



HAL
open science

Pandemic economics: Optimal dynamic confinement under uncertainty and learning

Christian Gollier

► **To cite this version:**

Christian Gollier. Pandemic economics: Optimal dynamic confinement under uncertainty and learning. Covid Economics Papers, 2020, 34, pp.1-14. hal-02909470

HAL Id: hal-02909470

<https://hal.science/hal-02909470>

Submitted on 16 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Pandemic economics: Optimal dynamic confinement under uncertainty and learning¹

Christian Gollier²

Date submitted: 1 July 2020; Date accepted: 2 July 2020

Most integrated models of the Covid pandemic have been developed under the assumption that the policy-sensitive reproduction number is certain. The decision to exit from the lockdown has been made in most countries without knowing the reproduction number that would prevail after the deconfinement. In this paper, I explore the role of uncertainty and learning on the optimal dynamic lockdown policy. I limit the analysis to suppression strategies. In the absence of uncertainty, the optimal confinement policy is to impose a constant rate of lockdown until the suppression of the virus in the population. I show that introducing uncertainty about the reproduction number of deconfined people reduces the optimal initial rate of confinement.

Covid Economics 34, 3 July 2020: 1-14

- 1 I thank Stefan Pollinger and Daniel Spiro for helpful comments. The research leading to these results has received the support from the ANR grants Covid-Metrics and ANR-17-EURE-0010 (Investissements d'Avenir program).
- 2 Toulouse School of Economics, University of Toulouse-Capitole.

Copyright: Christian Gollier

1 Introduction

Academic economists have recently spent a huge amount of energy to better understand the science of pandemic dynamics in the face of the emergence of the covid-19. Economists are contributing to the analysis of the covid-19 crisis by integrating economic dimensions to the models, such as the economic cost of social distancing and the statistical value of lives lost. These are key elements necessary for public and private decision-makers interested in shaping strategies and policies that minimize the welfare cost of the crisis. My preferred reading list on this issue as I write this paper is composed of papers by Acemoglu, Chernozhukov, Werning and Whinston (2020), Alvarez, Argente and Lippi (2020), Brotherhood, Kircher, Santos and Tertilt (2020), Favero, Ichino and Rustichini (2020), Fischer (2020), Greenstone and Nigam (2020), Miclo, Spiro and Weibull (2020), Pindyck (2020) and Pollinger (2020). This investment by the profession is impressive and highly policy-relevant. It raised critical debates about, for example, when and how much to deconfine people, who should remain confined longer, the value of testing and tracing, or whether the individual freedom of movement should be limited.

One of the most striking feature of the crisis is the deep uncertainties that surrounded most parameters of the model at the initial stage of the pandemic. To illustrate, here is a short list of the sources of covid-19 uncertainties: The mortality rate, the rate of asymptomatic sick people, the rate of prevalence, the duration of immunity, the impact of various policies (lockdown, social distancing, compulsory masks,...) on the reproduction numbers, the proportion of people who could telework efficiently, and the possibility of cross-immunization from similar viruses. Still, all models that have been built over such a short period of time by economists assumed no parameter uncertainty, and I am not an exception (Gollier, 2020). This is amazing. Large discrepancies between the predictions of these models and their associated "optimal" policies do not illustrate deep disagreements about the dynamics of the pandemic, but rather deep uncertainties about the true values of its parameters. This parameter uncertainty should be recognized and integrated in the modeling. Economists are well aware that uncertainty is typically a key component to explain observed behaviors and to shape efficient policies. Precautionary savings, option value to wait before investing, risk premia on financial markets, insurance demand, risk-sharing and solidarity mechanisms, and preventive efforts are obvious examples to demonstrate that risk and uncertainty are at the heart of the functioning of our society. But in the cases of climate change and covid-19, we most often assume no uncertainty to make policy recommendations in spite of the fact that uncertainty is everywhere in these contexts. I feel this fact as an impressive failure of our profession to be useful to make the world better.

In this paper, I go one step towards including risk in the shaping of efficient pandemic policies. Suppose that a virus has contaminated a small fraction of the population, and that no treatment or vaccine is available. Because of the high lethality of the virus, I suppose that the only feasible strategy is to "crush the (infection) curve" by imposing a partial lockdown. The intensity of the confinement can be adapted in continuous-time to the evolution of the pandemic in order to minimize the total cost of the confinement. Following Pollinger (2020), I show that in the absence of uncertainty, the optimal intensity of the lockdown should be constant over time until the eradication of the virus in the population. The optimal confinement intensity is the best compromise between the short-term cost of increasing the confinement and the long-term benefit of reducing the duration of the confinement. Confining

people modifies the reproduction number. Under the standard SIR pandemic model (Kermack and McKendrick, 1927), there is a quadratic relation between the instantaneous intensity of the confinement and the instantaneous reproduction number.

Consider the situation prevailing in the western world in April 2020, after a partial lockdown was imposed. In this context, suppose that the reproduction number under full lockdown is known, but the reproduction number under full deconfinement is uncertain. This uncertainty will evaporate within a few weeks by observing the propagation of the virus under the partial lockdown. How should this uncertainty with learning affect the initial intensity of the lockdown? Surprisingly, I show that it tends to reduce it. To obtain this result, I assume that the representative agent is risk-neutral. However, risk plays a role in this model because of two non-linear interactions: the quadratic relation between the cost of confinement and the instantaneous reproduction number, and the hyperbolic relation between the reproduction number and the duration of the pandemic. This double non-linearity makes the analysis quite complex, and I have been able to prove the main result only in the case of small risk. The calibration exercise suggests that my result holds for large risks too.

There is a long tradition in decision theory and finance on optimal choice under uncertainty and learning to which this paper is related. It is closest to the literature on the option value to wait introduced by McDonald and Siegel (1984) and popularized by Dixit and Pindyck (1994). An important message from this literature is that risk-neutral agents could optimally reduce their initial effort to achieve a long-term goal in order to obtain additional information about the welfare impact of this effort. I obtain a similar result in this pandemic model.

2 The model

My model is based on the classical SIR model developed by Kermack and McKendrick (1927) to describe the dynamics of a pandemic. Each person is either Susceptible, Infected or Recovered, i.e., the health status of a person belongs to $\{S, I, R\}$. This implies that $S_t + I_t + R_t = N$ at all dates $t \geq 0$. A susceptible person can be infected by meeting an infected person. Following the key assumption of all SIR models, this number of new infections is assumed to be proportional to the product of the densities of infected and susceptible persons in the population, weighted by the intensity of their social interaction. With no further justification, this is quantified as follows:

$$\frac{dS_t}{dt} = -\beta_t I_t S_t. \quad (1)$$

I will soon describe how β_t , which measures the intensity of the risk of contagion of a susceptible person by an infected person at date t , is related to the social interactions between these two groups and by the confinement policy. Once infected, a person quits this health state at rate γ , so that the dynamics of the infection satisfies the following equations:

$$\frac{dI_t}{dt} = \beta_t I_t S_t - \gamma I_t. \quad (2)$$

$$\frac{dR_t}{dt} = \gamma I_t \quad (3)$$

The pandemic starts at date $t = 0$ with I_0 infected persons and $N - I_0$ susceptible persons. I assume that the virus is eradicated when the number I_t of infected persons goes below I_{min} , in which case an aggressive tracing-and-testing strategy is implemented to eliminate the last clusters of the epidemic.

Because on average an infected person remains contagious for a duration $1/\gamma$, and because the instantaneous number of susceptible persons infected by a sick person is $\beta_t S_t$, the reproduction number at date t equals

$$r_t = \frac{\beta_t S_t}{\gamma}. \quad (4)$$

Herd immunity is obtained when the number of infected persons start to decrease over time. From equation (2), this is obtained when the number of susceptible persons goes below the herd immunity threshold $S^* = \gamma/\beta_t$, i.e., when the reproduction number goes below 1. In this paper, I focus on policies aimed at "crushing the curve", where r_t remains permanently below unity. Other policies, such as the laissez-faire policy or policies aimed at "flattening the curve", consist in building herd immunity through a rapid or gradual infection of a large fraction of the population, implying a large health cost but a limited economic cost. When crushing the curve, a sufficiently strong confinement is imposed to the population to maintain the reproduction number permanently below 1, so that the virus is eradicated relatively quickly. Under this family of scenarios, the number S_t of susceptible persons remain close to unity, very far from herd immunity under the laissez-faire policy. This implies that the changes in $I_t S_t$ in equation (2) mostly comes from changes in I_t . Following Pollinger (2020), I therefore simplifies the SIR dynamic described above into a single differential equation:

$$\frac{dI_t}{dt} = (\beta_t I_t S - \gamma) I_t, \quad (5)$$

where S is the average number of susceptible persons during the pandemic. This approximation of the SIR model is exact when the ratio of infected to susceptible is close to zero.

I examine policies of social distancing and lockdown. Let x denote the intensity of this policy. One can interpret x as a measure of the fraction of people that are confined. For simplicity, I assume that infected people are asymptomatic and that there is no PCR test, so that one cannot discriminate the intensity of confinement on the basis of the health status. This means that x is the fraction of people, both infected or susceptible, who are confined. A free infected person has a reproduction number $r_f = \beta_f S/\gamma$. I assume that there is no herd immunity at the start of the pandemic, i.e., that r_f is larger than unity, or $\beta_f S > \gamma$. The confinement reduces this number to $r_c = \beta_c S/\gamma$, with $\beta_c \leq \beta_f$. I assume that the full confinement of the population crushes the curve in the sense that $r_c < 1$, or $\beta_c S \leq \gamma$.

As said earlier, a crucial element of the SIR model is that the speed of infection is proportional to the product of the numbers of people infected and susceptible. Confining people reduces both the number of infected people and the number of susceptible persons, implying a quadratic relation between the intensity x of the confinement and propagation of the virus in the population (Acemoglu, Chernozhukov, Werning and Whinston (2020)). From this observation, the pandemic parameter β_t takes the following form:

$$\beta_t = \beta(x_t) = (\beta_c x_t + \beta_f (1 - x_t))(1 - x_t). \quad (6)$$

The true contagion rate $\beta_c x_t + \beta_f(1 - x_t)$ of infected people is a weighted average of the contagion rates β_c and β_f of infected people who are respectively confined and let free to live their life. They meet a reduced fraction $1 - x$ of susceptible people, because the remaining fraction x is lockdown. The quadratic nature of this relation plays a crucial role in this paper. The lockdown has also an economic cost. I assume that the instantaneous cost of confining a fraction x of the population at date t is equal to $w x$, where $w > 0$ can be interpreted as the sum of the wage and psychological costs of confinement. Abstracting from discounting given the short duration of the pandemic when crushing the curve, the objective of the policy is to minimize the total cost of the health crisis. This yields the following value function:

$$V(I) = \min_{x(\cdot)} w \int_0^T x(t) dt \quad \text{s.t. } I_0 = I \text{ and } I_T = I_{min}, \tag{7}$$

where I is the current rate of prevalence of the virus in the population. The termination date corresponds to the time when the rate of prevalence of the virus attains the eradication threshold I_{min} . Observe that I assume an objective that ignores the potential lethality of the virus. But even when the virus is lethal, policies aimed at crushing the curve typically yields economic costs that are at least one order of magnitude larger than the value of lives lost (Gollier (2020)), thereby justifying this objective of minimizing costs.

3 Optimal suppression under certainty

Pollinger (2020) derives the solution of a more general version of this dynamic problem under certainty. Using backward induction, problem (7) can be rewritten as follows:

$$\begin{aligned} V(I) &= \min_x wx\Delta t + V(I + (\beta(x)S - \gamma)I\Delta t) \\ &\approx \min_x wx\Delta t + V(I) + (\beta(x)S - \gamma)IV'(I)\Delta t, \end{aligned}$$

or, equivalently,

$$0 = \min_x wx + (\beta(x)S - \gamma)IV'(I). \tag{8}$$

The first-order condition of this problem is

$$w = -\beta_x(x^*)SIV'(I), \tag{9}$$

Under this notation, β_x is the derivative of β with respect to x . Equation (9) expresses the optimal intensity $x^*(I)$ of confinement as a function of the rate of prevalence of the virus. However, let us guess a constant solution x^* independent of I . From equation (9), this would be the case if $IV'(I)$ is a constant. In that case, the duration T of the pandemic will be such that

$$I_{min} = I \exp((\beta(x^*)S - \gamma)T). \tag{10}$$

This equation tells us that there is an hyperbolic relation between the reproduction number and the duration of the pandemic. The total cost under such a constant strategy is

$$V(I) = wx^*T = \frac{-wx^*}{\beta(x^*)S - \gamma} \ln\left(\frac{I}{I_{min}}\right). \tag{11}$$

This implies that $IV'(I)$ is a constant, thereby confirming the guess that it is optimal to maintain a constant intensity of lockdown until the eradication of the virus. Combining equations (9) and (11) yields the following optimality condition for x^* :

$$x^* = \frac{\beta(x^*)S - \gamma}{\beta_x(x^*)S}. \tag{12}$$

The optimal intensity of lockdown is a best compromise between the short-term benefit of easing the lockdown and the long-term cost of a longer duration of the pandemic. Under the quadratic specification (6) for beta, equation (9) simplifies to

$$x^* = \sqrt{\frac{\beta_f S - \gamma}{\beta_f S - \beta_c S}} = \sqrt{\frac{r_f - 1}{r_f - r_c}}. \tag{13}$$

Because $r_c < 1 < r_f$, the optimal intensity of confinement is between 0 and 1. For example, if the reproduction number goes from 2 to 0.5 when moving from the laissez-faire to the 100% lockdown, the optimal intensity of confinement is $\sqrt{2/3} = 81\%$. I summarize my results under certainty in the following proposition. Its first part is a special case of Pollinger (2020).

Proposition 1. *Under certainty, the optimal suppression strategy is to impose a constant intensity of confinement until the virus is eradicated. In the quadratic case (6), the optimal intensity of confinement is $\sqrt{(r_f - 1)/(r_f - r_c)}$, where r_f and r_c are the reproduction numbers under respectively the laissez-faire and the full lockdown.*

4 Optimal suppression under uncertainty

Suppose that some parameters of the pandemic are unknown at date 0. Suppose also that the only way to learn the true value of these parameters is to observe its dynamics over time. How should this parameter uncertainty affect the optimal effort to fight the virus in the population? I have not been able to solve the continuous-time version of this dynamic learning problem. I therefore simplified the problem as follows. I assume that parameter β_f is unknown. At date 0, a decision must be made for an intensity x_0 of confinement under uncertainty about β_f . This intensity of confinement will be maintained until date τ .¹ Between dates 0 and τ , the observation of the propagation of the virus will inform us about β_f . Therefore, at date τ , β_f is known and the intensity of confinement is adapted to the information. My objective is to compare the optimal x_0 under uncertainty to the x_0 that would be optimal when ignoring the fact that β_f is uncertain.

This is thus a two-stage optimization problem that I solve by backward induction. From date τ on, there is no more uncertainty. As observed in the previous section, it is optimal to revise the confinement policy to the information about the true β_f . We know from the previous section that the optimal contingent policy $x^*(\beta_f)$ is constant until the eradication of the virus. The minimal total cost of this policy is denoted $V(I_\tau, \beta_f)$. Combining equations (11) and (12), it is equal to

$$V(I_\tau, \beta_f) = \frac{-w}{\beta_x(x^*(\beta_f))S} \ln \left(\frac{I_\tau}{I_{min}} \right). \tag{14}$$

¹I assume that τ is small enough so that I_τ is larger than I_{min} with probability 1.

It is a function of the rate of prevalence I_τ of the virus observed at date τ and of the pandemic parameter β_f observed during the first stage of the pandemic.

The first stage of the pandemic takes place under uncertainty about β_f . I assume risk neutrality, so that the objective is to minimize the expected total cost of the suppression strategy:

$$W_0 = \min_{x_0} wx_0\tau + EV(I_\tau, \beta_f), \tag{15}$$

where $I_\tau = I_0 \exp((\beta(x_0, \beta_f)S - \gamma)\tau)$ is also a function of random variable β_f . The first-order condition of this stage-1 problem can be written as follows:

$$E [F(x_0^*, \beta_f)] = 1, \tag{16}$$

with

$$F(x_0, \beta_f) = \frac{\beta_x(x_0, \beta_f)}{\beta_x(x^*(\beta_f), \beta_f)}. \tag{17}$$

In the absence of uncertainty, i.e., when β_f takes value β_{f0} with probability 1, the optimal solution is the solution of equation (16) in that particular case, which implies

$$x_0^* = x^*(\beta_{f0}). \tag{18}$$

How does the uncertainty and learning about β_f affect the optimal effort to mitigate the pandemic? Because β is a convex function of the mitigation effort x , function F is increasing in x_0 . By Jensen's inequality, equation (16) implies that the uncertainty affecting β_f reduces the optimal initial mitigation effort if and only if F is convex in its second argument. I have not been able to demonstrate a general result of this nature. I therefore limited my analysis to the case of a small risk surrounding β_f . More precisely, suppose that β_f is distributed as $\beta_{f0} + h\epsilon$, where β_{f0} is a known constant, ϵ is a zero-mean random variable and h is an uncertainty-intensity parameter. I examine the sensitivity of the optimal confinement x_0^* as a function of the intensity h in the neighborhood of $h = 0$. In the Appendix, I demonstrate that F is locally convex in its second argument, i.e., that $x_0^*(h)$ is decreasing in h in the neighborhood of $h = 0$. More precisely, I show that $x_0^{*'}(0) = 0$ and $x_0^{*''}(0) < 0$. This yields the following main result of the paper.

Proposition 2. *Consider the quadratic case (6). Introducing a small risk about the transmission rate β_f reduces the optimal initial intensity of confinement.*

Proof: See Appendix.

5 Calibration exercise

In this section, I quantify the negative impact of uncertainty on the optimal confinement in the learning stage 1. I solve numerically the optimality condition (16) in the quadratic context. This equation takes the following form in that case:

$$E \left[\frac{(2\beta_f - \beta_c)S - 2(\beta_f - \beta_c)Sx_0^*}{(2\beta_f - \beta_c)S - 2\sqrt{(\beta_f - \beta_c)S(\beta_f S - \gamma)}} \right] = 1 \tag{19}$$




Figure 1: Optimal confinement x_0^* in stage 1 as a function of the intensity h of the uncertainty. I assume that $r_c = 0.5$ and $r_f = 1.5 + h\epsilon$, with $\epsilon \sim (-1, \pi; \pi/(1 - \pi), 1 - \pi)$.

It yields the following solution:

$$x_0^* = \frac{E \left[\frac{\sqrt{(r_f - r_c)(r_f - 1)}}{2r_f - r_c - 2\sqrt{(r_f - r_c)(r_f - 1)}} \right]}{E \left[\frac{r_f - r_c}{2r_f - r_c - 2\sqrt{(r_f - r_c)(r_f - 1)}} \right]}, \tag{20}$$

where $r_f = \beta_f S / \gamma$ and $r_c = \beta_c S / \gamma$ are the reproduction numbers in the laissez-faire and total lockdown respectively. I first describe a simulation in the spirit of the covid-19. There has been much debate about the reproduction number under the laissez-faire policy. Ferguson et al. (2020) assumed that it was between 2 and 2.6 at the beginning of the pandemic. However, I focus in this paper on a post-lockdown situation in which people have learned the benefit of washing hands, bearing masks and basic social distancing behaviors. Therefore, the expected reproduction number under the laissez-faire in this new situation is probably smaller than 2. I assume an expected value of $E r_f = 1.5$. For France, Santé Publique France² has estimated the reproduction number at different stages of the pandemic. It was estimated at 0.8 at the end of the strong confinement period in May. Because the confinement was partial, this observation is compatible with a r_c equaling 0.5.

In Figure 1, I describe the optimal intensity x_0^* in stage 1 as a function of the intensity h of the uncertainty surrounding r_f , with $r_f = 1.5 + h\epsilon$, with $E\epsilon = 0$. More specifically, I consider binary distribution with $\epsilon \sim (-1, \pi; \pi/(1 - \pi), 1 - \pi)$. In order to keep r_f above 1 with probability 1, I consider risk intensities h between 0 and 0.5. Under certainty ($r_f = 1.5$ with certainty, or $h = 0$), the optimal intensity of confinement is a constant $\sqrt{0.5} = 70.7\%$.

²<https://www.santepubliquefrance.fr/maladies-et-traumatismes/maladies-et-infections-respiratoires/infection-a-coronavirus/documents/bulletin-national/covid-19-point-epidemiologique-du-11-juin-2020>




Figure 2: Percentage reduction in the optimal confinement x_0^* in stage 1 due to uncertainty for different values of (r_c, \bar{r}_f) . I assume that r_f is distributed as $(1, 1/2; 2\bar{r}_f - 1, 1/2)$.

Suppose alternatively that r_f is either 1 or 2 with equal probabilities. In that case, the optimal confinement goes down to 66.2%. If our beliefs about the reproduction number r_f are distributed as 1 with probability 0.9 and 6 with probability 0.1, then the optimal initial confinement goes down to 61.4%.

In Figure 2, I describe the percentage reduction in the optimal initial confinement for different r_c and $r_f \sim (1, 1/2; 2\bar{r}_f - 1, 1/2)$. We see that the impact of uncertainty on the optimal confinement is largest when the reproduction numbers in the pre- and post-confinement are close to unity. Suppose for example that $r_c = 0.9$ and $r_f = 1.1$. In this context of certainty, the optimal confinement is 70.7%. If r_f is distributed as $(1, 1/2; 1.2, 1/2)$, the optimal initial confinement goes down to 34.7%, a 51% reduction in the initial mitigation effort.

6 Concluding remarks

The uncertainty surrounding the reproduction number when reducing the strength of the lockdown is an argument in favor of lowering the intensity of this lockdown in the learning phase of the pandemic. This rather surprising result is the outcome of two non-linearities of the model. First, the duration of the pandemic is an hyperbolic function of the reproduction number. Second, the reproduction number is a quadratic function of the cost of confinement. These two non-linearities explain why one should be sensitive to the uncertainty when shaping the confinement policy, but I confess that these observations do not explain why this uncertainty should reduce the optimal confinement at the first stage of the pandemic. More work should be done to explain this result.

This research opens a new agenda of research that I am glad to share with the readers of this paper. For example, shame on me, I assume here risk neutrality, in spite of the large size

of the risk and its correlation with aggregate consumption. Could there be a precautionary motive for a larger initial intensity of the confinement? No doubt that my result should be refined in that direction. Also, I limited the analysis to suppression policies. This restriction was necessary to simplify the dynamic equations of the generic SIR model, so that the assumption of an almost constant number of susceptible people in the population is a reasonable approximation. This excludes the possibility to compare the optimal solution among this family of policies to other plausible policies, in particular policies aimed at attaining a high rate of herd immunity. Introducing uncertainty in the generic SIR model and measuring its impact on the optimal policy is another promising and useful road for research. In my to-do list, I also have the exploration of other sources of uncertainty, such as not knowing the rate of prevalence, the fraction of the population already immunized, or the time of arrival of a vaccine. Finally, because the value of lives lost associated to most suppression strategies is typically one or two orders of magnitude smaller than the direct economic cost of the lockdown, I assumed that the objective of the social planner is to minimize the economic cost incurred to eradicate the virus in the population. It would be useful, as in Pollinger (2020), to incorporate the value of lives lost in the objective function.

Bibliography

Acemoglu, D., V. Chernozhukov, Ivan Werning and Michael Whinston, (2020), A multi-risk SIR model with optimally targeted lockdown, NBER WP 27102.

Alvarez, F., D. Argente and F. Lippi, (2020), A simple planning problem for COVID-19 lockdown, *CEPR Covid Economics* 14, 1-32.

Brotherhood, L., P. Kircher, C. Santos and M. Tertilt, (2020), An economic model of covid-19 epidemic: The importance of testing and age-specific policies, mimeo.

Dixit, A.K., and R.S. Pindyck, (1994), *Investment under uncertainty*, Princeton University Press, Princeton.

Favero, C., A. Ichino and A. Rustichini, (2020), Restarting the economy while saving lives under covid-19, WP Bocconi University.

Ferguson, N.M. et al., (2020), Impact of non-pharmaceutical interventions (npis) to reduce covid-19 mortality and healthcare demand, *CEPR Covid Economics* 2, 60-66.

Fischer, C., (2020), External costs and benefits of policies to address COVID-19, mimeo.

Gollier, C., (2020), Cost-benefit analysis of age-specific deconfinement strategies, *Covid Economics* 24, 1-31.

Greenstone, M., and V. Nigam, (2020), Does social distancing matter?, *Covid Economics* 7, 1-22 (May 2020).

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

McDonald, R. and D. Siegel, (1984), The value of waiting to invest, *Quarterly Journal of Economics*, 101, 707-728.

Miclo, L., D. Spiro and J. Weibull, (2020), Optimal epidemic suppression under an ICU constraint, mimeo, TSE.

Pindyck, R.S., (2020), Covid-19 and the welfare effects of reducing contagion, mimeo, MIT.

Pollinger, S., (2020), Optimal tracing and social distancing policies to suppress COVID-19, mimeo, TSE.

Salje, H., C. Tran Kiem, N. Lefrancq, N. Courtejoie, P. Bosetti, et al., (2020), Estimating the burden of SARS-CoV-2 in France. <https://hal-pasteur.archives-ouvertes.fr/pasteur-02548181>

Appendix: Proof of Proposition 2

In the quadratic case (6), we have that

$$\beta(x)S - \gamma = (\beta_c x + \beta_f(1-x))(1-x)S - \gamma = ax^2 - bx + c,$$

with

$$\begin{aligned} a &= (\beta_f - \beta_c)S > 0 \\ b &= (2\beta_f - \beta_c)S > 0 \\ c &= \beta_f S - \gamma > 0 \end{aligned}$$

Remember that we assume that $\beta_f S > \gamma$ and that $\beta_c S < \gamma$, so the signs of coefficients (a, b, c) , which are functions of β_f . This also implies that $\beta(x)S - \gamma$ alternates in sign, implying $b^2 - 4ac > 0$. We have that

$$x^*(\beta_f) = \sqrt{\frac{c}{a}} = \sqrt{\frac{\beta_f S - \gamma}{(\beta_f - \beta_c)S}}$$

Observe that $\beta_c S < \gamma$ implies that the optimal stage-2 confinement is smaller than unity. Stage-1 optimality condition (16) is now rewritten as follows:

$$E \left[\frac{2ax_0^* - b}{2\sqrt{ac} - b} \right] = 1. \tag{21}$$

As stated in the main part of the paper, let me parametrize the uncertainty by assuming that β_f is distributed as $\beta_{f0} + h\epsilon$, where β_{f0} is a known constant, ϵ is a zero-mean random variable and h is a measure of the uncertainty. The optimal stage-1 confinement is a function of h , and is denoted $x_0^*(h)$. I examine the properties of this function in the neighborhood of $h = 0$. When h equals zero, the above equation is solved with

$$x_0^*(0) = x^*(\beta_{f0}) = \sqrt{\frac{\beta_{f0}S - \gamma}{(\beta_{f0} - \beta_c)S}}.$$

I now estimate $x_0^{*'} = \partial x_0^*/\partial h$. To do this, I fully differentiate the optimality condition (21) with respect to h , taking account of the fact that (a, b, c) are functions of $\beta_f = \beta_{f0} + h\epsilon$. Let d be equal to ac . I obtain

$$E \left[\frac{(2a'x_0^* - b')\epsilon + 2ax_0^{*'}}{2\sqrt{d} - b} - \frac{(2ax_0^* - b)\epsilon}{(2\sqrt{d} - b)^2} \left(\frac{d'}{\sqrt{d}} - b' \right) \right] = 0. \tag{22}$$

When h equals zero, coefficients a, b, c and d are constant. Because $E\epsilon$ equals zero, the above equation has a single solution

$$x_0^{*'}(0) = 0 \tag{23}$$

when evaluating it at $h = 0$. At the margin, introducing a zero mean risk for the reproduction number has no effect on the optimal mitigation effort in stage 1.

Let me now turn to $x_0^{*''} = \partial^2 x_0^* / \partial h^2$. Let me specifically evaluate this second derivative at $h = 0$. Fully differentiating equation (22) with respect to h , and using property (23) yields

$$0 = \frac{(2a''x_0^* - b'')\sigma^2 + 2ax_0^{*''}}{2\sqrt{d} - b} - 2\frac{(2a'x_0^* - b')\sigma^2}{(2\sqrt{d} - b)^2} \left(\frac{d'}{\sqrt{d}} - b'\right) - \frac{(2ax_0^* - b)\sigma^2}{(2\sqrt{d} - b)^2} \left(\frac{d''}{\sqrt{d}} - \frac{1}{2}\frac{d'^2}{d^{\frac{3}{2}}} - b''\right) + 2\frac{(2ax_0^* - b)\sigma^2}{(2\sqrt{d} - b)^3} \left(\frac{d'}{\sqrt{d}} - b'\right)^2,$$

where $\sigma^2 = E\epsilon^2$ is the variance of ϵ . This is equivalent to

$$2a\sigma^{-2}x_0^{*''}(0) = -\left(2a''\sqrt{\frac{c}{a}} - b''\right) + 2\frac{(2a'\sqrt{\frac{c}{a}} - b')}{2\sqrt{d} - b} \left(\frac{d'}{\sqrt{d}} - b'\right) + \left(\frac{d''}{\sqrt{d}} - \frac{1}{2}\frac{d'^2}{d^{\frac{3}{2}}} - b''\right) - \frac{2}{2\sqrt{d} - b} \left(\frac{d'}{\sqrt{d}} - b'\right)^2.$$

We have that $a' = S$, $b' = 2S$, $d' = S(b - \gamma)$ and $d'' = 2S^2$. This allows me to rewrite the above equation as

$$\frac{2a\sigma^{-2}}{S^2}x_0^{*''}(0) = \frac{4\left(\sqrt{\frac{c}{a}} - 1\right)}{2ac - b\sqrt{ac}} \left((b - \gamma) - 2\sqrt{ac}\right) + \frac{2ac - \frac{1}{2}(b - \gamma)^2}{(ac)^{3/2}} - \frac{2}{2(ac)^{3/2} - abc} \left((b - \gamma) - 2\sqrt{ac}\right)^2.$$

Because $b^2 - 4ac$ is positive, we obtain that $x_0^{*''}(0)$ is negative if and only if

$$\left(3(b - \gamma)^2 - 4c(b - \gamma) - 4ac\right) (ac)^{1/2} - 2abc + 4ac\gamma + 8ac^2 \leq \frac{1}{2}b(b - \gamma)^2, \quad (24)$$

or, equivalently,

$$\left(3S^2\beta_c^2 - 2S(2S\beta_f + \gamma)\beta_c + (4\gamma S\beta_f - \gamma^2)\right)^2 (\beta_f(\beta_f S - \gamma) - (\beta_f S - \gamma)\beta_c) - S\left(-\frac{1}{2}S^2\beta_c^3 + (5S^2\beta_f - 3S\gamma)\beta_c^2 + \left(-4S^2\beta_f^2 - 2\gamma S\beta_f + \frac{7}{2}\gamma^2\right)\beta_c + 4S\gamma\beta_f^2 - 3\gamma^2\beta_f\right)^2 \leq 0.$$

Let me use the following notation:

$$v = (\beta_c S / \gamma) - 1 \\ z = \beta_f S / \gamma.$$

After tedious manipulations, the above inequality is true if and only if

$$H(v, z) = 4(1 - z)^2 + v(8.75 - 13z + 4z^2) + v^2(4.5 - 4z) - 0.25v^3 \quad (25)$$

is positive in the relevant domain of (v, z) , i.e., $v \in [-1, 0]$ and $z \geq 1$. Notice that H is clearly non-negative at the boundaries of the relevant domain:

$$H(0, z) = 4(1 - z)^2 \geq 0 \\ H(-1, z) = z \geq 0$$

$$H(v, 1) = 0.25v(1 - v)^2 \geq 0$$

$$\lim_{z \rightarrow +\infty} H(v, z) = \lim_{z \rightarrow +\infty} 4(1 + v)z^2 = +\infty.$$

More generally, H is non-negative in the relevant domain $D = \{(v, z) \mid (v, z) \in [-1, 0] \times [1, +\infty[)\}$. This implies that $x_0^{*''}(0)$ is negative, or that $x_0^*(h)$ is smaller than $x_0^*(0)$ in the neighborhood of $h = 0$. In other words, any small zero-mean risk surrounding β_f reduces the optimal confinement at stage 1. This concludes the proof of Proposition 2. ■