaccination, we assume that a constant number F_0 of susceptible individuals are vaccinated each day until $S(t)$ falls to a given value $\bar{S}$. At this point vaccination stops because there is no-one left who is willing to be vaccinated. This requires modifying equations (2) and (3) as follows:

$$\frac{dS(t)}{dt} = -\beta(t)S(t)I(t) - F(t)$$

$$\frac{dR(t)}{dt} = \gamma I(t) + F(t)$$

$$F(t) = \begin{cases} F_0 & \text{for } S(t) > \bar{S} \\ 0 & \text{otherwise} \end{cases}$$

We investigate the cases $\bar{S} = 0.5$ and $\bar{S} = 0.2$.

We also assume that vaccination reduces the average death rate amongst Covid victims. This is because vulnerable individuals at high risk of death are vaccinated first, thereby altering the risk profile of the remaining susceptible population. This
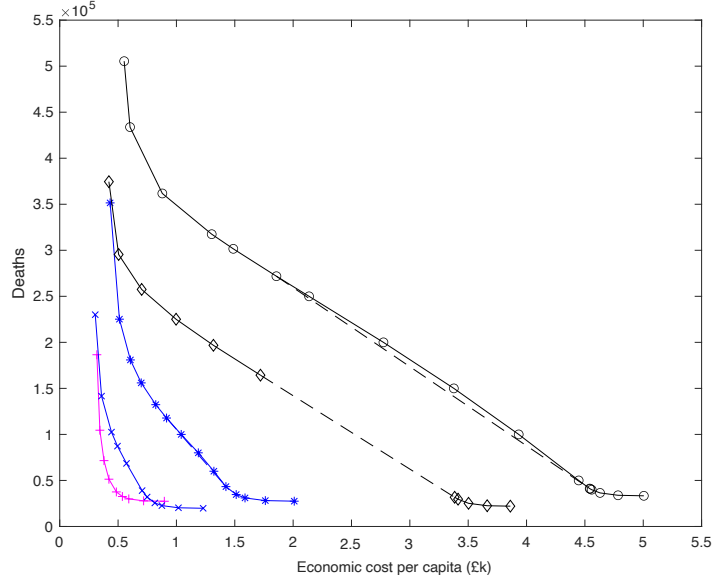
11

Figure 6: Frontier curves for various cases. Black curves relate to absence of vaccination. Blue curves relate to 50% vaccination (see section 6). The magenta curve relates to 80% vaccination. Diamonds ⋄ and crosses × mark the curves when Test and Trace is additionally present (see section 5). All the markers indicate solutions obtained for the optimisation problem. Broken lines span the discontinuity 'dents'.

effect is in addition to the assumed decline in the death rate due to medical advances (see (12)). Specifically, we assume that

$$\delta(t) = \begin{cases} \delta_0 e^{-\eta_2 t} & \text{for } t \leq t^* \\ \delta(t^*)e^{-\eta_1(t-t^*)} & \text{for } t > t^* \end{cases} \qquad (19)$$

with $\eta_2 > \eta_1$.

Figure 6 illustrates how vaccination affects the choices available to the government. It is derived from Figure 2 by inserting the equivalent curves when 50% and 80% of the population are vaccinated ($\bar{S} = 0.5$ and $\bar{S} = 0.2$, respectively). The main effect of vaccination is to shift the curves inwards. Not surprisingly vaccination greatly improves the trade-off between economic cost and deaths. It also reduces the size of the dent in the efficiency frontier or eliminates it altogether.

Costs of vaccination programmes have been estimated from the information available in [10].

7 Discussion of Results

Figures 7–11 show the optimal trajectories obtained with $\pi_D = 2000$ for the various cases shown in Figure 6. That is, each figure shows the trajectories corresponding to a single point on one of the curves in Figure 6. Note that $\pi_D = 2000$ is equivalent to £2 million for the value of life (as used by the UK Treasury for project evaluation). Table 1 shows, for each of these trajectories, the length of the initial lockdown and the 'total' economic cost, by which we mean the economic cost $\mathscr{E}$ as defined in (14), together with the cost of vaccination and of test-and-trace where appropriate. All the costs which appear in the table are *per capita*.

Comparing the optimal paths in Figures 7 and 8, we find that test-and-trace on its own, without vaccination, reduces the time spent in full-lockdown from 105 days to 55 days and the number of deaths from 36,000 to 25,000. Comparing Figures 9 and 10, we find that test-and-trace has a similar add-on effect if combined with vaccination.

Vaccination on its own, without test-and-trace, has a considerable effect on economic cost. This is partly because it shortens the length of full lockdown. It is also because it allows a rapid return to virtual normality once the lockdown is over. Under the other scenarios, without vaccination, costly government intervention continues long after full lockdown ends.

A striking feature of Table 1 is the fact that all of the optimal paths involve an initial full lockdown of more than 4 weeks, sometimes considerably more. This is for two reasons. The first reason is the high assumed value of life (£2 million). With a much lower value of life, it would be optimal to avoid a lengthy lockdown and accept a much higher death toll in return for less economic damage. A second reason is the large number of infectious people in circulation on day zero (1.9% of the population). Given a high assumed value of life, the priority is to prevent these people from infecting large numbers of others. At this stage, economic considerations are of secondary importance. Hence the initial lockdown.

It is worth explaining one feature of the results depicted in Figures 7– 9, something which was also evident in Figure 5. In most of the simulations shown, control is relaxed towards the very end of the time period, and the numbers of exposed and infectious are allowed to rise. This is simply because the optimisation exercise is not concerned with anything that happens after 730 days. As a result, it is desirable to relax control towards the end of the time period, thereby reducing the economic cost of the policy, because the full cost of doing this will only be evident after the time period under consideration.

A comparison of trajectories reveals some important differences between test-and-trace and vaccination. Test-and-trace removes people at a proportionate rate from the infected stock, and hence has most of its effect at the beginning of the period when prevalence is high. In contrast, vaccination is a gradual process which takes some time before it significantly affects the course of the disease. Eventually, as more people are vaccinated and prevalence declines, the effects of test-and-trace decline and those of vaccination come to the fore.

Another contrast between test-and-trace and vaccination concerns government intervention. When there is test-and-trace alone (i.e. without vaccination), government intervention continues at high intensity for a long time in order to prevent an explosion of the disease. In contrast, when there is vaccination alone, the number of susceptible
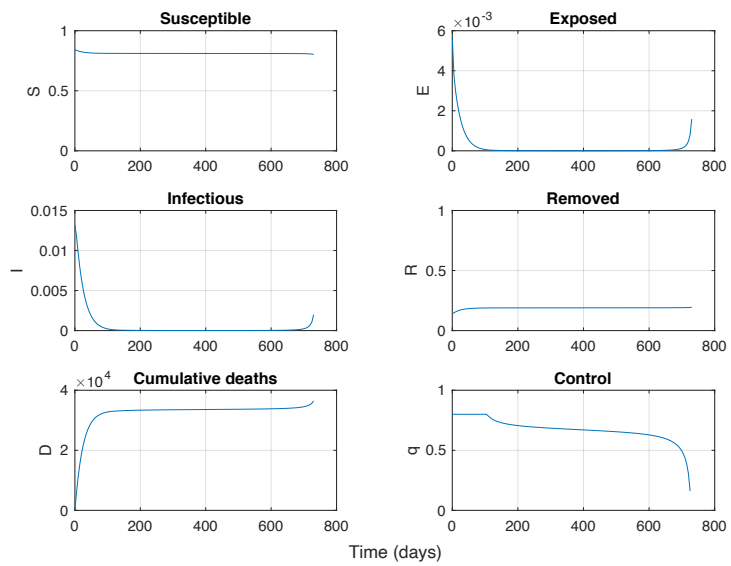
13

Figure 7: Trajectories of the optimal solution with $\pi_D = 2000$. No Vaccination or Test and Trace.

individuals shrinks steadily as the vaccine is rolled out, allowing intervention to be relaxed relatively fast (Figure 9). As a result, in our examples, economic cost is lower under vaccination alone than under test-and-trace alone.
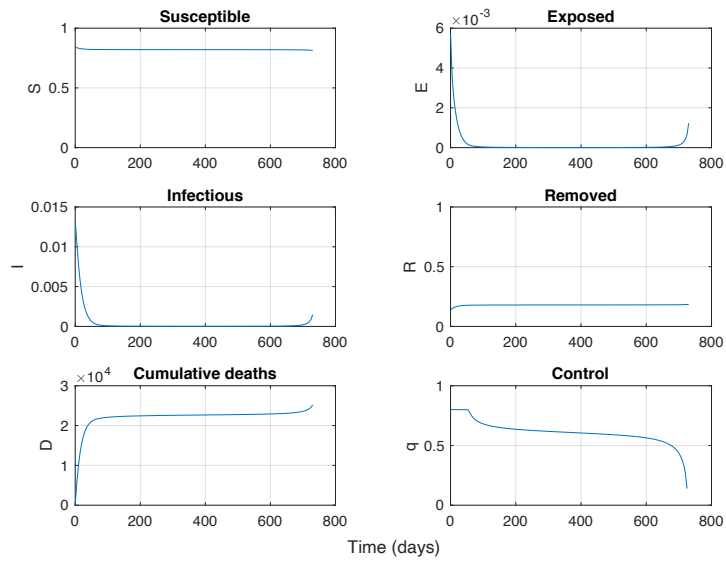
Figure 8: Trajectories of the optimal solution with $\pi_D = 2000$. With Test and Trace but no Vaccination.

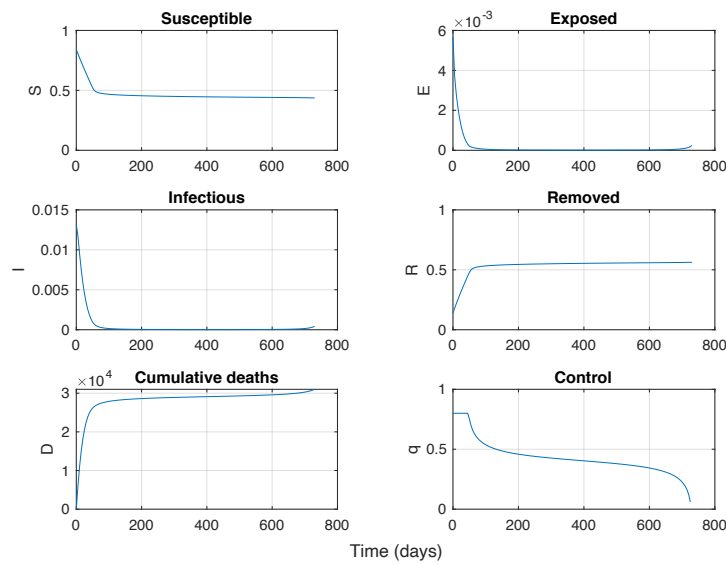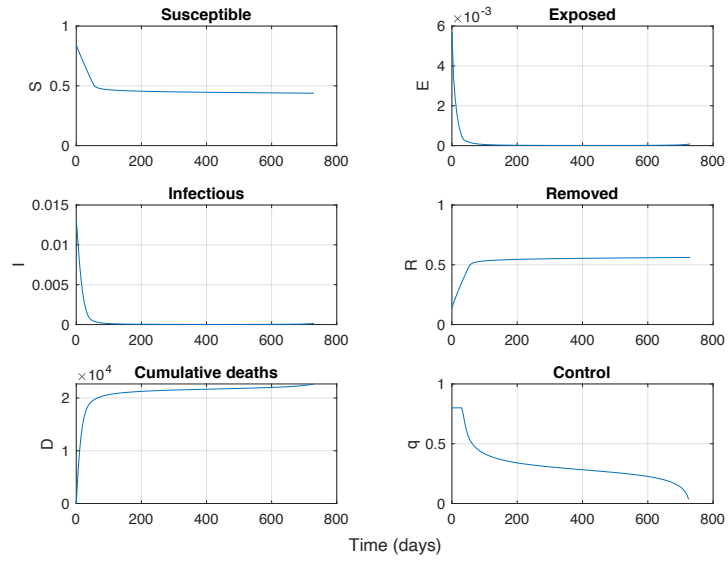Figure 9: Trajectories of the optimal solution with $\pi_D = 2000$. 50% vaccination but no Test and Trace.

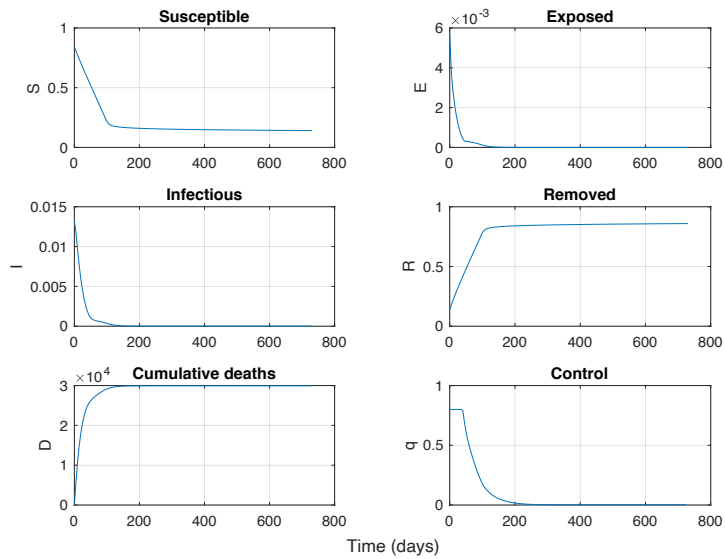Figure 10: Trajectories of the optimal solution with $\pi_D = 2000$. 50% vaccination and Test and Trace.

Figure 11: Trajectories of the optimal solution with $\pi_D = 2000$. 80% vaccination but no Test and Trace.

Trajectory	Initial lockdown (days)	Basic cost $\mathscr{E}$ (£)	Test and Trace cost (£)	Vaccination cost (£)	Total cost (£)	Deaths D_T (thousands)
Baseline (Figure 7)	105	4630			4630	36
Test & Trace (Figure 8)	55	3503	550		4053	25
50% Vaccination (Figure 9)	45	1588		83	1671	31
50% Vaccination with Test and Trace (Figure 10)	30	875	550	83	1508	23
80% Vaccination (Figure 11)	35	592		133	725	30

Table 1: Summary of optimal solutions with $\pi_D = 2000$. All costs are shown *per capita*.

8   Conclusions

We draw three important conclusions from this work.

(i) Test-and-trace policies and vaccination should be viewed as complementary policies which operate in different time frames. The former exerts its main effect in the opening weeks, whilst the latter acts more gradually and takes some months to achieve its full impact.

(ii) The optimal duration and severity of lockdown depends crucially on the effectiveness of these key disease control mechanisms.

(iii) Inclusion of the 'cost of death' into cost-benefit modelling, is both feasible and useful. However, it does have one potential drawback. Under certain conditions, a small change in the value of life can radically alter the optimal length of lockdown, and severity of subsequent policy. This reinforces the skepticism shown by Mark Carney, in his third Reith Lecture [5], as to whether the value of life constitutes a good guide for policy. It may be more appropriate, as he suggests, to decide on a socially acceptable number of deaths, and then to seek the optimum length of lockdown, and severity of subsequent policy, that would bring about the this outcome in the least costly way. We have shown how this could be done.

9   Acknowledgements

10   Funding

Appendix

A   Test and Trace — detailed model

We assume that the whole population is sampled randomly for testing, and that infected individuals are over-represented in the sample by a factor. Thus the proportion of Exposed and Infectious individuals in the sample (on day t) is $g_E E(t)$ and $g_I I(t)$, respectively.

Suppose that the proportion of false negatives for the Exposed and Infectious individuals in the sample is $1 - p_E$ and $1 - p_I$, respectively. Then the proportion of the sample that is correctly tested positive is $g_E p_E E(t) + g_I p_I I(t)$. Let the number of individuals tested on day t be $N(t)$. Then the number of Exposed individuals correctly tested positive is

$$E^+(t) = g_E p_E E(t) N(t)$$

and the number of Infectious individuals correctly tested positive is

$$I^+(t) = g_I p_I I(t) N(t)$$

Then the total number of individuals correctly tested positive is $[g_E p_E E(t) + g_I p_I I(t)]N(t)$.

Now suppose that each correctly-tested positive individual has c close contacts. Then the number of such contacts is $c[g_E p_E E(t) + g_I p_I I(t)]N(t)$, and if the proportion of such contacts that is positive is a then the number of such positive close contacts is $ac[g_E p_E E(t) + g_I p_I I(t)]N(t)$. Assume that none of the contacts of false positives are infected. Then the number of positive close contacts who are in the Exposed population is

$$C_E^+(t) = ac[g_E p_E E(t) + g_I p_I I(t)]N(t)\frac{E(t)}{E(t) + I(t)}$$

and the number of positive close contacts who are in the Infectious population is

$$C_I^+(t) = ac[g_E p_E E(t) + g_I p_I I(t)]N(t)\frac{I(t)}{E(t) + I(t)}$$

Thus the total number of Exposed individuals told to self-isolate is

$$E^+(t) + C_E^+(t) = N(t)\left[g_E p_E E(t) + ac[g_E p_E E(t) + g_I p_I I(t)]\frac{E(t)}{E(t) + I(t)}\right]$$

and the total number of Infectious individuals told to self-isolate is

$$I^+(t) + C_I^+(t) = N(t)\left[g_I p_I I(t) + ac[g_E p_E E(t) + g_I p_I I(t)]\frac{I(t)}{E(t) + I(t)}\right]$$

Not all of those told to self-isolate will comply with the instruction. Let the average compliance be b ($0 < b < 1$). Then on average the number of Exposed individuals who self-isolate on day t is

$$bN(t)\left[g_E p_E + \frac{ac[g_E p_E E(t) + g_I p_I I(t)]}{E(t) + I(t)}\right]E(t)$$

and the number of Infectious individuals who self-isolate is

$$bN(t)\left[g_I p_I + \frac{ac[g_E p_E E(t) + g_I p_I I(t)]}{E(t) + I(t)}\right]I(t)$$

Thus (recalling that we have normalised the total population to be 1) the fraction of the Exposed population that is effectively removed by test-and-trace is

$$r_E = bN(t)\left[g_E p_E + \frac{ac[g_E p_E E(t) + g_I p_I I(t)]}{E(t) + I(t)}\right]$$

and the fraction of the Infectious population that is effectively removed is

$$r_I = bN(t)\left[g_I p_I + \frac{ac[g_E p_E E(t) + g_I p_I I(t)]}{E(t) + I(t)}\right]$$

19

This leads us to modify equations (2) and (3) to (16) and (17). Since these individuals are removed from further interaction with the Susceptible population we modify equation (4) to (18). Note that this results in the total population remaining constant, ie (11) still holds.

Note there is a minor problem of timing with deaths but the ultimate total $D(T)$ is hardly affected.

B Solving the optimal control problem

The problem of minimising the total cost $J(\mathscr{E}, D)$ (see (15)), subject to the constraints imposed by the SEIR model (1)–(12), and by the constraint $0 \le q(t) \le q_{max}$, is an optimal control problem. The problem of minimising $J(\mathscr{E}, D)$ when either $\mathscr{E}$ or D is fixed (which arises when computing points on the 'dents' in Figures 2 or 6) is a closely-related optimal control problem.

We solve these problems by discretising time into intervals of length T_s, thus converting the problem into a finite-dimensional optimisation problem (with T/T_s decision variables), and applying numerical optimisation methods. The resulting problem is non-convex, despite $J(\mathscr{E}, D)$ being convex in $q(t)$, $I(t)$ and D_T, because the SEIR equations are nonlinear. We found the CasADi software framework [4], in which the optimisation solver was IPOPT [26], to be very effective for solving the problems that arose while writing this paper[2].

In all the scenarios considered in this paper we have $T = 730$ days (ie 2 years) and we used $T_s = 5$ days. We thus had 146 decision variables, and the same number of upper and lower bounds on q.

C Parameter values

We have attempted to choose parameter values which reflect the position of the UK in early January 2021.

Dynamic Parameters

$\gamma = 0.142 (= 1/7), \sigma = 0.333 (= 1/3)$. These values imply that on average a person becomes infectious 3 days after being infected and then remains infectious for a further 7 days. These values are consistent with the information in [18]. Phipps, Grafton and Kompas (2020).

The basic reproduction number is $\beta_0 / \gamma = 3.75$ and hence $\beta_0 = 0.536$.

Initial Conditions

$S(0) = 0.8419, E(0) = 0.0057, I(0) = 0.0132, R(0) = 0.1382, D(0) = 0$. The value of $1 - S(0)$ is derived from the ONS Antibody Survey [11] by summing over the four nations of the UK and averaging over the two months 10 December 2020 to 6 January 2021 and. 7 January to 3 February 2021. The sum $E(0) + I(0)$ is taken from the ONS Infection Survey [1]. It is assumed that $\sigma E(0) = \gamma I(0)$. The proportion in the removed

---

[2]The code supporting this article will be uploaded as part of the supplementary material.

category is obtained from the formula $R(0) = 1 - S(0) - E(0) - I(0)$. The initial value $D(0)$ is set equal to zero because optimisation is forward-looking.

Deaths
$\delta_0 = 0.01, \eta_1 = 0, \eta_2 = 0.0027$. The estimate for δ_0 is obtained by dividing the cumulative number of deaths [12] up to mid-January 2021 by the estimated number in the removed category (see above) at the beginning of the year. The parameter η_2 is a notional amount to allow for the fact that elderly and vulnerable people with a relatively high death rate are vaccinated first. It implies that $\delta(t) = 0.005$ after 250 days (assuming that vaccination continues that long).

Costs
$C_{max} = 0.01$. This is based on the deviation of daily UK per capita GDP in January 2021 below the pre-2020 trend [13].
$q_{max} = 0.8$. This value is consistent with an effective reproduction number $(1-q)\beta_0 S/\gamma = 0.75$ when $q = q_{max}, S = 1.0$ and $\beta_0/\gamma = 3.75$. In March 2021 the UK effective reproduction number was estimated to lie within the range 0.7-0.9 [14].
$\phi = 2$. This is the assumed measure of convexity. The higher the value of ϕ the more is policy biased away from lockdown (see [20]).
$\pi_I = 1$.This is a notional allowance for medical costs and lost output due to the absence from work of Covid-infected individuals and their close contacts. It is additional to the loss caused by economy-wide restrictions.
π_D varies in our investigations.

Test and Trace
$p_E = 0.35, p_I = 0.70$.These values are chosen in the light of the updated Cochrane review [17].
$g_E = 7, g_I = 14$. These values are chosen to ensure that $g_E p_E E_0 + g_I p_I I_0 = 0.143$. The quantity 0.143 is the proportion of cases testing positive in England in early January 2021 [15], after subtracting 0.008 for false positives (see [17]).
$n = 0.006$. This corresponds to the assumed number of tests per day (400,000).
$a = 0.1$. This is a conservative figure to allow for the possibility that relatively few close contacts are infected.
$b = 0.4$. This is a conservative figure to allow for the possibility that fewer people conform to self-isolation instructions than is reported in the UCL study [25].
$c = 1.76$. This is the ratio of close contacts reached to cases testing positive in England in early January 2021 [15].
$r_E = 0.01, r_I = 0.025$. These are the rounded values implied by the above assumptions.

Vaccination
$\bar{S} = 0.5, 0.2$. Illustrative assumptions.
$F_0 = 0.006$. This corresponds to an assumed rate of 400,000 vaccinations per day.

Simulation
Scenario duration $T = 730$ days.
Time discretisation $T_s = 5$ days.

References

[1]
https://www.ons.gov.uk/peoplepopulationandcommunity/
healthandsocialcare/conditionsanddiseases/bulletins/
coronaviruscovid19infectionsurveypilot/26march2021/
#number-of-people-who-had-covid-19-in-england-wales-
northern-ireland-and-scotland. Accessed 31 March 2021.

[2] D. Acemoglu et al. *A multi-risk SIR model with optimally targeted lockdown*. Tech. rep. National Bureau of Economic Research, 2020.

[3] F.E. Alvarez, D. Argente, and F. Lippi. *A Simple Planning Problem for COVID-19 Lockdown*. NBER Working Paper 26981. 2020.

[4] Joel A. E. Andersson et al. "CasADi – A software framework for nonlinear optimization and optimal control". In: *Mathematical Programming Computation* 11 (2019), pp. 1–36.

[5] House of Commons Public Accounts Committee. *COVID-19: Test, track and trace (part 1): 47th report of session 2019–2021*. Mar. 2021.

[6] M. Eichenbaum, S. Rebelo, and M. Trabandt. *The Macroeconomics of Epidemics*. Tech. rep. NBER Working Paper 26882, 2020.

[7] A. Eilersen and K. Sneppen. "Cost–benefit of limited isolation and testing in COVID-19 mitigation". In: *Scientific Reports* 10 (2020). https://doi.org/10.1038/s41598-020-75640-2, p. 18543.

[8] N. Ferguson and et al. *Report 9 – Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand*. https://www.imperial.ac.uk/mrc-global-infectious-disease-analysis/covid-19/report-9-impact-of-npis-on-covid-19/. 2020.

[9] C. Gros et al. "Containment efficiency and control strategies for the corona pandemic costs". In: *Sci.Rep.* 11 (2021), p. 6848. URL: https://doi.org/10.1038/s41598-021-86072-x.

[10] Department of Health and Social Security. *UK COVID-19 vaccines delivery plan*. Jan. 2021.

[11] Office for National Statistics. *Coronavirus (COVID-19) Infection Survey, antibody data for the UK: 16 March 2021*.
https://www.ons.gov.uk/peoplepopulationandcommunity/
healthandsocialcare/conditionsanddiseases/articles/\\
coronaviruscovid19infectionsurveyantibodydatafortheuk/
16march2021. Accessed 31 March 2021.

[12] Office for National Statistics. *Deaths in United Kingdom*.
https://coronavirus.data.gov.uk/details/deaths. Accessed 31 March 2021.

[13] Office for National Statistics. *GDP monthly estimate, UK: January 2021.* `https://www.ons.gov.uk/economy/grossdomesticproductgdp/bulletins/gdpmonthlyestimateuk/january2021`. Accessed 31 March 2021.

[14] Office for National Statistics. *The R value and growth rate in the UK.* `https://www.gov.uk/guidance/the-r-number-in-the-uk#latest-r-and-growth-rate`. Accessed 31 March 2021.

[15] Office for National Statistics. *Weekly statistics for NHS Test and Trace (England): 31 December 2020 to 6 January 2021.* `https://www.gov.uk/government/publications/nhs-test-and-trace-england-statistics-31-december-2020-to-6-january-2021/weekly-statistics-for-nhs-test-and-trace-england-31-december-2020-to-6-january-2021`. Accessed 31 March 2021.

[16] OECD. *Mortality Risk Valuation in Environment, Health and Transport Policies.* 2012, p. 140. DOI: `https://doi.org/https://doi.org/10.1787/9789264130807-en`. URL: `https://www.oecd-ilibrary.org/content/publication/9789264130807-en`.

[17] Brimingham Health Partners. *Updated Cochrane review assesses how accurate rapid tests are for detecting COVID-19.* `https://www.birminghamhealthpartners.co.uk/updated-cochrane-review-assesses-how-accurate-rapid-tests-are-for-detecting-covid-19/`. Accessed 31 March 2021. 2021.

[18] S.J. Phipps, R.Q. Grafton, and T. Kompas. "Robust estimates of the true (population) infection rate for COVID-19: a backcasting approach". In: *R. Soc. open sci.* 7.200909 (2020). URL: `http://doi.org/10.1098/rsos.200909`.

[19] Robert Rowthorn. "A Cost–Benefit Analysis of the Covid-19 Disease". In: *Covid Economics* 9 (2020).

[20] Robert Rowthorn and Jan Maciejowski. "A cost-benefit analysis of the Covid-19 disease". In: *Oxford Review of Economic Policy* 36.S1 (2020), S38–S55.

[21] T. Sasse, C. Haddon, and A. Nice. *Science advice in a crisis.* Institute for Government. `https://www.instituteforgovernment.org.uk/sites/default/files/publications/science-advice-crisis_0.pdf`. Dec. 2020.

[22] A. Scherbina. "Could the United States Benefit from a Lockdown? A Cost-Benefit Analysis". In: *Covid Economics* 65 (2021), pp. 78–107.

[23] S. Sheffield et al. *Strict Physical Distancing May Be More Ecient: A mathematical argument for making the lockdowns count.* MedRxiv `https://doi.org/10.1101/2020.05.19.20107045`. 2020.

[24] Scott Sheffield. *Best and worst policy control in low-prevalence SEIR.* 2020. arXiv: `2009.07792 [math.OC]`.

[25] UCL. *Lockdown compliance improving but low take up of Covid tests 'worrying'*. `https : / / www . ucl . ac . uk / news / 2021 / jan / lockdown - compliance - improving - low - take - covid - tests - worrying`. Accessed 31 March 2021. 2021.

[26] A. Wächter and L.T. Biegler. "On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming". In: *Mathematical Programming* 106.1 (2006), pp. 25–57.