

On Fast Multi-Shot Epidemic Interventions for Post Lock-Down Mitigation: Implications for Simple Covid-19 Models

M. Bin, P. Cheung, E. Crisostomi, P. Ferraro, H. Lhachemi, R. Murray-Smith, C. Myant, T. Parisini, R. Shorten, S. Stein, L. Stone

Version 4, April 8th, 2020

Also added more supporting literature and summary of mitigation measures. New co-authors included (Lhachemi, Murray-Smith, Stein, Stone).

Executive Summary: The initial UK government strategy of a timed intervention as a means of combating Covid-19 is in stark contrast to policies adopted in many other countries which have embraced more severe social distancing policies. Our objective in this report is to enunciate the differences between these policies and to suggest modified policies, for post lock-down, that may allow management of Covid-19, while at the same time enabling some degree of reduced social and economic activity. The suggested policies can also be used to augment and compliment other post lock-down strategies to provide additional levels of robustness in the post lock-down period.

Disclaimer : Our results are based on elementary SIR, SIQR and SIDARTHE models. We are also not epidemiologists. More extensive validation is absolutely necessary on accurate Covid-19 models. Our intention is simply to make the community aware of such policies. All the authors are available for discussion via email addresses or by contacting the corresponding author R. Shorten.

Change-log : Version 2: edited to include some of the related literature. Also mitigation strategy is further verified on a recent Italian model [1]. All results presented here are qualitatively consistent with this model.

Version 3: more refined numerical results are now included using hybrid systems integration solver [2].

Version 4: More extensive report with: (i) detailed simulations; (ii) description of outer loop; and (iii) sensitivity analyses.

M. Bin and T. Parisini are with the Department of Electrical and Electronic Engineering, Imperial College London. Email: {m.bin, t.parisini}@imperial.ac.uk

P. Cheung, P. Ferraro, C. Myant, R. Shorten are with the Dyson School of Design Engineering, Imperial College London. Email: {p.cheung, p.ferraro, conor.myant, r.shorten}@imperial.ac.uk

H. Lhachemi is with the School of Electrical & Electronic Engineering, University College Dublin. Email: hugo.lhachemi@ucd.ie

R. Murray-Smith and S. Stein are with the School of Computing Science, University of Glasgow. Email: {roderick.murray-smith@glasgow.ac.uk, Sebastian.Stein@glasgow.ac.uk}@imperial.ac.uk

T. Parisini is also affiliated with the Department of Engineering and Architecture, University of Trieste, Italy.

R. Shorten is also affiliated with the Department of Electrical and Electronic Engineering, University College Dublin

L. Stone is with The George S. Wise Faculty of Life Sciences, Tel Aviv University. Email: lewi@tauex.tau.ac.il

I. Background

Presently, governments worldwide are struggling to contain the Covid-19 epidemic. Most countries, such as Italy, China, USA, Germany and France, have adopted severe social distancing policies, which amount to a total *lock-down* of their populations, in an attempt to combat the virus [4]. In contrast, other governments have attempted to control the effect of the virus through *timed* interventions [3]. We call the first policy, the *lock-down policy* (LDP), and the second policy, the *timed intervention policy*¹ (TIP).

Roughly speaking, the LDP attempts to stop the virus in its track, and in doing so, buys society some time to find effective mitigation strategies such as a vaccine, or to build healthcare capability. While such measures are likely to be effective in reducing the spread of the virus (c.f. China) they do come at a heavy economic cost. For example, in India in the first two weeks of lock-down it is estimated that 100 Million people got unemployed and in the first two days of lock-down in Ireland, it is estimated that 140,000 people were made redundant (approx 6% of the workforce).² These statistics are likely to become more grim, and may be followed by even more severe economic consequences, as personal/mortgage loan defaults emerge, which may spread to the banking sector. Despite these costs, the LDP makes sense if we are able to utilise the time gained to develop a vaccine for the general population. On the other hand, if we are unable to develop a vaccine quickly, this policy gives no clear exit strategy from the current crisis, and the virus may simply re-emerge once the policy of strict social distancing is relaxed.

The TIP can be considered as a demand-management policy. The key consideration here is the capacity of the healthcare system to absorb and treat new illnesses that arise as a result of Covid-19. As interventions such as social distancing place difficult burdens on society, the argument is that these

¹<https://www.technologyreview.com/s/615375/what-is-herd-immunity-and-can-it-stop-the-coronavirus/>

²<https://www.irishpost.com/news/140000-people-ireland-lose-jobs-due-coronavirus-crisis-forcing-businesses-close-181717>

interventions must be timed for maximum benefit to the healthcare system by reducing the number of infected people to a manageable level. The difficulty with this approach is timing. Intervene too early, and one simply shifts the peak of ill people to a later date, whereas too late an intervention will not limit the peak of infections at all. The issue of timing is exacerbated by the virus (apparently) having up to a 14-day incubation period³, as well as an initial exponential growth rate. Thus, the problem of observing the true state of the epidemic, in the face of exponential growth, makes the effectiveness of this policy very sensitive to the timing of intervention.

Our suggestion in this note, as in [5], is to use *multi-shot* interventions to manage the epidemic, once the current lock-down policies have brought the epidemic under control. The principle is to allow some level of social interaction, followed by social distancing. This, as suggested in the recent report by a team at Imperial College [5], is in fact the basis of TIPs. The policy suggested therein is developed from the perspective of intermittent social distancing policies with a view to manage the amount of infected people at a given time, so that the healthcare system is able to cope. A consequence of this policy, which is based on measurements from the healthcare system, is to control the spread of the virus at a rate to ensure a level of infected at any time is below the healthcare capacity. In this note, but in contrast to [5], we argue that *open-loop* interventions over very short time-scales, **rather than interventions based on measurements over long time-scales**, may also be good as a strategy. We also propose an *outer supervisory control loop* that takes explicitly into account the significant delays inherent in the measurements to enhance – at a slow rate – the switching characteristics of the open-loop intervention policy. This is not only to control the number of infections, but to also suppress the virus at a lower cost to society. The possible advantages of this approach are that, as an exit from the current lock-down strategy, a multi-shot policy may allow some level of economic activity, as well as reducing the sensitivity to the timing of interventions and mitigating the risk of a new wave of the epidemic.

II. Modelling and Control of the Covid-19 Disease: Overview

A. Epidemiological Models

While time series analysis for modelling infectious diseases has a long history [16], it was only in the first half of the last century [15] that *compartmental models* were first proposed to describe the relationships between *susceptible*, *infected*, and *recovered* people in a population. Accurate epidemiological models are fundamental to take measures to mitigate the effects of a disease. For example, the predictive ability of a model can be used for better planning of health services [17], and to compare the expected impact of different mitigation interventions, including non-pharmaceutical interventions such as wearing masks and social distancing, or full isolation

of positive cases (e.g., quarantines).

The SIR model [8] is the classic model that is widely adopted to describe epidemiological dynamics in a well-mixed population (see, for example [3], [4], [9]):

$$\begin{aligned}\frac{dS}{dt} &= -\frac{\beta SI}{N} \\ \frac{dI}{dt} &= \frac{\beta SI}{N} - \gamma I \\ \frac{dR}{dt} &= \gamma I.\end{aligned}\tag{1}$$

In this model S , I , and R denote the aforementioned susceptible, infected, and recovered people in a total population of N individuals. All quantities vary in time, but for simplicity of notation the dependency on time has been dropped in the previous equations (and the time dependency will not be written explicitly in the remainder of the section either). Note that at any instant in time one individual can only belong to one of such three classes, as at any time instant $S+I+R = N$. Finally, note that R may be interpreted as the number of *resistant* people as well, as they are supposed to have acquired immunity after recovering from the disease [4], which is usually true for most epidemics. The parameter β represents the rate of effective contacts between infected and susceptible individuals, and γ represents the disease-specific rate at which infected individuals recover. An important quantity characterising an epidemic is the *basic reproduction number*, defined as $R_0 = \beta/\gamma$ [3].

B. SIR-like Models

Many variants have been derived from the original SIR model, to capture peculiar aspects of specific diseases. Some interesting generalisations include the following.

- **Forced SIR Models**, where the contact rate β is assumed to be a time-varying quantity that is affected by a seasonal forcing term that takes into account the seasonally changing contact rates in recurrent epidemics [18].
- **SEIR Models**, where a further class of exposed people is added to the states of the SIR model to take into account that fact that in some cases the infected people may become infectious only after a latent period of time.
- **SIQR Models**, where a further class of people who are in quarantine is added to the three classes of the SIR model [11];
- **Meta Population Models**, where the population is sub-divided into sub-populations [3], to take into account that the spreading of the disease will in general be different in sub-populations that may correspond to people of different ages, or to people living in different areas where the homogeneous assumption of SIR-like

³<https://www.healthline.com/health/coronavirus-incubation-period>

models does not hold.

The SIQR model appears to be particularly convenient for modelling the Covid-19 disease [9], as it has the advantage of considering infectious people and those who are in quarantine. This is convenient because it models the fact that many governments, including the Italian one, are forcing individuals tested positive (positive individuals) to self-isolate from the community, and also because it distinguishes between the infectious people who do self-isolate, and/or those who do not (mostly likely because they have not developed the symptoms of the disease and are not aware of actually being infectious). For this reason, we now provide in more detail the equations underlying an SIQR model:

$$\begin{aligned}
 \frac{dS}{dt} &= -\frac{\beta SI}{N} \\
 \frac{dI}{dt} &= \frac{\beta SI}{N} - (\alpha + \eta)I \\
 \frac{dQ}{dt} &= \eta I - \delta Q \\
 \frac{dR}{dt} &= \delta Q + \alpha I.
 \end{aligned} \tag{2}$$

In this SIQR model, the I state actually includes positive individuals who will never develop symptoms; positive individuals who have not developed symptoms yet; and positive individuals who have symptoms, but have not been tested positive and isolated (i.e., they may believe that they are experiencing flu-like symptoms) [9]. The parameters in (2) have the same meaning as the classic SIR model, i.e., $\alpha + \eta$ playing the role of γ in (1). In addition, we have parameter δ , whose inverse δ^{-1} can be estimated considering the average number of days after which isolated and hospitalized patients recover or die (in both cases, they are assumed to pass to the R state).

Finally, we conclude this section by noting one further SIR-like model, which has been introduced with the purpose to describe the specific behaviour of the COVID-19 model. This is called the SIDARTHE model [1], and it is convenient as it introduces a large number of classes to partition the infected people according to the degree of severity of their symptoms. In particular, eight states are identified that correspond to: ‘ S ’, the susceptible people as usual; ‘ I ’, the asymptomatic undetected infected people; ‘ D ’, the diagnosed people, corresponding to asymptomatic detected cases; ‘ A ’, the ailing people, corresponding to the symptomatic undetected cases; ‘ R ’, the recognized people, corresponding to the symptomatic detected cases; ‘ T ’, the threatened people, corresponding to the detected cases with life-threatening symptoms; ‘ H ’, the healed people; and the ‘ E ’, the extinct or dead people. Obviously, this model is more detailed than the previous ones, and it takes into account that, for instance, the transition from the ‘infected’ state to a ‘recovered’ state depends on the severity of the infection. Also, this model is convenient as not all states of an SIR model may be regarded as observable, or in other words, directly measurable. For

instance, it may be simpler to quantify the number of recovered people with life-threatening symptoms, than the number of overall infected people, that actually includes asymptomatic cases as well.

The eight ordinary differential equations of a SIDARTHE dynamical system are reported below (from [1]. Note that parameter names have been adapted for consistency with the previously introduced models):

$$\begin{aligned}
 \frac{dS}{dt} &= -\frac{\beta S}{N} \cdot (\sigma_1 I + \sigma_2 D + \sigma_3 A + \sigma_4 R) \\
 \frac{dI}{dt} &= \frac{\beta S}{N} \cdot (\sigma_1 I + \sigma_2 D + \sigma_3 A + \sigma_4 R) - (\sigma_5 + \sigma_6 + \sigma_7) I \\
 \frac{dD}{dt} &= \sigma_5 I - (\sigma_8 + \sigma_9) D \\
 \frac{dA}{dt} &= \sigma_6 I - (\sigma_{10} + \sigma_{11} + \sigma_{12}) A \\
 \frac{dR}{dt} &= \sigma_8 D + \sigma_{10} A - (\sigma_{13} + \sigma_{14}) R \\
 \frac{dT}{dt} &= \sigma_{11} A + \sigma_{13} R - (\sigma_{15} + \sigma_{16}) T \\
 \frac{dH}{dt} &= \sigma_7 I + \sigma_9 D + \sigma_{12} A + \sigma_{14} R + \sigma_{15} T \\
 \frac{dE}{dt} &= \sigma_{16} T.
 \end{aligned} \tag{3}$$

The interpretation of the parameters is as follows (from [1]):

- Parameters σ_1 , σ_2 , σ_3 and σ_4 denote the transmission rate from the susceptible state to any of the four infected states;
- Parameters σ_5 and σ_{10} denote the rate of detection of asymptomatic and mildly symptomatic cases;
- Parameters σ_6 and σ_8 denote the probability rate with which (asymptomatic and symptomatic) infected subjects develop clinically relevant symptoms;
- Parameters σ_{11} and σ_{13} denote the probability rate with which (undetected and detected) infected subjects develop life-threatening symptoms;
- Parameter σ_{16} denotes the mortality rate for people who have already developed life-threatening symptoms;
- Parameters σ_7 , σ_9 , σ_{12} , σ_{14} and σ_{15} denote the rate of recovery for the five classes of infected subjects (including those in life-threatening conditions).

In the remainder of the report we shall describe results obtained using these models. As expected, qualitatively similar results have been obtained in general, since all SIR-like models contain, at their core, the same SIR dynamics.

C. Tuning of the COVID-19 Parameters

The COVID-19 outbreak has given rise to an unprecedented world-wide epidemic in terms of cases of infections and deaths. Consequently, research communities all over the world have joined efforts to improve the mathematical modelling of the disease, as a means to investigate and compare possible countermeasures. Here we briefly discuss some of the challenges in estimating the model parameters:

- **Undetected cases:** Since a consistent number of cases is asymptomatic, it is very hard to estimate the correct number of infected people, especially in the initial phase of the outbreak when Covid-19 test kits were very few in number limiting meaningful disease surveillance. For instance, it is estimated in [19] that during the initial spread of the virus in China (in the period 10-23 January, 2020), 86% of the infections went undocumented and that of these 55% were contagious. While such number can be reduced with different sampling and testing strategies, recalibrated models for the Italian case similarly estimate a number of 40 – 50% of non-diagnosed infected individuals ([1], [9]). Such numbers may be even higher should the findings of [23] be confirmed.
- **Time-varying parameters:** Models are usually calibrated under the assumption that parameters are constant in time (at least for a significant number of days). Obviously, most parameters actually change in time due to changes in the environment such as changes in human behaviour (e.g., washing hands frequently and wearing protective masks), and evolutions of the disease. Clearly, parameters are expected to change even more dramatically when total lock-down measures are taken.
- **Geographical differences:** While similar patterns of the evolution of infected people have been observed in different countries, possibly with some delays, models calibrated upon data from one country cannot be used *as-is* for other countries. This is also true even within countries, from one region to another [10].
- **Heterogeneity of the population:** Demographic aspects (e.g., the density of the population, see for instance the particular case of the Diamond Princess Ship [20]), age of the population, and to a lesser extent other factors like religion, ethnicity or socio-economic status are known to also have an impact on the spreading of a disease [3].
- **Time delays:** While many quantities may be measurable, it should be noted that often they can only be measured with systemic delays (for instance, an infected individual takes some time before eventually showing symptoms; then it takes some time before he/she gets tested; plus a further few hours before a diagnosis is received [4]). This is a critical aspect to consider any time a control action is designed on the basis of a measured quantity.

While qualitative analyses can provide an intuitive understanding and notionally predict the evolution of an epidemic, particular care is required when making quantitative predictions as estimates of parameters are inherently uncertain. One approach to dealing with this uncertainty is via sensitivity analyses, whereby effects of changes in estimated parameters on dependent variables of interest are investigated. Bayesian inference is a principled approach to quantify uncertainty by estimating posterior distributions over model parameters from prior knowledge and observed data. This has been applied to infer the impact of non-medical interventions on R_0 and deaths avoided through interventions, and for sensitivity analysis with respect to subsets of the data and prior assumptions [23].

D. Control and Mitigation Actions

In most countries, governments have gradually increased measures to limit social mixing to abate the course of the epidemic. Basic measures include maintaining a safe distance from other people, frequent washing of hands, and encouraging other general behaviours (e.g., not touching ones own face, or how to properly sneeze or cough). As the epidemic progressed governments have started taking more severe measures (often in a gradual fashion), which include closing schools and universities, restaurants and bars/pubs, theaters and cinemas, and eventually forcing a lock-down. While lock-downs themselves have been implemented in different ways in different countries, and sometimes with a different level of social compliance, the lock-down policies appear to be the most effective and heavily implemented worldwide. For example, the authors in [22] conclude that in some plausible scenarios, case isolation alone would be unlikely to control transmission within 3 months. Yet China appears to show that isolation of infected populations can contain the epidemic [21] within such a time-scale.

There is general agreement that a lock-down phase is necessary to abate the number of infected people to negligible values. However, as soon as normal activities are restored, there is a significant risk of a new wave of the epidemic. Thus, several exit strategies are being discussed as a means to allowing economic activity after an initial period of lock-down. Roughly speaking, these can be classified as follows.

- (i) **Data driven intermittent lock-downs:** Some authors have suggested using feedback to trigger lock-down and release-from-lock-down interventions [5]. As we have suggested, the use of feedback in the context of Covid-19 is potentially problematic due to the speed at which the virus grows, and measurement delays. In addition, such policies are likely to be highly localised, meaning that policies would be implemented non-uniformly across geographic areas, and thus be non-trivial to implement.
- (ii) **Contact tracing with/without testing:** Contact tracing, however interesting an option, is also not without

flaws. Apart from the obvious ethical, privacy and technical challenges (Bluetooth reliability, GPS errors, data management, scalability issues), other issues arise, including that large demographic groups may not find it easy to access this technology, and thus would act as un-observables in any nationwide strategy. For example, over 70's have limited access to smart-phones in many countries. Furthermore, asymptomatic Covid-19 carriers would also be difficult to detect. Finally, contact tracing is also open to cyber-physical attacks that can, in this case, be very-harmful to society.

- (iii) **One-shot interventions:** One-shot interventions have also been discussed in [3]. Such interventions are also problematic as any optimal strategy would require the number of infectious in society to be large. Such policies would not cope with localised disease dynamics and demographic heterogeneity.
- (iv) **Population scheduling:** Another approach is to split the population into a number of bins, and for members of these bins to take turns in lock-down. Apart from the complexities of organising society around such a strategy, significant cross-bin leakage is likely, due to social interactions in households.
- (v) **Immunity passports:** Finally, some governments suggest the use of immunity passports to identify and bestow rights on citizens who are immune to the disease. Again, such policies are problematic. Apart from stratifying society, they are open to being exploited, and may in fact encourage individuals in financial need to deliberately contract the virus.

We are interested in developing multi-shot epidemic interventions, as an alternative, or perhaps more likely, to complement to the above strategies, and as a generalisation of the one-shot interventions previously described. Our objective is to demonstrate that after an initial lock-down period, a frequently alternating sequence of lock-down and working days is expected to provide positive results in terms of overall expected infected people, while at the same time mitigating the huge economic impact of the epidemic. It should be noted that such an idea is not totally new, for instance [5] had already investigated the impact of an 'ON'-'OFF' triggered quarantine policy (i.e., where quarantine is enforced when then number of cases gets above a given threshold) in terms of the overall number of infected people (i.e., summing the number of infections at every open window of time). However, the frequency was on average very low (i.e., months) and driven by a feedback signal. Other switched, and more specifically, open-loop periodic strategies have been suggested in the context of other epidemics. In particular, periodic vaccination is suggested in [12], and periodic quarantines for combating computer worms, and viral epidemics, are suggested in [13], [14], [26]. Note however, that most of the research findings in these papers appear to relate to impulsive strategies; namely,

states of the viral dynamics are reset periodically to reflect the effect of an intervention policy. In contrast, our approach differs in that we assume to have limited or no room for intervention on the people that are already infected. Our policy consists of adjusting the infection rate of the disease by introducing a periodic suppression based on switching between the transmission rates of lock-down and not-lock-down. Notwithstanding this difference, we believe these works are closely related to our approach and may be consistent with the *hypothesis* that fast periodic switching may be useful, and consequently that such a policy may be a viable exit strategy to the current lock-down situation.

III. The Fast Periodic Switching Policy (FPSP)

The strategy proposed in [5] – and advocated in this document – gives rise to what is known as a switched system [6]. In the proposed strategy, we simply switch between allowing society to return to normality and accept virus to spread slowly, and enforced strict social isolation. Switched systems have been studied extensively since the mid-1990's and give rise to many interesting phenomena. Among these, it is well known that the choice of switching strategy fundamentally affects the behaviour of the system being influenced by the switching, and that sometimes – rather counter-intuitively – fast switching can be better than slow switching. More specifically, in the language of switching systems, the policy suggested in [5] is a slow switching strategy based upon a feedback signal (here hospitalised patients), and in fact resembles closely the multiple-Lyapunov function ideas developed by Michael Branicky [7] in the late 1990's. Furthermore, as we have already mentioned, control of systems growing exponentially fast with large time delays is very difficult. Our suggestion, on the other hand, is to use an **open-loop fast switching strategy** to control and suppress the growth of the virus in society.

In this policy, we simply allow society to function as normal for X days, followed by social isolation of Y days. This is then repeated (hence the *periodic* nature of the switching policy). As we shall see, policies for which X and Y are small, can be developed for which the virus is suppressed rapidly, and for which the peak level of infections is (relatively) low. To understand the effect of this policy in terms of the SIR and SIQR models described in Section II, the effect of the open-loop FPSP policy is to adjust the parameter β in accordance to the policy adopted:

$$\beta = \begin{cases} \beta^+ & \text{during inactive lock-down} \\ & \text{(society functioning as normal)} \\ \beta^- & \text{during lock-down and social isolation} \end{cases} \quad (4)$$

where β^+ and β^- are values corresponding to each situation, as described in Section II. Figure 1 shows a possible FPSP instance characterised by a *period* $T = 7$ days and a *duty-cycle* $D = 28.6\%$.

IV. Illustrative Simulations

In this section, we present some simulations of the proposed FPSP policy applied to the SIQR and SIDARTHE models




Fig. 1: Example of a FPSP policy with $T = 7$ days and $D = 28.6\%$.

described in Section II. The parameters for the different sets of equations can be found in Table I (for a full discussion on these models and their parameters, please refer to [1] and [9]). For all models, the starting number of susceptibles and infected and the total population are chosen as $(S(0), I(0)) = (N - \varepsilon, \varepsilon)$, with $\varepsilon = 83.333$, while the initial conditions for all the other state variables are set to 0. Moreover all the simulations follow the pattern below:

- **Phase 1:** The virus spreads with no containment attempts. This happens for $t < 20$ days and in this phase $\beta = \beta^+$.
- **Phase 2:** A strict lock-down is enforced to contain the spread of the virus. This can be considered analogous to the policies that some European governments are enforcing at the moment. We assume that the value β^+ switches to the value $\beta^- = q\beta^+$ with $q = 0.15$ in the SIQR model, and $q = 0.175$ in the SIDARTHE model. This happens for $20 \leq t < 50$ days.
- **Phase 3:** Once the number of infected people has decreased, the FPSP policy is enforced. This happens for $t \geq 50$ days.

Different simulations are obtained for different values of the period and the duty cycle. In particular, each FPSP is identified by a specific pair (X, Y) . The period of a (X, Y) -FPSP policy is defined as $T = X + Y$ days, while the duty cycle as $D = 100 \cdot X/(X + Y)$. Figures 2-3 show the distribution of the maximum peak-values (in percentage) of infected people obtained for each model by each (X, Y) -FPSP policy with X and Y ranging from 0 to 14 (i.e., the value of $(100/N) \cdot \sup_{t \geq 50} [I(t) + Q(t)]$ for the SIQR model and $(100/N) \cdot \sup_{t \geq 50} [I(t) + D(t) + A(t) + R(t) + T(t)]$ for the SIDARTHE model). Figures 4-5, instead, show the time instants at which such peaks are attained. Finally, only for the SIQR model, Figures 6 and 7, show the distribution of the peaks and peak-times for the FPSP obtained with X and Y ranging from 0 to 98 days with a resolution of 7 days.

For each model, we notice the following:

- There is a *stability region*, painted in light blue and located in the bottom-left part of the images, in which the peak values are similar to the one we would observe with a complete lock down, i.e. with any policy in which




Fig. 2: Percentage of peak infections parametrised by (X, Y) in a population of 10^7 individuals in the SIQR model.




Fig. 3: Percentage of peak infections parametrised by (X, Y) in a population of 10^7 individuals in the SIDARTHE model.

the number X of work days is set to zero, and the peak times are close to the time ($t = 50$ days) in which the policies are started. More precisely, with reference to Figures 4-5 and 7, we observe that a policy (X, Y) belonging to the stability region attains the peak at time $t \leq 50 + X$. This, in turn, implies that the trajectory of the total infected population obtained under these policies starts to decay after the first X days of the first period. We further underline that, although two policies belonging to the stability region have a similar peak value, they may show quite different behaviours. This is shown, for instance for the SIQR model, in Figure 8, in which the $(1, 6)$ and $(1, 3)$ policies are compared.

- There is an *instability region*, painted in dark blue and located in the top-right part of the images, in which the peak values are similar to the one we would observe without lock down, i.e. with any policy in which the number Y of quarantine days is set to zero. As the peak-time distributions in Figures 4-5 and 7 show, the policies belonging to this instability region do not necessarily lead to the same time evolution. In fact, policies with more days of quarantine (i.e., larger Y) are associated with larger peak times. This, in turn, implies that more days of quarantine still have the positive effect of delaying the peak.
- There is a *compromise region*, which contains the

TABLE I: Parameters of the different models

Model	Parameters	Values
SIQR	$(\beta; \alpha; \eta; \delta; q; N)$	$(0.373; 0.067; 0.067; 0.036; 0.175; 10^7)$
SIDARHTE	$(\beta; \sigma_i, i = 1, \dots, 16; q; N)$	$(1, 0.570, 0.011, 0.456, 0.011, 0.171,$ $0.371, 0.125, 0.125, 0.012, 0.027, 0.003,$ $0.034, 0.034, 0.017, 0.017, 0.017, 10^7)$




Fig. 4: Distribution of the peak times corresponding to Figure 2.




Fig. 5: Distribution of the peak times corresponding to Figure 3.

remaining policies, and which is located in the central band going from the top-left to the bottom-right corner. The policies of this region yield a peak of the number of infected people which is considerably larger than the value attained after the initial lock-down phase. However, they are associated with a larger duty cycle (i.e. a large fraction of work days X) than the policies belonging to the stability region, thus allowing a larger number of work days. Figure 9 compares the two policies, in the SIQR model, obtained with $(X, Y) = (2, 5)$ and $(X, Y) = (14, 35)$. Both the policies have the same duty cycle $D = 100 \cdot X / (X + Y) = 28, 6\%$.

- All the figures show a *main growth direction* going from the bottom-left to the top-right corner. This growth direction is associated with a reduction of the duty cycle $D = 100 \cdot X / (X + Y)$ of the policies, and reflects the fact that a higher number of work days leads to higher



Fig. 6: Percentage of peak infections parametrised by (X, Y) in a population of 10^7 individuals in the SIQR model. In this instance, the distribution is obtained with X and Y ranging from 0 to 98 days with a resolution of 7 days.

Fig. 7: Distribution of the peak times corresponding to Figure 6.

peaks in the infected population.

- All the figures show a *secondary growth direction*, which is orthogonal to the previous one, from the top-left to the bottom-right corner. This growth direction concerns only the stability and compromise regions, and it is associated with an increment in the period $T = X + Y$ of the policies. As shown in Figures 10-11, indeed, the simulations suggest that, for similar values of the duty cycle, **higher frequencies are associated with smaller and delayed peaks** of the infected population. This aspect is also shown in Figure 9, in which the two compared policies on the SIQR model, one obtained with $(X, Y) = (2, 5)$, the other with $(X, Y) = (14, 35)$, have the same duty cycle but different periods.




Fig. 8: Time evolution of the percentage of infected people $(100/N) \cdot (I(t) + Q(t))$ in the SIQR model for the FPSP policies obtained with $(X, Y) = (1, 6)$ and $(X, Y) = (1, 3)$ respectively. The red dashed line marks time $t = 20$ days, in which the lock-down phase starts. The blue dashed line marks time $t = 50$ days, when the FPSP policy commences.




Fig. 9: Time evolution of the percentage of infected people $(100/N) \cdot (I(t) + Q(t))$ in the SIQR model for the FPSP policies obtained with $(X, Y) = (2, 5)$ and $(X, Y) = (14, 35)$ respectively. The two policies have equal duty cycle $D = 100 \cdot X / (X + Y) = 28.6\%$. The red dashed line marks time $t = 20$ days, in which the lock-down phase starts. The blue dashed line marks time $t = 50$ days, when the FPSP policy commences.

Finally, notice that the peak level of infections (see Figures 2, 3) depends on the level of infected citizens at $t = 50$ (i.e., when the FPSP policy is enforced). However, if this number can be driven low enough (e.g., by prolonging the lock-down period), then this policy appears to be an effective quarantine exit strategy, that avoids a second increase in infected individuals while at the same time allowing a certain level of economic activity.




Fig. 10: Value of the peak (in percentage) of the infected population for different values of the duty cycle and period of the FPSP policies applied to the SIQR model. The plotted lines are obtained as linear interpolations of the peak values shown in Figures 2 and 6.




Fig. 11: Value of the peak times of the infected population for different values of the duty cycle and period of the FPSP policies applied to the SIQR model. The plotted lines are obtained as linear interpolations of the peak values shown in Figures 4 and 7, by removing the peak times of the policies belonging to the stability region (i.e., the peak times less than 55 days).

V. SENSITIVITY ANALYSIS

This section explores the sensitivity of the number of infected individuals in the SIQR model ($I + Q$) and in the SIDARTHE model ($I + D + A + R + T$) with respect to quarantine effectiveness, anticipatory and compensatory population behavior, and uncertainty in the model parameters. Unless stated otherwise, parameters are defined as in Table I.

1) *Quarantine Effectiveness*: The effectiveness of quarantine on a reduction q in the rate of infectious contacts cannot




Fig. 12: Sensitivity analysis of the quantity $Q + I$ in the SIQR model on quarantine effectiveness q . The FPSP-(1, 6) policy remains stable for $q \leq 0.255$, corresponding to a 74.5% reduction in infectious contacts during periodic quarantine days. In comparison to the assumed effectiveness $q = 0.175$, this leaves a safe error margin of 8%.




Fig. 14: Sensitivity analysis of the quantity $Q + I$ in the SIQR model on increased infection rate $(1 + d)\beta^+$ during periodic working days. The FPSP-(1, 6) policy remains stable for $d \leq 0.40$, corresponding to a 40% increase in infectious contacts during periodic working days due to compensatory and anticipatory population behavior.




Fig. 13: Sensitivity analysis of the quantity $I + D + A + R + T$ in the SIDARTHE model on quarantine effectiveness q . The FPSP-(1, 6) policy remains stable for $q \leq 0.33$, corresponding to a 66% reduction in infectious contacts during periodic quarantine days. In comparison to the assumed effectiveness $q = 0.175$, this leaves a safe error margin of 15.5%.




Fig. 15: Sensitivity analysis of the quantity $I + D + A + R + T$ in the SIDARTHE model on increased infection rate $(1 + d)\beta^+$ during periodic working days. The FPSP-(1, 6) policy remains stable for $d \leq 0.80$, corresponding to a 80% increase in infectious contacts during periodic working days due to compensatory and anticipatory population behavior.

be observed directly and can only be estimated with some delay from the start date of the initial quarantine period. It is commonly assumed that this effect is constant over time. The population's compliance with quarantine measures may reduce once a less stringent policy as proposed here is introduced. Exemplary simulation results are presented showing the effect of varying quarantine effectiveness q and an FPSP policy with $(X=1, Y=6)$ working days and quarantine days, respectively.

In simulation with the SIQR model (Fig. 12), the policy remains stable for $q \leq 0.255$, corresponding to a 74.5% reduction in infectious contacts during periodic quarantine days. In comparison to the assumed effectiveness $q = 0.175$, this leaves a safe error margin of 8%. In simulation with the SIDARTHE model (Fig. 13), the policy remains stable for $q \leq 0.33$, corresponding to a 66% reduction in infectious contacts during periodic quarantine days. In comparison to

the assumed effectiveness $q = 0.175$, this leaves a safe error margin of 15.5%.

2) Anticipatory and Compensatory Population Behavior:

Thus far, it has been assumed that the rate of infection during working days β^+ would revert back to the rate observed pre-quarantine. However, it is possible that it may increase above pre-quarantine level if the population mixes with higher frequency during working days. This may occur initially in response to prolonged quarantine, and subsequently in anticipation of future quarantine periods. Exemplary simulation results are presented modelling increased mixing with an infection rate $(1 + d)\beta^+$ during periodic working days with an FPSP-(1,6) policy. The SIQR model (Fig. 14) remains stable for $d \leq 0.40$, corresponding to a 40% increase in infectious contacts during periodic working days. The SIDARTHE model (Fig. 15) remains stable for $d \leq 0.80$, corresponding to a 80% increase in infectious contacts during periodic working days due to compensatory and anticipatory population behavior.

3) *Uncertainty in Model Parameters:* Deterministic epidemic models are highly sensitive to their parameters, which in practice cannot be observed directly but are instead inferred from test results or prior knowledge. It is important to quantify uncertainty in the model parameters and explore the range of possible effects of policy decisions under uncertainty.

a) *SIQR:* We consider uncertainty in the following parameters. The basic reproduction rate $R_0 \sim \mathcal{N}(\mu = 2.676, \sigma = 0.572)$ was sampled from the consensus distribution in [25]. The probability of an infected individual to show symptoms $p_s = 0.5$, the rate with which individuals develop symptoms and get tested $r_t = 0.2$, the probability of an infected symptomatic individual to test positive $p_d = 0.67$, the recovery rate of non-quarantined individuals $r_I = 0.1$, and the recovery rate of quarantined individuals $\delta = 0.036$ were taken from [9] and were each sampled from a truncated Normal distribution with mean as above and 10% standard deviation. The remaining parameters of the SIQR model were derived as in (5),

$$\begin{aligned} p_q &= p_s p_d \\ \alpha &= (1 - p_q) r_I \\ \eta &= p_q r_t \\ \beta^+ &= R_0 (\alpha + \eta). \end{aligned} \quad (5)$$

Monte Carlo simulations with 1000 draws from the joint distribution were performed with varying FPSP policy parameters. In order to isolate the effect of (X, Y) , the median number of quarantined individuals after 50 days was estimated as $Q(50) = 3251$ or 0.0325% of the total population. Initial quarantine periods were extended to $\text{argmin}_{t \geq 50} Q(t) \leq 3251$. The median, 75-percentile and 95-percentile of infected individuals under example policies are shown in Fig. 16. Note that the equi-percentile graphs do not correspond to individual Monte Carlo samples. The depicted examples illustrate that fast switching policies exist that are stable under a wide range of basic reproduction rates

and considerable uncertainty in all other model parameters.

b) *SIDARTHE:* We consider uncertainty in model parameters represented as zero-truncated Normal distributions with means $\sigma_1, \dots, \sigma_{16}$ and 10% standard deviation, and estimate β from samples of the basic reproduction rate $R_0 \sim \mathcal{N}(\mu = 2.676, \sigma = 0.572)$. Analogously to Monte Carlo simulations with the SIQR model, 1000 samples were drawn from the joint distribution and initial quarantine periods were extended to $\text{argmin}_{t \geq 50} D(t) + R(t) + T(t) \leq 212433$. The median, 75-percentile and 95-percentile of infected individuals under example policies are shown in Fig. 17. Stable FPSP parameterisations avoiding a second peak of infections with 95% probability exist also under the SIDARTHE model, both for weekly ($X \leq 2$) and biweekly ($X \leq 4$) switching periods.

c) *Remark:* SIR-like models assume a fraction of each compartment moving from one compartment to another. For example, if γ^{-1} is the average recovery time, then γI is the rate at which infectious individuals leave the I class and enter the R class. Of course, this is an approximation, and recoveries, quarantines, and other quantities, are in reality governed by a distribution. More accurate modelling will be explored to incorporate such effects, and establish their impact on the fidelity of the control schemes proposed in this document.

VI. Slow Outer Supervisory Control Loop

In the previous section, we showed how a fixed switching policy between quarantine and work days can be effectively used to reduce the number of infected individuals in a population without resorting to a complete lock down. While these results are promising and show that the several FPSP policies are feasible, it remains to establish how to select a specific pair (X, Y) . In fact, even though it is possible to specifically select one such policy from the examples above, we are still left with the question of how reliably we can estimate a model and its parameters in order to make such a choice with the caveat that standard feedback mechanisms may fail dramatically due to the high-level of uncertainty and delay in the measurements [4].

Due to the critical nature of this issue, we investigate the design of a slow outer supervisory control loop to compensate for model mismatch, that *does not depend on the specific choice of the model, nor its parameters*, and that satisfies the following basic requirements:

- (i) The supervisory control loop needs to find a policy (X, Y) such that the infection is suppressed.
- (ii) The supervisory control loop needs to be robust with respect to the uncertainties on the parameters and independent on the choice of the model.
- (iii) The supervisory control loop needs to take into account the delays in the observed measurements and inherent in




Fig. 16: Sensitivity analysis of the quantity $Q + I$ in the SIQR model on uncertainty in epidemiological model parameters. The basic reproduction rate $R_0 \sim \mathcal{N}(\mu = 2.676, \sigma = 0.572)$ was sampled from the consensus distribution in [25]. The probability of an infected individual to show symptoms $p_s = 0.5$, the rate with which individuals develop symptoms and get tested $r_t = 0.2$, the probability of an infected symptomatic individual to test positive $p_d = 0.67$, the recovery rate of non-quarantined individuals $r_I = 0.1$, and the recovery rate of quarantined individuals $\delta = 0.036$ were taken from [9] and were each sampled from a truncated Normal distribution with mean as above and 10% standard deviation. The remaining parameters of the SIQR model were derived as in (5).

the system's dynamics (e.g., the incubation period of the disease).

Owing to Requirements (i)-(iii), we propose to design an hysteresis-based supervisory control loop which is characterized by the simplicity of implementation and by its inherent robustness due to the independence on the structure of the model of the infection.

The supervisory outer control policy can be described as follows. Denote by t_0 the time-instant when the control action starts (i.e., the end of a prolonged lock-down) and set $X(t_0) = 0, Y(t_0) = c$. We consider the set of integers $T_c(t) = \{X, Y \in \mathbb{N} : X + Y = c\}$, where c is the time period during which the pair (X, Y) remains constant. Then, by considering the half-closed intervals $(t_k, t_{k+1}]$, with $t_{k+1} - t_k = c$, and using the shorthand $X(k+1)$ to denote $X(t_{k+1})$, the hysteresis-based supervisory outer control law can be expressed as follows:

$$X(k+1) = \text{mid}(0, X(k) + \text{sign}(\psi_X(k+1)), c), \quad (6)$$

$$Y(k+1) = \text{mid}(0, Y(k) - \text{sign}(\psi_Y(k+1)), c), \quad (7)$$

where

$$\text{mid}(a, b, c) = \begin{cases} a, & \text{if } b \leq a \\ b, & \text{if } a < b < c \\ c, & \text{otherwise} \end{cases}$$

In (6) and (7), functions ψ_X and ψ_Y are given by

$$\psi_X(k+1) = (1 - \alpha_X) \int_{t_{k-1}}^{t_k} [O(s) - O(t_{k-1})] ds - \int_{t_k}^{t_{k+1}} [O(s) - O(t_k)] ds,$$

$$\psi_Y(k+1) = -(1 + \alpha_Y) \int_{t_{k-1}}^{t_k} [O(s) - O(t_{k-1})] ds + \int_{t_k}^{t_{k+1}} [O(s) - O(t_k)] ds.$$

where α_X, α_Y represent two positive design constants and $O(t)$ denotes the observed amount of infected people (recall that $O(t)$ is affected by significant uncertainty and delay).

To show the effectiveness of the open-loop FPSP with the duty-cycle tuned over time by the hysteresis-based supervisory outer control law, the following simulation analysis using the SIDARTHE model is reported. The control parameters are $c = 14, \alpha_Y = 0, O(t) = D(t) + R(t) + T(t) + E(t)$. Specifically, we consider two possible scenarios.

- *Scenario 1:* During work days the infection behaves as if there were no social distancing measures and people could behave exactly as if there was no pandemic ($R_0 = 2.38$). We consider $\alpha_X = 0.4$. Results are




Fig. 17: Sensitivity analysis of the quantity $I + D + A + R + T$ in the SIDARTHE model on uncertainty in epidemiological model parameters. Parameters were sampled from zero-truncated Normal distributions with means $\sigma_1, \dots, \sigma_{16}$ and 10% standard deviation. β was estimated from samples of the basic reproduction rate $R_0 \sim \mathcal{N}(\mu = 2.676, \sigma = 0.572)$ representing the consensus distribution in [25].

shown in Figure 18: a rather conservative policy $(X, Y) = (3, 11)$ is reached.

- *Scenario 2:* During work days we assume mild social measures to diminish the effect of the spread of the disease ($R_0 = 1.66$). We consider $\alpha_X = 0.4$. Results are shown in Figure 19: a reasonably non-conservative policy $(X, Y) = (6, 8)$ is reached.

In both scenarios, the controller is able to find a suitable FPSP policy such that the disease is suppressed and the social cost to contain the pandemic is reduced. Of course, as expected, the control action in Scenario 1 leads to a more conservative policy due to the lack of social measures taken during the work days.

Notice that, unlike a control action that regulates a pulsed quarantine period using a threshold feedback (i.e., based on the amount of population being infected), the proposed system fixes a time window c (effectively the frequency of the control action) and slowly varies the duty cycle of the policy in order to have as many working days as possible, while at the same time maintaining the virus suppression. The main difference between the two strategies, lies in the use of a slow varying duty cycle with a fast switching policy instead of a threshold-based one. Intuitively this means that, unlike the threshold-based approach. In our protocol, if the control action is slower than the maximum incubation time of the disease (i.e., if $c \geq 14$) then time delays will not affect the performance of the controller.




Fig. 18: Outer Loop effects for the Scenario 1. The upper panel shows the total amount of infected, the middle panel shows the observed state, $O(t)$, the lower panel shows the control action. Notice that $O(t)$ is delayed with respect to the number of infected people. In this simulation $\alpha_X = 0.4$. The two vertical lines in the upper and middle panel represent, respectively, the beginning of the full lock-down and of the FPSP policy. In the lower panel, we show only the vertical line corresponding to the beginning of the FPSP policy.

As a final important consideration, it is worth noting that the combined open-loop FPSP/closed-loop outer supervisor paves




Fig. 19: Outer Loop effects for the Scenario 2. The upper panel shows the total amount of infected, the middle panel shows the observed state, $O(t)$, the lower panel shows the control action. Notice that $O(t)$ is delayed with respect to the number of infected people. In this simulation $\alpha_X = 0.3$. The two vertical lines in the upper and middle panel represent, respectively, the beginning of the full lock-down and of the FPSP policy. In the lower panel, we show only the vertical line corresponding to the beginning of the FPSP policy.

the way to the future possibility of pursuing an optimization-based design with the aim of minimizing an expected “total cost” $J(t)$, measuring how the control strategy is believed to impact on the epidemic growth and on the society. For example, with reference to the models described in Section II, a possible total cost may be written as

$$\underbrace{J(t)}_{\text{total cost}} = \underbrace{\rho(S(t), I(t))}_{\text{epidemic growth}} + \underbrace{\kappa \cdot C(t)^2}_{\text{societal cost}}$$

in which the “epidemic growth” term $\rho(S(t), I(t))$ measures how bad is the current epidemic state, the “social cost” $C(t)$ measures the cumulative social cost due to the overall lock-down period already imposed, and κ is a free parameter defining the relative importance of the two terms in the sum.

VII. Findings

In this note we consider strategies that may mitigate the effect of Covid-19. Such strategies currently include: (i) complete lock-down for a long duration; (ii) managed strategies in a manner that does not overwhelm the healthcare system. Our findings are as follows.

- (i) Fast switching between two societal modes appears to be an interesting mitigation strategy. These modes are *normal behaviour* and *social isolation*.
- (ii) The fast switching policy may allow a predictable (X days on, Y days off) and continued (albeit reduced) economic activity.

- (iii) Fast switching may suppress the virus propagation, mitigate secondary virus waves, and may be a viable alternative to sustained lock-down (LDP) and timed intervention (TIP) policies.
- (iv) The fast switching policy may be a viable exit strategy from current lock-down policies when the number of infected individuals reduces to a lower level.
- (v) The fast switching policy can be implemented through the aid of a outer loop whose aim is to slowly increase the duty cycle of the policy, given a fixed frequency.

Finally we emphasize that the fast switching policy should not necessarily be viewed as a stand-alone policy, and can also be used to augment and compliment other post lock-down strategies to provide additional levels of robustness in the post lock-down period. For example, it may be worth considering in combination with other strategies such as using contact tracing, face-masks, and reduced social distancing. In combination with these, or as the number of susceptible people decreases, the policy may allow a gradual return to normality over time.

REFERENCES

- [1] Giordano, G., Blanchini, F., Bruno, R., Colaneri, P., Di Filippo, A., Di Matteo, A., Colaneri, M., and the COVID19 IRCCS San Matteo Pavia Task Force, A SIDARTHE Model of COVID-19 Epidemic in Italy, 2020, <https://arxiv.org/abs/2003.09861>.
- [2] Sanfelice, R., A Hybrid Systems Simulation Toolbox for Matlab/Simulink (HyEQ), <https://hybrid.soe.ucsc.edu/software>.
- [3] Di Lauro, F. and Kiss, I. Z and Miller, J., The timing of one-shot interventions for epidemic control, 2020, <https://www.medrxiv.org/content/early/2020/03/06/2020.03.02.20030007>, medRxiv
- [4] Casella, F., Can the COVID-19 epidemic be managed on the basis of daily data?, March 2020, <https://arxiv.org/abs/2003.06967>
- [5] Ferguson, N., et. al., Impact of non-pharmaceutical interventions (NPIs) to reduce COVID- 19 mortality and healthcare demand, <https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/IR>
- [6] Shorten, R., Wirth, F., Mason, O., Wulff, K., and King, C., Stability Criteria for Switched and Hybrid Systems, SIAM Review, 49(4), 545-592, 2007.
- [7] Branicky, M., Multiple Lyapunov functions and other analysis tools for switched and hybrid systems, IEEE Transactions on Automatic Control, vol. 43, no. 4, pp. 475-482, April 1998.
- [8] Anderson, R.M. and May, R.M., Infectious Diseases of Humans. Oxford University Press. Oxford, 1991.
- [9] Pedersen, M.G. and Meneghini, M., Quantifying undetected COVID-19 cases and effects of containment measures in Italy, available online at https://www.researchgate.net/publication/339915690_Quantifying_undetected_COVID-19_cases_and_effects_of_containment_measures_in_Italy
- [10] Peng, L., Yang, W., Zhang, D., Zhuge, C. and Hong, L., Epidemic analysis of COVID-19 in China by dynamical modeling, available online at <https://www.medrxiv.org/content/10.1101/2020.02.16.20023465v1>.
- [11] Hethcote, H., Zhien, M. and Shengbing, L., Effects of quarantine in six endemic models for infectious diseases. Mathematical biosciences, vol. 180, no. 1-2, pp. 141-160, 2002.
- [12] X., Liu and P., Stechliniski, The Switched SIR Model. Infectious Disease Modeling, Nonlinear Systems and Complexity, vol 19. Springer, Cham, 2017.
- [13] L., Zhu, and Y., Zhou, The Dynamics of an SIQS Epidemic Model with Pulse Quarantine. IEEE Chinese Control and Decision Conference, pp. 3546-3551, 2008.
- [14] J., Lai, S., Gao, Y., Liu, and X., Meng, Impulsive Switching Epidemic Model with Benign Worm Defense and Quarantine Strategy. Complexity, 2020.

- [15] Kermack, W. and McKendrick, A., A contribution to the mathematical theory of epidemics, Proceedings of the Royal Society of London, Series A, vol. 115., pp. 700-721, 1927.
- [16] Graunt, J., Natural and Political Observations Made Upon the Bills of Mortality, 1662.
- [17] Wu, J.T., Leung, K., Bushman, M., Kishore, N., Niehus, R., de Salazar, P.M., Cowling, B.J., Lipsitch, M., and Leung, G.M., Estimating clinical severity of COVID-19 from the transmission dynamics in Wuhan, China, Nature Medicine Letters, 2020.
- [18] Stone, L., Olinky, R., and Huppert, A., Seasonal dynamics of recurrent epidemics, Nature Letters, vol. 446, 2007.
- [19] Li, R., Pei, S., Chen, B., Song, Y., Zhang, T., Yang, W., and Shaman, J., Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV2), Science, 10.1126/science.abb3221, 2020.
- [20] Mizumoto, K., and Chowell, G., Transmission potential of the novel coronavirus (COVID-19) onboard the diamond Princess Cruises Ship, 2020, Infectious Disease Modelling, vol. 5, pp. 264-270, 2020.
- [21] Anderson, R.M., Heesterbeek, H., Klinkenberg, D., and Hollingsworth, T.D., How will country-based mitigation measures influence the course of the COVID-19 epidemic, The Lancet, vol. 395, pp. 931-934, 2020.
- [22] Hellewell, J., Abbott, S., Gimma, A., Bosse, N.I., Jarvis, C.I., Russell, T.W., Munday, J.D., Kucharski, A.J., Edmunds, W.J., Centre for the Mathematical Modelling of Infectious Diseases COVID-19 Working Group, Funk, S., and Eggo, R.M., Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts, The Lancet, Vol. 8, No. 4, 2020.
- [23] Flaxman, S., et al., Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries, Imperial College COVID-19 Response Team, 2020.
- [24] Pike, W.T., and Saini, V., An international comparison of the second derivative of COVID-19 deaths after implementation of social distancing measures, 2020.
- [25] N. G. Davies, et al., The effect of non-pharmaceutical interventions on COVID-19 cases, deaths and demand for hospital services in the UK: a modelling study, 1st April 2020, https://cmmid.github.io/topics/covid19/control-measures/report/uk_scenario_modelling_preprint_2020_04_01.pdf
- [26] Liu, X., and Stechlinkski, P., Infectious Disease Modeling - A Hybrid System Approach, Nonlinear Systems and Complexity, Series Editor: Albert C.J. Luo, Springer Nature, 2017.

Acknowledgements:

R.M-S. and S.S. acknowledge funding support from EPSRC grant EP/R018634/1, *Closed-loop Data Science*. M.B. and T.P. acknowledge funding support from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No 739551 (KIOS CoE). H.L. and R.S. acknowledge the support of Science Foundation Ireland.