

A cost–benefit analysis of the COVID-19 disease

Robert Rowthorn* and Jan Maciejowski**

Abstract: The British government has been debating how to escape from the lockdown without provoking a resurgence of the COVID-19 disease. There is a growing recognition of the damage the lockdown has caused to economic and social life. This paper presents a simple cost–benefit analysis inspired by optimal control theory and incorporating the SIR model of disease propagation. It also reports simulations informed by the theoretical discussion. The optimal path for government intervention is computed under a variety of conditions. These include a cap on the permitted level of infection to avoid overload of the health system, and the introduction of a test and trace system. We quantify the benefits of early intervention to control the disease. We also examine how the government’s valuation of life influences the optimal path. A 10-week lockdown is only optimal if the value of life for COVID-19 victims exceeds £10m. The study is based on a standard but simple epidemiological model, and should therefore be regarded as presenting a methodological framework rather than giving policy prescriptions.

Keywords: COVID-19, benefit–cost, lockdown

JEL classification: C61, H12, J17

I. Introduction

There has been a debate in Britain about the best policy for dealing with the COVID-19 virus. The official policy was originally to proceed step by step and intensify, as required, the measures that encourage hygiene and social distancing. Such measures range from careful hand-washing through to the banning of large public gatherings, the closing of pubs, restaurants, and many shops, quarantine or near quarantine of vulnerable people, and restrictions on national and international travel. The gradualist approach of the government was attacked by critics who called for vigorous action of the type observed in Italy and Spain. The government responded by implementing an unprecedented lockdown on economic and social life. A factor behind this change of heart was concern about the potential shortage of intensive care beds if the disease was

*Emeritus Professor of Economics, University of Cambridge; e-mail: rer3@cam.ac.uk

**Emeritus Professor of Control Engineering, University of Cambridge; e-mail: jmml@cam.ac.uk

We are grateful to Bill Martin, Pal Ormerod, Flavio Toxvaerd, Martin Wolf, and especially Nicholas Rau and David Vines, for their comments. We should also like to thank an anonymous referee for providing a valuable review at short notice. The final version of this paper was submitted for publication on 9 July 2020. doi:10.1093/oxrep/graa030

© The Author(s) 2020. Published by Oxford University Press.

For permissions please e-mail: journals.permissions@oup.com

not brought quickly under control. At the time of writing, the government was still searching for a way to exit the lockdown without provoking a surge in the disease.

The measures required to inhibit disease transmission can be very costly in economic and social terms, including depression and other ‘diseases of despair’ among the millions who lose their jobs. These costs must be weighed against the medical benefits of intervention. The decision when to intervene and on what scale is a classic optimal control problem. This paper explores the choices facing the government using a simple mathematical model that is inspired by optimal control theory.¹ For clarity we omit details of the full optimal control model which are to be found in [Rowthorn \(2020\)](#). The paper complements the theoretical analysis with some illustrative simulations. These simulations should not be taken literally but they indicate some of the issues and orders of magnitude involved.

The economic literature on the optimal control of disease is sparse and its models mostly deal with individual behaviour and the externalities of individual decision-making with regard to treatment, vaccination, or social distancing.² These are not our concern here. Our interest is in the cost–benefit analysis of large-scale interventions such as lockdowns. This involves an approach that is unusual in the existing optimal control literature on disease. Costs and benefits in existing optimal control models are typically functions of the health status of individuals, computed by assigning values or weights to individuals according to their health status. This is a procedure followed here. However, unlike these models we also make an explicit allowance for the more general costs of comprehensive interventions such as lockdowns. Such costs depend on the scale and type of intervention but they are not linked in a direct way to the health status of individuals. These costs are given a central role in this paper.

Since the outbreak of the epidemic there has been a spate of thought-provoking articles on economic aspects of COVID-19. Two, in particular, deserve special mention. [Acemoglu *et al.* \(2020\)](#) examine targeted lockdowns in a multi-group SIR model where infection, hospitalization, and fatality rates vary between groups—in particular between the ‘young’, ‘the middle-aged’, and the ‘old’. They also allow for the fact that lockdown damages the economy and reduces the productivity of non-infected members of the workforce. Their paper, incidentally, contains a good review of the recent literature. [Giordano *et al.* \(2020\)](#) draw on the experience of the Italian epidemic. Their model distinguishes between detected and undetected infection cases, and between cases with different severity of illness. They argue that social-distancing measures are necessary and effective, and should be promptly enforced at the earliest stage. They also argue that lockdown measures can only be relieved safely when an effective system of testing and contact tracing is in place. These are both excellent articles, and nothing in the present article contradicts their findings.

A system of testing and tracing is most effective when the number of people to be tested or contacted is relatively small. It may be feasible to test small subgroups of the population on a frequent basis and trace their contacts if they test positive

¹ The term ‘optimal control theory’ is conventionally restricted to models that utilize Pontryagin’s maximum principle.

² [Toxvaerd \(2020\)](#) for a brief survey of this literature. Among the articles worthy of note are [Sethi \(1978\)](#), [Gersovitz \(2010\)](#), [Reluga \(2010\)](#), [Chen *et al.* \(2011\)](#), [Chen \(2012\)](#), [Fenichel \(2013\)](#), [Rowthorn and Toxvaerd \(2015\)](#), [Toxvaerd \(2019\)](#), [Toxvaerd and Rowthorn \(2020\)](#).

(Cleevely *et al.*, 2020). Care home workers are an example. However, a policy of targeted testing is of limited use as a means of infection control if the disease is widespread, since most of the infected population will not be in the groups selected for testing. The alternative is universal and frequent random testing, but this is likely to be prohibitively expensive, as Cleevely *et al.* point out. If the scale of infection is too large for the system of testing and tracing to handle unaided, and if there is currently no treatment or vaccine, some form of social distancing will be required. This is the case in the present article. Indeed, our basic model goes further. It assumes that a perfect vaccine will become available on a known date in the future and that prior to this date there exists no testing and tracing regime at all. There is also no currently available treatment for the disease. Hence social distancing is the only feasible means of disease control. However, in one simulation we consider a scenario in which a test and trace regime is established in advance of vaccination.

The analysis assumes that the scale of social distancing is determined by government fiat alone. In reality, as the disease spreads and people become aware of the risks involved, there will be a degree of voluntary social distancing. As a result, the more apocalyptic predictions of what would happen without draconian intervention may be wide of the mark. The implications of endogenous behaviour are not explored here, but are the subject of another paper (Ormerod *et al.*, 2020).

The theoretical section of this paper was written the day after Prime Minister Johnson announced a full-scale lockdown. The first batch of simulations was completed shortly thereafter with the aim of influencing the ongoing policy debate. The paper including simulations was published in mid-April in the CEPR real-time online journal *Covid Economics* (Rowthorn, 2020). These simulations were comprehensively revised in May and June for this issue of the *Oxford Review of Economic Policy*. By the time the journal appears, the die will have been cast and the actual policy choices of the government will be there for all to see. However, we hope that this paper will continue to provide a useful framework for thinking about the cost–benefit analysis of disease control. Our study is based on a standard but simple epidemiological model, and should therefore be regarded as presenting a methodological framework rather than giving policy prescriptions.

II. The model

The analysis in this paper uses a modified version of the standard SIR model of disease propagation. Ignoring births and deaths from non-COVID-19 causes, the initial population will divide in the future into three groups of people: susceptible, infected, and removed—denoted, respectively, by $S(t)$, $I(t)$, and $R(t)$. The removed group includes people who have died from the disease. They are denoted by $D(t)$. The population at the start of the epidemic is normalized to 1, so these various quantities can be interpreted as shares. Individuals who are infected remain infectious until they recover or die. Infected individuals who recover acquire complete immunity, so the journey from $S(t)$ via $I(t)$ to $R(t)$ is in one direction only.

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

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

$$\frac{dI(t)}{dt} = \beta(t)S(t)I(t) - \gamma I(t) \quad (2)$$

$$\frac{dR(t)}{dt} = \gamma I(t) \quad (3)$$

$$\frac{dD(t)}{dt} = \delta \gamma I(t) \quad (4)$$

$$S(0) = S_0 \geq 0 \quad (5)$$

$$I(0) = I_0 \geq 0 \quad (6)$$

$$R(0) = R_0 \geq 0 \quad (7)$$

$$D(0) = D_0 \geq 0 \quad (8)$$

$$S(t) + I(t) + R(t) = 1 \quad (9)$$

where γ and δ are constant. These constants indicate, respectively, the rate at which infected individuals cease to be infectious, and the probability that an infected individual will die. Note that there are only two genuinely independent state variables in this model. For example, if the trajectories of $I(t)$ and $R(t)$ are known, the trajectories of $S(t)$ and $D(t)$ are uniquely determined by equations (1) and (4).

Equation (1) indicates how the pool of susceptible individuals 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. Equation (2) indicates how the pool of currently infected individuals is augmented by the inflow of newly infected individuals and depleted by the outflow of infected individuals who recover or die. The rate of outflow is $\gamma I(t)$ of whom a fraction δ are dead. Equation (3) indicates how the removed category is augmented by the inflow of newly recovered or dead individuals.

The coefficient $\beta(t)$ in equation (1) is a variable which depends on the current intensity of social interaction. The intensity of social interaction depends, in turn, on the measures that the government puts in place to inhibit the spread of the disease. Specifically, it is assumed that:

$$\beta(t) = (1 - q(t))\beta_0 \quad (10)$$

where $q(t)$ is an index of policy severity. The effective reproduction rate of the disease is

$$r(t) = (1 - q(t))S(t)r_0 \quad (11)$$

where

$$r_0 = \frac{\beta_0}{\gamma}$$

The number r_0 indicates how many people the average infected person would infect in a situation where everyone was susceptible to the disease and there was no government intervention to control its spread. The number $r(t)$ indicates how many people are infected if there is government intervention and some people are immune. If $r(t) < 1$, the prevalence of the disease will be diminishing through time.

Government intervention comes at a cost $C(q(t))$ in the form of damage to the economy. This cost is independent of the number of people currently infected and is the result of society-wide measures to control the disease. It is in addition to the various costs arising directly from infection. The function $C(\cdot)$ is assumed to be twice differentiable and such that

$$\begin{aligned} C(0) = 0, C(q_{\max}) = C_{\max} < \infty \\ C'(q) \geq 0, C''(q) > 0 \text{ for } q \in [0, q_{\max}] \end{aligned} \quad (12)$$

where $q_{\max} < 1$ is an upper limit beyond which it is not feasible to increase q . Thus, $C(q)$ is strictly convex over the relevant range. Examples are shown in [Figure 1](#) which plots the function $C(q) = C_{\max} \left(\frac{q}{q_{\max}}\right)^{1+\phi}$ for various values of $\phi > 0$. When q is close to zero, the marginal cost of intervention is low but rises steeply at higher values of q . These are realistic assumptions. Think of hand-washing at one end of the scale and the closure of shops, pubs, cafés, and restaurants at the other.

The government is assumed to have perfect foresight. Thus, the entire control trajectory is decided at the very outset. The system is therefore open loop, whereas in a closed loop system the control is modified in the light of new information. We assume that an effective vaccine will become available at time T at negligible cost.³ For simplicity we also assume that a cure will become available at the same time as the vaccine at zero cost.

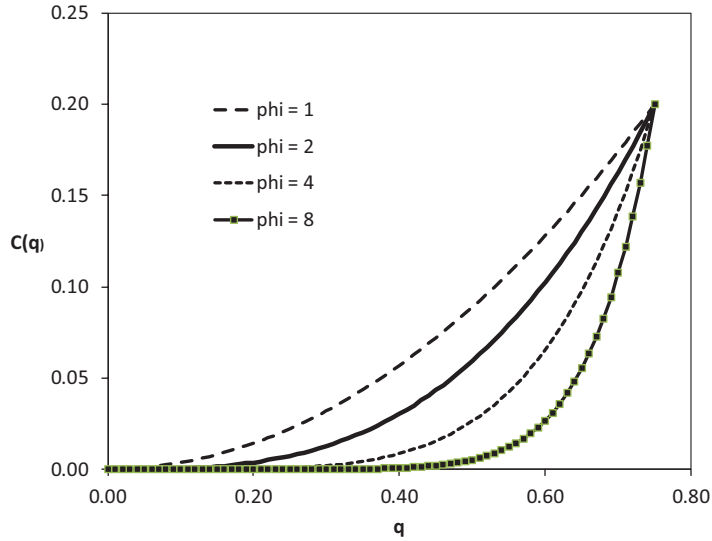
The government chooses the trajectory $q(t)$ so as to minimize the following quantity subject to the foregoing equations:

$$J = \int_0^T [\pi_A I(t) + C(q(t))] dt + \pi_D [D(T) - D(0)] \quad (13)$$

where π_A is the monetary value that planners assign to each person who is currently alive and infected and π_D the additional value they assign to those who die.

Total economic cost refers to the cost of economy-wide measures such as lockdown, plus the various costs arising directly from the disease, such as the cost of treatment and the loss of output due to withdrawal from production of infected individuals and their close associates. It is defined as follows:

³ In a game theoretic paper on social distancing [Reluga \(2010\)](#) also assumes that vaccination will occur on a fixed date in the future. In their recent paper, [Acemoglu et al. \(2020\)](#) assume that a vaccine and a cure become simultaneously available.

Figure 1: Weekly cost $C(q)$: £'000 per capita

$$E = \int_0^T [\pi_A I(t) + C(q(t))] dt \quad (14)$$

Thus,

$$J = E + \pi_D [D(T) - D(0)] \quad (15)$$

The monetary allowance for death π_D is not included in economic cost since most of the people who die from the disease are not economically active, so their death does not have a significant effect on output. Their cost of treatment prior to death is included in the π_A term which is an average for all infected individuals, including those who die and those who are asymptomatic or require no treatment.

III. Simulation

The optimization problem defined above has no explicit solution. In the absence of such a solution, the obvious procedure is to explore the properties of the system by means of numerical simulation. We solved the optimization problem by posing it as a nonlinear programming problem.⁴

Assumptions

The simulations assume that the cost of intervention is given by the function:

⁴ This involved discretizing time into steps of tenths of a week, then minimizing a function of 520 variables (in one case 1,040 variables) under constraints. We used the Model Predictive Control Toolbox in Matlab.

$$C(q) = C_{\max} \left(\frac{q}{q_{\max}} \right)^{1+\phi} \quad (16)$$

where C_{\max} is the cost of the maximum feasible lockdown and $\phi > 0$. The larger is the value of ϕ the lower is the cost of the other interventions relative to lockdown and the greater is the economic benefit of moving to less draconian forms of intervention (see [Figure 1](#)).

Our simulations use parameter values that we hope are realistic, although given the paucity of reliable data, a fair amount of guesswork is involved. The simulations take 1 April 2020 as their notional starting point for optimization, although the epidemic is assumed to have started some weeks earlier. The lockdown was officially announced on 23 March, but it was not until 1 April that it had a clear effect on the number of people infected ([King’s College, 2020](#)). The unit of time is a week and the time horizon is $T=52$. The monetary unit of account is thousands of UK pounds. There are initially 2m people currently infected and therefore infectious. In addition a further 1.4m have had the disease and recovered or died. The initial conditions are thus $I_0=0.030$, $R_0=0.021$. The death rate is $\delta = 0.7$ per cent. The UK population is assumed to be 66.8m.

The parameters in the baseline scenario have the following values: $\beta_0 = 4.8$, $\gamma = 1.6$, $C_{\max} = 0.20$, $\pi_A = 2$, $\pi_D = 2,000$. Infected individuals cease to be infectious at an exponential rate of -1.6 per week, which implies that after 2 weeks 96 per cent are no longer infectious. They have either recovered or died. In the absence of intervention the net reproduction rate $r_0 = 3$. The *per capita* weekly cost of full lockdown is £200 which is approximately 35 per cent of *per capita* GDP at factor cost, in line with the Office for Budget Responsibility’s prediction of what the lockdown might do to the UK economy ([OBR, 2020](#)). The values $\pi_A = 2$ and $\pi_D = 2,000$ assume that planners assign a monetary value of £2,000 per week to the average currently infected person, plus a further £2m to each fatality. The latter figure is what the UK Treasury assumes in project evaluation as the value of a prevented fatality ([Dolan and Jenkins, 2020](#)).

To derive the path before 1 April, we assume that 4.7 weeks previously the state of the system was $S_{-4.7} = 1 - 10^{-8}$, $I_{-4.7} = 10^{-8}$, $R_{-4.7} = 0$, $D_{-4.7} = 0$. From this starting point the system is assumed to grow freely with parameters $\beta = 4.8$, $\gamma = 1.6$, $\delta = 0.007$ until 1 April, when government intervention in our simulations begins. We ignore the limited interventions of the government before 1 April.

IV. Results

[Tables 1](#) and [2](#) provide information about the optimum path under various scenarios. The numbers for deaths and total economic cost in these tables have been adjusted to include the pre-intervention weeks. This is a small adjustment which does not materially affect the results. It makes it easier to compare scenarios with different starting dates for intervention.

[Figure 2](#) shows what happens if the government does nothing to control the disease and restricts itself to the medical treatment of those infected. Within a few weeks, 90 per cent of the population has been infected and the cumulative death toll by the end of the year is 440,000 ([Table 1](#)). At the peak of the epidemic 20m people are currently infected and hence infectious.

Table 1: Optimal paths compared

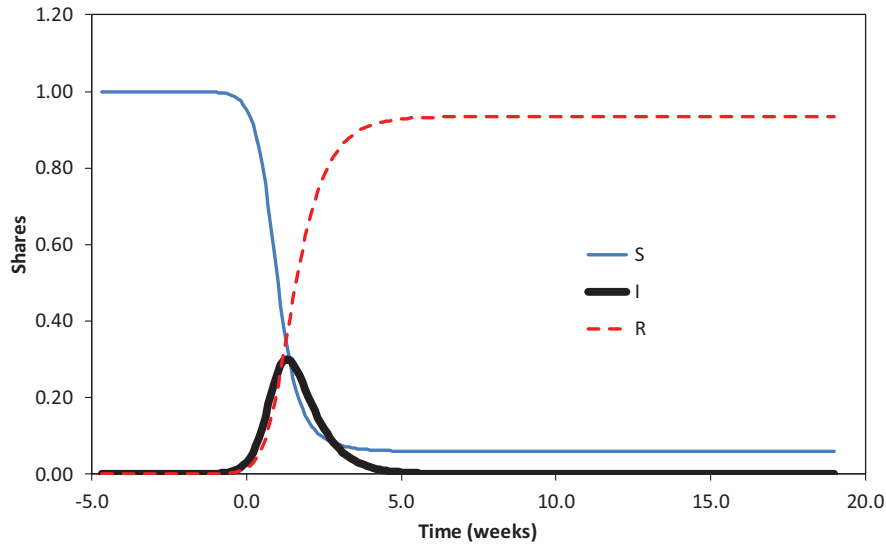
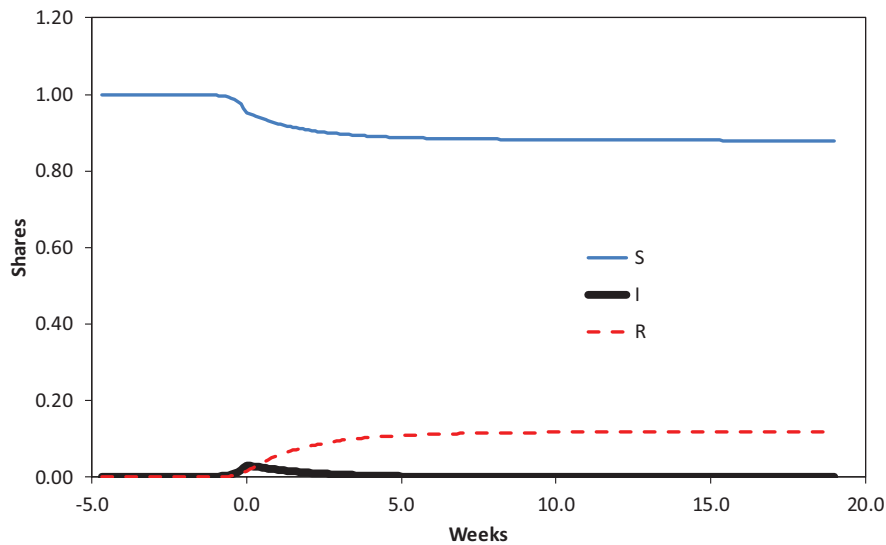
	ϕ	Value of life (£m)	Duration of lockdown (weeks)	Peak infection (m)	Total deaths (thousands)	Total economic cost (£ per capita)
Do nothing	2	2.0	0	20.1	439.8	14,342
Baseline	2	2.0	5.3	2.0	59.9	6,589
Low relative cost	4	2.0	1.8	2.0	67.1	4,811
High relative cost	1	2.0	7.9	2.0	57.6	7,660
Long time horizon:						
unconstrained	2	2.0	0	7.0	270.5	1,916
constrained	2	2.0	0	3.3	268.1	2,093
Early start	2	2.0	0.9	0.3	8.3	7,360
Test & trace	2	2.0	6.0	2.0	60.1	3,551

Table 2: Optimal paths and the value of life

	ϕ	Value of life (£m)	Duration of lockdown (weeks)	Peak infection (m)	Total deaths (thousands)	Total economic cost (£ per capita)
Baseline	2	2.0	5.3	2.0	59.9	6,589
High value of life	2	5.0	8.4	2.0	55.9	6,768
Low value of life:						
unconstrained	2	1.0	0	7.1	275.3	1,582
constrained	2	1.0	0	3.3	269.5	1,776
Nil value of life:						
unconstrained	2	0	0	8.8	333.0	1,139
constrained	2	0	0	3.3	335.1	1,327

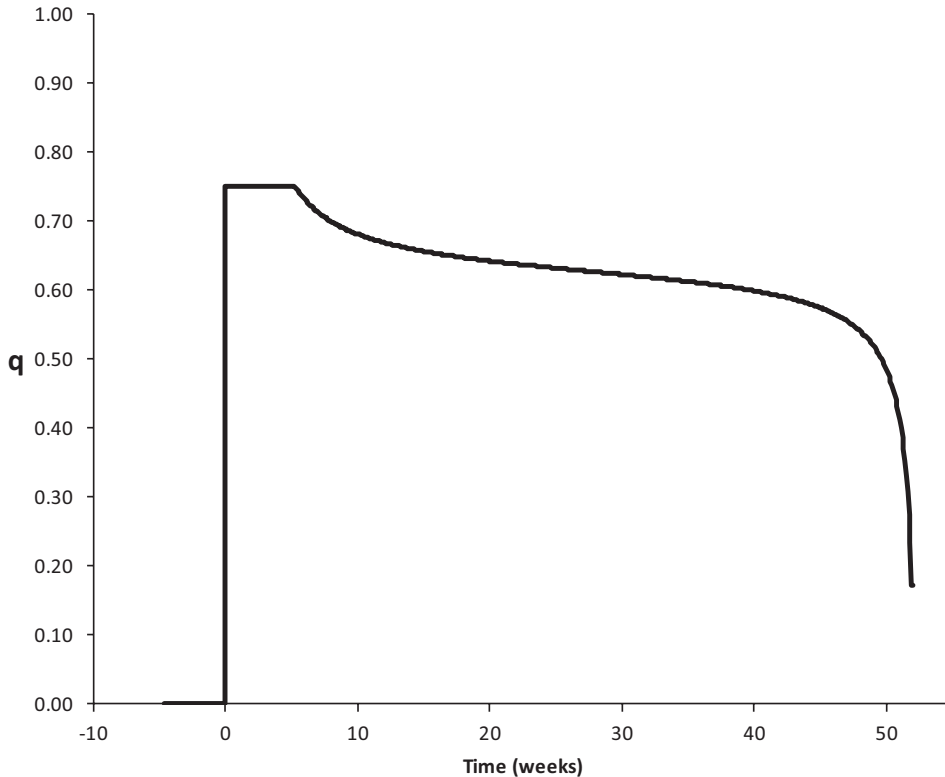
Under the Baseline scenario, the optimum response of the government is to impose a tight lockdown at the very beginning of the planning year. The lockdown lasts 5.3 weeks and brings the disease under control quite soon, although not before millions of people have been infected and many thousands have died (Figure 3). The eventual death toll is 60,000. The death toll is so high because the lockdown is not complete. Lockdown reduces the transmission of the disease but does not entirely prevent it. As a result there is inertia in the system. If the level of infection is already high when the lockdown is imposed, this will continue to be the case for some time thereafter. This is a good reason for acting swiftly before the disease has really taken hold. Once the lockdown is relaxed there is a prolonged period when it is optimal to maintain restrictions close to the minimal level required to contain the disease (Figure 4). During this period the effective reproduction rate r , although rising, is close to 1 (Figure 5). As the vaccination date draws near, restrictions are lifted at an accelerating pace until eventually they are largely abandoned. The result is a brief resurgence of infection which is halted by vaccination or treatment.

Figure 6 compares the course of infection under various scenarios. Under the Early Start scenario, the lockdown is imposed a week earlier, with the result that infection and deaths are much lower. The eventual death toll is around 8,000 as compared to 60,000 under the Baseline scenario. The lockdown is also much shorter: 0.9 weeks as compared

Figure 2: Disease trajectory: no intervention**Figure 3:** Disease trajectory: baseline scenario

to 5.3 weeks. This comparison illustrates clearly the harm that may arise from even a short delay.

In Rowthorn (2020), it was argued that extending the planning horizon does not greatly affect the results. This conclusion is not supported by the more sophisticated simulations reported here. Suppose the vaccine comes on stream after 2 years instead of one. The effect on the optimal path is dramatic. There is no lockdown and the final death toll is 271,000. Peak infection is 7m and eventually over 40 per cent

Figure 4: Optimal path for q : baseline scenario

of the population catches the disease. Infection on the peak scale would impose an intolerable burden on the health system. To avoid such an eventuality, we repeat the simulation with a ceiling of 3.3m on the permitted level of infection. This is just over 50 per cent more than the initial level of infection (2m). The existence of this constraint has little impact on the eventual death toll, although it does reduce the peak load on the health system.

Relative costs

The parameter ϕ conveys information about the relative cost of various interventions. When ϕ is small the economic benefit from a partial relaxation of the lockdown is also small. This creates an incentive to extend the duration of lockdown. Why relax an effective policy for so little economic gain? Conversely, if ϕ is large, the economic gain from a partial relaxation is large. The duration of lockdown is therefore short. Under the Baseline scenario $\phi = 2$ and the lockdown lasts for 5.3 weeks. If $\phi = 1$, the lockdown lasts for 7.9 weeks. If $\phi = 4$, it lasts for 1.8 weeks.

Test and trace

A test and trace system is designed to isolate infectious individuals and their contacts, so that they cannot infect the general population. Within the framework of the present

Figure 5: Optimal path for r : baseline scenario

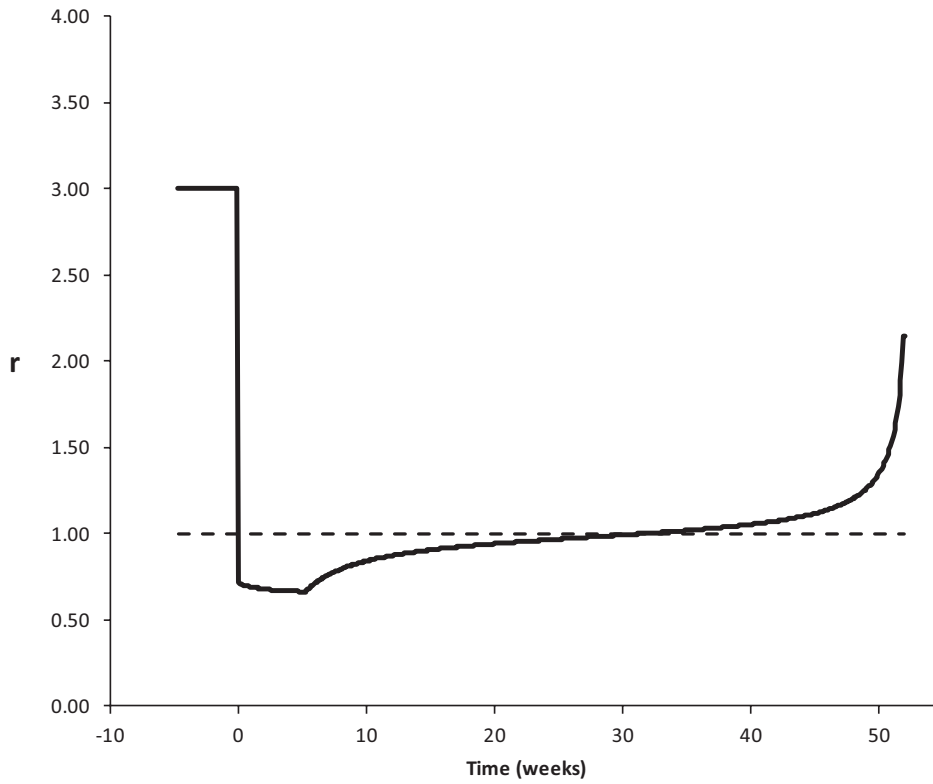


Figure 6: Infection scenarios compared, millions

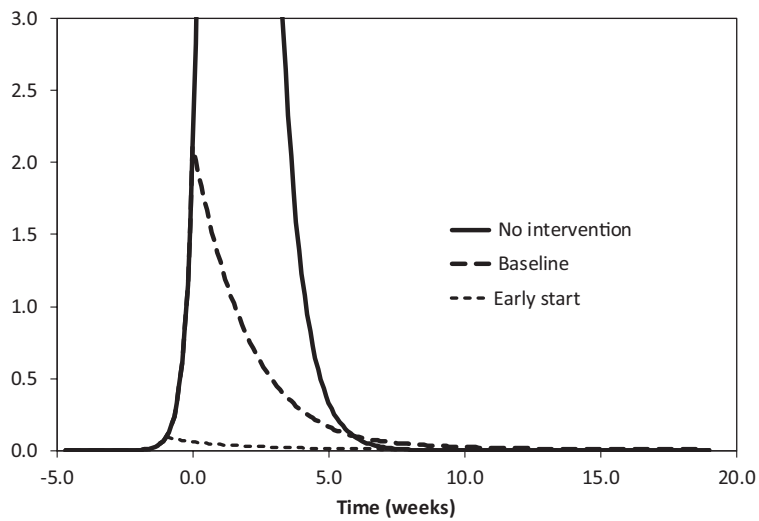
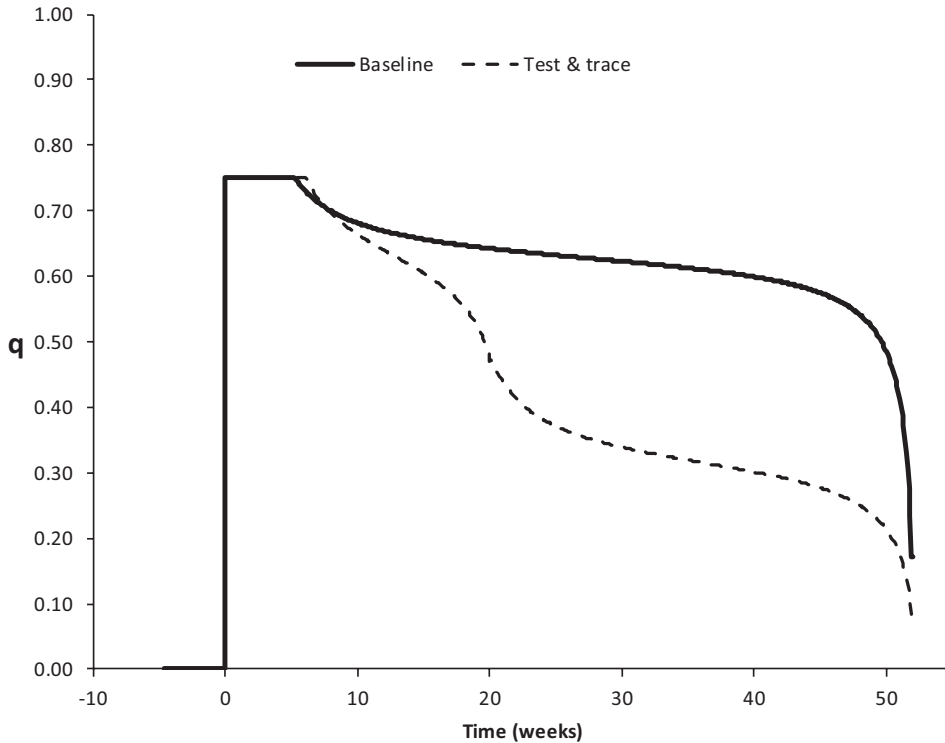


Figure 7: Optimal paths for q : scenarios compared

model it is equivalent to either a reduction in the transmission coefficient β_0 or else an increase in γ , the rate at which infected people cease to be infectious. To explore the implications of the system introduced by the UK government, we assume that it becomes fully operational in week 20. This is later than the government's initial target, but we allow for teething problems. The system has a capacity of 200,000 tests per day. We assume it has negligible cost. The effectiveness of a test and trace system depends on the following factors: (i) the number of tests carried out, (ii) the share of infected individuals in the tested population, (iii) the fraction of infected individuals who are available for testing, and (iv) the number of infected contacts who self-isolate following a positive diagnosis. The roles of these various factors are discussed in the Appendix. The parameters we use for our simulation are somewhat arbitrary, but the results illustrate clearly the impact of test and trace on the optimum path. [Figure 7](#) plots the optimum paths with and without test and trace. The effect of test and trace is to lower the trajectory of the control variable q . The reason for this is as follows. The existence of a test and trace system reduces the impact of present interventions on the future course of infection. Planners therefore have less need to be concerned about the future. They can afford to relax since test and trace will help deal with the outcome. This is true both before and after test and trace comes into operation. The test and trace system in our simulation is not perfect, so some degree of social distancing is still required after this system becomes operational.

The value of life

Any cost–benefit analysis of optimal policy towards COVID-19 requires some assumption about the value of human life (Social Value UK, 2016; Dolan and Jenkins, 2020). This assumption may be explicit or it may be implicit. Governments may reject the whole idea of valuing life in the context of disease control, but to the extent that their actions are consistent, they imply some tacit valuation of life. In other policy areas, such as transport and drug evaluation, it is normal for government agencies to put a value on life. In our simulations a reduction in the value of life implies a shorter lockdown or maybe no lockdown at all (Table 2). This is true even if we impose a ceiling on the permitted scale of infection. Under the Baseline scenario, the value of life is £2m and the optimal lockdown lasts for 5.3 weeks. Holding other parameters constant, it becomes optimal to dispense with the lockdown altogether once the value of life drops below £1.68m. If we impose the condition that peak infection must not exceed what the health service can handle, it is optimal to dispense with lockdown when the value of life is below £1.56m. At the other end of the spectrum, the optimal duration of lockdown becomes rather insensitive to further increases in the value of life. The optimal lockdown is not much different if the value of life is £10m or £20m (Figure 8).

Figure 9 plots the relationship between total deaths and total economic cost. Through its impact on optimal policy, the value of life affects both the economic cost of the disease and the number of people who die from it. Each point on the curve corresponds to a certain value of life, and the variables shown are calculated on the assumption that the government behaves optimally given this value of life.⁵

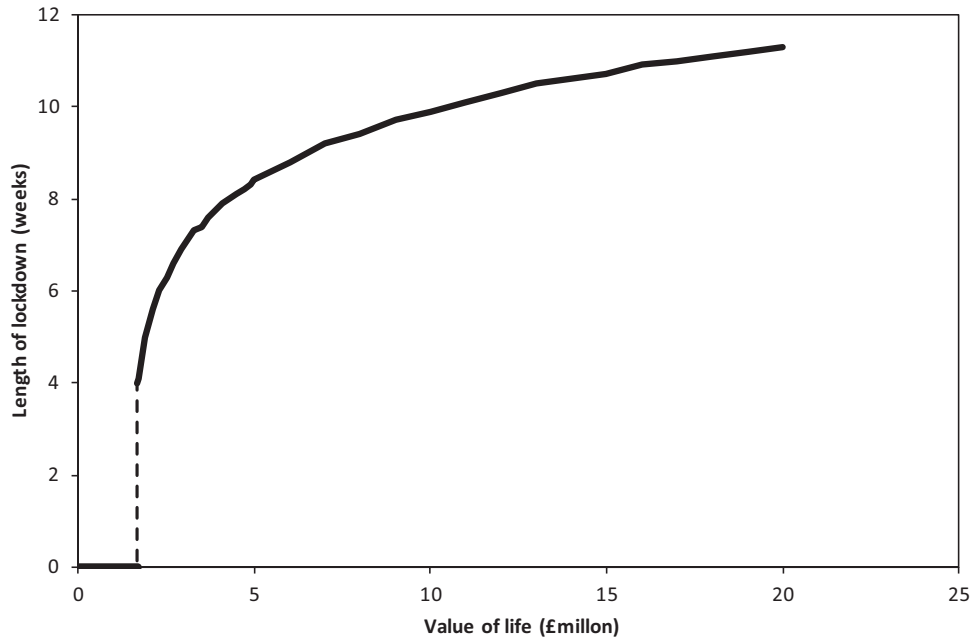
The curve is downward sloping, as expected. If the government assigns a low value to life it will optimally choose a trajectory that involves a very short lockdown or no lockdown at all. This will ensure a low economic cost but will also involve a large number of deaths. Conversely, if the government assigns a high value to life it will opt for a long lockdown, thereby saving lives, although at much greater economic cost.

A striking feature of Figure 9 is the discontinuity indicated by the broken line. This was unexpected, but appears genuine. We checked it using two different programs. This break in the curve marks the transition between two radically different types of policy. To the right of the break, the optimal policy is lockdown with a low death rate. To the left, the optimal policy is no lockdown and a high death rate. This transition occurs abruptly when the value of life is around £1.68m. It is clearly visible in Figure 8.

What light does this discussion throw on the actual policy of the UK government? The period of maximum lockdown lasted approximately 10 weeks. With the baseline cost structure ($\phi = 2, \pi_A = 2$), a lockdown of this length is only optimal when the value of life exceeds £10m. If $\phi = 1, \pi_A = 3$ the figure is £4m. These numbers are much larger than the value of life implied by the official guidelines for drug evaluation (£200,000 to £300,000).⁶ To the extent that the government is behaving optimally, these comparisons imply that it values the lives of potential COVID-19 victims a lot more highly than those of other types of victim.

⁵ Technically speaking, the curve is parameterized by π_D .

⁶ The National Institute for Health Care and Clinical Excellence assumes £20,000–£30,000 per quality-adjusted year of life. Office for National Statistics life tables and statistics on the age, sex, and underlying health condition of COVID-19 fatalities suggest that the average person dying from the disease lost about 10 years of life.

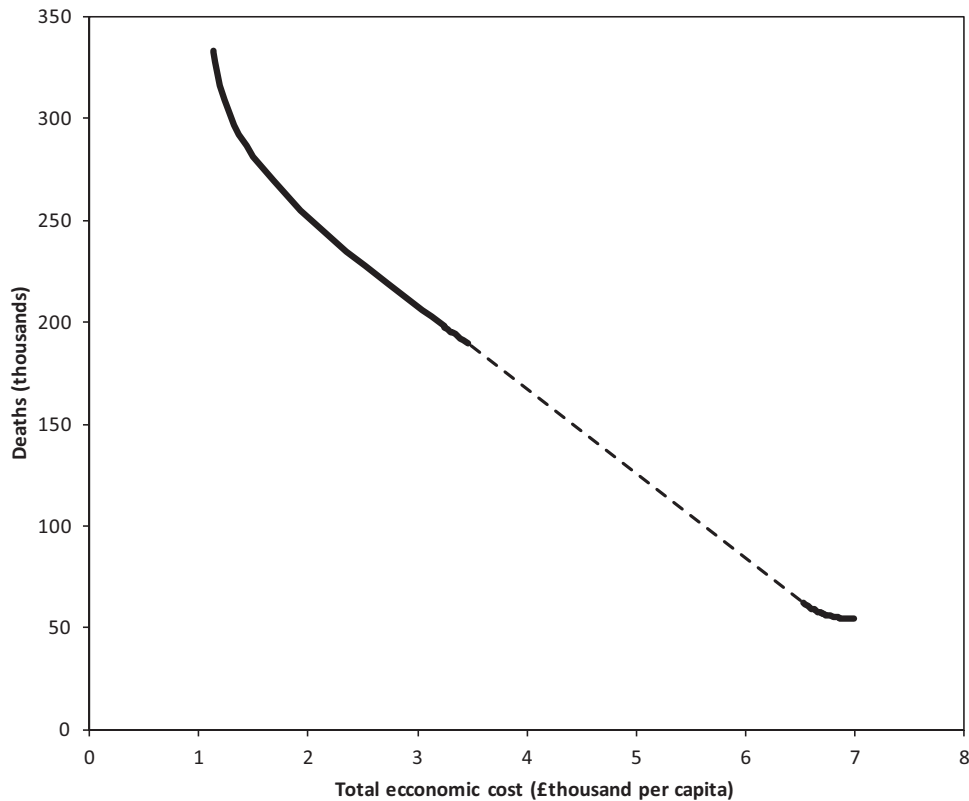
Figure 8: Optimal lockdown: baseline parameters (unconstrained simulation)

V. Concluding remarks

Soon after the implications of lockdown became evident people began to ask the obvious question: ‘Is the cure worse than the disease?’ (Miles *et al.*, 2020). Governments began to seek cost-effective policies that would enable them to exit the lockdown without setting off a renewed surge of infection. Although they are speculative in nature and limited in their methodology, the simulations presented here and their underlying theory may throw some light on government policy.

The original motivation for the lockdown was a fear that the health system would be overwhelmed if the disease were to get out of hand. However, this does not explain why the lockdown continued for such a long time. The explanation may be inertia or excessive caution. Or it may be that the government (and the public) values the lives of potential COVID-19 victims far more highly than those of certain other types of victim. Whatever the explanation, it is clear that government policy towards the COVID-19 disease has not been subject to the same forensic cost–benefit analysis that is applied in other areas of health policy.

In his *Covid Economics* paper, Rowthorn (2020) argued that, if a relatively inexpensive way can be found to maintain an r value close to 1, this is the policy to aim for in the medium term. A lockdown may (or may not) be necessary to halt the explosive spread of the disease, but once this aim has been achieved it would be a costly mistake to stick with expensive social distancing policies that aim to keep r well below 1. This conclusion is reinforced by our example of test and trace. If there is an effective test and trace system in the offing, it may even be optimal to let r exceed 1 during the weeks before this system becomes operational. This will cause infection to increase somewhat, but the potential explosion will be prevented when test and trace comes on stream. The same is true during the run-up to mass vaccination.

Figure 9: Deaths and economic cost: baseline parameters (unconstrained simulation)

One issue that this paper has not dealt with is that of ignorance. We have assumed that there is a menu of known policies, with known effects, from which the government can choose at will. In fact governments and their advisors may have very limited knowledge about the disease and potential policies, and they may be reluctant to experiment because they are concerned about the risk of a mistake.

Despite these caveats, we believe that the approach adopted in this paper provides a useful framework for thinking about policy choices and their timing.

Appendix: Test and trace

Throughout this appendix the symbol I refers to infected individuals who are not isolated and can therefore infect the susceptible population. Isolated individuals who are infectious are classified as removed.

Suppose that a fraction a of the infected population I is currently available for testing. The rest are either asymptomatic or unwilling to undergo testing. For those available for testing; the probability of not being tested positive in a period of length s is equal to e^{-ps} where p is constant. The probability that an infected individual will cease to be

infectious in the small time interval Δs is $[-\frac{d}{ds}(e^{-\gamma s})] \Delta s = \gamma e^{-\gamma s} \Delta s$. Thus, the probability of recovering or dying without testing positive is:

$$\int_0^{\infty} \gamma e^{-(p+\gamma)s} ds = \frac{\gamma}{p+\gamma} \quad (\text{A1})$$

The probability of being tested positive at some time is therefore

$$1 - \frac{\gamma}{p+\gamma} = \frac{p}{p+\gamma} \quad (\text{A2})$$

The average length of time that an individual remains infected is $\frac{1}{\gamma}$. Thus, the probability that he or she will test positive during a small time interval of length Δt , is equal to:

$$\frac{p}{p+\gamma} \left(\frac{\Delta t}{\frac{1}{\gamma}} \right) = \left(\frac{\gamma p}{p+\gamma} \right) \Delta t \quad (\text{A3})$$

The number of infected individuals who are available for testing is aI . The number of such individuals who test positive in the time interval Δt is

$$aI \left(\frac{\gamma p}{p+\gamma} \right) \Delta t$$

The rate per unit of time at which they are tested positive is therefore

$$\left(\frac{\gamma p}{p+\gamma} \right) aI$$

Suppose there is no constraint on testing. Then $p = \infty$ and the rate of testing infected persons is:

$$A = \gamma aI \quad (\text{A4})$$

In the constrained case assume that M is the maximum number of tests per week. Assume also that a constant fraction b of these tests is directed at infected persons. Then access to testing will be capacity constrained if $bM < \gamma aI$. In this case

$$A = \left(\frac{\gamma p}{p+\gamma} \right) aI = bM \quad (\text{A5})$$

Thus,

$$p = \frac{\gamma aI}{\gamma aI - bM} \quad (\text{A6})$$

Assume that for each person who tests positive the number of infected persons who self-isolate (including the tested person) is c . Then infected persons are isolated at the rate cA . They are classified as removed.

Suppose the test and trace system comes on stream at time T^* . Define the following function:

$$\begin{aligned} Q(t, I) &= 0 \text{ for } t < T^* \\ &= c \times \min(bM, \gamma aI) \text{ for } t \geq T^* \end{aligned} \quad (\text{A7})$$

The equations of motion become:

$$\begin{aligned} \frac{dI}{dt} &= (1 - q)\beta_0 SI - \gamma I - Q(t, I) \\ \frac{dR}{dt} &= \gamma I + Q(t, I) \end{aligned} \quad (\text{A8})$$

Our simulation assumes a daily capacity of 200,000 for test and trace. This amounts to 1,400,000 per week, which is equal to a fraction 0.021 of the population. Thus, $M = 0.021$. It is also assumed that the test and trace system becomes fully operational in week 20 and that $a = 0.5, b = 0.5, c = 1.6$.

References

- Acemoglu, D., Chernozhukov, V., Werning, I., and Whinston, M. D. (2020), ‘Optimal Targeted Lockdowns in a Multi-Group SIR Model’, NBER Working Paper No. 27102.
- Chen, F. (2012), ‘A Mathematical Analysis of Public Avoidance Behavior During Epidemics Using Game Theory’, *Journal of Theoretical Biology*, **302**, 18–28.
- Jiang, M., Rabidoux, S., and Robinson, S. (2011), ‘Public Avoidance and Epidemics: Insights from an Economic Model’, *Journal of Theoretical Biology*, **278**, 107–19.
- Cleavelly, M., Susskind, D., Vines, D., Vines, L., and Wills, S. (2020), ‘A Workable Strategy for Covid-19 Testing: Stratified Periodic Testing rather than Universal Random Testing’, *Oxford Review of Economic Policy*, **36**(Supplement).
- Dolan, P., and Jenkins, P. (2020), ‘Estimating the Monetary Value of the Deaths Prevented from the UK Covid-19 Lockdown when it was Decided Upon—and the Value of “Flattening the Curve”’, LSE, 18 April, <https://pauldolan.co.uk/wp-content/uploads/2020/04/Estimating-the-monetary-value-of-the-deaths-prevented-from-the-UK-Covid-19-lockdown.pdf>
- Fenichel, E. P. (2013), ‘Economic Considerations for Social Distancing and Behavioral Based Policies During an Epidemic’, *Journal of Health Economics*, **32**(2), 440–51.
- Gersovitz, M. (2010), ‘Disinhibition and Immunization in a Model of Susceptible-Infected-Susceptible (SIS) Diseases’, mimeo.
- Giordano, G., Blanchini, F., Bruno, R., et al. (2020), ‘Modelling the COVID-19 Epidemic and Implementation of Population-wide Interventions in Italy’, *Nat Med*, <https://doi.org/10.1038/s41591-020-0883-7>
- King’s College (2020), <https://covid.joinzoe.com/>, King’s College, London
- Miles, D., Stedman, S., and Heald, A. (2020), ‘Living with COVID-19: Balancing Costs Against Benefits in the Face of the Virus’, mimeo.
- OBR (2020), ‘Coronavirus: Reference Scenario’, Office for Budget Responsibility, <https://obr.uk/coronavirus-reference-scenario/>
- Ormerod, P., Rowthorn, R., and Nyman, R. (2020), ‘Why Epidemiological Models Exaggerate the Risks of a Second Wave of Covid-19’, mimeo.
- Reluga, T. C. (2010), ‘Game Theory of Social Distancing in Response to an Epidemic’, *PLoS Computational Biology*, **6**(5), e1000793.
- Rowthorn, R. (2020), ‘A Cost–Benefit Analysis of the Covid-19 Disease’, *Covid Economics*, **9**, 24 April, London, CEPR.
- Toxvaerd, F. (2015), ‘The Optimal Control of Infectious Diseases via Prevention and Treatment’, mimeo.

-
- Sethi, S. P. (1978), 'Optimal Quarantine Programmes for Controlling an Epidemic Spread', *Journal of the Operational Research Society*, **29**(3), 265–8.
- Social Value UK (2016), <http://www.socialvalueuk.org/resource/valuation-of-a-life/>
- Toxvaerd, F. (2019), 'Rational Disinhibition and Externalities in Prevention', *International Economic Review*, **60**(4), 1737–55.
- (2020), 'Equilibrium Social Distancing', Cambridge–INET Working Paper Series No. 2020/08.
- Rowthorn, R. (2020), 'On the Management of Population Immunity', mimeo.