

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, C. Myant, T. Parisini, R. Shorten

version 1.4, March 27, 2020

Executive Summary: The initial UK government strategy of a timed intervention as a means of combatting 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 note is to enunciate the differences between these policies and to suggest modified policies, for post lockdown, that may allow management of Covid-19, while at the same time enabling some degree of reduced social and economic activity.

Disclaimer : Our results are based on elementary SIR and SIQR models. We are also not epidemiologists. While promising in simulation, more extensive validation is absolutely necessary on realistic Covid-19 models. Our intention in this note 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 1.0 is the base version.

Version 1.1: edited to make it clear that we are speaking about exit strategies from current lockdown.

Version 1.2: edited to make it clear that we are speaking about exit strategies from current lockdown. Figures have been modified to reflect this.

Version 1.3: 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 1.4: more refined numerical results are now included using hybrid systems integration solver [2].

I. Background

Presently, governments worldwide are struggling to contain the Covid-19 epidemic. Most governments, 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

M. Bin is with the Department of Electrical and Electronic Engineering, Imperial College London: m.bin@imperial.ac.uk

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

E. Crisostomi is with the Department of Energy, Systems, Territory and Constructions Engineering, University of Pisa: emanuele.crisostomi@unipi.it

T. Parisini is with the Department of Electrical and Electronic Engineering, Imperial College London and he is also affiliated with the Department of Engineering and Architecture, University of Trieste, Italy: t.parisini@imperial.ac.uk

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

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

a heavy economic cost. For example, 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 in the coming days, 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 interventions must be timed for maximum benefit to the healthcare system by reducing the number of ill 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, and to repeat this policy while waiting for a vaccine or suitable antibody tests to become viable. This, as suggested in the recent report by a team at Imperial College [5], is in fact the basis of the TIP policies. 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 infections 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. 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 is 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.

II. Models

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 number of *susceptible*, *infected*, and *recovered* people in a total population of N individuals. In particular, note that at any instant in time one individual can only

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

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

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]. The parameters β and γ depend on the specific considered disease, and how contagious an individual is early in the epidemic in the absence of intervention and depletion of susceptibles. By definition, $\mathcal{R}_0 = \beta/\gamma$ is the characteristic *basic reproduction number* of a disease [3]. Current estimates of \mathcal{R}_0 in the Covid-19 disease lie in the range between 2 and 4, with most of them between 2.5 and 3 [9].

Many variants have been derived from the original SIR model. While these have been used to include more sophisticated and realistic stochastic, time dependent or spatial aspects of epidemic dynamics, they all maintain, at their core, the SIR dynamics. Among the most popular of these are: (i) the SEIR model [10], in which a further class of exposed people is added to the classic three classes of the SIR model to take into account the infected people who become infectious only after a latent period of time; and (ii) the SIQR model [11] in which a further class of people who are in quarantine is added to the classic three classes of the SIR model.

The SIQR model appears to be particularly convenient to model the Covid-19 disease [9], as it has the advantage of considering two separate states for the infectious people and those who are in quarantine. On the one side, 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). The SIQR model may be described as follows:

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

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, the 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). The results presented in [9] provide estimates of the SIQR model parameters (based on publicly available data) for the Italian region; $\alpha = \eta = 0.067$ (which implies that half of the infectious individuals are asymptomatic, while half show symptoms and get quarantined); $\delta = 0.036$ (which corresponds to an approximate recovery or death of 25 days); and $\beta = 0.373$.

Also, it is noted in [9] that during isolation periods, for a strategy to be effective, β should be decreased by about 65% at least, so that an $\mathcal{R}_0 < 1$ is achieved and the spreading of the virus decreases at least during such isolation periods. However, they also point out that β should be reduced by 90% at least, to reduce the number of infectious individuals in a reasonable timeframe. Following these recommendations, in this work, we assume that β can in fact be reduced by 90% during isolation periods. In the remainder of this note, we shall denote by β^+ the value of β when no lock-down measures are

taken, and by β^- the value of β when social isolation is enforced. While it is hard to evaluate the impact of a given containment strategy on the corresponding value of β , here we shall assume that β^- will be about 10 to 15% of the value of β^+ .

III. The Proposed Policy

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 a **open-loop fast switching strategy** to control and suppress the growth of the virus in society.

The Fast Periodic Switching Policy (FPSP)

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:

$$\beta = \begin{cases} \beta^+ & \text{during inactive lock-down (society functioning as normal)} \\ \beta^- & \text{during lock-down and social isolation} \end{cases} \quad (3)$$

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* $D = 28.6\%$.




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

IV. Illustrative Simulations

To illustrate the performance of the FPSP policy, some preliminary simulations are presented. In what follows, we use the SIQR model described in Section II, with parameters $(\beta, \alpha, \eta, \delta, N) = (0.373, 0.067, 0.067, 0.036, 6 \cdot 10^7)$. Each FPSP policy is characterised by the pair (X, Y) where X denotes the normal working time in days followed by Y lock-down days. An example of such a policy is shown in Figure 2 and compared to a complete lock-down.

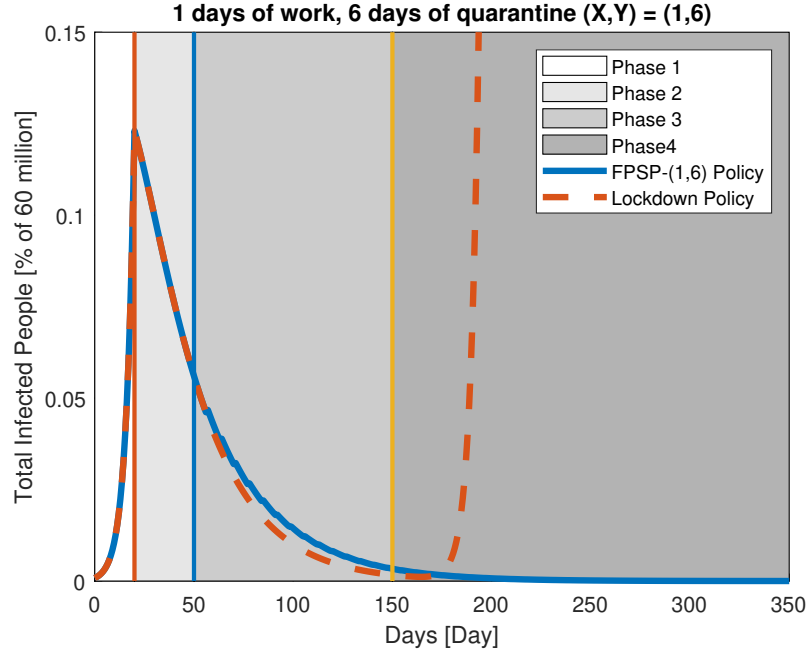


Fig. 2: Time evolution of the quantity $Q + I$. The red line marks the beginning of Phase 2, the blue line marks the beginning of Phase 3 and the yellow line marks the beginning of phase 4. The policy chosen in phase 3 is FPSP-(1,6).

In this scenario the parameters (X, Y) are set to $(2, 6)$. The simulation shows a population of $N = 6 \cdot 10^7$ individuals with an initial amount of infected $I(0) = 500$ individual and it is divided into four phases:

- *Phase 1*: The virus spreads with no containment attempts. This happens for $t < 20$ days.
- *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 $\beta^- = c \cdot \beta^+$ with $c = 0.15$ (which corresponds to the effect of a lock-down). This happens for $20 \leq t < 50$ days.
- *Phase 3*: Once the number of infected people has decreased, the FPSP-(2, 6) policy is enforced and, as it is possible to see from Figure 2, this policy successfully suppresses the virus. This happens for $50 \leq t < 150$ days.
- *Phase 4*: In order to show that the effects of a prolonged quarantine are not sufficient to eliminate the disease, after the virus has been reduced, the lock-down is lifted and the system goes back very fast to the behaviour of phase 1. In contrast the FPSP-(2, 6) policy is able to contain the spread of the virus indefinitely. This happens for $t \geq 150$ days.

The peak level of infections (using the FPSP policy) reached in the third phase of the epidemic 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 (by prolonging the lock-down period, for example), 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. For completeness, in Figure 3 we show the peak levels of the infected individuals for various choices of (X, Y) ranging in the interval $[0, 14]$, after the initial lock-down period has expired (i.e, for $t > 50$, $I(50)$).




Fig. 3: Percentage of peak infections parameterised by (X, Y) in a population of $6 \cdot 10^7$ individuals.

Remark 1: Notice that in Figure 3 the peak values for several choices of (X, Y) are the same. This means that after phase 3 starts, for these values the number of infected does not rise any more but keeps decreasing. This, of course, does not imply an equivalence in terms of dynamic behaviour. Moreover, notice that some choices of (X, Y) might result into an increase of the amount of infected individuals. As an example consider Figure 4, where the pair (X, Y) was set to $(2, 5)$.

Remark 2: We want to stress, again, that the FPSP policy is intended to be used, **exclusively** as an exit strategy.

V. 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.

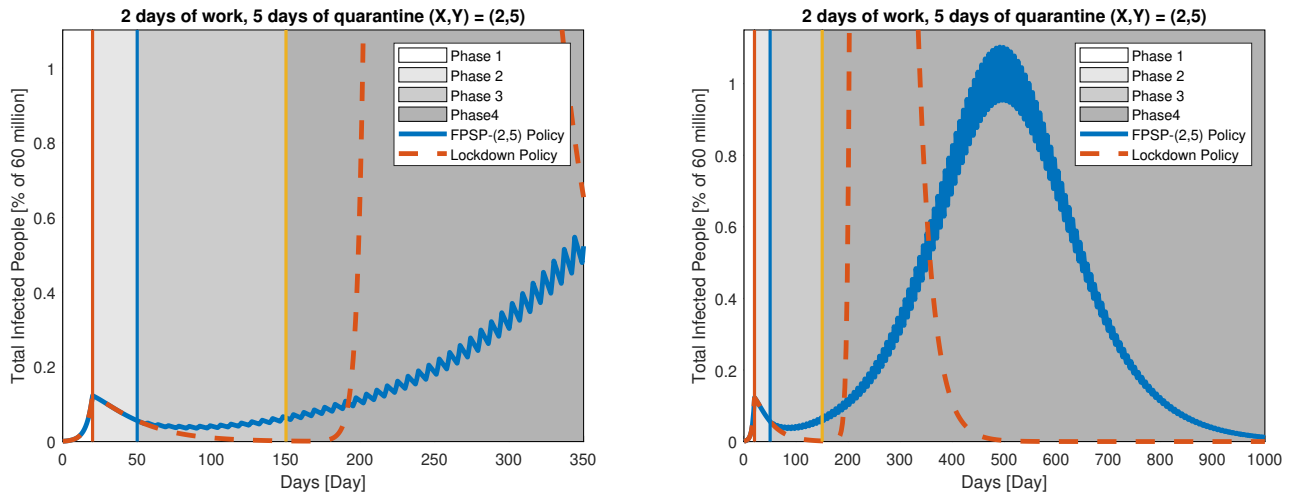


Fig. 4: Time evolution of the quantity $Q + I$. The red line marks the beginning of Phase 2, the blue line marks the beginning of Phase 3 and the yellow line marks the beginning of phase 4. The policy chosen in phase 3 is FPSP-(2, 5). The time scale of the right hand Figure is larger in order to show the entire evolution of the system.

- (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.
- (iv) The fast switching policy may be worth considering in combination with other strategies. For example, 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.

Supporting literature : Since writing version 1.2 of this report, we have become aware of a number of studies that propose similar strategies to the one in this report. Switched, and more particularly, open-loop periodic strategies have also been suggested in the context of other epidemics. In particular, periodic vaccination is suggested in [12], and periodic quarantines for combatting computer worms, and viral epidemics, are suggested in [13], [14]. 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. The approach in this document differs from the ones listed above in that we assume to have limited or no room for intervention on the people that are already infected. Our policy consists in adjusting the infection rate of the disease by introducing a periodic suppression based on switching between essentially the transmission rates of lockdown and not-lockdown. 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 an exit strategy to the current lock-down situation. As we become aware of other related strategies we shall add reference to them.

REFERENCES

- [1] Giordano, C. et al., A SIDARTHE Model of COVID-19 Epidemic in Italy, <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, <https://www.medrxiv.org/content/early/2020/03/06/2020.03.02.20030007>, medRxiv, 2020.
- [4] Casella, F., Can the COVID-19 epidemic be managed on the basis of daily data?, <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/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf>
- [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., Zhen, 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.

VI. Appendix - Future Research

Future research efforts will be devoted to investigate aperiodic switching policies based on suitable outer feedback strategies even though care should be exercised since standard feedback mechanisms may fail dramatically due to the high-level of uncertainty and delay in the measurements [4].

Over a longer time-horizon, different phases and scenarios of the epidemic may suggest an aperiodic switching strategy. For example, during the initial period the virus typically shows an exponential growth hence requiring more severe mitigation actions whereas less aggressive actions may be needed in a terminal phase to mitigate growth. An example of such an aperiodic policy is depicted in Figure 5.

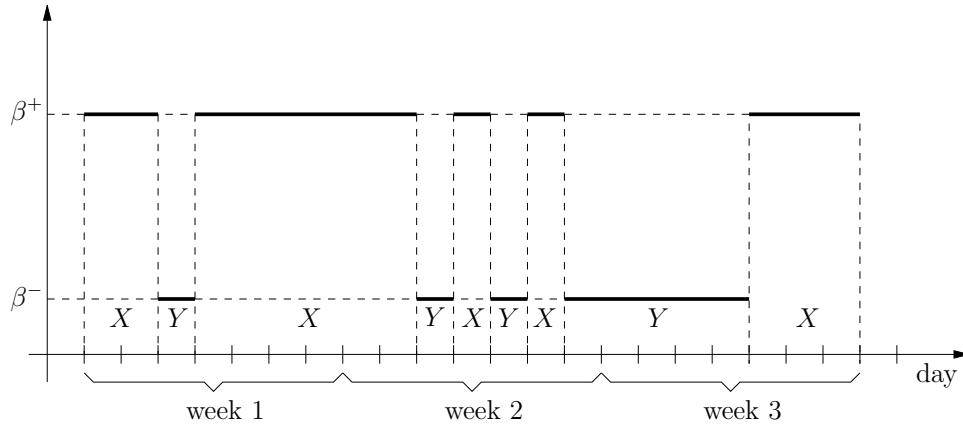


Fig. 5: Example of a FASP policy (normal and lock-down phases are not reported for clarity).

The switching action may be driven by an **optimization policy**. Specifically, with reference to either the SIR model (1) or the SIQR model (2), the a “total cost” $J(t)$ can be introduced aiming at quantifying the impact on the epidemic growth as well as the societal costs:

$$\underbrace{J(t)}_{\text{total cost}} = \underbrace{\rho[S(t), I(t)]}_{\text{epidemic growth}} + \underbrace{\kappa \cdot C(t)^2}_{\text{societal cost}} \quad (4)$$

In cost (4), the “epidemic growth” term $\rho[S(t), I(t)]$ weights the badness of the current epidemic state, whereas the “social cost” $C(t)$ measures the cumulative social cost due to the overall lock-down period already imposed. The design parameter κ quantifies the relative importance of the two terms in the sum

$J(t)$. A switching policy that aims at minimizing the cost (4) can be devised. Clearly, this policy would be a feedback one since it would depend on the current values of the state variables $S(t), I(t)$. Early theoretical and simulation results are very promising and the authors are currently devoting efforts to validate this feedback switching policy.