



Epidemic Models for Personalised COVID-19 Isolation and Exit Policies Using Clinical Risk Predictions

Theodoros Evgeniou, Mathilde Fekom, Anton Ovchinnikov, Raphaël Porcher, Camille Pouchol, Nicolas Vayatis

► To cite this version:

Theodoros Evgeniou, Mathilde Fekom, Anton Ovchinnikov, Raphaël Porcher, Camille Pouchol, et al.. Epidemic Models for Personalised COVID-19 Isolation and Exit Policies Using Clinical Risk Predictions. 2022, Production and Operations Management, Special Issue on Managing Pandemics: A POM Perspective, 10.2139/ssrn.3588401 . hal-02559446v2

HAL Id: hal-02559446

<https://hal.science/hal-02559446v2>

Submitted on 8 May 2020

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.

Epidemic Models for Personalised COVID-19 Isolation and Exit Policies Using Clinical Risk Predictions

Theodoros Evgeniou^{a,*}, Mathilde Fekom^b, Anton Ovchinnikov^c, Raphael Porcher^d, Camille Pouchol^e, Nicolas Vayatis^f

^aINSEAD, Bd de Constance, 77300 Fontainebleau, France

^bUniversit Paris-Saclay, ENS Paris-Saclay, CNRS, Centre Borelli, 94235 Cachan, France

^cSmith School of Business, Queen's University, Kingston, ON, K7L3N6, Canada

^dUniversit de Paris CRESS, INSERM, INRA, 75004 Paris, France

^eMAP5 Laboratory, FP2M, CNRS FR 2036, Universit de Paris, 75006 Paris, France

^fUniversit Paris-Saclay, ENS Paris-Saclay, CNRS, Centre Borelli, 94235 Cachan, France

Abstract

Background: In early May 2020, following social distancing measures due to COVID-19, governments consider relaxing lock-down. We combined individual clinical risk predictions with epidemic modelling to examine simulations of risk-based differential isolation and exit policies.

Methods: We extended a standard susceptible-exposed-infected-removed (SEIR) model to account for personalised predictions of severity, defined by the risk of an individual needing intensive care if infected, and simulated differential isolation policies using COVID-19 data and estimates in France as of early May 2020. We also performed sensitivity analyses. The framework may be used with other epidemic models, with other risk predictions, and for other epidemic outbreaks.

Findings: Simulations indicated that, assuming everything else the same, an exit policy considering clinical risk predictions starting on May 11, as planned by the French government, could enable to immediately relax restrictions for an extra 10% (6 700 000 people) or more of the lowest-risk population, and con-

*Corresponding author

Email addresses: theodoros.evgeniou@insead.edu (Theodoros Evgeniou), mathilde.fekom@ens-paris-saclay.fr (Mathilde Fekom), anton.ovchinnikov@queensu.ca (Anton Ovchinnikov), raphael.porcher@aphp.fr (Raphael Porcher), camille.pouchol@parisdescartes.fr (Camille Pouchol), nicolas.vayatis@ens-paris-saclay.fr (Nicolas Vayatis)

sequently relax the restrictions on the remaining population significantly faster – while abiding to the current ICU capacity. Similar exit policies without risk predictions would exceed the ICU capacity by a multiple. Sensitivity analyses showed that when the assumed percentage of severe patients among the population decreased, or the prediction model discrimination improved, or the ICU capacity increased, policies based on risk models had a greater impact on the results of epidemic simulations. At the same time, sensitivity analyses also showed that differential isolation policies require the higher risk individuals to comply with recommended restrictions. In general, our simulations demonstrated that risk prediction models could improve policy effectiveness, keeping everything else constant.

Interpretation: Clinical risk prediction models can inform new personalised isolation and exit policies, which may lead to both safer and faster outcomes than what can be achieved without such prediction models.

Funding: No funding was used for this research.

Research in context

Evidence before this study

Several countries have implemented non-pharmaceutical interventions based on social distancing and isolation measures in order to limit the spread of COVID-
5 19. There has been limited differentiation in the degree of isolation measures, except for those critical for the functioning of the healthcare system and other key services. There is limited evidence about the impact of relaxing these measures as this has happened only recently and in only a few countries. Investigating the potential impact of differential restrictions depending on medical factors, such
10 as the risk of severe symptoms if infected by Sars-Cov-2, may inform policies for imposing or relaxing isolation policies when these are considered.

Added value of this study

This study investigates incorporating clinical risk predictions in epidemic models, allowing to explore isolation policies that consider individual clinical risks
15 using simulations.

Implications of all the available evidence

Epidemic simulations of isolation policies that consider predicted clinical risks in order to differentiate restrictions indicate the feasibility of policies that may be otherwise impossible to implement without undertaking this type of risk-based
20 approach.

1. Introduction

As of May 2020, many countries have adopted non-pharmaceutical interventions, such as isolation restrictions,¹ to control the spread of COVID-19. Epidemic models have been used to inform such policies.^{2,3} Governments now consider
25 relaxing these restrictions. Immunity tests⁴ and technology⁵ may support such policy decisions. We considered an application of predictive technologies, such as machine learning, that can be used to better understand outbreaks using epidemic simulations: using personalised predictions of severity risk, defined as requiring ICU if infected, in epidemic simulations to examine differential risk
30 based isolation policies. Other clinical risks can be considered similarly.

We extended standard epidemic models, namely a version of SEIR,⁶ to incorporate personalised risk predictions. Using simulations, we investigated how prediction models for patient severity may inform policy in two scenarios. First, when there is an ongoing outbreak as it was the case in France on the 17th
35 of March, 2020, when lock-down started. Second, when the outbreak has been curbed by lock-down and progressive loosening of isolation policies ("exit") may take place, as planned in France starting from the 11th of May, 2020.

There has already been research indicating differential impact of COVID-19 across patients, for example depending on hypertension, diabetes or other factors,⁷ but there is currently no standard risk prediction model considering all
40 factors, although some early versions are available.^{8,9} Therefore we assumed hypothetical risk prediction models and studied the sensitivity of the simulation results with respect to model discrimination.

To populate the models, we used available COVID-19 estimates and data from
45 France as of early May 2020.¹⁰ At the time, there were about 4 000 beds occupied by people with COVID-19, with a peak at 7 148 people in intensive care, compared to the total roughly 10 000 capacity recently reached by the French health system. We used current estimates with a reproduction number value of $\mathcal{R}_0 = 3 \cdot 3$ prior to lock-down, and an estimate of 2 · 5 million people
50 who had been immune or infected when it started on March 17.¹¹

Numerous containment/lock-down strategies have been proposed to flatten the predicted curve of the number of severe cases by means of non-pharmaceutical interventions such as strict quarantine,¹² social distancing, strict hygiene, population screening, etc.^{3,13,14,15} Although social distancing might be applied to
55 each individual in a similar way, a different policy consists of adjusting it to the seriousness of the symptoms one has if infected.^{16,17,18} However, severity of symptoms upon infection is not known a priori and, at best, can only be predicted. Such predictions may be possible using predictive models, for example using rules, data science, and machine learning methods and principles, based
60 on data related to infection symptoms.^{8,9}

2. Methods

2.1. Risk-Extended SEAIR Model

The model extends a standard compartmentalised SEAIR model with an added compartment for people in ICU by introducing four subcategories for each com-
65 partment. The risk model classifies each individual into the “high” versus “low” risk group, where risk refers to an individual experiencing severe symptoms requiring ICU.

Any prediction model makes mistakes. As a result, each individual would fall into exactly one of four groups which form a so-called “confusion matrix”: (a)
70 he/she would actually need ICU bed upon infection and was classified as high-risk (true positive), (b) would need ICU bed upon infection but classified low-risk (false negative), (c) would not need ICU bed upon infection but classified high-risk (false positive), (d) would not need ICU bed upon infection and classified low-risk (true negative). As a whole, this led to 22 possible compartments
75 in the extended SEAIR model.

For the control parameter of a policy applying differentiated isolation restriction w.r.t. to the predicted risk of developing severe symptoms we considered the proportion ρ of individuals subject to low isolation restrictions, i.e. those classified by the risk-model as the low-risk group. We denote with p the fraction

of individuals who are actually in the low-risk group (that is, those who would actually have mild symptoms and not require ICU if infected.) Depending on the characteristics of the risk prediction model and on p and ρ , the false positive and false negative rates, denoted by q^{FP} and q^{FN} , vary per the following implicit function equation:

$$(1 - q^{FP})p + q^{FN}(1 - p) = \rho. \quad (1)$$

We assumed that the differentiated isolation restrictions change people's behaviour per two behavioural parameters δ_r (for the group with low isolation restrictions, i.e. released) and δ_c (for the group with high isolation restrictions, i.e. confined); $0 \leq \delta_r < \delta_c \leq 1$. Variables δ_r and δ_c capture a level of "protection", which may aggregate several factors such as respiratory and hand hygiene, how much a person has lowered the number of exits from home and social interactions, etc. How individuals of a given category reduce their contacts not only depends on δ_c and δ_r but also on the proportion ρ of people with less strict restrictions. The decrease in contact rates, c_r and c_c for the two groups, may be obtained with the following formula:

$$1 - c_a = (1 - \delta_a) ((1 - \delta_r)\rho + (1 - \delta_c)(1 - \rho)), \quad a \in \{r, c\}.$$

These parameters have been used in the literature modelling the current lock-down.^{19,10} We also denote by c and δ the containment and protection parameters after a complete lock-down as it happened in France on March 17 ($c = 1 - (1 - \delta)^2$).

⁸⁰ The full description of the ODE model is given in the Appendix, a simplified version of which is presented in Figure 1.

2.2. Estimation of key parameters

When fitting available data, the key parameters driving the simulation results are: the fraction p of individuals with mild symptoms if infected, the
⁸⁵ reduction of contact rates during lock-down c , and the numbers of people exposed/asymptomatic/infected at the beginning of lock-down, on March 17.




Figure 1: Simplified compartmental model used. Corresponding notations are given in the Appendix.

These parameters were estimated by comparing model predictions to the actual data of ICU occupation from March 17 to May 2 in France.

We used the Approximate Bayesian Computation method (ABC)²⁰ to derive
 90 estimates and confidence intervals for the three parameters of interest. The ABC method was implemented with the root mean standard error as a distance function,²¹ with a maximum error set at 1 000 beds on average over the 47 data points, which corresponds to an acceptance rate of about 13%. We assumed the prior distributions to be independent with the following choices:
 95 Beta distribution for p with parameters $2\,390$, $11 \cdot 29$, fitting the mean and 95% confidence interval observed,¹¹ uniform distribution for the total number of people exposed/asymptomatic/infected on March 17 with range between 1 and 2 million, and uniform distribution for c between 65% and 75%.

2.3. Using personalised Risk Predictions

100 A policy based on simulations with a risk-extended epidemic model requires identifying individuals at highest risk of severity and correspondingly advising them to remain in strict isolation, while relaxing isolation restrictions for individuals at lower risk. Such identification is done by risk-“scoring” models using common data science and machine learning techniques: logistic regression, random forest, and the likes.
 105 A standard metric to assess the discriminating power of such models is the Area Under the Curve (AUC) of the Receiver Operating

Characteristic (ROC) curve.²²

Combining the model’s ROC curve, with the more strictly isolated fraction, $1 - \rho$, and the prevalence of severe symptoms in the population, p , results in the false positive and false negative error rates q^{FP}, q^{FN} . When ρ is small, the number of negative predictions is small (most people are in strict isolation), but the fraction of false negatives is also small, meaning that the vast majority of people in low isolation conditions will not experience severe symptoms. Increasing ρ not only releases more people from isolation, but also increases the false negative rate. Technical details are provided in the Appendix.

A key question is to select a tolerable fraction of people (ρ) being submitted to lower isolation as a function of the performance of a risk prediction model so that to subject only few people to stricter isolation, yet not to violate the ICU bed capacity due to the errors made by the model in identifying such people.

2.4. Summary of Key Parameters and Data

All parameters used in simulations are listed in Table 1.

For the purpose of illustration, the class-conditional distributions w.r.t. to high/low risk were modelled with Beta distributions. Risk predictions, ROC curves and AUC were derived accordingly.

Parameters for the initial conditions S_0, E_0, A_0, I_0, U_0 and R_0 were taken depending on the investigated scenario: in the first one the “day 0” is set on March 17 2020 – the first day of country-wide lock-down in France, while in the second one “day 0” is set on May 11 2020 – the announced day for the beginning of the post lock-down exit.

3. Results

Figure 2 displays the number of individuals requiring an ICU bed w.r.t. time t . The March 17 scenario is in the left column, the May 11 scenario is in the right. Two risk models are considered: one with AUC of $95 \cdot 99\%$ (top row)

Table 1: Simulation parameters used with relevant 95% confidence intervals.

Symbol	Description	Value(s)	Reference
N_0	total initial number of people in the population	$6 \cdot 7 \cdot 10^7$	
S_0	total initial number of infected people in the population	computed	
E_0	total initial number of exposed people in the population	case-dependent	estimated
A_0	total initial number of asymptomatic people in the population	case-dependent	estimated
I_0	total initial number of infected people in the population	case-dependent	estimated
U_0	total initial number of people in ICU	case-dependent	known/estimated
R_0	total initial number of immune people in the population	case-dependent	¹¹ /estimated
I_{max}	hospital capacity for COVID-19 ICU beds	7 250	assumed
p	proportion with mild symptoms (prior)	$0 \cdot 9953 [0 \cdot 9918 - 0 \cdot 9975]$	¹¹
β	transmission rate	computed	¹⁹
\mathcal{R}_0	basic reproduction number	$3 \cdot 3$	¹⁰
ε	waiting rate to viral shedding	$1/3 \cdot 7 \text{ day}^{-1}$	¹⁰
σ	waiting rate to symptom onset	$1/1.5 \text{ day}^{-1}$	¹⁰
η	waiting rate from symptom onset to ICU	$1/7 \text{ day}^{-1}$	¹¹
γ_m	recovery rate from mild symptoms	$1/2 \cdot 3 \text{ day}^{-1}$	¹⁰
γ_s	recovery rate for people in ICU	$1/17 \text{ day}^{-1}$	¹¹
α	mortality rate for people in ICU	$1/11 \cdot 7 \text{ day}^{-1}$	¹¹

and another, with AUC of $75 \cdot 71\%$ (bottom row), bracketing the performance
of initial risk models developed for COVID-19.^{8,9}

In each plot, ρ represents the maximal percentage of the population that is
submitted to lighter restrictions ($\delta_r = 0 \cdot 1$) in such a way that the 95% confidence
interval of the number of individuals requiring an ICU bed when using the
risk prediction model (green and orange curves) remains below the maximum
number of ICU beds assumed (7 250). In these first simulations, the rest of
the population is confined with stricter restrictions, $\delta_c = 0 \cdot 9$. Finally, the red
curves show the number of individuals requiring an ICU bed w.r.t. time if the
same ρ of population is in lower isolation, but selected at random without any
risk prediction model.

Figure 2 shows that a high-AUC model (green curve) allows for having 60% in
low isolation ($\delta_r = 0 \cdot 1$, corresponding to a decrease of social interaction by
47%) from March 17 on, while a low-AUC model (orange curve) enables only
51%. In France, with a population of 67 million, these percentage differences
correspond to 6 700 000 people. Plots for lock-down exit strategies (May 11)

150 investigate the effect of the same risk prediction accuracies. All differences (60% vs 51% for March 17, 70% vs 59% for May 11) are statistically significant at the 5% level. Lastly, without a risk prediction model, the ICU beds demand greatly exceeds the current capacity at either ρ .




Figure 2: Number of individuals requiring an ICU bed w.r.t. time t (days). Left column starts on March 17 (the day of the initial lock-down in France), right column starts on May 11 (the day when lock-down ends). The dotted line on the left column shows the actual data for France from March 17 to May 2. Top row uses a risk prediction model with AUC $95 \cdot 99\%$, bottom row uses a risk prediction model with AUC $75 \cdot 71\%$.


Figure 3 presents sensitivity analyses of the difference between the maximal per-

centage of people which may be in low isolation without exceeding ICU capacity for several risk-prediction models, relative to the same maximal percentage, but with no risk-prediction model. The results are shown for both the March 17 and the May 11 scenarios. Sensitivity is tested with respect to the discrimination performance of the risk prediction models used for the risk-extended SEAIR simulations and the degree of isolation of the more strictly isolated population (δ_c). We also alter the degree of isolation for the less strictly isolated population (δ_r) across different plots.


As expected, the higher the discrimination of the prediction model, the bigger the difference. The degree of isolation restrictions has different effects depending on who is considered: for the more strictly isolated population (the higher risk one) the stricter the isolation (parameter δ_c) the larger the impact of the risk prediction model. For the less strictly isolated population (the lower risk one – $\delta_r = 0 \cdot 1$ or $0 \cdot 2$ in Figure 3), the results are more intricate. It is also better to isolate more strictly, except when the risk prediction model is of very high quality and people of high risk are in stricter isolation. It is therefore important to both assume in models and encourage in practice (for example by focusing distribution of masks and other resources, strictly isolating nursing homes, etc.) realistic isolation practices for the high risk population.

Figure 4 shows sensitivity analyses regarding the percentage p of the population with severe symptoms upon infection, if we were to only decrease (from $0 \cdot 995$ to $0 \cdot 98$) the value for p while following the same analysis as before. The results are shown with a different ICU constraint (15 000 instead of 7 500) than in Figure 3 to also study how the ICU constraint affects the conclusions.


First, as expected, the lower p the less the impact of a risk prediction model keeping AUC constant: given the limited ICU – and possibly other – resources, larger p allows for a smaller range of percentages of the population being under less strict isolation restrictions, making all differences between policies smaller in absolute terms. Second, when we compare Figures 3 (a) and (c) with Figures 4 (a) and (c), where the only difference is in the total ICU assumed (7 500 vs 15 000), we see that the more ICUs available the larger the impact of using a risk-




(a) March 17, $\delta_r = 0 \cdot 1$



(b) March 17, $\delta_r = 0 \cdot 2$




(c) May 11, $\delta_r = 0 \cdot 1$




(d) May 11, $\delta_r = 0 \cdot 2$

Figure 3: Difference in maximum possible percentage of people in low isolation without hospital saturation. Maximal number of people in the low isolation group without exceeding the limit of 7 250 beds, with a margin of 2 000 beds as imposed by typical 95% confidence intervals, compared to the case of not using a risk prediction model. Plotted as a function of the AUC of a risk prediction model and the protection level δ_c for people recommended to be in isolation with stricter restrictions. $p \sim 0 \cdot 995$ for all figures.


prediction model, keeping everything else constant. More available resources allow for a larger range of percentage of people in less strict isolation making the differences between policies – risk based vs not – larger in absolute terms. Note that in all cases a risk prediction model approach allows for fewer people in strict isolation: this is consistent with value of information related arguments, as any test provides information which can be beneficial assuming everything




(a) March 17, $p \sim 0.995$



(b) March 17, $p = 0.98$



(c) May 11, $p \sim 0.995$



(d) May 11, $p = 0.98$

Figure 4: Difference in maximum possible percentage of people in low isolation without hospital saturation. Maximal number of people in the low isolation group without exceeding the limit of 15 000 beds, with a margin of 2 000 beds as imposed by typical 95% confidence intervals, compared to the case of not using a risk prediction model. Plotted as a function of the AUC of a risk prediction model and the protection level δ_c for people recommended to be in isolation with stricter restrictions. $\delta_r = 0.1$ for all figures.

else (including behavioural aspects) kept constant.

Finally, we explored the implications of our simulation for gradual exit strategies. To consider practical and realistic scenarios we solved the resultant optimisation (dynamic program) allowing releases every 30 days at multiples of 5% of the population, while ensuring that the maximum number of ICU beds needed was at most 5 250, being at least 2 000 beds (the typical maximum confidence

Four waves				Three waves			
	High AUC	Low AUC	No		High AUC	Low AUC	No
	model	model	model		model	model	model
$\delta_c = 0.9$	5	7	>12	$\delta_c = 0.9$	5	9	>12
$\delta_c = 0.8$	6	9	>12	$\delta_c = 0.8$	7	>12	>12
$\delta_c = 0.7$	8	9	>12	$\delta_c = 0.7$	>12	>12	>12

Four waves				Three waves			
	High AUC	Low AUC	No		High AUC	Low AUC	No
	model	model	model		model	model	model
$\delta_c = 0.9$	4	7	11	$\delta_c = 0.9$	4	7	>12
$\delta_c = 0.8$	5	8	11	$\delta_c = 0.8$	6	11	>12
$\delta_c = 0.7$	6	8	11	$\delta_c = 0.7$	6	9	>12

Table 2: Minimal time (in months) required for all people to exit high isolation restrictions, starting from March 17 (upper tables) or May 11 (lower tables), depending on δ_c , model quality and the number of waves of gradual relaxations. $\delta_r = 0.1$

range in Figure 2) below the maximum ICU capacity assumed.

Table 2 shows the total number of months to release the entire population for different scenarios, keeping all other parameters constant: no model, low-AUC
200 model, high-AUC model, and for 3 different values of δ_c . We considered only gradual releases in 3 or 4 waves – the ICU system was overwhelmed when using only 2 waves for most simulations. The main insight is that using no risk model would require more than a year in all scenarios, while exit based on risk-based
205 models would lead to relaxing restrictions for the entire population as fast as over 5 (March 17) or 4 (May 11) months.

Figure 5 shows example gradual policies for May 11 corresponding to the four waves part of Table 2, and assuming $\delta_c = 0.9$ as in Figure 2. The insights complement those for single release policies: with risk-prediction models, a smaller
210 percentage of the population may need to be subject to stricter isolation policies. Consequently, one could also reach the moment when isolation measures could be lifted sooner.

For example, using the high-AUC model and without exceeding the ICU capacity at any point, 70% of the lowest-risk population could be released on May 11 (“day 0”), followed by another 5% on July 10 (“day 60”), and yet another 15% on August 9 (“day 90”), and finally releasing the remaining 10% on September 8 (“day 120”). The resultant ICU demand is shown as a green line on Figure 5a. Implementing the same exit schedule without a model would lead to ICU demand of over 25 000 beds (red line). In contrast, a capacity-abiding exit strategy without a model (blue line) would put only 50% of population in low isolation on “day 0” (May 11), another 10% on “day 90” (August 9), additional 10% on “day 180” (November 7), and the last 30% only on “day 330” (April 6, 2021), – 7 months later than the similar risk-model-based strategy. Such an extended isolation would also apply to many more people: 10% with the model versus 30% without; for France this means the additional ~13.5 million people in isolation for the additional 7 months.

For both scenarios, Figure 5b shows the percentage of the population that becomes immune over time. Because model-based policies release larger portions of the low-risk population and do so faster, they also achieve the levels of “herd immunity” ($1 - \frac{1}{\mathcal{R}_0} \approx 0.7$), allowing for the ultimate protection against the disease. In contrast, herd immunity is not achieved by a policy without the risk-model: the disease is suppressed, but could explode again. In other words, assuming everything else constant, the simulations indicated that, using risk prediction models, isolation restriction may be relaxed faster and safer.

4. Discussion

Simulations indicated that considering differential relaxation of isolation restrictions depending on predicted severity risk can decrease the immediate percentage of the population in France under strict restrictions by 10% or more relative to not using such risk predictions, and fasten the complete exit by several months. This result was robust to changes in risk prediction accuracy, percentage of severe-if-infected cases in the population, availability of resources




Figure 5: Examples of gradual schedules of relaxing isolation restrictions with and without model-based risk predictions. High AUC model (green) and no model (red) $\rho = [0 \cdot 7, 0 \cdot 75, 0 \cdot 9, 1]$ and $\mathbf{t} = [60, 90, 120, 600]$, no model (blue): $\rho = [0.5, 0 \cdot 6, 0 \cdot 7, 1]$ and $\mathbf{t} = [90, 180, 330, 600]$. Vectors $\mathbf{t} = [t_1, t_2]$ give the population release schedules $\rho = [\rho_1, \rho_2]$ as follows: $\rho_1 \cdot 100\%$ of the population is released on day 0, then $(\rho_2 - \rho_1) \cdot 100\%$ are released on day t_1 , etc.

(such as ICUs), and social distancing. Benefits increased when risk prediction accuracy increased, percentage of severe-if-infected cases in the population decreased, availability of resources (such as ICUs) increased, and social distancing
245 increased. All results were developed using hypothetical risk prediction models for COVID-19, with discrimination ranges in line with early indications from initial models developed as of early May 2020.^{8,9}

The proposed approach can also be adopted for other epidemic models, and personalisation can further be explored using this approach for policies other
250 than isolation restrictions. Moreover, predicted risk based isolation restrictions can be combined with other policies such as test-based ones,^{4,5} possibly also using other relevant prediction models, to limit the impact of outbreaks such as COVID-19. Finally, the same analysis can be done focusing on mortality or other outcomes instead of ICU demand.

255 Several caveats should be noted. First, epidemic models – and the conclusions they may support – rely on a number of parameters, for example virus incuba-

tion and recovery times and the basic reproduction number \mathcal{R}_0 , while the effects
 of policies also depend on healthcare system factors such as the availability of
 relevant resources (e.g., trained personnel). Second, these parameters are uncer-
 260 tain and evolve dynamically;²³ the resultant policies are therefore contingent.
 Observing an ICU demand that is closer to an upper boundary of the confidence
 interval may require the next wave to be delayed or involve a smaller release
 percentage than our current simulations built from day zero suggest. Third, pol-
 icy decisions require careful context-specific robustness analysis; however, using
 265 risk prediction models can at worst make no significant difference while at best
 improve policies by a significant margin, fixing all other conditions. Fourth, risk
 prediction models cannot be used when, or for people, for whom the necessary
 data is unavailable. In this case, simple models (e.g., based only on age and
 some reliable chronic disease data) may need to be used, which may limit the
 270 benefits of the approach. Finally, risk-predictions based policies using epidemic
 simulations should be developed taking into account behavioural aspects that
 may prove any model predictions and policy actions wrong; ethical issues, fear,
 widespread non-compliance to isolation measures, and the likes.
 In conclusion, combining prediction models using data science and machine
 275 learning principles may improve outbreak management policies and should be
 considered when developing isolation and exit policies.

Appendix A. Extended SEAIR model

Appendix A.1. Complete version

We first break down the 5 compartments S , E , A , I and R in 2 subcategories, depending on whether people are going to have mild ("m") or severe ("s") symptoms requiring ICU, upon infection. An added compartment for people in the "s" category is that of people in ICU, denoted U .

These people are also labelled depending on whether they are considered low-risk and submitted to low isolation ("r"), or are considered high-risk and recommended to be in high isolation ("c").

The notations are as follows: for a given category $Q \in \{S, E, A, I, U, R\}$, we use superscripts for risk prediction, and subscripts for actual status. In other words:

- $Q_m^{(r)}$: would not require ICU if infected, submitted to low isolation restrictions, *i.e.* true negative,
- $Q_m^{(c)}$: would not require ICU if infected, submitted to high isolation recommendations, *i.e.* false positive,
- $Q_s^{(r)}$: would require ICU if infected, submitted to low isolation restrictions, *i.e.* false negative,
- $Q_s^{(c)}$: would require ICU if infected, submitted to high isolation recommendations, *i.e.* true positive.

We also denote in category $Q \in \{S, E, A, I, U, R\}$

- $Q^{(r)} = Q_m^{(r)} + Q_s^{(r)}$: submitted to low isolation restrictions,
- $Q^{(c)} = Q_s^{(c)} + Q_m^{(c)}$: submitted to high isolation recommendations,
- $Q_m = Q_m^{(r)} + Q_m^{(c)}$: having mild symptoms if infected,
- $Q_s = Q_s^{(c)} + Q_s^{(r)}$: requiring ICU if infected,
- $Q = Q_m + Q_s = Q^{(r)} + Q^{(c)}$: total number of people in category Q .

We define the effective number of contagious people as

$$I^e = (1 - \delta_r)(A^{(r)} + I^{(r)}) + (1 - \delta_c)(A^{(c)} + I^{(c)}).$$

The equations read

$$\begin{aligned}\dot{S}_m^{(r)} &= -(1 - \delta_r)\beta I^e S_m^{(r)} \\ \dot{E}_m^{(r)} &= (1 - \delta_r)\beta I^e S_m^{(r)} - \varepsilon E_m^{(r)} \\ \dot{A}_m^{(r)} &= \varepsilon E_m^{(r)} - \sigma A_m^{(r)} \\ \dot{I}_m^{(r)} &= \sigma A_m^{(r)} - \gamma_m I_m^{(r)} \\ \dot{R}_m^{(r)} &= \gamma_m I_m^{(r)} \\ \dot{S}_m^{(c)} &= -(1 - \delta_c)\beta I^e S_m^{(c)} \\ \dot{E}_m^{(c)} &= (1 - \delta_c)\beta I^e S_m^{(c)} - \varepsilon E_m^{(c)} \\ \dot{A}_m^{(c)} &= \varepsilon E_m^{(c)} - \sigma A_m^{(c)} \\ \dot{I}_m^{(c)} &= \sigma A_m^{(c)} - \gamma_m I_m^{(c)} \\ \dot{R}_m^{(c)} &= \gamma_m I_m^{(c)} \\ \dot{S}_s^{(c)} &= -(1 - \delta_c)\beta I^e S_s^{(c)} \\ \dot{E}_s^{(c)} &= (1 - \delta_c)\beta I^e S_s^{(c)} - \varepsilon E_s^{(c)} \\ \dot{A}_s^{(c)} &= \varepsilon E_s^{(c)} - \sigma A_s^{(c)} \\ \dot{I}_s^{(c)} &= \sigma A_s^{(c)} - \eta I_s^{(c)} \\ \dot{U}^{(c)} &= \eta I_s^{(c)} - (\gamma_s + \alpha)U^{(c)} \\ \dot{R}_s^{(c)} &= \gamma_s I_s^{(c)} \\ \dot{S}_s^{(r)} &= -(1 - \delta_r)\beta I^e S_s^{(r)} \\ \dot{E}_s^{(r)} &= (1 - \delta_r)\beta I^e S_s^{(r)} - \varepsilon E_s^{(r)} \\ \dot{A}_s^{(r)} &= \varepsilon E_s^{(r)} - \sigma A_s^{(r)} \\ \dot{I}_s^{(r)} &= \sigma A_s^{(r)} - \eta I_s^{(r)} \\ \dot{U}^{(r)} &= \eta I_s^{(r)} - (\gamma_s + \alpha)U^{(r)} \\ \dot{R}_s^{(r)} &= \gamma_s U_s^{(r)}.\end{aligned}$$

The number of deaths is obtained by

$$\dot{D} = \alpha (U^{(c)} + U^{(r)}).$$

Appendix A.2. Compact version

Although all compartments are needed, whether it is for the purpose of computations (see the definition of I^e) or for tracking the numbers in each category, the model may be written in a more compact form for convenience:

$$\begin{aligned}\dot{S}^{(r)} &= -(1 - \delta_r)\beta I^e S(r) \\ \dot{E}^{(r)} &= (1 - \delta_r)\beta I^e S(r) - \varepsilon E^{(r)} \\ \dot{S}^{(c)} &= -(1 - \delta_c)\beta I^e S(c) \\ \dot{E}^{(c)} &= (1 - \delta_c)\beta I^e S(c) - \varepsilon E^{(c)} \\ \dot{A} &= \varepsilon E - \sigma A \\ \dot{I}_m &= \sigma A_m - \gamma_m I_m \\ \dot{I}_s &= \sigma A_s - \eta I_s \\ \dot{U} &= \eta I_s - (\gamma_s + \alpha)U \\ \dot{R}_m &= \gamma_m I_m \\ \dot{R}_s &= \gamma_s I_s\end{aligned}$$

The number of deaths is obtained by

$$\dot{D} = \alpha U.$$

Computed parameters. The initial number of susceptibles S_0 is computed as follows:

$$S_0 = N_0 - (E_0 + A_0 + I_0 + R_0 + U_0),$$

and the transmission rate β is given by the following formula:

$$\beta = \frac{\mathcal{R}_0}{S_{\text{init}}} \frac{1}{p(\sigma^{-1} + \gamma_m^{-1}) + (1 - p)(\sigma^{-1} + \eta^{-1})},$$

which can be obtained by stability analysis.¹⁹ Here S_{init} stands for the initial
 305 number of susceptible people for the period during which \mathcal{R}_0 has been estimated.
 As all estimates for \mathcal{R}_0 have been obtained for the period prior to lock-down,
 we chose $S_{\text{init}} = N_0$ where N_0 is the total number of people in the population
 of interest.

Appendix A.3. Simulations

310 All simulations were run using Python. The number of time discretisation points
 per day was fixed at 500.

A given strategy of gradually relaxing restrictions is defined by N fractions of
 individuals put in the low-risk group

$$(\rho_0, \rho_1, \dots, \rho_{N-1}),$$

together with the times at which the policy changes

$$(T_0, T_1, \dots, T_{N-1}, T_N)$$

with $T_0 = 0$ and $T = T_N$ a final horizon of interest.

The corresponding solution

$$y = (S_m^{(r)}, E_m^{(r)}, A_m^{(r)}, I_m^{(r)}, R_m^{(r)}, S_m^{(c)}, E_m^{(c)}, A_m^{(c)}, I_m^{(c)}, R_m^{(c)}, \\ S_s^{(c)}, E_s^{(c)}, A_s^{(c)}, I_s^{(c)}, U^{(c)}, R_s^{(c)}, S_s^{(r)}, E_s^{(r)}, A_s^{(r)}, I_s^{(r)}, U^{(r)}, R_s^{(r)}),$$

was computed up until the final time T - further details below.

Each ρ_i uniquely determines how many people will be considered to be high or
 low risk, through

$$(1 - q_i^{FP})p + q_i^{FN}(1 - p) = \rho_i.$$

Assume that the solution at time T_i has been computed for a given i . Since the
 315 number of individuals in the low-risk group changes at time T_i , the ODE must
 be integrated for a corrected initial condition on the interval $[T_i, T_{i+1}]$.

This initial condition was obtained from reallocating people depending on their
 new labelling. Denoting y^{old} and y^{new} the value of $y(T_i)$ before and after rela-

bellling, the new value is given as a function of the previous one by

$$\begin{aligned}
y_j^{\text{new}} &= (1 - q_i^{FP})(y_j^{\text{old}} + y_j^{\text{old}}), & j = 1, \dots, 5 \\
y_{j+5}^{\text{new}} &= q_i^{FP}(y_j^{\text{old}} + y_j^{\text{old}}), & j = 1, \dots, 5 \\
y_j^{\text{new}} &= (1 - q_i^{FN})(y_j^{\text{old}} + y_j^{\text{old}}), & j = 11, \dots, 16 \\
y_{j+6}^{\text{new}} &= q_i^{FN}(y_j^{\text{old}} + y_j^{\text{old}}), & j = 11, \dots, 16.
\end{aligned}$$

Appendix A.4. Further details for Figure 2

The number of samples for the prior distribution for p and the initial number of exposed/asymptomatic/infected at the beginning of lock-down $E_0 + A_0 + I_0$ was set at $n = 10\,000$ (resp. $100\,000$ for robustness when computing means). This led to around $1\,300$ (resp. $13\,000$) posterior samples as the acceptance rate was at 13%. In order to reduce the parameter space, we have estimated the total number of people exposed/asymptomatic/infected and inferred the number in each state by using the fractions of the mean time spent in each category, as given in Table 1.

In all figures showing the evolution of the number of people in ICU, the initial condition $E_0 + A_0 + I_0$ and the proportion of people p not requiring ICU were sampled according to their posterior distribution. Subsequent 95% confidence intervals were derived by removing the $2 \cdot 5\%$ and $2 \cdot 5\%$ upper and lower values for the computed number of ICU beds, respectively.

Appendix A.5. Synthetic risk prediction distributions

For all figures, the risk prediction was obtained from synthetic data. Distributions for people having critical symptoms requiring ICU and people having milder symptoms were assumed to follow Beta-distributions. More precisely, denoting a_m, b_m and a_s, b_s the parameters for the respective distributions for people with mild symptoms and for people with severe symptoms, we fixed $b_s = a_m = 2$ and made the other two parameters vary from $a_s = b_m = 2 \cdot 5$, with a maximum at $a_s = b_m = 6 \cdot 5$ used for Figures 3 and 4.

In Figures 2, 5 and Table 3, *no model* refers to $a_m = b_m = a_s = b_s = 1$, *low*
 340 *AUC model* refers to $b_s = a_m = 2$, $a_s = b_m = 3$ and *high AUC model* refers to
 $b_s = a_m = 2$, $a_s = b_m = 5$.

Appendix A.6. Further details on Figures 3, 4, 5 and Table 3

Since Figures 3, 4, T and Table 3 require grid searches, we did not sample
 according to the posterior distribution for each scenario, but rather computed
 345 mean values in order to ease the computational burden.

March 17: For March 17, simulations are obtained as follows:

- the initial condition for ICU beds is known, and we used the estimate for
 the total number of infected people at that date. With our notations, the
 number $E_0 + A_0 + I_0 + R_0$ is about $2.6 \cdot 10^6$ people.¹¹ We only needed an
 350 estimate for $E_0 + A_0 + I_0$ and we did so by averaging over the posterior
 for this variable,
- p was taken as the mean along posterior samples,
- the ODE was then integrated on $T = 200$ days.

May 11: For May 11, simulations were obtained as follows:

- 355 • sampling according to the posterior for p and $E_0 + A_0 + I_0$, and integrating
 the ODE system from March 17 to May 11, we obtained a sample of initial
 conditions for May 11, of which we took the average,
- the ODE was then integrated on $T = 200$ days.

Figure 5 and Table 3: Table 3 computes the minimal possible time the whole
 360 population may be in low isolation, without breaking the ICU constraint of
 7 250 beds with a 2 000 bed margin taking into account typical sizes of 95%
 confidence intervals. The minimal time was computed by searching for strategies
 where times of (gradual) relaxations are chosen every month, and percentages
 of people in low isolation are multiples of 5%. The initial percentage of people
 365 in low isolation is always chosen to be the highest possible.

Appendix B. Risk Prediction Models

The purpose of this section is to provide details about the practical use of ROC curves to calibrate the discrimination threshold in the personalised risk predictions. The key objective is to target the individuals at highest risk and
370 apply differential policies to protect them from the risk of infection. Assuming that there is a constraint on the resources (for instance the number of ICU beds) and that even without constraints intensive care may not prevent death of all patients admitted (mortality rates for patients with mechanical ventilation may be higher than 50%), the questions underlying the implementation of large-scale
375 risk predictions are the following: (i) how to rank the population with respect to their risk (e.g., of being eventually transferred to ICU), (ii) how to select the threshold value on this risk in order to further prevent individuals deemed at high risk being in contact with other individuals who may infect them, (iii) how sensitive will these estimates (ranking, threshold) be w.r.t. to sampling bias and
380 to prediction accuracy. Here we only discuss the practical aspects related to (i) and (ii), and briefly comment on the sources of uncertainty (iii).

Building a personalised risk predictions. Personalised risk predictions are build upon past data which correlate individually a vector X of known factors with the outcome Y materialising the risk (e.g. the patient needing to be transferred
385 to the intensive care unit, ICU) which is represented as a binary event ($Y = +1$ is a positive instance vs. $Y = -1$ is a negative instance). Such data are usually obtained after clinical studies. The fraction of high risk individuals in the population is denoted by $p = \mathbf{P}(Y = -1)$. Typical examples of methods to estimate personalised risk predictions are, for example, logistic regression
390 in parametric statistics, or random forests, regularisation methods, or bipartite ranking algorithms in machine learning. In order to compare different estimation strategies, it is standard to assess performance using the ROC curve,²² defined as the parametric curve which maps, for a given risk prediction level s , each threshold value t to a point in the unit square with coordinates $(\mathbf{P}(s(X) > t \mid$
395 $Y = -1), \mathbf{P}(s(X) > t \mid Y = +1))$ (plot of the true positive rate against the false

positive rate).

Discriminating high risk individuals. Personalised risk predictions can be used to discriminate high versus low risk individuals based on the knowledge of individual values of relevant factors such as comorbidities. The decision is
400 binary and it is taken after having calibrated the threshold t with respect to control parameters of a policy which are essentially driven by the constraints on the resources. Such constraints are for example: (a) the ICU capacity, which induces a constraint on the precision $\mathbf{P}(Y = 1 \mid s(X) > t)$ and such quantity is directly related to the true positive rate thanks to the Precision-Recall curve, or
405 (b) the economic and psychological consequences of isolation measures on the population, which gives a constraint on $\mathbf{P}(s(X) > t)$.²⁴

Sources of uncertainty. The previous estimates (risk predictions, decision threshold) are subject to uncertainties which can be due to sampling bias (e.g. clinical trial data used are not reflecting population data or electronic medical records
410 failing to account for the part of the population who was never admitted to the hospital) or methodological bias (model misspecification, suboptimal machine learning/statistical method used). In order to provide statistical guarantees on the estimators obtained, it is necessary to compute confidence bands on the estimated ROC curve which will then lead to explicit confidence bands on the
415 decision parameters of the risk prediction strategy. Typical approaches to derive confidence bands are to perform error propagation on distribution parameters (in a parametric framework), or to generate several ROC curves and Precision-Recall curves through resampling and provide some bootstrap estimate of the confidence band.²⁵ Resampling strategies may include label flipping (prediction
420 uncertainty), sample perturbation or shifting (sampling bias). Lastly, the quality of a prediction model was assessed using AUC; other measures may be used for ranking based on risk.²⁶

Acknowledgements

The authors are grateful to Ramss Djidjou-Demassee for detailed exchanges
425 about the model of his team,¹⁹ and to Amaury Lambert and Pierre-Yves Mass
for the interesting discussions. We also thank Olivier Boulant for his help in
making the code available on GitHub.

Funding:

430 No funding was used for this research.

Competing interests

No competing interests.

435 Data availability

Not Applicable.

Code availability

Code for this research is available at <https://reine.cmla.ens-cachan.fr/boulant/seair>
440

Correspondence

Correspondence should be addressed to TE (theodoros.evgeniou@insead.edu)
and AO (anton.ovchinnikov@insead.edu) for general questions and to CP
(camille.pouchol@parisdescartes.fr) for questions regarding code and methods.

445

Contributions

All authors contributed equally for the research and analysis. Camille Pouchol
and Mathilde Fekom developed the code.

¹ Koo JR, Cook AR, Park M, Sun Y, Sun H, Lim JT, et al.. Interventions to
450 mitigate early spread of SARS-CoV-2 in Singapore: a modelling study; 2020.
Available from: <http://www.sciencedirect.com/science/article/pii/S1473309920301626>.

- ² Flaxman S, Mishra S, Gandy A, Unwin H, Coupland H, Mellan T, et al..
Report 13: Estimating the number of infections and the impact of non-
pharmaceutical interventions on COVID-19 in 11 European countries; 2020.
- ³ Ferguson N, Laydon D, Nedjati Gilani G, Imai N, Ainslie K, Baguelin M,
et al.. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce
COVID19 mortality and healthcare demand; 2020.
- ⁴ Petherick A. Developing antibody tests for SARS-CoV-2; 2020. Available
from: [https://www.thelancet.com/journals/lancet/article/PIIS0140-6736\(20\)30788-1/fulltext](https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(20)30788-1/fulltext).
- ⁵ Wang CJ, Ng CY, Brook RH. Response to COVID-19 in Taiwan: Big Data
Analytics, New Technology, and Proactive Testing; 2020. Available from:
<https://doi.org/10.1001/jama.2020.3151>.
- ⁶ Kucharski AJ, Russell TW, Diamond C, Liu Y, Edmunds J, Funk S, et al..
Early dynamics of transmission and control of COVID-19: a mathematical
modelling study; 2020. Available from: <http://www.sciencedirect.com/science/article/pii/S1473309920301444>.
- ⁷ Guan W, Ni Z, Hu Y, Liang W, Ou C, He J, et al.. Clinical Characteristics
of Coronavirus Disease 2019 in China; 2020. Available from: <https://www.nejm.org/doi/full/10.1056/NEJMoA2002032>.
- ⁸ Mortality Risk Calculatory; 2020. Available from: <https://www.covidanalytics.io/calculator>.
- ⁹ Urwin SG, Kandola G, Graziadio S. What prognostic clinical risk prediction
scores for COVID-19 are currently available for use in the community setting?;
2020. Available from: <https://www.cebm.net/covid-19/what-prognostic-clinical-risk-prediction-scores-for-covid-19-are-currently-available-for-use-in-the-community-setting/>.
- ¹⁰ Di Domenico L, Pullano G, Sabbatini CE, Boëlle PY, Colizza V. Expected
impact of lockdown in Île-de-France and possible exit strategies. Cold Spring

Harbor Laboratory Press; 2020. Available from: <https://www.medrxiv.org/content/early/2020/04/17/2020.04.13.20063933>.

- ¹¹ Salje H, Tran Kiem C, Lefrancq N, Courtejoie N, Bosetti P, Paireau J, et al.. Estimating the burden of SARS-CoV-2 in France; 2020. Working paper or preprint. Available from: <https://hal-pasteur.archives-ouvertes.fr/pasteur-02548181>.
485
- ¹² Volpert V, Banerjee M, Petrovskii S. On a quarantine model of coronavirus infection and data analysis. EDP Sciences; 2020.
- ¹³ Chen X, Qiu Z. Scenario analysis of non-pharmaceutical interventions on global COVID-19 transmissions; 2020.
490
- ¹⁴ Manchén C, Brugnago EL, da Silva RM, Mendes CFO, Beims MW. Strong correlations between power-law growth of COVID-19 in four continents and the inefficiency of soft quarantine strategies; 2020.
- ¹⁵ Leung K, Wu JT, Liu D, Leung GM. First-wave COVID-19 transmissibility and severity in China outside Hubei after control measures, and second-wave scenario planning: a modelling impact assessment; 2020. Available from: <http://www.sciencedirect.com/science/article/pii/S0140673620307467>.
495
- ¹⁶ Zhang AZ, Enns E. Optimal timing and effectiveness of COVID-19 outbreak responses in China: a modelling study; 2020.
500
- ¹⁷ Zhigljavsky A, Whitaker R, Fesenko I, Kremnizer K, Noonan J, Harper P, et al.. Generic probabilistic modelling and non-homogeneity issues for the UK epidemic of COVID-19; 2020.
- ¹⁸ Fineberg HV. Ten Weeks to Crush the Curve; 2020.
- ¹⁹ Djidjou-Demasse R, Michalakakis Y, Choisy M, Sofonea MT, Alizon S. Optimal COVID-19 epidemic control until vaccine deployment. Cold Spring Harbor
505

Laboratory Press; 2020. Available from: <https://www.medrxiv.org/content/early/2020/04/14/2020.04.02.20049189>.

²⁰ Marjoram P, Molitor J, Plagnol V, Tavaré S. Markov chain Monte Carlo
510 without likelihoods. National Acad Sciences; 2003.

²¹ Britton T. Stochastic epidemic models: a survey. Elsevier; 2010.

²² Fawcett T. An introduction to ROC analysis; 2006. ROC Analysis in Pattern
Recognition. Available from: <http://www.sciencedirect.com/science/article/pii/S016786550500303X>.

²³ Alban A, Chick S, Dongelmans D, Sluijs A, Wiersinga W, Vlaar A, et al.. ICU
515 Capacity Management During the COVID-19 Pandemic Using a Stochastic
Process Simulation; 2020.

²⁴ Cléménçon S, Vayatis N. Ranking the Best Instances. JMLR.org; 2007.

²⁵ Macskassy SA, Provost F, Rosset S. ROC Confidence Bands: An Empirical
520 Evaluation. In: Proceedings of the 22nd International Conference on Ma-
chine Learning. ICML 05. New York, NY, USA: Association for Computing
Machinery; 2005. p. 537544. Available from: <https://doi.org/10.1145/1102351.1102419>.

²⁶ Reid MD, Williamson RC. Information, Divergence and Risk for Binary Ex-
525 periments; 2011. Available from: <http://jmlr.org/papers/v12/reid11a.html>.