UNIVERSITY *of York*

Discussion Papers in Economics

No. 20/05

Epidemic policy under uncertainty and information

Martin Forster, Emanuela Randon

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

Epidemic policy under uncertainty and information

Martin Forster[*] Emanuela Randon[†]

May 20, 2020

Abstract

We present a model of infectious disease control which incorporates uncertainty and information. A policy-maker possesses beliefs about the value of a key parameter – we choose the level of herd immunity in the population – and seeks the welfare-maximising level of intervention, accounting for both the public health benefit and economic cost. An approximation to the optimal rule shows that it accounts for interactions between beliefs, the policy-maker's attitude to risk, the production technology and costs, and the weights in the welfare function. We consider the role of information, in the form of expert opinion and scientific advice, in influencing the policy-maker's beliefs and the optimal policy. We assess the framework's potential for advancing the economic modelling of epidemic control.

JEL codes: I18, H12, H51
Keywords: Economics; Epidemics; COVID-19.

1 Introduction

Policy-makers reacted in a variety of ways to the initial stages of the COVID-19 outbreak in their countries. Some introduced stringent control measures, others did not (Frey et al., 2020). Differences in first-responses can be attributed, at least in part, to the perceived trade-off that exists between the public health benefits that intervention brings (reduced transmission rates, lower mortality, reduced strain on health care systems, and so on) and the economic costs that it imposes (reduced economic activity, the closure of firms, unemployment, and so on). See, for example, contributions by Jenny (2020), Thunstrom et al. (2020) and Jones et al. (2020).

In the United Kingdom, Ferguson et al. (2020) modelled two principal intervention strategies for controlling viral spread in the absence of a vaccine: 'suppression', in which measures are taken to stop human-to-human transmission and reduce the reproduction number to below 1;

[*]Department of Economics and Related Studies, University of York, York, U.K. mf8@york.ac.uk. This is preliminary and incomplete work in progress that has not been peer reviewed and is issued for discussion. Nothing in this paper constitutes a policy recommendation. Mistakes are ours. We thank Simon Smith and Luigi Franzoni for comments on previous drafts.

[†]Department of Economics, University of Bologna, Bologna, Italy. emanuela.randon@unibo.it

'mitigation', in which targeted measures are taken to limit the health impact, building population immunity over time, and eventually reducing the number of cases and the rate of transmission. Ferguson et al. (2020) recommended suppression as the only viable strategy, but noted that there was no easy policy decision to be made. They stressed that 'very large uncertainties' existed around transmission, the effectiveness of different interventions, and the degree to which populations will adopt behaviours which reduce risk. Uncertainty, in various guises, has been recognised by economists too. For example, Jones et al. (2020, page 39), in an economic model of contagion dynamics, note there is 'much uncertainty' about the parameters governing progression of the disease; Rowthorn (2020, page 101), in a cost-benefit analysis of policy-making for COVID-19, notes that 'considerable guesswork' is involved in fixing parameter values; Atkeson (2020) discusses the difficulties of estimating the fatality rate in a Susceptible-Infected-Removed (SIR) model of COVID-19; Avery et al. (2020, page 22) note the 'great uncertainty' associated with using predictive models to make forecasts in real time during the early days of the pandemic.

Notwithstanding such uncertainties, the policy-maker (PM) faced with having to determine its optimal response to the outbreak must process information about the pandemic – including expert opinion, scientific papers, advisory board advice and so on – and act upon it as best it can. But such information can, at times, be contradictory (Jenny, 2020), and concerns have been raised about the increasing use of preprint archives to release research which has yet to be peer reviewed (Lakëns, 2020). Jenny (2020) concludes his review of governance and resilience during the COVID-19 crisis by calling for better integration of risk factors into economic analysis and policy-making.

In this paper, we consider whether it is possible to make a degree of headway in addressing these concerns by modelling uncertainties as 'measurable risks' (Knight, 1921) within a model of optimal policy-making for infectious disease control. We consider a policy-maker who must decide on the strength of an intervention to tackle an infectious disease outbreak in the presence of uncertainty over the value of a key parameter – we choose the level of herd immunity in the population, but the idea extends to other parameters – and the arrival of information, in the form of scientific evidence and advice from experts. We account for both the public health benefits of intervention and the economic costs imposed on citizens and firms. Ours is not the first piece of economic work to study uncertainty and its effects. For example, Kozlowski et al. (2020) model the 'scarring' of beliefs about the true distribution of shocks to the economy following the COVID-19 pandemic. And Bu et al. (2020) compare the risk-taking behaviour of a sample of subjects quarantined in Wuhan, China, and exposed to the COVID-19 outbreak, with subjects in regions with lower exposure and less severe quarantine restrictions. However, we believe that ours is the first to model explicitly a PM's beliefs about a key parameter value within an optimisation framework, and to investigate the way these beliefs impact the optimal intervention.

Our specific focus is on how the expected value and variance of a random variable describing the PM's beliefs about the level of herd immunity interact with the PM's appetite for risk-taking to determine its optimal level of intervention. In contrast to the dynamic economic models which have developed the classical SIR model of disease transmission of Kermack and McKendrick (1927),[1] we take a single period approach to the problem. This is deliberate, and reflects our

[1] Recent contributions include Farboodi et al. (2020); Rowthorn (2020); Gonzalez-Eiras and Niepelt (2020);

interest in PM behaviour under uncertainty and the role played by scientific evidence and expert opinion in influencing it. However, it should be possible to extend the ideas to a multi-period setting using, e.g., the 'continuation value' approach of Moser and Yared (2020) or Bayesian models of learning and sequential decision-making. We discuss such matters in section 3.

We derive an approximation to the rule defining the PM's optimal level of intervention, showing that it involves interactions between the expected value and variance of beliefs, the PM's appetite for risk-taking (notably, the PM's degree of absolute risk aversion and prudence), the shape of the production technology, interactions between the marginal product of the intervention and the level of immunity, and marginal costs, all appropriately weighted. We consider how information impacts the optimal rule and illustrate how it can change beliefs and move the optimal level of intervention up and down, causing it to diverge from, or converge towards, the optimal deterministic intervention (the level based on the expected value of beliefs alone, ignoring the variance).

2 Model

A policy-maker seeks the optimal level of intervention to limit the damage done by an infectious disease outbreak in their country, or a city or region within it. The PM may be confronting the start of an outbreak, may be mid-outbreak or may have experienced one or multiple waves of infection and be anticipating a subsequent wave. Welfare is determined by two factors: 1. that which results from limiting the damage from the spread of the disease among citizens (the 'public health' welfare) and 2. that which results from limiting the economic shock caused by the intervention, as it impacts citizens and firms (the 'economic' welfare). This is a stylised representation, and reflects the approach taken in a number of recent contributions (examples include Rowthorn (2020) and Gonzalez-Eiras and Niepelt (2020)).

The PM faces a situation of uncertainty over a range of dimensions. They receive information, in the form of scientific evidence and advice from groups of experts and other sources, some of which may be contradictory (Jenny, 2020, section 1). The PM may assess the quality of this information and use it when determining its response, working within the social norms and legislature of its jurisdiction, which can vary across nations (Jenny, 2020, section 2). Uncertainties facing the PM include, but are not limited to:[2]

1. the level of herd immunity in the population and/or the degree to which infection of an individual confers immunity. Both scenarios have characterised the COVID-19 pandemic (see, e.g., the policy debate in the United Kingdom over herd immunity, to which Lourenço et al. (2020) contributed, and World Health Organisation (2020));

2. the effectiveness of the measures which may comprise the intervention. For example, the PM may be unsure about the degree to which the population will comply with social-

Jones et al. (2020).

[2]Other uncertainties include the case-fatality rate, whether patients are infectious before the onset of symptoms, the proportion of asymptomatic cases in the population and the duration of the infectious period (Anderson et al., 2020).

distancing policies (recent discussion and empirical contributions on this matter include Avery et al. (2020); Painter and Qiu (2020) and Engle et al. (2020));

3. the economic damage to citizens and firms in the economy as a result of the intervention. Many studies now present estimates of the cost of the outbreak, some of which also assess the degree of uncertainty surrounding them (e.g. Ludvigson et al. (2020) and Fadinger and Schymik (2020))

We focus on 1, noting that it is straightforward to extend the methods to 2 and 3.

2.1 Uncertainty and attitude to risk

The PM's choice is made in the presence of uncertainty over the proportion of the population that is herd immune, $p_A \in [0, 1]$. One way to model this would be to treat it as a fixed parameter in the optimisation problem, obtain a rule for an optimal level of intervention, and conduct a comparative static analysis which varies that parameter using analytical methods or simulation. We take a different approach, and incorporate the PM's uncertainty about p_A into the optimisation problem. Assume that the PM's beliefs about p_A may be modelled as a random variable with expected value μ_{p_A}, variance $\sigma^2_{p_A} \geq 0$, and density function f_{p_A}. We may consider higher central moments, but limit analysis to the mean and variance for the purposes of exposition. This approach fits naturally with a Bayesian take on modelling beliefs, one which has been successfully applied in other health-related fields (e.g. Spiegelhalter et al. (1994)).

To borrow the language of the classical SIR models, the proportion of the population that has not acquired immunity at the time the PM is choosing its intervention, $1 - p_A$, is classed as 'susceptible'. The higher is $\sigma^2_{p_A}$, the greater the PM's uncertainty about the true level of herd immunity. The PM expects the proportion of the population that is susceptible to be $1 - \mu_{p_A}$. Setting $\sigma^2_{p_A} = 0$ reduces the problem to the case where the PM believes herd immunity is known to equal μ_{p_A}. We refer this scenario as the 'deterministic' case.

We adopt an expected utility framework. Assume that, if the PM does not intervene to control the spread of the virus, all susceptible individuals eventually become infected. The most optimistic scenario for the PM is that 100% of the population possesses immunity, in which case it need take no action; both public health welfare and economic welfare, as defined in the introduction to this section, are maximised. The most pessimistic scenario is that there is zero immunity. If the PM acts in a risk-neutral manner, public health (PH) welfare, U, is increasing and linear in p_A. If it acts in a risk-averse manner, welfare decreases at an increasing rate as the proportion of susceptible individuals increases (the proportion with immunity decreases); U is then an increasing, strictly concave function of p_A.

We assume that the PM's beliefs about p_A are such that $0 < \mu_{p_A} < 1$. Using a Taylor series expansion to the second order around $U(\mu_{p_A})$, expected PH welfare is approximated as follows:

$$\mathbb{E}[U(p_A)] \approx U(\mu_{p_A}) + \frac{\sigma^2_{p_A}}{2} \frac{\partial^2 U(p_A)}{\partial p_A^2}\Big|_{\mu_{p_A}},$$

where we assume that we are considering a small perceived risk (Pratt, 1964). For a PM acting in a risk-averse manner, $\mathbb{E}[U(p_A)] < U(\mu_{p_A})$ by Jensen's inequality. If, instead, the PM acts in a risk-neutral manner, $\mathbb{E}[U(p_A)] = U(\mu_{p_A})$.

Define $\hat{p}^A$ as the 'certainty equivalent' such that, when the PM acts in a risk-averse manner, $U(\hat{p}^A) = \mathbb{E}[U(p_A)]$ for a value $\hat{p}_A < \mu_{p_A}$. If the PM's intervention can guarantee that a defined proportion of the population, $x > \hat{p}^A$, is not exposed to the virus which causes the disease, the PM guarantees that PH welfare under the intervention is greater than that which is expected under no intervention ($U(x) \geq \mathbb{E}[U(p_A)]$ for $x > \hat{p}^A$). One such intervention would be to create a 'cordon sanitaire' around an area where infection has occurred, preventing travel to other areas. This is what happened in Wuhan, China, on 23rd January 2020 (Pan et al., 2020). In taking such action, the PM may bear a level of infection over and above that which it expects if it does not intervene. The cost of bearing the risk is the risk premium, $\rho = \mu_{p_A} - \hat{p}^A$. Continuing the analysis of small perceived risks, approximating $U(\hat{p}^A)$ by a first order Taylor series expansion, $U(\hat{p}^A) \approx U(\mu_{p_A}) + (\partial U(\mu_{p_A})/\partial p_A)(-\rho)$:

$$\rho = \frac{\sigma_{p_A}^2}{2}\mathrm{C}_a > 0, \tag{1}$$

where $\mathrm{C}_a = -(\partial^2 U/\partial p_A^2)(\partial U/\partial p_A)^{-1}$ is the coefficient of absolute risk aversion (Arrow, 1963; Pratt, 1964). When the PM acts in a risk-neutral manner, $\rho = 0$.

If $\partial \mathrm{C}_a/\partial p_A < 0$ for all $p_A \in [0,1]$, the risk premium falls as the level of herd immunity increases, and the lower is the cost of being uncertain about p_A. A necessary and sufficient condition for absolute risk aversion to be strictly decreasing in p_A is for the PM's degree of absolute prudence,

$$\mathrm{P}_a = -(\partial^3 U/\partial p_A^3)(\partial^2 U/\partial p_A^2)^{-1}, \tag{2}$$

to be greater than C_a for all $p_A \in [0,1]$ (Eeckhoudt et al., 2005, Prop 1.6).

2.2 Optimal intervention

We consider how the PM may intervene to protect susceptibles from becoming infected, but at a cost. We do so by splitting the proportion of susceptibles into those who remain uninfected by the end of the outbreak, owing to the intervention, and those who become infected. We introduce into the welfare function additional terms for the damage done to citizens and firms.

Define α as the PM's intervention. α may include social distancing policies, the closure of schools and non-essential services, and so on.[3] When $\alpha = 0$, the PM does not intervene. The more stringent the intervention, the higher is α. α increases p_C, the proportion of the population who are susceptible and who are uninfected post-intervention ($\partial p_C/\partial\alpha > 0$). We refer to $p_C(\alpha; p_A)$ as the 'production technology' and assume that $p_C(0; p_A) = 0$. Assuming

[3] Ferguson et al. (2020, Table 2) model the following strategies, actioned either alone or in combination: case isolation in the home; voluntary home quarantine; social distancing for over-70s: social distancing; closure of schools and universities. Chen and Qiu (2020, pages 47–48) provide details of country-level interventions. Hale et al. (2020) use a policy stringency index.

		Immune?		
		Yes	No	Total
Infected?	Yes	p_A	$p_B(\alpha; p_A)$	$p_A + p_B(\alpha; p_A)$
	No	0	$p_C(\alpha; p_A)$	$p_C(\alpha; p_A)$
Total		p_A	$p_B(\alpha; p_A) + p_C(\alpha; p_A)$	1

Table 1: Contingency table of immune and infected post-intervention.

that immunity is only acquired by having been infected, the classifications 'immune', 'susceptible and infected' and 'susceptible and not infected' are mutually exclusive and collectively exhaustive. Hence the proportion of susceptible and infected by the end of the outbreak is $p_B(\alpha; p_A) = 1 - p_A - p_C(\alpha; p_A)$ (Table 1).

Welfare, W, is a function of both the total proportion of the population that is not infected post-intervention, $T = p_A + p_C(\alpha; p_A)$, and the degree of economic damage imposed on citizens and firms as a result of the intervention. U remains the utility from the total proportion of the population that is not infected post-intervention, so is a function of T. Define V_1 and V_2 as the welfare costs imposed on citizens and firms, respectively. We assume that these are increasing, convex, functions of the amount of the nation's wealth, Y, that is spent on the intervention and that they are independent of p_A. We assume that, when there is no intervention, no costs are incurred, but that $\partial Y / \partial \alpha > 0$.

The PM chooses α to maximise overall expected welfare, defined as the weighted average of public health and economic welfare:

$$W = (1 - \beta_1 - \beta_2) \int_0^1 U(p_A + p_C(\alpha; p_A)) \, f_{p_A}(p_A) \, \mathrm{d}p_A - \beta_1 \, V_1(Y(\alpha)) - \beta_2 \, V_2(Y(\alpha)). \quad (3)$$

Define $\phi \equiv 1 - \beta_1 - \beta_2$ and α^* as the optimal level of intervention, which we assume to hold at an interior solution for the problem (we assume that $\partial^2 W / \partial \alpha^2 < 0$ at α^*). Then α^* satisfies:

$$\phi \int_0^1 \frac{\partial U(p_A + p_C(\alpha^*; p_A))}{\partial T} \frac{\partial p_C(\alpha^*; p_A)}{\partial \alpha} f_{p_A}(p_A) \mathrm{d}p_A - \frac{\partial Y}{\partial \alpha} \sum_{i=1}^{2} \beta_i \frac{\partial V_i(Y)}{\partial Y} = 0, \quad (4)$$

where the first term on the LHS of Eq. (4) is the weighted expected marginal benefit of the intervention and the second is the weighted marginal cost. We study the qualitative nature of Eq. (4) by approximating the expected value using a Taylor series expansion to the second order, and obtain the following results, in which α^* is an approximation to the welfare-maximising level of the intervention. Proofs are presented in Appendix A.

Proposition 1.A. *The welfare-maximising PM chooses the optimal level of intervention, α^*, to* *satisfy:*

$$
\begin{aligned}
&\frac{\partial Y}{\partial \alpha}\left(\frac{\beta_1}{\phi}\frac{\partial V_1}{\partial Y}+\frac{\beta_2}{\phi}\frac{\partial V_2}{\partial Y}\right)\Big|_{\alpha^*} \approx \frac{\partial U}{\partial T}\frac{\partial p_C}{\partial \alpha}\left\{1+\frac{\sigma^2_{p_A}}{2}\Omega_1\Omega_2+\rho\,\Upsilon\right\}\Big|_{\mu_{p_A},\alpha^*},\\
&\quad \textit{where } \Upsilon = P_a\left(\frac{\partial T}{\partial p_A}\right)^2 - 2\,\Omega_1\frac{\partial T}{\partial p_A}-\frac{\partial^2 p_C}{\partial p_A^2},\\
&\quad 1+\frac{\sigma^2_{p_A}}{2}\Omega_1\Omega_2+\rho\,\Upsilon\;>\;0,
\end{aligned}
\tag{5}
$$

$\Omega_1 = (\partial^2 p_C/\partial\alpha\partial p_A)(\partial p_C/\partial\alpha)^{-1}$, $\Omega_2 = (\partial^3 p_C/\partial\alpha\partial p_A^2)(\partial^2 p_C/\partial p_A\partial\alpha)^{-1}$, *$\rho$ is the risk premium defined in Eq. (1), with T replacing p_A, and P_a is the degree of absolute prudence, from Eq. (2).*

Proposition 1.B. *Other things equal, the greater the weight placed on public health welfare relative to economic welfare, the more aggressive is the intervention.*

Proposition 1.C. *Other things equal, when the PM acts in a risk-neutral manner, the level of intervention is more/the same/less aggressive when it acts in a deterministic manner ($\sigma^2_{p_A} > 0$) compared with when $\sigma^2_{p_A} = 0$, according to whether $\Omega_1\Omega_2 \gtreqless 0$ at μ_{p_A}.*

Proposition 1.D. *Other things equal, when $\sigma^2_{p_A} > 0$, when the PM acts in a risk-averse manner, the intervention is more/the same/less aggressive than when it acts in a risk-neutral manner according to:*

$$
P_r \gtreqless \varepsilon^{-1}_{T,p_A}(2\,\varepsilon_{\partial p_C/\partial\alpha,p_A}+\varepsilon_{\partial T/\partial p_A,p_A}), \tag{6}
$$

at μ_{p_A}, where $P_r = P_a\,T$ is the degree of relative prudence and $\varepsilon_{a,b}$ the elasticity of a with respect to b.

Eq. (5) of Proposition 1.A shows how the variance used to model the PM's beliefs about herd immunity adjusts the standard, deterministic, rule which equates the marginal cost of the intervention (LHS) with the marginal benefit, $(\partial U/\partial T)(\partial p_C/\partial\alpha)$. Failure to account for the variance imposes a welfare loss. In the discussion that follows, we speak about the PM operating in three contexts: 1. a deterministic manner, ignoring the variance of beliefs and basing the optimal intervention decision on the expected value alone; 2. a risk-neutral manner, accounting for both the expected value and the variance of beliefs; 3. a risk-averse manner, accounting for both the expected value and variance of beliefs.

Two effects on the RHS of Eq. (5) drive the results covered by the propositions:

1. the 'pure technology' effect, $(\sigma^2_{p_A}/2)\Omega_1\Omega_2$. This measures the adjustment to the marginal product of the intervention, $\partial p_C/\partial\alpha$, which results from the interaction between the variance of beliefs and the production technology. The term is independent of the PM's attitude to risk (Proposition 1.C). Since the terms $\partial^2 p_C/\partial\alpha\partial p_A$ in the product $\Omega_1\Omega_2$ cancel, the effect may be written as $(\sigma^2_{p_A}/2)(\partial^3 p_C/\partial\alpha\partial p_A^2)(\partial p_C/\partial\alpha)^{-1}$. Under the assumption that

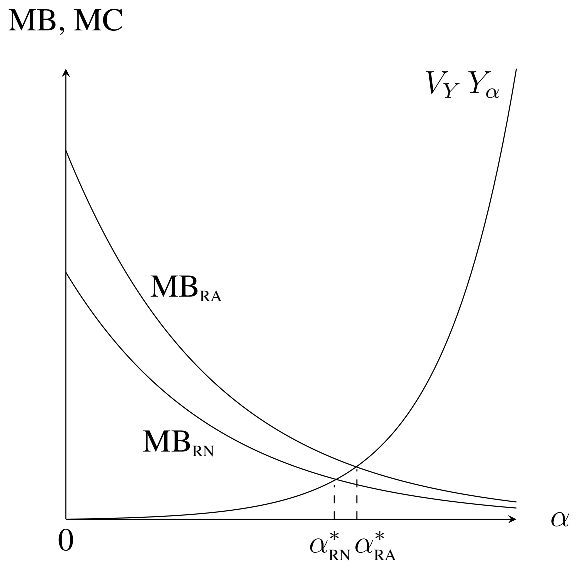

Figure 1: Marginal benefit and costs for the PM when they act in risk-averse and a risk-neutral manner, together with optimal levels of intervention, α^*, for the case of $\text{P}_r > \varepsilon_{T,p_A}^{-1}(2\,\varepsilon_{\partial p_C/\partial\alpha,p_A} + \varepsilon_{\partial T/\partial p_A,p_A})$.

$\partial p_C/\partial\alpha > 0$, the sign of the effect is equal to the sign of the second derivative of marginal product with respect to herd immunity;[4]

2. the 'welfare/technology' effect. This measures the adjustment to marginal benefit which results from the interaction between the welfare function and the production technology, expressed as a rescaling of the risk premium, ρ. The term applies when the PM acts in a risk-averse manner ($\rho > 0$), but not in a risk-neutral manner. Proposition 1.D shows how the adjustment is positive/zero/negative according to the relationship between the PM's relative prudence (LHS) and elasticities measuring the response of p_C to changes in the level of herd immunity. Figure 1 provides an illustration of a positive adjustment when the PM acts in a risk-averse manner, compared with when they act in a risk-neutral manner.

We state a corollary which deals with the case of a specific functional form for p_C, which we use in the numerical simulation of section 2.3.

Corollary 1. *When $p_C = (1 - p_A)\psi(\alpha)$, and ψ is a continuous, twice-differentiable, function of α such that $\psi(0) = 0$, the PM will pursue a more/the same/a less aggressive intervention when they are risk-averse compared with when they are risk-neutral according to whether, at μ_{p_A}:*

$$P_r \gtreqless 2\,\varepsilon_{T,p_A}^{-1}\varepsilon_{\partial p_C/\partial\alpha,p_A}.$$

[4]If marginal product is linear in p_A, so that the absolute value of gain/loss to marginal product is equal for a given small deviation of p_A from its expected value, the term is equal to zero and the optimal intervention of the PM who acts in a deterministic manner (ignoring the variance) coincides with the optimal intervention when they act in a risk-neutral manner and account for the variance. Otherwise, under strict concavity, $\partial^3 p_C/\partial\alpha\partial p_A^2 < 0$, the term is strictly negative (and decreases expected marginal benefit relative to marginal cost; optimal intervention decreases); under strict convexity, it is strictly positive (and increases expected marginal benefit relative to marginal cost; optimal intervention increases).

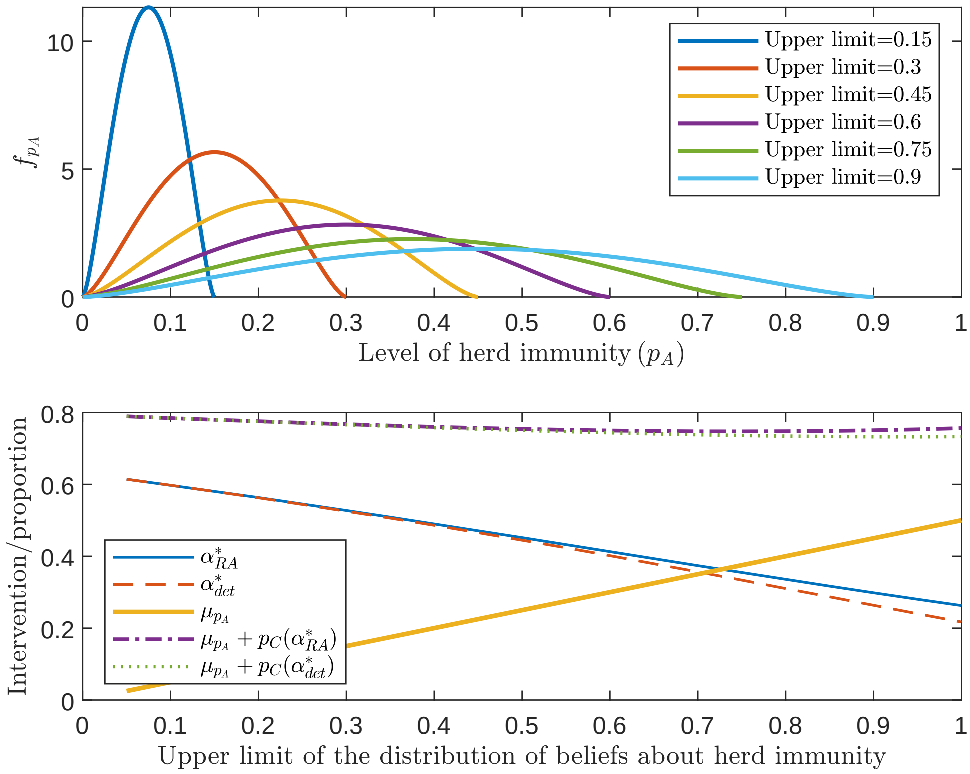

Figure 2: Upper: changing beliefs about herd immunity. Lower: expected level of immunity, μ_{p_A}; optimal intervention when the PM acts in a deterministic manner ($\sigma^2_{p_A} = 0$) and when the PM acts in a risk-averse manner ($\sigma^2_{p_A} > 0$); proportions uninfected, $p_A + p_C(\alpha^*)$, under each policy.

Finally, we note that it is not necessarily the case that, when the PM acts in a risk-averse manner, it will pursue a more aggressive intervention when it accounts for the variance of beliefs, compared with when it does not (the deterministic case). The result is driven by whether or not the pure technology effect and the welfare/technology effect are reinforcing or opposing.

Proposition 2. *When the PM acts in a risk-averse manner, it pursues a more/the same/a less aggressive intervention when $\sigma^2_{p_A} > 0$ compared with when it acts in a deterministic manner ($\sigma^2_{p_A} = 0$) according to whether*

$$P_r - \Theta \gtreqless \varepsilon^{-1}_{T,p_A}(2\,\varepsilon_{\partial p_C/\partial\alpha,p_A} + \varepsilon_{\partial T/\partial p_A,p_A}) \tag{7}$$

at μ_{p_A}, where $\Theta = \varepsilon^{-2}_{T,p_A} C_r^{-1} \varepsilon_{\partial p_C/\partial\alpha,p_A} \varepsilon_{\partial^2 p_C/\partial\alpha\partial p_A,p_A}$ and $C_r = C_a T$ is the coefficient of relative risk aversion.

2.3 The role of information

Sources of information, such as expert opinion and/or scientific evidence, can influence the PM's beliefs and therefore one or both of μ_{p_A} and $\sigma^2_{p_A}$. These effects, in turn, can impact the optimal level of intervention. Define the information set available to the PM as $\mathcal{S}$, so that $\mu_{p_A}(\mathcal{S})$ and $\sigma^2_{p_A}(\mathcal{S})$. Figure 2 presents an illustrative example of what happens to beliefs (upper figure), the optimal level of intervention, and expected cases not infected post-intervention, $T = p_A + p_C$ (lower figure), as $\mathcal{S}$ changes, for a PM acting in a risk-averse manner. We use a four parameter beta distribution to represent beliefs about herd immunity, making the density function symmetric and setting a progressively higher maximum value for the level of immunity, holding the minimum constant at zero. The lowest maximum value is 0.05 and the highest is 1, and selected density functions representing the PM's beliefs about p_A are drawn for the cases of an upper limit equal to 0.15, 0.30, ..., 0.90. Increasing the upper limit of the distribution increases both the expected value and variance. The production technology is assumed to be $p_C = (1 - p_A)\alpha^{1/2}$ (Corollary 1). The lower part of Figure 2 shows the optimal level of intervention, contrasting the level when the PM acts in a deterministic manner, ignoring the variance of beliefs (red, dashed, line) with the level when the variance is not ignored and Eq. (5) is used (blue, continuous, line).[5]

Reading the upper part of Figure 2 from left to right, the leftmost density function shows the PM believes the proportion of the population that is immune to be low, and not higher than 0.15. The expected level of immunity is equal to 0.075 and the variance is 3/3200. The lower part of Figure 2 shows that, in this case, the optimal level of α is approximately equal to 0.56 when the PM acts in a deterministic manner, as well as in a risk-averse manner.

As the upper limit on the PM's beliefs about the maximum level of immunity increases, holding the minimum level constant at zero, the density functions shift rightwards, as shown. These distributions have successively higher expected values, but also higher variances. The lower part of Figure 2 shows that the optimal level of intervention falls as the upper limit increases, in both the deterministic ($\sigma^2_{p_A} = 0$) and stochastic ($\sigma^2_{p_A} > 0$) cases. This reflects the increasing expected value of beliefs about herd immunity. It also shows that the difference between the optimal level of intervention for the two cases grows, owing to the increasing variance. In particular, when the upper limit on the distribution is equal to 1 (distribution not shown in Figure 2), Figure 2 (lower) shows that the optimal level of intervention in the deterministic scenario is 0.217 and that in the stochastic case is 0.263 (+21%). Also shown are the proportions remaining uninfected post-intervention, T: 0.756 (stochastic) and 0.733 (deterministic), up from $\mu_{p_A} = 0.5$.

3 Discussion

The framework shows how an approximation to the optimal level of intervention for infectious disease control may be obtained and interpreted for a PM whose beliefs about the value of a key parameter – we have chosen to work with the level of herd immunity in the population – are modelled using a random variable. By working in the neighbourhood of the expected value,

[5] We cross-checked the error obtained from using Eq. (5) to obtain the optimal level of intervention with the value resulting from a discretised maximisation of Eq. (3) and found good agreement (see Appendix B).

Eq. (5) shows that the intervention is determined by interactions between functions describing beliefs, attitude to risk, the production technology and the economic cost functions, weighted appropriately. Failure to account for uncertainty when setting the optimal level of intervention imposes a welfare loss.

Compared with the case of a PM who acts in a deterministic manner, ignoring the variance of beliefs, Propositions 1.A to 1.D show that the optimal intervention differs when the PM acts in a risk-neutral or risk-averse manner, accounting for the variance, owing to the pure technology effect (risk-neutral) and both the pure technology and the welfare/technology effect (risk-averse). In particular, when the PM acts in a risk-averse manner, the optimal level of intervention is higher than when it acts in a risk-neutral manner if the degree of relative prudence is high enough (Proposition 2).

This paper has sought to address some of the concerns raised by, e.g., Jenny (2020) and Avery et al. (2020) about the design of economic policy when there exists uncertainty over the value or values of key parameters in an economic model of an epidemic. In order to have quantitative guidance for policy, specific functional forms on the economic cost function and the production technology, suitably calibrated, are required. It may be possible to estimate these from the rich and rapidly-expanding data sets on the outbreak (a list of sources is provided by Avery et al. (2020, pages 67–68)). The PM's PH welfare function must also be specified, together with specification of a prior distribution on the parameter value of interest (methods could draw on random sampling of the population and/or techniques described in, for example, Kadane and Wolfson (2002), Gill and Walker (2005), Oakley and O'Hagan (2007) and Johnson et al. (2010)).

The analysis highlights the importance to the PM of assessing the provenance and quality of the sources of information which comprise its information set, $\mathcal{S}$. The simulation of section 2.3 shows how changes in $\mathcal{S}$ can increase both the expected value and the variance of beliefs, leading to a reduction in the optimal level of intervention for both the deterministic and the stochastic case. It also shows that there is increasing divergence between the levels of intervention as the expected value and variance increase. Other scenarios are, of course, possible: evidence/advice could increase μ_{p_A} and reduce $\sigma^2_{p_A}$ or vice versa, or reduce both μ_{p_A} and $\sigma^2_{p_A}$ (read the sequence of distributions in Figure 2 from right to left).

Reviewing five models of the current COVID-19 outbreak, Avery et al. (2020) note that there appears to be a 'general tendency for researchers to report a greater degree of confidence than is warranted for an existing model, in part because it is not straightforward to quantify parameter uncertainty or to trace the effects of those uncertainties in a non-linear model' and that 'the language of these papers suggests a degree of certainty that is simply not justified.' Similar concerns have been expressed by Lakëns (2020). One way to assess the quality of information would be to use tools from checklists and guidelines that are used to assess scientific research and evaluate the degree to which such sources of information have, themselves, taken into consideration uncertainties. For example, the following questions from the checklist of Drummond et al. (2005), often used to assess economic evaluations of medical technologies, could be used: 'Was allowance made for uncertainty in the estimates of costs and consequences? ...Were the conclusions of the study sensitive to the uncertainty in the results, as quantified by the statistical and/or sensitivity analysis?'

We list ideas for improvements and extensions, referencing the COVID-19 pandemic.

1. *Incorporation of uncertainties on other parameters*. As noted in section 2, the COVID-19 outbreak is characterised by uncertainty over a range of dimensions on which policy-makers are likely to form beliefs. It should be reasonably straightforward to incorporate additional uncertainties, accounting for correlations between random variables so introduced. For example, beliefs about the economic cost functions, as well as the effectiveness of the production technology, could be modelled (for an example of the former, see Rheinberger and Treich (2017, section 4.2)).

2. *Sequential modelling using Bayes' theorem*. The PM's information set, $\mathcal{S}$, can be updated over time. A dynamic version of the model could use Bayes' theorem to update beliefs, permitting the optimal level of intervention to evolve as $\mathcal{S}$ changes:

   (a) *A model of optimal testing for immunity within the population.* Starting with a prior distribution on p_A, the results of section 2 could be used within a two stage model of optimal testing for immunity, followed by intervention. In Stage 1, the optimal sample size for random testing of the population (or a sub-group within it) could be carried out, at a given fixed marginal cost, c, per test, with the optimal level of intervention being decided in Stage 2, conditional upon the result of testing. Bayes' theorem could be used to update the prior distribution using the result of the random sampling. The optimal size of the test could then be obtained recursively. Define $v_2^*(\alpha^*(\mu_{p_A}, \tilde{\sigma}^2_{p_A}(n); \mathcal{S}_1))$ as the Stage 2 maximised value which results from the solution to Eq. (5) given the information set at the start of Stage 1 and the result of the test. Assuming that n individuals are sampled in Stage 1, the optimal sample size solves:

   $$\text{argmax}_n \int_0^1 v_2^*(\alpha^*(\mu_{p_A}, \tilde{\sigma}^2_{p_A}(n); \mathcal{S}_1))\, \mathrm{d}g_{\mu_{p_A}} - c\,n,$$

   where $g_{\mu_{p_A}}$ is the density function for the predictive distribution of the posterior expected level of immunity prior to carrying out the test.

   (b) *Incorporation of Bayesian updating into existing SIR models.* Avery et al. (2020) note a range of heterogeneities that could be accounted for in the standard SIR models of viral spread (including those relating to exposure, responses across people, medical capacity, dosage and the nature of exposure and whether multiple strains of the virus exist). Bayesian learning about such parameters could be incorporated into these models and solved using dynamic programming methods. For example, beliefs about the economic damage functions could use results from observational analyses such as Fadinger and Schymik (2020), who estimate that the damage of confinement for the German economy is around 1.6% of GDP per week and/or Inoue and Todo (2020), who estimate that the lockdown will lead to a production loss that is equivalent to 5.3% of Japan's annual GDP.

3. *Alternative approaches to modelling risk and uncertainty*. An alternative approach would be to solve stochastic optimal control versions of the SIR models, such as those that are

often applied in the financial and environmental economics literature (see, e.g. Dixit and Pindyck (1994)). Focus would then be on how the variances and covariances of stochastic processes affect the optimal level of intervention, or the optimal stopping time for, e.g., the relaxation of a lockdown.

A Proofs

Proof of Proposition 1.A. The Taylor series expansion of the expected value in Eq. (4) is:

$$\begin{aligned}\mathbb{E}\left[\frac{\partial U(p_A+p_C(\alpha;p_A))}{\partial T}\frac{\partial p_C(\alpha;p_A)}{\partial \alpha}\right] = \\ \left[\frac{\partial U(p_A+p_C(\alpha;p_A))}{\partial T}\frac{\partial p_C(\alpha;p_A)}{\partial \alpha}+\right. \\ \left.\frac{\sigma^2_{p_A}}{2}\frac{\partial^2}{\partial p_A^2}\left(\frac{\partial U(p_A+p_C(\alpha;p_A))}{\partial T}\frac{\partial p_C(\alpha;p_A)}{\partial \alpha}\right)\right] + h.o.t.\Big|_{\mu_{p_A}}.\end{aligned} \tag{8}$$

We obtain an approximation to the optimal level of intervention by working with the second order expansion. Extension to higher orders is straightforward. Dropping the dependence of the functions on their arguments, evaluate the second derivative on the RHS of Eq. (8):

$$\begin{aligned}\frac{\partial^2}{\partial p_A^2}\left[\frac{\partial U}{\partial T}\frac{\partial p_C}{\partial \alpha}\right]\Big|_{\mu_{p_A}} = \left[\frac{\partial^3 U}{\partial T^3}\left(\frac{\partial T}{\partial p_A}\right)^2\frac{\partial p_C}{\partial \alpha}+\frac{\partial^2 U}{\partial T^2}\frac{\partial^2 T}{\partial p_A^2}\frac{\partial p_C}{\partial \alpha}+2\frac{\partial^2 U}{\partial T^2}\frac{\partial T}{\partial p_A}\frac{\partial^2 p_C}{\partial \alpha\partial p_A}+\right. \\ \left.\frac{\partial U}{\partial T}\frac{\partial^3 p_C}{\partial \alpha \partial p_A^2}\right]\Big|_{\mu_{p_A}}.\end{aligned} \tag{9}$$

Using Eq. (9) in Eq. (8) gives Eq. (5). □

Proof of Proposition 1.B. The result is clear from inspection of Eq. (5): an increase in ϕ reduces the marginal cost of the intervention and so increases the optimal level of the intervention. □

Proof of Proposition 1.C. The risk premium ρ is equal to zero when the PM acts in a risk-neutral manner. Under the assumption that $\rho = 0$, Eq. (5) simplifies to:

$$\frac{\partial Y}{\partial \alpha}\left(\frac{\beta_1}{\phi}\frac{\partial V_1}{\partial Y}+\frac{\beta_2}{\phi}\frac{\partial V_2}{\partial Y}\right)\Big|_{\alpha^*} \approx \frac{\partial U}{\partial T}\frac{\partial p_C}{\partial \alpha}\left\{1+\frac{\sigma^2_{p_A}}{2}\Omega_1\Omega_2\right\}\Big|_{\mu_{p_A},\alpha^*}. \tag{10}$$

Hence, when $\sigma^2_{p_A} > 0$, the marginal benefit (RHS) is greater/equal to/less than that when $\sigma^2_{p_A} = 0$ according to whether $\Omega_1\Omega_2 \gtreqless 0$ at μ_{p_A} and α^*. □

Proof of Proposition 1.D. When $\sigma^2_{p_A} > 0$ and the PM acts in a risk-averse manner, the optimal level of intervention is determined by the sign of $\rho\Upsilon$ in Eq. (5). We solve for the value of P_r which makes this term greater than/equal to/less than zero, which gives Eq. (6). □

Proof of Corollary 1. The result follows from making the appropriate simplification to Eq. (6) under the assumption of the corollary. □

Proof of Proposition 2. When $\sigma^2_{p_A} > 0$ and the PM acts in a risk-averse manner, the optimal level of intervention is determined by the sign of $(\sigma^2_{p_A}/2)\Omega_1\Omega_2 + \rho\Upsilon$ in Eq. (5). Repeating the approach taken for the proof of Proposition 1.D gives Eq. (7). □

B Simulation

For the illustrative simulation used to produce Figure 2, we used the following functional forms: $U = T^{1-\lambda}/(1-\lambda), \lambda = 2$, $p_C = (1-p_A)\alpha^{1/2}$ and we assumed a single economic cost function of the form $V = \alpha$. To model beliefs about herd immunity, we used a four parameter beta distribution (Hintze, J./NCSS, 2012, Chapter 551) with density function $f(p_A; 2.5, 2.5, 0, c)$, varying c in the interval $[0.05, 1]$. We used Matlab R2019a to derive Figure 2. We used Maple 18 to cross-check results and found good agreement, with approximations to the optimal α differing by between 0% and 1.47%. Matlab code here: matlabCodeBeliefs, calls detMax. Please replicate, check, improve.

References

Anderson, R. M., Heesterbeek, H., Klinkenberg, D., et al. (2020). How will country-based mitigation measures influence the course of the COVID-19 epidemic? *Lancet*, 395:931–934.

Arrow, K. (1963). Liquidity preference. Lecture VI in Lecture Notes for Economics 285, The Economics of Uncertainty, Stanford University.

Atkeson, A. (2020). How deadly is COVID-19? Understanding the difficulties with estimation of its fatality rate. NBER Working Paper No. 26965. Available from https://www.nber.org/papers/w26965.

Avery, C., Bossert, W., Clark, A., Ellison, G., and Ellison, S. (2020). Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists. *Covid Economics*, 12:21–68.

Bu, D., Hanspal, T., Lao, Y., and Liu, Y. (2020). Risk taking during a global crisis: evidence from Wuhan. *Covid Economics*, 5:106–146.

Chen, X. and Qiu, Z. (2020). Scenario analysis of non-pharmaceutical interventions on global Covid-19 transmissions. *Covid Economics*, 7:46–67.

Dixit, A. and Pindyck, R. (1994). *Investment under Uncertainty*. Princeton University Press, Princeton, New Jersey, First edition.

Drummond, M., Sculpher, M., Claxton, K., et al. (2005). *Methods for the Economic Evaluation of Health Care Programmes*. Oxford, London, Fourth edition.

Eeckhoudt, L., Gollier, C., and Schlesinger, H. (2005). *Economic and Financial Decisions under Risk*. Princeton University Press, Princeton, N. J.; Oxford, First edition.

Engle, S., Stromme, J., and Zhou, A. (2020). Staying at home: Mobility effects of Covid-19. *Covid Economics*, 4:86–102.

Fadinger, H. and Schymik, J. (2020). Impact of working from home on infection risk in German regions. *Covid Economics*, 9:107–134.

Farboodi, M., Jarosch, G., and Shimer, R. (2020). Internal and external effects of social distancing in a pandemic. *Covid Economics*, 9:22–58.

Ferguson, N., Laydon, D., Nedjati-Gilano, G., et al. (2020). Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Report 9, Imperial College COVID-19 Response Team. Available from: https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf.

Frey, C., Chen, C., and Presidente, G. (2020). Democracy, culture, and contagion: Political regimes and countries' responsiveness to Covid-19. *Covid Economics*, 18:222–238.

Gill, J. and Walker, L. (2005). Elicited priors for Bayesian model specifications in political science research. *The Journal of Politics*, 3:841–872.

Gonzalez-Eiras, M. and Niepelt, D. (2020). On the optimal 'lockdown' during an epidemic. *Covid Economics*, 7:68–87.

Hale, T., Webster, S., Petherick, A., Phillips, T., and Kira, B. (2020). Oxford COVID-19 Government Response Tracker. Blavatnik School of Government. Data use policy: Creative Commons Attribution CC BY standard. Available from: https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker.

Hintze, J./NCSS (2012). NCSS User's Guide V.

Inoue, H. and Todo, Y. (2020). The propagation of the economic impact through supply chains: The case of a mega-city lockdown to contain the spread of Covid-19. *Covid Economics*, 2:43–59.

Jenny, F. (2020). Economic resilience, globalisation and market governance: Facing the Covid-19 test. *Covid Economics*, 1:64–78.

Johnson, S., Tomlinson, G., Hawker, G., Granton, J., and Feldman, B. (2010). Methods to elicit beliefs for Bayesian priors: a systematic review. *Journal of Clinical Epidemiology*, 63(4):355–369.

Jones, C., Philippon, T., and Venkateswaran, V. (2020). A note on efficient mitigation policies. *Covid Economics*, 4:25–46.

Kadane, J. and Wolfson, L. (2002). Experiences in elicitation. *The Statistician*, 47(1):3–19.

Kermack, W. and McKendrick, A. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society, Series A*, 115:700–721.

Knight, F. (1921). *Risk, Uncertainty and Profit*. First edition.

Kozlowski, J., Veldkamp, L., and Venkateswaran, V. (2020). Scarring body and mind: the long-term belief-scarring effects of Covid-19. *Covid Economics*, 8:1–26.

Lakëns, D. (2020). Pandemic researchers – recruit your own best critics. *Nature*, 581:121.

Lourenço, J., Paton, R., Ghafari, M., et al. (2020). Fundamental principles of epidemic spread highlight the immediate need for large-scale serological surveys to assess the stage of the SARS-CoV-2 epidemic. Available from: https://www.medrxiv.org/content/10.1101/2020.03.24.20042291v1.article-info.

Ludvigson, S., Ma, S., and Ng, S. (2020). Covid-19 and the macroeconomic effects of costly disasters. *Covid Economics*, 9:1–21.

Moser, C. and Yared, P. (2020). Pandemic lockdown: The role of government commitment. *Covid Economics*, 17:1–23.

Oakley, J. and O'Hagan, A. (2007). Uncertainty in prior elicitations: a nonparametric approach. *Biometrika*, 94(2):427–441.

Painter, M. and Qiu, T. (2020). Political beliefs affect compliance with Covid-19 social distancing orders. *Covid Economics*, 4:103–123.

Pan, A., Liu, L., Wang, C., et al. (2020). Association of public health interventions with the epidemiology of the COVID-19 outbreak in Wuhan, China. *Journal of the American Medical Association*, doi:10.1001/jama.2020.6130.

Pratt, J. (1964). Risk aversion in the small and in the large. *Econometrica*, 32:122–136.

Rheinberger, C. and Treich, N. (2017). Attitudes toward catastrophe. *Environmental Resource Economics*, 67:609–636.

Rowthorn, R. (2020). A cost-benefit analysis of the Covid-19 disease. *Covid Economics*, 9:97–106.

Spiegelhalter, D. J., Abrams, K. R., and Myles, J. P. (1994). *Bayesian Approaches to Clinical Trials and Health-care Evaluation*. Wiley, Chichester, First edition.

Thunstrom, L., Nebold, S., Finnof, D., Ashworth, M., and Shogren, J. (2020). The benefits and costs of flattening the curve for COVID-19. Available from: https://ssrn.com/abstract=3561934orhttp://dx.doi.org/10.2139/ssrn.3561934.

World Health Organisation (2020). "Immunity passports" in the context of COVID-19. Scientific Brief, 24 April, 2020. Available from: https://www.who.int/news-room/commentaries/detail/immunity-passports-in-the-context-of-covid-19.